Planetary Rover Visual Motion Estimation improvement for Autonomous, Intelligent, and Robust Guidance, Navigation and Control

J. Nsasi Bakambu *, C. Langley *, G. Pushpanathan *, W. James MacLean *, R. Mukherji *, E. Dupuis **

* MDA Corporation, Brampton, Ontario, Canada
e-mail: Joseph.bakambu, chris.langley, giri.pushpanathan, raja.mukherji@mdacorporation.com; james.maclean@utoronto.ca

** Space Exploration, Canadian Space Agency, Saint-Hubert, Quebec, Canada
e-mail: Erick.dupuis@asc-csa.gc.ca

Abstract

This paper presents the Mojave Desert field test results of an improved planetary rover visual motion estimation technique for the Autonomous, Intelligent, and Robust Guidance, Navigation, and Control for Planetary Rovers (AIR-GNC). The main improvements include: optimal use of different features from stereo-pair images as visual landmarks, and the use of VME-based feedback to close the path tracking loop. As well, a long-range and wide FOV active 3D sensor was used to extract long-range fixed landmarks for enabling visual motion estimation observability, and thus improving the accuracy of the VME. The field test, conducted in relevant Mars-like terrains, under dramatically changing weather and lighting conditions, shows good localization accuracy on average. Moreover, the MDA developed Enhanced IMU-corrected odometry was reliable and had good accuracy in all test locations including in loose sand dunes. These results are based on data collected during 7.3 km of traverses, under both fully autonomous and tele-operated control.

1 Introduction

One of the continuing challenges for future unmanned exploration rover missions on the moon and Mars is the ability to accurately estimate the rover's position and orientation. In the absence of an equivalent to the Global Positioning System on earth, designers of future system must exploit a combination of vehicle odometry, inertial sensing and visual information to obtain this localization information. Future missions such as ESA's ExoMars and NASA's Mars Mobile Science Laboratory rover require the rover to autonomously traverse hundreds of meters to one km daily at speeds of up to 100 m/h [1]. Equally challenging is the need to localize the position of the rover to an accuracy of between one and four percent of distance traveled. Accurate localization is arguably the most fundamental competence required for long range autonomous navigation. In this context, MDA Space Missions and the Canadian Space Agency (CSA) have embarked on a project to further their capability for visual motion estimation (VME) of planetary rovers, which is described in this paper.

VME algorithms have recently seen a considerable amount of interest from the planetary exploration rover community as a solution for accurate localization. On the Mars Exploration Rovers (MERs), VME was not considered part of the main system for localization, but was shown to work relatively well. In [2], JPL reported that "Visual Odometry software has enabled precision drives over distances as long as 8 m on slopes greater than 20 degrees, and has made it possible to safely traverse the loose sandy plains of Meridiani". The algorithm works by tracking features (Harris corners) in a stereo image pair from one frame to the next (frame-to-frame). Thus, the problem is one of determining "the change in position and attitude for two pairs of stereo images by propagating uncertainty in a 3D to 3D pose estimation formulation using maximum likelihood estimation". The evaluation tests conducted at the JPL Marsyard and Johnson Valley, California showed that the absolute position errors were less than 2.5% over the 24 meter Marsyard course, and less than 1.5% over the 29 meter Johnson Valley course. The rotation error was less than 5.0 degrees in each case.

LAAS/CNRS (Laboratoire d'Architecture et d'Analyse des Systèmes/Centre National de la Recherche Scientifique) VME is based on the frame-to-frame pixel tracking method. Landmarks are extracted from images by finding points of interest identified by image intensity gradients. The test results using the Lama rover [3] showed an error of 4% on a 25 meter traverse. After improvement of the algorithm in [4], the overall error of 2% on a 70 m traverse was maintained.

A survey of the literature [5][6][7] shows that the

state-of-the-art in localization has yet to meet the ExoMars localization accuracy requirement of 1% of the traveled distance, without the extensive use of post-processing methods (for example, bundle adjustment).

Previous work at MDA has shown that a full simultaneous localization and mapping (SLAM) approach yields good accuracy, but has a high computational burden, forcing lower commanded speeds for the vehicle [8]. Conversely, if no map is maintained, a frame-to-frame technique has been shown to operate with acceptable speed but at the cost of decreased accuracy. MDA proposed a multi-frame VME algorithm that attempts to obtain a suitable balance between accuracy and speed. The results in hardware indicate that traverses of greater than 200 m are possible to an accuracy between one and four percent in planetary conditions. Our approach uses 3D odometry and a stereo-pair to identify visual landmarks by using the Scale Invariant Feature Transform (SIFT). While the result is commensurate with the ExoMars localization accuracy requirements, further development is required to improve the accuracy and achieve robustness under a variety of terrain conditions, and decreased computational cost for greater practicality for a flight mission.

The latest MDA VME approach is presented in this paper. This extended localization system was tested in a field campaign in the Mojave Desert using the robot shown in Figure 2. The new features being tested which yield improvements from the previous work include:

- Optimal use of different features from stereo-pair images as visual landmarks. SIFT features are invariant to image translation, scaling, rotation, and partially invariant to illumination changes and affine projection, making them suitable landmarks for use in motion estimation. See [9] for details. However, in testing to date of the VME SLAM system it has been found that in some cases SIFT features alone are insufficient for tracking. For example, on sandy terrain with smooth-shaped rocks, the SIFT detector often produces only small scale features, and these are often unstable. Other features extracted from the stereo-pair images are Maximally-Stable Extremal Regions (MSERs) [10] and Harris-Laplace [11] features.
- Investigation of the use of short-range active 3D sensors (lighting condition independent) as the visual front end to VME. A flash lidar-based, rigid transform invariant 3D feature extraction technique has been developed.
- Extension of the localization system to use inclinometry to improve roll/pitch angle estimation. Inclinometry maintains an observable estimate of the gravity vector, instead of relying on roll and pitch rate integration.
- Extension of the localization system to use a long-range and wide FOV active 3D sensor in order to extract long-range fixed landmarks for enforcing VME observability, and thus improving the accuracy of the approach.
- Development of an adaptive Extended Kalman Filter (EKF) based SLAM
- Use of the improved VME estimated pose as feedback for closed loop motion control. Prior work in this field has focused on VME as an open-loop observer, for providing accurate *a posteriori* localization when a destination is reached. In our novel approach, the low-rate and delayed VME measurements and the high-rate IMU-odometry measurements are fused to yield a high-rate and smooth position and attitude measurement signal. This technique is called *"frequency lifting"*. The frequency lifted output is then used as the feedback signal for closed loop path tracking

This new VME technology was integrated as part of a complete short-range GN&C solution for planetary rovers, including 3D terrain modeling and assessment, motion planning and tracking, hazard detection and avoidance algorithms, and high level supervisory control based on CORTEX, a tool developed by the CSA for design and real time execution of state machines [12].

The rest of this paper is organized as follow. Section 2 describes a typical field test scenario and provides the field test terrain landscape through pictures and maps reconstructed using sensor data. The challenges and the issues observed during the field test campaign are also presented. The experimental results are presented and discussed in section 3, and the lessons learned are also discussed. Section 4 concludes the paper.

2 Test Scenario and Field Test Terrain Landscapes

This section describes a typical field test scenario of AIR-GNC and provides the field test terrain landscapes through pictures and maps reconstructed using sensor data. The goal of AIR-GNC was to advance the capabilities of autonomous terrain modeling and assessment, and localization for future planetary missions. In order to demonstrate these capabilities, it was necessary to test them in a relevant planetary environment, which includes rocks, crevices and cracks, loose sand, and slope hazards.

The dry lake beds of the Mojave Desert in the Southern United States have been used as Martian

analogues in the past by NASA's Jet Propulsion Laboratory [13]. The terrain offers a variety of soil conditions, from dry fine grain sand, to hard packed clay, to loose gravel. Slopes and rock fields can also be found near the edges of the lake beds. The area is also relatively free of vegetation.

The testbed with the newly developed software modules can be fully exercised due to the presence of slope and rock hazards, the likelihood of wheel slippage, and the presence of representative natural image features.

2.1 The Challenges and Issues Observed

This subsection intends to address some of the general issues observed which affected the field campaign. Comments on the performance of the localization system itself will be addressed in the results section. General issues during the field campaign included:

- Weather: January is the rainy season in the Mojave Desert. Unfortunately, in this field campaign, it rained more frequently and heavily than usual. As a result, a few full days of testing were completely lost, and about four days had to be cut short due to rain. Rainfall could be heavy at times, flooding the dry lake beds. With any more than a light sprinkle, the playa becomes very slick and loose, quickly causing a risk to the vehicles. Further, the ground needs time to dry out after a heavy rainfall, meaning that more than one day can be lost at a time.
- Glare: The glare from the sunlight on the moistened, high-reflectivity sand frequently caused glare in the VME stereo camera images. These specular reflections had a negative impact on the performance of the VME localization. For a more complete discussion, see the results section.
- Rover Terrainability: The terrainability of the rover (See Figure 2) used in the test campaign is very limited. Having no suspension and limited ground clearance means that very small rocks can be hazardous. While the existing sensor suite certainly has the resolution to detect hazards for a planetary-representative chassis, with the Red Rover there are still some undetected rocks which, while not catastrophic, can cause the rover to stick and skid, or undergo abrupt pitch and roll motions.
- Shadows in the images: No effort was made to control the direction of travel with respect to the sun, and as such there are many instances in which the shadow of the rover appears in the stereo camera images. This ensures that the motion filtering in the feature matching algorithms is fully exercised.

2.2 Field Test Scenario

To better understand the sequence of operations during the test campaign, this subsection describes a typical long-range autonomous traverse scenario. This scenario can be classified as an autonomous navigation test in unknown terrain. The terrain is considered static, although the system is equipped with dynamic obstacle detection ability.

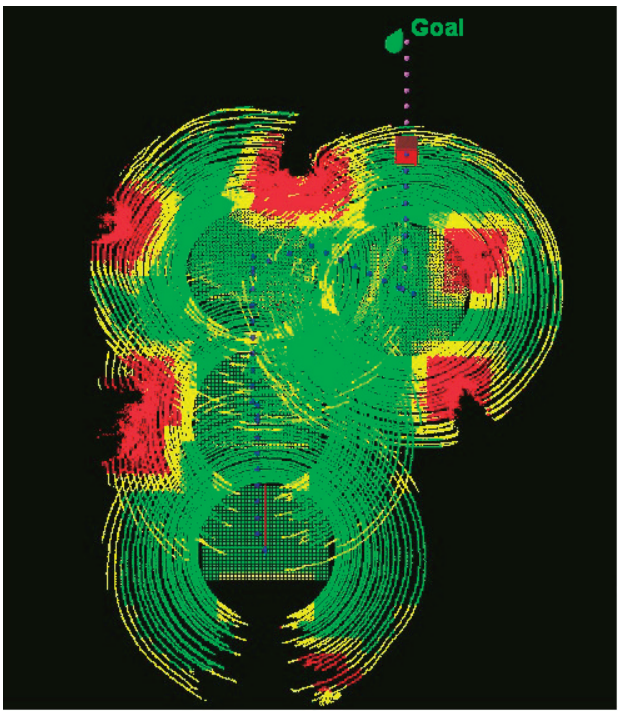

Figure 1: Typical Field test scenario, showing 3D terrain data from the sensor, colour coded according to traversability.

Typical operation (See Figure 1) starts with:

- A remotely located operator or scientist selects goal locations that are outside of the rover's on-board sensor ranging envelope, and/or obscured by an obstacle from a start location.
- The rover takes a high resolution medium range (10 to 15 m) scan of its surrounding environment using its integrated sensor suite.
- The rover constructs an internal representation of the surrounding terrain, conducts the terrain traversability assessment and plans a hazard free path toward the goal. This path is safe inside the scanned area, and aims directly at the goal outside the sensed area.
- The operator station displays the terrain map including traversability map overlays, and the planned path.
- The rover tracks the planned path up to the boundary of the medium range map and stops. Telemetry and environment data are collected during the motion and at the stop location.
- Mapping, terrain assessment, and path planning

are automatically repeated until the final destination is reached. If the goal is not reachable, the rover will move as close as possible to the defined goal and request a new goal or a new task.

All of the above steps are coordinated and monitored by a high level supervisory control and data acquisition module developed using the CSA's CORTEX autonomy framework.

2.3 Field Test Terrain Landscapes and Reconstructed Maps

The field tests were conducted in and around Silver Dry Lake, Silurian Dry Lake, the Mudflat, and Dumont Little Dunes in the Mojave Desert. The test environs offer many interesting terrains (in terms of rock, bush, and slope hazards, visual features, soil characteristics) all within a small geographical area.

Silver Dry Lake Test Locations

Three locations were selected for testing in the Silver dry lake bed: the "playa", the "boulder", and the "plateau" locations. The "playa" location is characterized by relatively flat, open, and hazard-free terrain. The soil is hard clay, which has cracked as it dries. This produces a fractal-like pattern of fissures on the surface. The fissures also make the surface bumpy, which can cause the rover to shake as it moves (See Figure 2).

The "boulder" field is characterized by relatively flat loose sand and playa, with large rock hazards randomly scattered throughout. (See Figure 3).

The "plateau" area is characterized by loose, small grain gravel with few rocks, but many medium-sized bushes (See Figure 4 and Figure 5).

Figure 2: Typical view of the Playa location

Silurian Dry Lake Test Locations

Three locations were selected for testing in the Silurian dry lake bed: the "playa", the "bush island", and the "shore" locations. The open playa of Silurian lake is virtually the same as the playa of Silver lake, with the exception that the fissures in the clay are slightly larger. The "bush island" location is in the middle of the Silurian lake playa, and contains a wide variety of soil conditions: hard cracked clay, soft cracked clay, gravel, and sand. There are several medium-sized bushes and a sparse scattering of rocks. There are also some mild gulleys to provide interesting slopes. The "shore" location is simply the shore of the Silurian dry lake bed, where the playa turns into loose gravel with some bushes (See Figure 6).

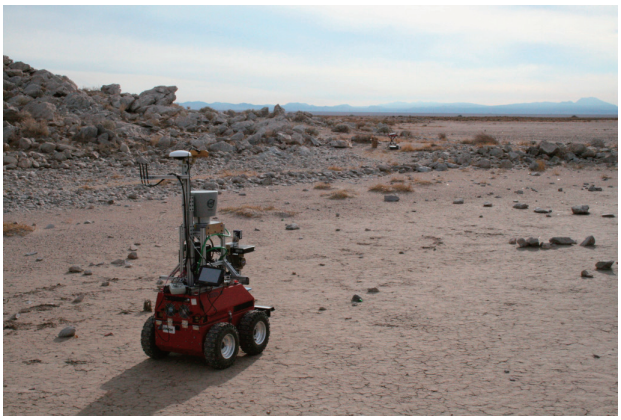

Figure 3: Typical view of the Boulder location

Mudflat Test Location

This area contains a lot of loose sand and gravel, with many ravines and gulleys. There are only a few bushes. Most of the hazards are due to slope rather than rocks (See Figure 7 and Figure 8).

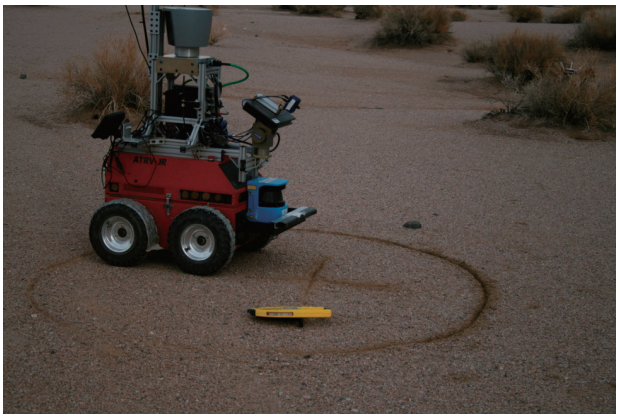

Figure 4: Typical view of the Plateau location

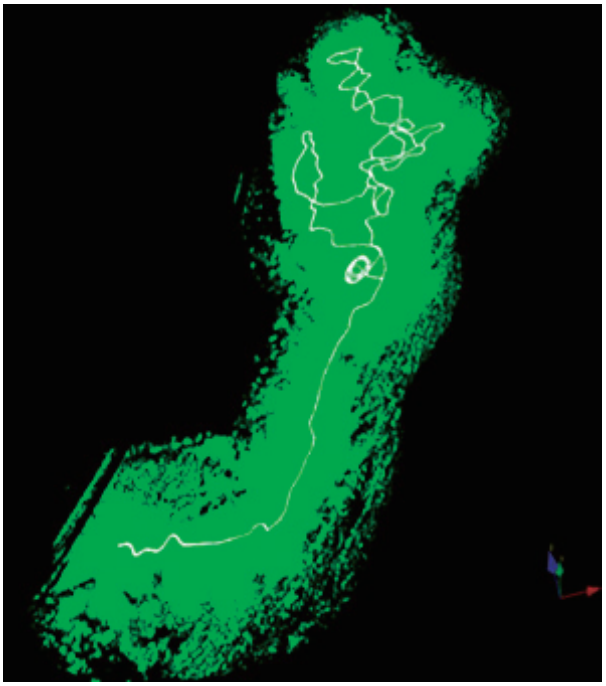

Figure 5: Reconstructed Digital Elevation Map of the plateau and the traveled path (200 x 255m, 960 m traveled distance).

Figure 6: Typical view of the Shore location

Figure 7: Typical view of the mudflat location

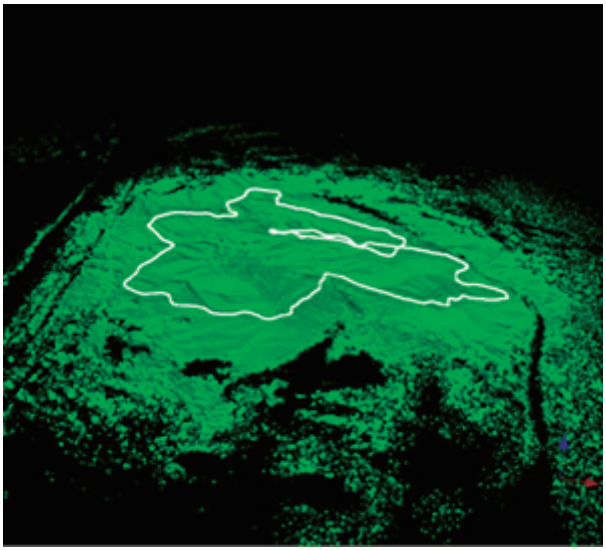

Figure 8: Reconstructed Mudflat Digital Elevation Map and traveled path overlays (160m x 140m, 445 m traveled distance)

Dumont Little Dunes Test Location

The terrain is loose sand dunes, with occasional sparse bushes as shown in Figure 9.

Figure 9: Typical view of the Dunes location

3 Experimental Results

This section describes the field trials of AIR-GNC performed in the Mojave Desert of Southern California in January, 2010.

3.1 Performance Metrics and Statistics

The Ground truth position was measured using the NavCom 3020M real-time kinematic differential global positioning system (RTK DGPS). The manufacturer's specification states an accuracy of 1 cm when operating

in RTK mode [14]. GPS latitude, longitude, and altitude data were converted to the local frame where positive X is East, positive Y is North, and positive Z is up (i.e., coincident with the local gravity direction).

Localization data was logged with respect to the initial rover body frame at the start of the traverse. As such, the ground truth must be expressed in the initial rover frame before accuracy analyses can begin. The position offset is easy to correct by simply subtracting the initial GPS position from all subsequent measurements. Pitch and roll are also compensated, since the VME initializes itself using the inclinometry from the IMU's accelerometers, which give very clean measurements since the rover is stationary. However, in the absence of an absolute heading sensor on the rover, the alignment of the ground truth becomes somewhat more challenging.

In this project, with so many different options to the VME algorithm, it was more desirable to align the ground truth with the IMU-corrected odometry. Note that no matter which signal is used for the alignment, the accuracy results will be biased (even if very slightly) in favour of that signal over the others. In the absence of absolute attitude ground truth, this problem cannot be avoided.

The distance between the end point of the estimate and the end point of the ground truth, referred to as end-point distance, is used to evaluate the localization accuracy.

3.2 Result and Analysis

VME localization accuracy results are compared based on the visual feature used for localization. Figure 10 shows the normalized localization percentage error by feature based on the data collected in all field test locations. The width of each bin of the histogram is 2% (of distance travelled), and for clarity, the axis is cut off at 10% error. All runs with errors greater than 10% (more likely outliers) are lumped into a single histogram bin, which is presented to the right of the regular histogram.

As can be seen in the figure above, on average MSER outperformed both Harris-Laplace and SIFT. This suggests that in these terrains the image regions had greater stability than the local (corner) features. On the average, Enhanced IMU-corrected odometry had slightly better accuracy than MSER. Figure 11 and Figure 12 show the normalized localization result of respectively the Enhanced IMU-corrected odometry and MSER across test locations.

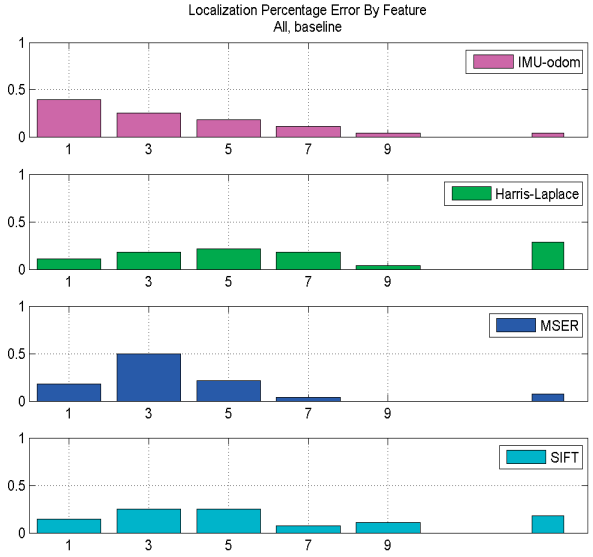

Figure 10: localization percentage error by feature

From Figure 11 and Figure 12 it can be seen that MSER outperformed Enhanced IMU-corrected odometry in the Playa area and the Dunes area, the latter certainly due to the high amount of wheel slippage.

The Mudflat location was particularly difficult for every localization method, including Enhanced IMU-corrected odometry. In the latter case, the terrain was quite loose, which caused considerable slippage. VME was primarily affected by the high reflectivity of the terrain (for example, the glare which was present in many runs). Moreover, the stereo cameras used in this project for VME have automatic brightness and contrast adjustment capability. One of the lessons learned from this field trial is that automatic exposure adjustment may actually hinder the performance of the stereo camera vision system. In many runs there is a significant and sometimes intense glare from the sunlight, which washes out large portions of the image and increases the contrast. Example runs in which the glare is very prominent include those on the Playa, and in the Mudflat location before the sun goes behind the mountains (See Figure 13). Further, this glare is specular, so it changes with the relative angle to the sun. As such, the glare patches themselves do not make very reliable features as they change with viewpoint. This effect was not noticed during the previous field campaign in 2008, which took place during the dry season [8].

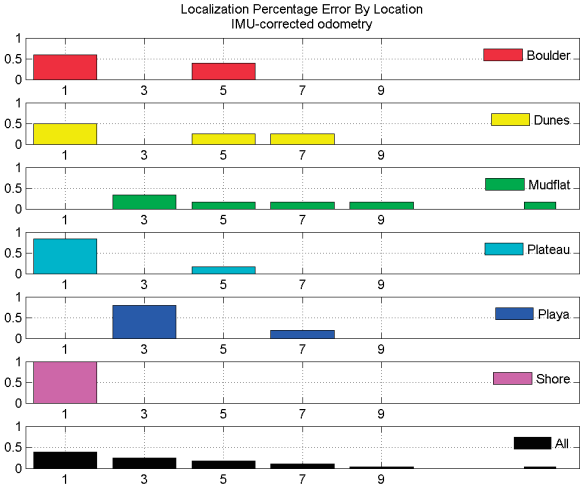

Figure 11: Enhanced IMU-corrected odometry localization percentage across all test locations

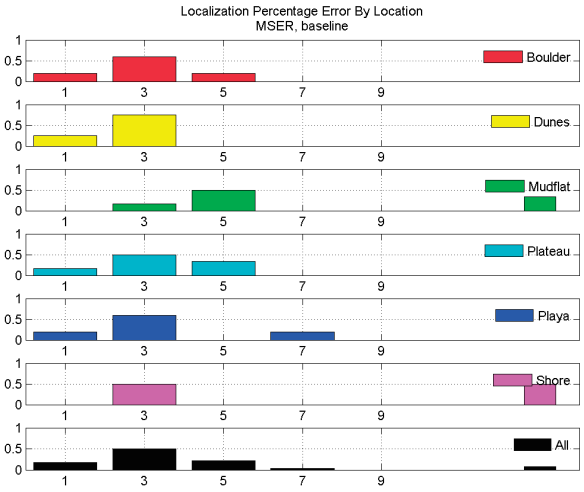

Figure 12: MSER localization percentage across all test locations.

3.3 Summary of the results

Table 1 summarizes the localization accuracy of stereo camera feature detectors shown in Figure 11.

Table 1: Localization accuracy summary for stereo camera feature detectors

Feature	Min % Error	Mean % Error	Std. Dev
Harris-Lap.	0.8	7.3	5.2
MSER	**0.4**	**3.9**	**2.8**
SIFT	1.2	8.8	11.9
IMU-Odo	**0.4**	**3.5**	**2.2**

MSER had better accuracy and lower standard deviation than both Harris-Laplace and SIFT. This suggests that in these types of terrains, region-based features are more useful than the local (corner) features. This is particularly noticeable in the Playa areas. MDA Enhanced IMU-corrected odometry did very well in most locations; slightly better than MSER. MSER outperformed Enhanced IMU-corrected odometry in the Playa and Dunes areas. Harris-Laplace performed nearly as well as Enhanced IMU-corrected odometry in the Shore area (the plot is not shown). The Mudflat location was particularly difficult for all methods of localization. As mentioned above, Enhanced IMU-corrected odometry was subject to slippage, and camera images were subject to glare.

Based on above results, an elegant localization solution will rely on Enhanced IMU-corrected odometry and continuous monitoring for wheel slippage. If the slippage starts, the localization will switch from Enhanced IMU-corrected odometry to VME and /or map-based localization and stay in this mode until the slippage ends. Another solution is to robustly extract and match long-range landmarks online to anable the observability of EKF-based SLAM. Offline results (See Table 2) using data collected during the field trials have shown very promising results.

Table 2: Offline Long-range landmark based observable VME (LRO).

Location	Distance	IMU-Odo	MSER	LRO
Mudflat, run 1	85 m	2.0%	3%	**0.6%**
Mudflat, run 2	21 m	3.0%	9.4%	**0.9%**
Plateau	34 m	1.5%	5.3%	**1.3%**

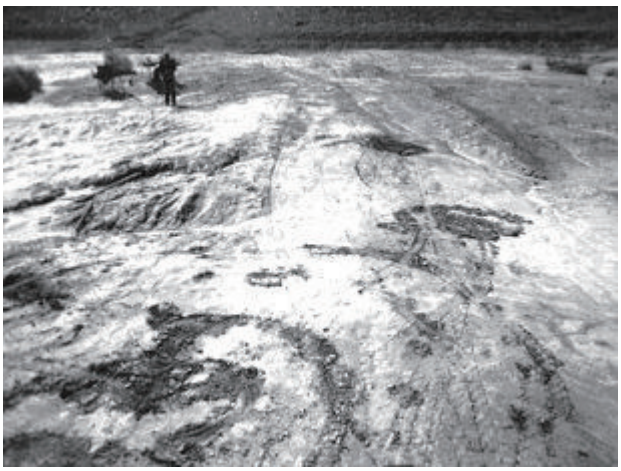

Figure 13: Mudflat terrain - High contrast due to glare off of the moist mudflat terrain

4 Conclusions and Future work

This paper has presented the field test results of the planetary VME for AIR-GNC project conducted by MDA in collaboration with and funded by the CSA.

The operational scenario of AIR-GNC, which includes terrain scanning, modeling and traversability assessment, and path planning and tracking, has been tested to a high level of maturity. CORTEX-based supervisory control was effective for monitoring and executing traverses, which resulted in a noticeable increase in the autonomy of the testbed. An added advantage is that experiments could be set up and run in a fraction of the time compared to the 2008 field campaign. VME-based feedback was successfully used to close the path tracking loop.

Enhanced IMU-corrected odometry performed very well; on average it performed slightly better than MSER based VME. MSER outperformed Harris-Laplace, and SIFT as a feature extraction front end to VME.

Future work includes wheel slippage monitoring and estimation using visual evidence, as discussed in section 3.3. This improved odometry could potentially be combined with state estimation using a dynamic model of the rover, and would be a wide-ranging benefit. For the visual front end, a fast and robust long range 3D-feature extraction method will be critical in order to take full advantage of observability.

5 Acknowledgment

This work was funded by the Canadian Space Agency. The authors would like to thank Régent L'Archeveque and David Gingras from the CSA for their valuable contribution in path planning and CORTEX based supervisory control. The authors also thank Manickam Umasuthan from MDA for his contributions to the 3D feature extraction implementation, and Tung Nguyen and Kelvin Tao, students at the University of Toronto, for integration of sensors and initial testing of the system.

References

[1] R. Volpe, Rover functional autonomy development for the Mars mobile science laboratory. In Proceedings of the IEEE Aerospace Conference, volume 2, pages 643–652, Big Sky, MT, USA, 2006.

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

[3] S. Lacroix, *et al.* "Autonomous Rover Navigation on Unknown Terrains. Demonstrations in the Space Museum Cité de l'Espace at Toulouse". In 7th Int. Symposium on Experimental Robotics. Honolulu, HI (USA), 2000.

[4] S. Lacroix S. Lacroix, *et al.*. "Autonomous Rover Navigation on Unknown Terrains: Functions and Integration", In International Journal of Robotics Research, 21(10-11), pages 917-942, 2002.

[5] P.I. Corke, *et al*. "Omnidirectional visual odometry for a planetary rover". In Proceedings of Intelligent Robot and System, 2004

[6] J.J. Biesiadecki,et al., "Mars exploration rover surface operations: Driving opportunity at meridiani planum". In IEEE Conference on Systems, Man and Cybernetics, The Big Island, Hawaii, USA, 2005.

[7] M.W. Maimone, Y. Cheng, L. Matthies, "Two years of visual odometry on the mars exploration rovers: Field reports", J. Field Robot., 24(3):169-186, 2007.

[8] J.N. Bakambu, C. Langley, and R. Mukherji, "Visual Motion Estimation: Localization Performance Evaluation Tool for Planetary Rovers", In International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), 2008.

[9] S. Se, T. Barfoot, P. Jasiobedzki, "Visual Motion Estimation and Terrain Modeling for Planetary Rovers", iSAIRAS 2005, Munich, Germany.

[10] F. Kristensen, W.J. MacLean, "Real-time Extraction of Maximally Stable Extremal Regions on an FPGA", IEEE Int. Symp. Circuits and Systems, 2007.

[11] K. Mikolajczyk, et al., "A Comparison of Affine Region Detectors", IJCV, 65(1-2): 43-72, 2005.

[12] E. Dupuis, P. Allard and R. L'Archevêque, "Autonomous Robotics Toolbox", iSAIRAS, Munich, 5-8 September 2005.

[13] T. Huntsberger, et al, "Rover autonomy for long range navigation and science data acquisition on planetary surfaces", IEEE ICRA 2002.

[14] NavCom Technology Inc., "RT-3020 GPS Products User Guide", 2008.