Field Testing of a Rover Guidance, Navigation, & Control Architecture to Support a Ground-Ice Prospecting Mission to Mars

Timothy Barfoot[a], Paul Furgale[a], Braden Stenning[a], Patrick Carle[a], Laura Thomson[b], Gordon Osinski[b], Michael Daly[c], Nadeem Ghafoor[d]

[a]University of Toronto Institute for Aerospace Studies
{tim.barfoot,paul.furgale,braden.stenning,patrick.carle}@utoronto.ca

[b]University of Western Ontario, Depts. of Earth Science, Physics and Astronomy
{lthoms2,gosinski}@uwo.ca

[c]York University, Dept. of Earth and Space Science and Engineering
dalym@yorku.ca

[d]MacDonald Dettwiler and Associates Ltd. (MDA), Space Missions
nadeem.ghafoor@mdacorporation.com

Abstract

This paper provides an overview of a rover guidance, navigation, and control (GN&C) architecture being developed to support a ground-ice prospecting mission to Mars. The main contribution of this paper is to detail an integrated field campaign that demonstrates the viability of the key rover GN&C techniques needed to carry out this mission. Tests were conducted on Devon Island in the Canadian High Arctic during the summer of 2009, wherein a large field robot was driven on real polygonal terrain (a landform of interest on Mars). Lessons learned and recommendations for future work are provided.

1. Introduction

Environmental conditions on Mars today are such that any water reserves will be in the form of ice, either in the polar caps or as ground ice at lower latitudes [6]. Deposits of ground ice may be key sites for future human exploration missions due to the possibility for in-situ resource utilization. Additionally, the climate history of Mars may be partly revealed through scientific study of the physiochemical characteristics of the detected ground ice. *Polygonal terrain* (a network of interconnected trough-like depressions in the ground) is a landform commonly found throughout the polar regions of both Earth [19, 25, 11] and Mars [27, 22]. In terrestrial environments, these features are formed by the response of a frozen substrate to thermal forcing mechanisms induced by winter freezing and subsequent warming later in the season and are often indicative of subsurface ice bodies [19]; on Mars, it is believed that such thermal forcing mechanisms may also be responsible for the observed formations [22, 30].

On Mars, the recent Phoenix lander mission [36] appears to have confirmed the presence of a frozen substrate in polygonal terrain, but the nature of any underlying massive ice bodies has not yet been determined. We have been investigating a rover mission concept, perhaps as a follow-on to Phoenix, to carry out ground-ice prospecting using a lidar (Light Detection and Ranging), a stereo camera, and a ground-penetrating radar (GPR) [2]. The main operational steps of our mission concept are shown in Figure 1. A landing site (e.g., containing polygonal terrain) is selected based on orbital imagery; the system lands. A large-scale initial site survey (out to a few hundred meters) is gathered using a lidar. A path is carefully planned (by a human) for the rover to traverse the terrain (e.g., crossing polygon boundaries perpendicularly to enable GPR scans). The rover drives the path, using visual odometry (VO) for motion feedback, while gathering information about the surface (using the stereo camera) and the subsurface (using the GPR). A coupled surface-subsurface model is created based on the VO motion estimate and the stereo/GPR data [12, 14]. Based on the stereo/GPR model, potential ground-ice deposits are identified for sampling. The rover uses a visual backtracking technique [13] to return along its path to the sample locations (and possibly all the way to the lander/ascent vehicle in the case of a sample return scenario).

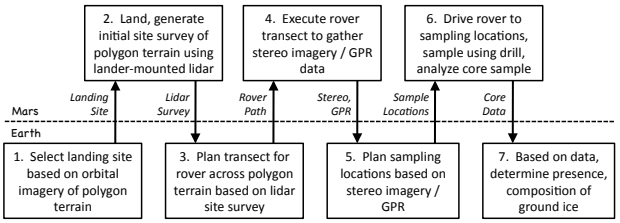

Figure 1: Flowchart of ground-ice prospecting mission concept introduced by Barfoot et al. [2].

The main contribution of this paper is to provide an overview of an integrated field campaign that demonstrates the viability of the key rover guidance, naviga-

tion, and control (GN&C) techniques required to carry out this ground-ice prospecting mission to Mars. Tests were conducted on Devon Island in the Canadian High Arctic during the summer of 2009. This paper extends the work of Barfoot et al. [2] in that a real field robot was used to test the engineering aspects of the mission concept (i.e., rover GN&C techniques), whereas the earlier paper focussed on showing the science instruments could produce the required data products to support the concept. The paper is organized as follows. Related work on rover GN&C for planetary exploration is discussed, followed by our specific GN&C architecture. Next, our experimental setup is presented, followed by our field test results. A discussion of the results ensues including an elaboration of our lessons learned and recommendations for future work, followed by a conclusion.

2. Related Work

Our mission concept employs lidar, stereo vision, and GPR as the main *scientific* instruments used for ground-ice prospecting. In this section we briefly introduce these instruments to motivate the required capabilities of our rover GN&C. We then summarize the state of the art in the areas of (i) traversing away from a lander, and (ii) backtracking to places along the outbound path.

2.1. Lidar, stereo camera, and GPR

Measuring surface properties and morphologies are of critical importance to understanding geological processes. This information can help to predict the presence of ground ice in, for example, ice-wedge polygons. Surface properties can also help in the interpretation of coupled subsurface geophysical data (e.g., GPR data). Two candidate instruments capable of measuring three-dimensional (3D) surface properties on future planetary missions are lidar and stereo camera. These instruments are complementary in that lidar actively illuminates a scene using laser light and works well from short to long range, while a stereo camera is passive, works only to short ranges, but can provide valuable colour information. Osinski et al. [31] shows the benefits of using a combination of lidar and stereo camera to build detailed 3D models of various geological features including polygonal terrain, terraced crater walls, impact breccias, and gullies in the Haughton impact structure. The Phoenix Lander had both a lidar, as part of its Meteorological Station [36], and a stereo camera as part of its Surface Stereo Imager [21]. The lidar was pointing upward and thus did not image the terrain around the lander. The stereo camera did produce 3D images of the terrain, but only out to a short range. The ability to move these instruments across the surface would be highly beneficial on future missions.

To prospect for ground ice, subsurface measurements are also required. One candidate instrument is ground-penetrating radar (GPR), which is widely used to determine subsurface structures and the distribution of ground ice based on differences in the dielectric properties of subsurface materials [1]. When a GPR survey is conducted along a surface traverse, individual traces can be combined to produce a two-dimensional profile showing continuous subsurface reflective layers, which allows for enhanced stratigraphic interpretation. Rover-based GPR has thus been proposed for a variety of planetary missions and will be included on the European Space Agency's ExoMars mission [37]. Fong et al. [9] investigated the use of lidar, stereo camera, and GPR to carry out site surveys on Devon Island; it is noteworthy that GPR traverses were precisely planned in advance based on orbital imagery and then driven using GPS for rover positioning. Furgale et al. [14] also detail a technique to build a coupled surface/subsurface model using a stereo camera and GPR, without the need for GPS.

2.2. Outbound traverse

There are typically several GN&C techniques required to drive a rover away from the lander towards some goal including (i) long-range (strategic) path planning, (ii) relative localization, (iii) path tracking, and (iv) short-range (tactical) path-planning to avoid local obstacles. For reasons related to the acquisition of our science data, we carry out long-range path planning manually and currently have not implemented short-range path planning, although the need for the latter is one of our lessons learned. Since path tracking is quite vehicle-specific and fairly well understood, we limit our discussion of the outbound traverse to the issue of relative localization (a.k.a., dead-reckoning), as we feel this is the main bottleneck in the process at the moment.

The MERs serve as a benchmark for comparison. On nominal terrain, they utilize wheel odometry measurements to estimate position changes. However, the MERs have frequently encountered steep slopes and sandy terrain, resulting in large amounts of wheel slip, thereby rendering the odometry measurements unreliable [26]. For these high-slip situations, they employ a technique called *visual odometry* (VO) over short distances. VO is based on earlier work by Matthies et al. [28], where updated vehicle poses are determined by tracking features between stereo camera image pairs. VO has proven to be a crucial tool in high-tilt terrains, and was used frequently in the Eagle and Endurance craters, as well as during climbing operations in the Columbia Hills [26]. With some improvements [17], the upcoming Mars Science Laboratory rover will also

make use of visual odometry. Note that VO is still a dead-reckoning technique and therefore estimation errors grow as a function of distance travelled. The MER implementations typically have errors of a few percent of distance travelled.

Over longer distances (e.g., a few hundred meters), the MERs periodically employ a technique called local bundle adjustment (BA) to improve their VO/odometric/inertial estimate [23]. Common tie points are found manually in a network of images along the rover's traverse and are then used in an optimization procedure to estimate all the associated rover poses. Typically, rocks are identified in a forward-looking PanCam image and then re-identified in a backwards-looking image at a new rover pose. Follow-on work has investigated the automation of the rock-matching step [24], but this has not yet been tested on Mars. The MER BA framework claims to have error growth as low as 0.1% of distance travelled over a 5 km traverse [23].

The other noteworthy method of improving VO/odometric/inertial estimates is matching of 3D lidar scans from different robot poses. This idea has been investigated by several research groups. Specifically for rover applications, the works of Se et al. [35], Rekleitis et al. [33], Dupuis et al. [8], and Wettergreen et al. [40] are representative examples of this concept. Variants of the popular Iterative Closest Point (ICP) algorithm [5] are typically used to align the lidar scans before and after a few hundred meters of travel. The framework of Se et al. [35] is most similar to our current work, as we employ VO to estimate position for a few hundred meters, then correct our pose by aligning before-and-after lidar scans.

Both local BA and lidar scan-to-scan matching still accumulate error growth as a function of distance travelled, albeit at a slower rate than VO/odometric/inertial estimates.

2.3. Inbound traverse

Traversing back to the lander/ascent vehicle is a less mature topic than the other aspects of rover GN&C. Both NASA and the European Space Agency are developing Mars-sample-return missions [34]. Landing with pinpoint accuracy at a site of scientific interest is either unlikely [7] or not possible (such as in our ice prospecting mission concept wherein in-situ data is needed to identify sampling locations). For sample-return missions, it would be extremely beneficial if the rover could perform an autonomous return to the lander/ascent vehicle in a single command cycle; this could allow the robot to spend most of the mission on the outbound traverse and then quickly return.

NASA has tested two GN&C technologies to allow a rover to return to a lander/ascent vehicle: a radio-frequency beacon and visual tracking. The visual tracking system identified the lander using the rover's mast camera from up to 125 m away. The rover tracked the lander in subsequent images as it approached [16, 3, 38]. This technique only works if the rover can see the lander, limiting its range of application.

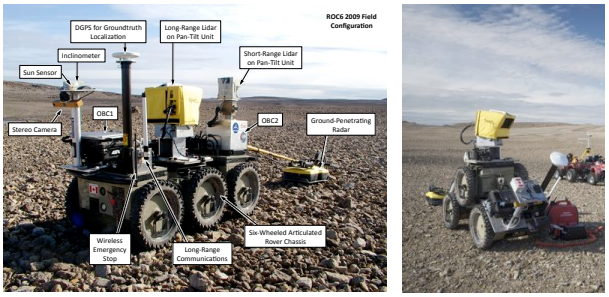

(a) Driving configuration. (b) Middle pod lifted.

Figure 2: ROC6 field robot at Lake Orbiter Site.

One concept that can work over the horizon is *visual homing*, which enables a return by 'memorizing' a route on the outgoing drive and then repeating it in reverse. Baumgartner and Skaar [4] develop a navigation system that uses visual cues (in this case artificial ring shaped markers), wheel odometry, and Kalman filter localization against a known map. The Kalman filter estimates are recorded and later used to retrace the path. While the map of the path is not globally consistent, it is locally consistent and allows the robot to navigate tight spaces and repeat its path.

The concept of a manifold of maps [15] was extended to the problem of visual homing for a mobile robot driving over rough outdoor terrain using a stereo camera for feedback by Furgale and Barfoot [13]. Here a sequence of locally-consistent overlapping maps of distinctive visual features is built up on the outbound pass and then used to reverse the path very accurately back to the start location. This was demonstrated over 32 km of driving on Devon Island during the summer of 2009. The rover drove without intervention for 99.6% of the distance. This serves as the visual homing implementation used in the current work.

3. Experimental Setup

3.1. Field robot

The ROC6 field robot, used as a proxy for a real planetary rover in this study, is a ruggedized mobile robot platform that may be used to carry out a variety of mission scenarios (see Figure 2(a)).

The chassis has six wheels on three main pods joined by active/passive hydraulic articulation joints. In passive mode, the base is able to keep all six wheels in good contact with the terrain ensuring excellent traction. In active mode, the center pod may be lifted

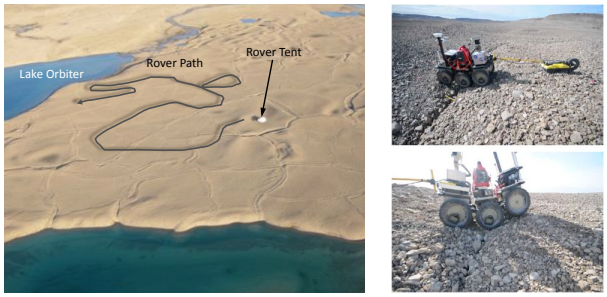

(a) Aerial photo showing Polygon Site with rover tent; 734 m rover path overlaid (Photo: Haughton-Mars Project).

(b) Rover in polygon troughs.

Figure 3: Lake Orbiter Polygon Site.

to, for example, gather a long-range lidar scan (see Figure 2(b)). The ROC6 has several engineering sensors/payloads including: long-range lidar (Optech ILRIS 3D), panoramic camera, stereo camera (Point-Gray Research XB3), ground-penetrating radar (Sensors & Software Noggin 250 MHz), and differential GPS (for benchmark localization).

3.2. Field test site

The experiments described in this paper were conducted on real polygonal terrain on Devon Island in the Canadian High Arctic, as part of the Canadian Space Agency's Canadian Analogue Research Network (CARN) program [32], and in conjunction with the Haughton-Mars Project Research Station (HMPRS) [20]. Our experiments were conducted approximately 10 km north of HMPRS near Lake Orbiter, at UTM 16X 419410 E 8380581 N.

Devon Island presents unique qualities for planetary analogue studies because it offers an unusually wide variety of geological features and microbiological attributes of strong planetary analogue value or potential. It has been used for rover testing in the past [39, 10, 9], as it presents real challenges to field exploration that are analogous in fundamental ways to those expected in planetary exploration. The lack of vegetation and variety of terrain make it well suited for rover GN&C field tests.

Generally, the Lake Orbiter site contains polygonal shapes that measure a few meters to tens of meters between subsequent troughs, with individual troughs averaging approximately 1-2 meters in width and tens of centimeters in depth. The entire surface is covered with rocks ranging from centimeters to tens of centimeters in size. Figure 3(a) shows an aerial view of the site (taken from a helicopter) with our rover traverse overlaid. Figure 3(b) shows the ROC6 field robot driving through two larger polygon troughs; the terrain was quite three-dimensional as can be seen by the degree to which the chassis articulated to keep its six wheels in contact with the ground.

4. Experimental Results

This section describes an end-to-end mission simulation we carried out using the ROC6 field robot at the Lake Orbiter site. The rover traversed 734 m away from the lander (broken down into three 'Legs') and then back. We describe our results in the following sections, organized by the steps in our mission architecture.

4.1. Steps 1-3: Select Landing Site, Land/Generate Initial Site Survey, Plan Transect

Our landing site was chosen based on a Quickbird satellite image of the site (viewable in Google Earth) that has a 60 cm/pixel resolution. This is similar in resolution to data that are available for portions of Mars from the HiRISE camera on the Mars Reconnaissance Orbiter [18, 29]. The primary polygon troughs can be easily identified at this scale. The site survey was conducted using our long-range lidar and a panoramic camera.

In our architecture, the outbound traverse path of the rover is planned manually. Figure 4 shows a screenshot of the graphical user interface (GUI) developed to carry out path planning. Both the lidar scan and the image panorama are displayed to the user. The lidar scan can be rendered from any viewpoint (i.e., it is 3D data). The user enters the rover's path as a series of waypoints joined by straight line segments. The GUI selects the lidar point closest to the selected screen point as the waypoint. Figure 4 shows the planned path for Leg 1 (based on Scan 1). Our science team, which carried out the planning, remained inside a large tent to ensure they could not augment their knowledge of the rover's situation by viewing it directly. Here we see the plan was chosen to cross several polygon boundaries in a perpendicular fashion to facilitate the GPR data collection and to ensure the rover would remain

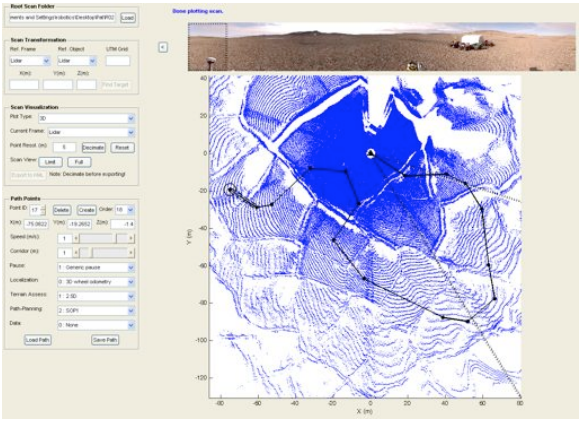

Figure 4: GUI used to plan rover traverse in lidar reference frame.

stable during the crossings. The path was planned as far as the polygon boundaries could be confidently identified in the lidar scan.

4.2. Step 4: Drive path

Once the rover's path had been planned, it attempted to drive this path using a vehicle-specific path-tracking controller and visual odometry for localization feedback. As with all VO implementations, there was a growth of localization error as a function of distance travelled. In other words, the 'Actual Outbound Path' (measured by DGPS) gradually drifts away from the 'Planned Path'. Figure 5 shows the path-following errors throughout the entire outbound traverse (i.e., all three legs). The error shown here is a combination of the path-tracking error and the VO localization error. However, the latter quickly dominates so that we can think of this as the VO localization error. Upon arriving at the end of a given leg, the localization error effectively drops back to zero because we plan a new path based on a new lidar scan from rover's current vantage point. Note, during the outbound traverse the rover is logging stereo imagery and GPR data for use in the next step of the GN&C architecture.

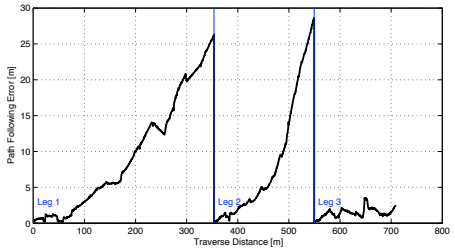

Figure 5: Growth of errors between planned and actual outbound path due to visual odometry drift.

At the end of each leg, when the rover stops, it is desirable to improve the VO estimate of rover pose using the next leg's lidar scan and/or image panorama. Figure 6 shows an example of aligning two lidar scans using a least-squares method with manually-selected tie points. Ten corresponding points are manually selected in each scan. Only three points are needed to solve for the pose between the scans, but using more than three effectively averages the errors associated with precisely picking identical points in the two scans. This reduced the VO localization error to approximately 1-2 m RMS position error (from 25-30 m RMS VO error).

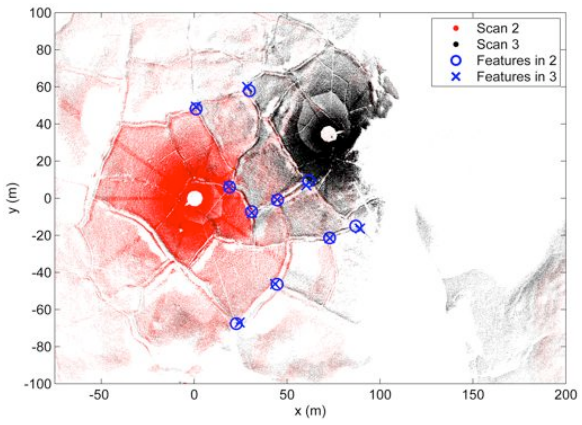

Figure 6: Manual tie point selection followed by least-squares alignment of lidar scans was carried out to correct VO localization at the end of each traverse leg.

4.3. Step 5: Choose subsurface sampling locations

In our architecture, at the end of each traverse leg, the rover builds a coupled surface/subsurface model using the logged stereo imagery and GPR data and sends this back to the science team to be used to select subsurface sampling locations. Although we did log the data during this field test, we did not create the coupled model because our focus is on the rover GN&C aspects not tested in previous work [2]. Figure 7 shows an example of a surface/subsurface model created using data from the same stereo camera and GPR units on a nearby polygon site [14].

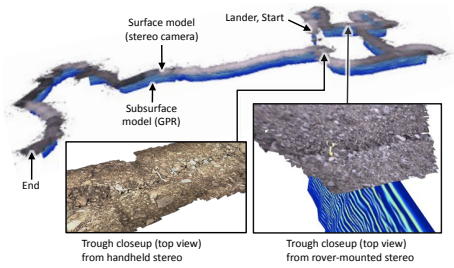

Figure 7: Example coupled surface/subsurface model from [14].

4.4. Step 6: Return to subsurface sampling locations

After completing all the outbound traverse legs, the mission scenario calls for the rover to return along its path to specific sampling stops, possibly all the way to the lander/ascent vehicle. To demonstrate this capability, we returned the rover along all three traverse legs to the start location using visual homing [13]. The sequence of overlapping locally-consistent maps used by visual homing was built automatically using the stereo imagery logged during the outbound pass. Note that the stereo camera was forward-looking during the outbound traverse and therefore backward-looking during the inbound traverse. We removed the GPR during the inbound traverse as we did not develop a means to reverse the route with it attached (this was not our focus).

Anecdotally, the rover often drove directly in its own tracks during visual homing; Figure 8 shows an example of this behaviour at another site on Devon Island (unfortunately, the tracks do not show up well in the photographs of the Lake

Figure 8: Rover driving in its own tracks during visual homing.

Orbiter site due to the nature of the terrain). Figure 9 shows the path-following errors for the inbound traverses of Legs 1, 2, and 3. The errors plotted here are between the actual outbound path (measured using DGPS) and the actual inbound path (measured using DGPS). During the inbound traverse of Leg 2, the GPS system dropped out of differential mode and thus only regular GPS information is available for this leg (i.e., the groundtruth system was not adequate to measure the performance of visual homing). For Legs 1 and 3, we see that the path-following error is generally less than 1 m, meaning the rover was able to very accurately drive along the same route back to the start without human intervention. Figure 10 also plots both the outbound and inbound traverses, although it is hard to tell them apart at this scale.

Pausing at any designated location along the inbound traverse in order to gather subsurface samples would be straightforward, although we did not demonstrate this. Actually gathering the samples from the subsurface is a more difficult task that we have not yet attempted to address in this research.

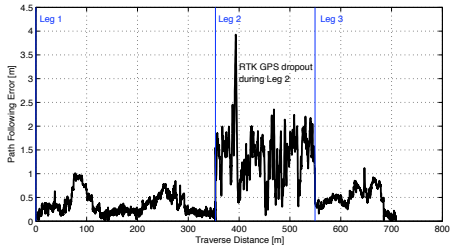

Figure 9: Bounded error between actual inbound and outbound paths due to visual backtracking.

5. Discussion

We have learned several lessons during our field test that may inform future research in this area. Most importantly, we learned that there is merit to conducting integrated field tests such as the one reported here. By bringing together all the requisite GN&C elements into a single test, we were able to identify the deficiencies in their ability to enable the overall mission concept. See Figure 10 for a summary of our results. Clearly, it would be nice to be able to report on a much larger number of tests, but running large-scale robotics experiments in this remote venue was logistically challenging. Still, we believe there are some specific findings worth detailing from our campaign.

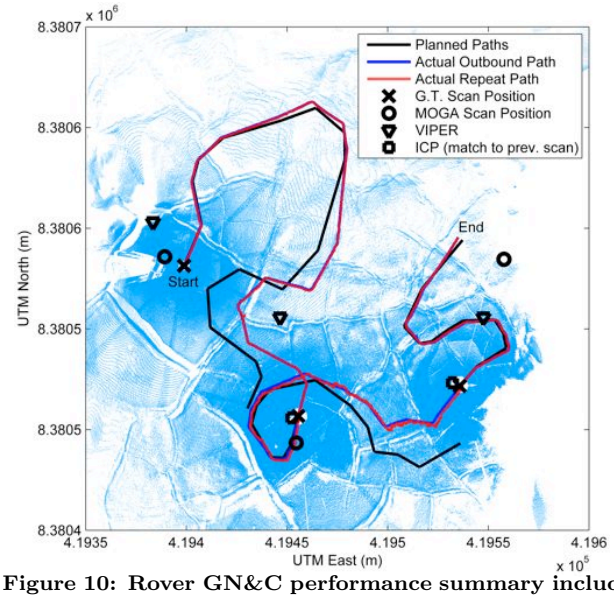

Figure 10: Rover GN&C performance summary including: (i) groundtruth scan positions and aligned lidar scans of polygon terrain, (ii) planned paths based on lidar scans, (iii) actual outbound path using VO for localization, (iv) improved localization (using VO as the initial guess) at the end of Legs 1 and 2, and (v) actual inbound path using visual homing.

The users of the graphical user interface during our test were planetary scientists, who were not involved in the GUI's development. They planned the rover's path with advice from a rover specialist. The main lesson learned in this step is that human users are very comfortable with image panoramas and less comfortable with lidar data. The GUI allowed both data sources to be viewed simultaneously and visually correlated using annotations on the screen. Although the GUI could certainly be refined, the capability to view images and lidar data together was found to be very successful. Generally, the idea of planning the rover's path manually based on the raw geometry of the polygon terrain (as measured by the lidar scan) did prove to be successful as it provided adequate situational awareness of the polygonal terrain's specific structure around the rover.

In terms of driving the outbound path, our main conclusion here is that we either need to improve the accuracy of visual odometry, or reduce the length of each traverse leg. A combination of these two changes is also possible. From Figure 5 we see that the position errors on our traverse legs grew as high as tens of meters. This is not acceptable for two reasons. First, the rover may not cross the polygon troughs perpendicularly, thereby compromising the science data collection; we can see the manifestation of this effect in Figure 10. Second, the rover may drive into unsafe terrain when it deviates from the planned path by a large amount. Knowing the limitations of VO, we can partially account for the second issue by ensuring there

is a safety corridor on either side of the planned path. However, as Figure 5 indicates, the VO localization errors grow super-linearly with distance travelled, meaning the safety corridor would need to grow impractically large. A quick solution is to reduce the traverse legs to less than 150 m each, whereby the localization errors would shrink to about 2-3% of distance travelled. A more permanent solution would be to switch to a VO implementation that maintains good accuracy over longer distances.

Another lesson learned is that a simple hazard-detection system should be implemented as there will inevitably be some deviation between the planned and driven paths. As the stereo camera is forward-looking during the outbound pass, a basic stereo reconstruction could serve in this regard. We did not include such a monitoring system due to lack of development time, but will do so in future work.

We found that visual homing was able to very accurately return the rover along its outbound path back to the start location. It is worth stressing that the performance of visual homing is unaffected by any path-following errors accumulated during the outbound traverse; visual homing simply drives the robot back along the actual outbound path. The biggest concern with

(a) Image from the outbound pass.

(b) Image from the failed inbound traverse 4 hours later.

Figure 11: Images from the start of Leg 2 visual homing (all taken at Scan 3 location) showing that lighting can cause drastic changes to a scene's appearance.

visual homing is lighting sensitivity. Because the technique is appearance-based, drastic changes in lighting can cause the algorithm to be unable to recognize places it has visited before. During the summer in the Arctic, the sun moves rapidly around the horizon and thereby quickly changes the shadow directions. The Lake Orbiter site is also littered with many similar-sized rocks. The combination of these effects caused visual homing to be unable to report the rover's position at the start of Leg 2 (see Figure 11). A period of a several hours had elapsed between the outbound and inbound traverses, changing the appearance of the scene. Visual homing failed to recognize the location. With no changes to the algorithm, a second attempt the next morning allowed the inbound traverse to be completed successfully.

6. Conclusion

We have presented the results of an integrated field test of a rover GN&C architecture designed to support a ground-ice prospecting mission to Mars. The experiments were conducted at a Mars analogue site on Devon Island in the Canadian High Arctic on real polygonal terrain using a real field robot. The novel contribution of this paper is the combination and critique of several rover GN&C techniques in a single test. We believe that the mission architecture presented in this paper is viable; however, we identified a number of improvements that need to be made to the individual rover GN&C components and their interactions to reach adequate performance levels for a real mission.

Acknowledgements

We would like to thank the Canadian Space Agency for funding this field campaign under its Canadian Analogue Research Network Grant Program. Additional funding was provided by the Northern Scientific Training Program, Natural Sciences and Engineering Research Council of Canada, and Canada Foundation for Innovation. We would also like to thank the Haughton-Mars Project, the Polar Continental Shelf Project, and Technical Field Support Services for logistics and field support. Optech Inc. and Marwan Hussein generously provided support of the lidar system during our test. Matt Izawa of the University of Western Ontario also provided invaluable field support at the Lake Orbiter site. Sensors & Software Inc. generously provided equipment loans/donations in the form of the GPR.

References

[1] Arcone, S., Prentice, M., Delaney, A., 2002. Stratigraphic profiling with ground-penetrating radar in permafrost: A review of possible analogs for mars. Journal of Geophysical Research 107 (E11), 5108.

[2] Barfoot, T. D., Furgale, P. T., Osinski, G. R., Ghafoor, N., Williams, K., March 2010. Field testing of robotic technologies to support ground-ice prospecting in martian polygonal terrain. Planetary and Space Science, 58 (4), 671–681.

[3] Baumgartner, E. T., Leger, P. C., Schenker, P. S., Huntsberger, T. L., 1998. Sensor-fused navigation and manipulation for a planetary rover. In: Proc. of SPIE Vol 3523.

[4] Baumgartner, E. T., Skaar, S. B., 1994. An autonomous vision-based mobile robot. Automatic Control, IEEE Transactions on 39 (3), 493–502.

[5] Besl, P., McKay, H., 1992. A method for registration of 3-D shapes. IEEE Transactions on pattern analysis and machine intelligence 14 (2), 239–256.

[6] Carr, M. H., 1996. Water on Mars. Oxford University Press, New York.

[7] Clark, B. C., 2007. Mars sample return: The critical next step. Acta Astronautica 61 (1-6), 95–100.

[8] Dupuis, E., Rekleitis, I., Bedwani, J. L., Lamarche, T., Allard, P., Zhu, W. H., 2008. Over-the-horizon autonomous rover navigation: Experimental results. In: Proc. of the 9th

Int. Symp. on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS). Los Angeles, CA.

[9] Fong, T., Allan, M., Bouyssounouse, X., Bualat, M., Deans, M., Edwards, L., Fluckiger, L., Keely, L., Lee, S., Lees, D., To, V., Utz, H., February 26-29 2008. Robotics site survey at haughton crater. In: Proceedings of the 9th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS). Los Angeles, CA.

[10] Fong, T., Deans, M., Lee, P., Bualat, M., March 12-16 2007. Simulated lunar robotic survey at terrestrial analog sites. In: Proceedings of the 38th Lunar and Planetary Science Conference. League City, Texas.

[11] Fortier, D., Allard, M., 2004. Late holocene syngenetic ice wedge polygons development, bylot island, canadian arctic archipelago. Canadian J. of Earth Sciences 41, 997–1012.

[12] Furgale, P., Barfoot, T. D., Ghafoor, N., 14-16 July 2009. Rover-based surface and subsurface modeling for planetary exploration. In: Proceedings of Field and Service Robotics (FSR). Cambridge, MA.

[13] Furgale, P. T., Barfoot, T. D., 2010. Visual teach and repeat for long-range rover autonomy. Journal of Field Robotics, *special issue on "Visual mapping and navigation outdoors"*. Manuscript # ROB-09-0081, accepted on March 31, 2010.

[14] Furgale, P. T., Barfoot, T. D., Osinski, G. R., Williams, K., Ghafoor, N., 2010. Field testing of an integrated surface/subsurface modeling technique for planetary exploration. International Journal of Robotics Research, *special issue on "Field and Service Robotics"*. Manuscript # IJR-09-0636, accepted on March 23, 2010.

[15] Howard, A., Sukhatme, G. S., Mataric, M. J., 2006. Multirobot simultaneous localization and mapping using manifold representations. In: Proc. of the IEEE. Vol. 94. pp. 1360–1369.

[16] Huntsberger, T. L., Baumgartner, E. T., Aghazarian, H., Cheng, Y., Schenker, P. S., Leger, P. C., Iagnemma, K. D., Dubowsky, S., September 1999. Sensor fused autonomous guidance of a mobile robot and applications to mars sample return operations. In: Proc. SPIE Symp. on Sensor Fusion and Decentralized Control in Robotic Systems II, Vol. 3839.

[17] Johnson, A. E., Goldberg, S. B., Cheng, Y., Matthies, L. H., 2008. Robust and efficient stereo feature tracking for visual odometry. In: IEEE Int. Conf. on Robotics and Automation. pp. 39–46.

[18] Johnston, M., Graf, J., Zurek, R., Eisen, H., 2005. The Mars Reconnaissance Orbiter mission. In: IEEE Conference on Aerospace. pp. 1–18.

[19] Lachenbruch, A., 1962. Mechanics of thermal contraction cracks and ice-wedge polygons in permafrost. Special paper to the Geological Society of America 70, 69.

[20] Lee, P., et al., 2007. Haughton-mars project: 10 years of science operations and exploration systems development at a moon/mars analog site on devon island, high arctic. In: Proceedings of the 38th Lunar and Planetary Science Conference. League City, Texas, pp. 2426–2427.

[21] Lemmon, M. T., et al., 2008. The phoenix surface stereo imager (ssi) investigation. In: Proceedings of the 39th Lunar and Planetary Science Conference. No. 2156.

[22] Levy, J., Head, J., Marchant, D., 2009. Thermal contraction crack polygons on mars: Classification, distribution, and climate implications from hirise observations. Journal of Geophysical Research 114 (E01007).

[23] Li, R., Arvidson, R. E., Di, K., et al., February 2007. Opportunity rover localization and topographic mapping at the landing site of Meridiani Planum, Mars. Journal of Geophysical Research 112 (E2).

[24] Li, R., Di, K., Howard, A. B., March 2007. Rock modeling and matching for autonomous long-range Mars rover localization. Journal of Field Robotics 24 (3), 187–203.

[25] Mackay, J. R., Burn, C. R., 2002. The first 20 years (1978-1979 to 1998-1999) of ice wedge growth at the illisarvik experimental drained lake site, western arctic coast, canada. Canadian Journal of Earth Sciences 39 (1), 95–111.

[26] Maimone, M., Cheng, Y., Matthies, L., 2007. Two years of visual odometry on the mars exploration rovers. Journal of Field Robotics 24 (3), 169–186.

[27] Mangold, N. A., 2005. High latitude patterned grounds on mars: Classification, distribution, and climatic control. Icarus 174, 336–359.

[28] Matthies, L., Maimone, M., Johnson, A., Cheng, Y., Willson, R., Villalpando, C., Goldberg, S., Huertas, A., Stein, A., Angelova, A., 2007. Computer vision on mars. International Journal of Computer Vision 75 (1), 67–92.

[29] McEwen, A., Delamere, W., Eliason, E., et al., March 2002. HiRISE - The High Resolution Imaging Science Experiment for Mars Reconnaissance Orbiter. In: Lunar and Planetary Science XXXIII, Houston, TX.

[30] Mellon, M. T., Arvidson, R. E., Marlow, J. J., Phillips, R. J., Asphaug, E., 2008. Periglacial landforms at the phoenix landing site and the northern plains of mars. Journal of Geophysical Research 113 (E00A23).

[31] Osinski, G. R., Barfoot, T. D., Ghafoor, N., Izawa, M., Bangerjee, N., Jasiobedzki, P., Tripp, J., Richards, R., Auclair, S., Sapers, H., Thomson, L., Flemming, R., March 2010. Lidar and the mobile scene modeler (mSM) as scientific tools for planetary exploration. Planetary and Space Science, 58 (4), 691–700.

[32] Osinski, G. R., Leveille, R., Berinstain, A., Lebeuf, M., Bamsey, M., 2007. Terrestrial analogues to mars and the moon: Canada's role. Geoscience Canada 33 (4), 175–188.

[33] Rekleitis, I., Bedwani, J.-L., Dupuis, E., 2007. Over-the-horizon, autonomous navigation for planetary exploration. In: Bedwani, J.-L. (Ed.), Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems 2007. pp. 2248–2255.

[34] Schenker, P. S., Huntsberger, T. L., Pirjanian, P., Baumgartner, E. T., Tunstel, E., 2003. Planetary rover developments supporting mars exploration, sample return and future human-robotic colonization. Autonomous Robots 14 (2), 103–126.

[35] Se, S., Ng, H.-K., Jasiobedzki, P., Moyung, T.-J., 2004. Vision based modeling and localization for planetary exploration rovers. In: Proceedings of the 55th International Astronautical Congress. Vancouver, Canada.

[36] Smith, P. H. et al., 2008. Introduction to special section on the phoenix mission: landing site characterization experiments, mission overviews, and expected science. Journal of Geophysical Research 113 (E00A18).

[37] Vago, J., Gardini, B., Kminek, G., Baglioni, P., Gianfiglio, G., Santovincenzo, A., Bayón, S., van Winnendael, M., May 2006. ExoMars - searching for life on the Red Planet. ESA Bulletin 126, 16–23.

[38] Volpe, R., Baumgartner, E., Scheaker, P., Hayati, S., 2000. Technology development and testing for enhanced mars rover sample return operations. In: Aerospace Conference Proceedings, 2000 IEEE. Vol. 7. pp. 247–257 vol.7.

[39] Wettergreen, D., Dias, M., Shamah, B., Teza, J., Tompkins, P., Urmson, C., Wagner, M., Whittaker, W., May 2002. First experiment in sun-synchronous exploration. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA). Washington, DC, pp. 3501–3507.

[40] Wettergreen, D., Jonak, D., Kohanbash, D., Moreland, S., Spiker, S., Teza, J., 14-16 July 2009. Field experiments in mobility and navigation with a lunar rover prototype. In: Proceedings of the 7th International Conference on Field and Service Robotics. Cambridge, Massachusetts, USA.