PLANETARY SURFACE EXPLORATION USING A NETWORK OF REUSABLE PATHS: A PARADIGM FOR PARALLEL SCIENCE INVESTIGATIONS

B.E. Stenning[1], G.R. Osinski[2], T.D. Barfoot[1], G. Basic[1], M. Beauchamp[2], M. Daly[3], R. Francis[2], P. Furgale[1], J. Gammell[1], N. Ghafoor[4], P. Jasiobedzki[4], A. Lambert[1], K. Leung[1], M. Mader[2], C. Marion[2], E. McCullough[2], C. McManus[1], J. Moores[2], and L.J. Preston[2]

[1]*University of Toronto Institute for Aerospace Studies, Canada,* {braden.stenning, tim.barfoot}@utoronto.ca,
[2]*University of Western Ontario, Depts. of Earth Science, Physics and Astronomy, Canada,* gosinski@uwo.ca,
[3]*York University, Dept. of Earth and Space Science and Engineering, Canada*
[4]*MacDonald Dettwiler and Associates Ltd. (MDA), Space Missions, Canada*

ABSTRACT

A *network of reusable paths* (NRP) allows for a new approach to planetary surface exploration using a mobile robot. NRP gives the robot the ability to accurately return to any previously visited point. This allows mission-level improvements by enabling parallel exploration of scientific targets. NRP would be particularly useful for sample-return missions to the Moon or Mars. The approach was tested in a mock Lunar sample-return mission near the impact crater located in Sudbury, Ontario, Canada. There, NRP enabled nearly twice as many sites to be investigated as compared to a serial approach to exploration. In this mock mission, the robot drove more than 3.9 km, allowing for in situ analysis and sample collection and return at many sites.

Figure 1. A robot operating on a network of reusable paths. The robot is repeating a previously driven path (in its own tracks) to return to a previously visited position.

1. INTRODUCTION

The Mars Exploration Rovers (MERs) have driven over 40 kilometres, visiting many sites of scientific interest along the way. The exploration strategy for each rover was serial in the sense that scientific objectives were completed at one site before departing for the next [1]. This means the robot remained in place while mission controllers decided which measurements to collect. Fig. 2 shows the traverse map for the Spirit MER as of sol 2555. Note that many of the sites were near each other, and on several occasions the rover would roughly follow its previous track, over the course of several sols and many command cycles, to return near to a previous position.

The coming decades will see sample-return missions to both Mars and the Moon. Here, we advocate for a planetary exploration strategy that allows sites of interest to be studied in parallel, rather than in series. We believe this better supports the overarching aims of sample-return missions, as a methodical down-selection process may be employed to identify the key specimens to be returned to Earth. We show that by using a *network of reusable paths* (NRP) [2] a rover can revisit places of scientific interest and thus allow the study of sites in parallel. This new approach was field tested (see Fig. 1) in a mock Lunar sample-return mission conducted near the Sudbury impact crater in Canada [3, 4, 5, 6]. In this paper, we emphasize three points:

1. *NRP allows a robot to return to a previously visited position with a single command,*
2. *this allows for parallel exploration, and*
3. *parallel exploration allows for an efficient down-selection process to identify key samples for return.*

We discuss NRP in Section 2, and in Section 3 we present the mock sample-return mission. Section 4 identifies other uses of NRP, and in Section 5 we identify challenges, along with some future works.

2. A NETWORK OF REUSABLE PATHS FOR PLANETARY SURFACE EXPLORATION

Next, we give a background on the network-of-reusable-paths approach and its enabling technology. We then present how NRP can be used in a sample-return mission.

2.1. Extending Visual Teach and Repeat to NRP

A network of reusable paths is an extension of *visual teach and repeat* (VT&R) [7, 8]. A VT&R system allows a robot to drive arbitrarily long distances, without the use of GPS, along previously established routes. In

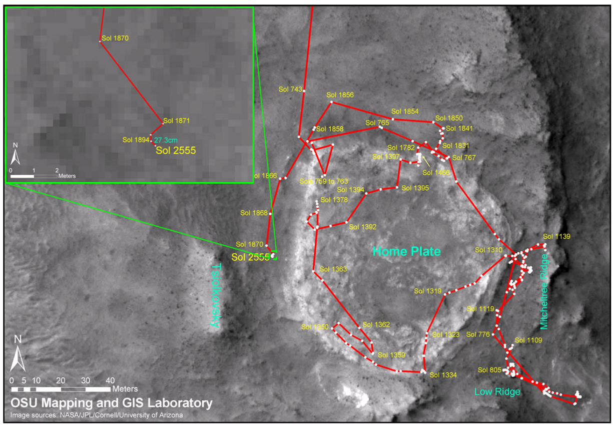

Figure 2. Traverse map at sol 2555 for the Mars Exploration Rover, Spirit. Credit: OSU Mapping and GIS Laboratory, NASA/JPL/Cornell/University of Arizona.

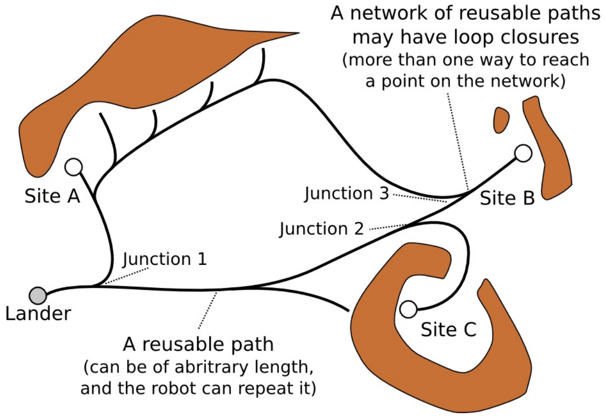

Figure 3. A simple network of reusable paths is shown in black. The robot can return to any point on the network and can grow the network into new areas. To go from site B to C, the rover reuses the previous paths by traveling through junction 3 and then 2, before going to site C.

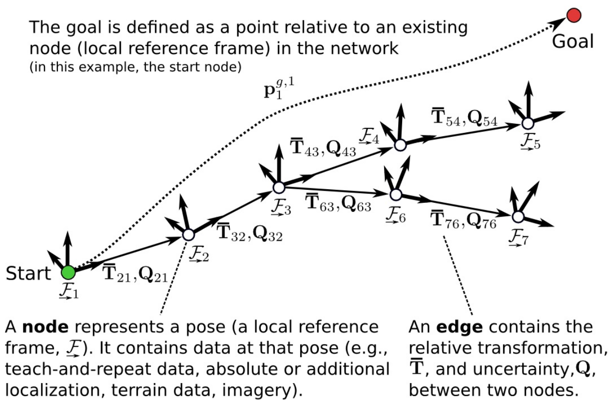

Figure 4. Definition of the key components of the graph structure used in a network of reusable paths. There is no privileged coordinate frame, everything is relative.

these systems, a chain of small maps is attached along the robot's path (estimated using visual odometry [9]) during a teaching phase; to repeat the path, the robot localizes against each small map in sequence as it drives. At any time the robot can return to a previous position on the path. Visual teach and repeat remembers the path that was traveled, and the case for doing so is strong; *knowing that the robot has already successfully driven the path is strong evidence that the path is traversable.*

The capability of the system of Furgale and Barfoot [7] has been demonstrated through 32 km of autonomous driving in both an urban setting and a planetary analogue environment in the Canadian High Arctic. The system, making use of a stereo camera, often performed so well that it repeated the path in its own tracks. It autonomously drove all but a few tens of meters of the desired 32 km of paths (< 0.4 % of distance traveled). A lighting-invariant extension to this work, using a high-framerate lidar [8], has been developed to address one of the major challenges of the stereo-camera-based system [10]. Namely, appearance changes due to changing lighting conditions can make it challenging, or impossible, to localize against the map when revisiting places. The lidar-based system allows for operation in complete darkness, making it suitable for exploration of permanently-shadowed regions such as those at the Lunar South Pole.

There are other approaches to teach and repeat. One makes use of a planar laser rangefinder for underground mining applications [11]. Another example uses an omnidirectional camera [12], and still others use different techniques that provide a similar teach-and-repeat function [13, 14, 15, 16, 17].

The NRP concept arises when we recognize that VT&R systems can be extended, from using a simple chain of local maps, to an arbitrary network of local maps [2]. Consider the example shown in Fig. 3. The robot can return to any point on the network (shown in black), and by driving into new areas, the network can be extended.

The NRP is represented by a graph, $\mathcal{G}$, that consists of a set of nodes, V, and a set of edges, E (see Fig. 4).

Each node represents a previous pose, and the node contains the local visual-landmark-map (used by VT&R) associated with that pose. All poses are relative. The estimated mean transformation, $\bar{\mathbf{T}}$, and the associated covariance matrix representing uncertainty, $\mathbf{Q}$, are stored at an edge connecting two nodes. The actual transformation, $\mathbf{T}$, is unknown. The relative transformation and uncertainty between any two connected nodes can be found by compounding the relative transforms (and uncertainties) along a chain joining the two. The result is a graph similar to the paradigm of Sibley et al. [18], in that there is no privileged coordinate frame; everything is relative. The nodes and edges are the paths the robot has previously taken; based on existing VT&R capabilities, we assume that these paths (and subsets) can be repeated exactly, in either direction, with the robot in the same orientation as the initial (teach, mapping) pass. Other information can be stored at the edges or the nodes (e.g., terrain assessment data, absolute localization such as that from orbital observation, imagery, science data).

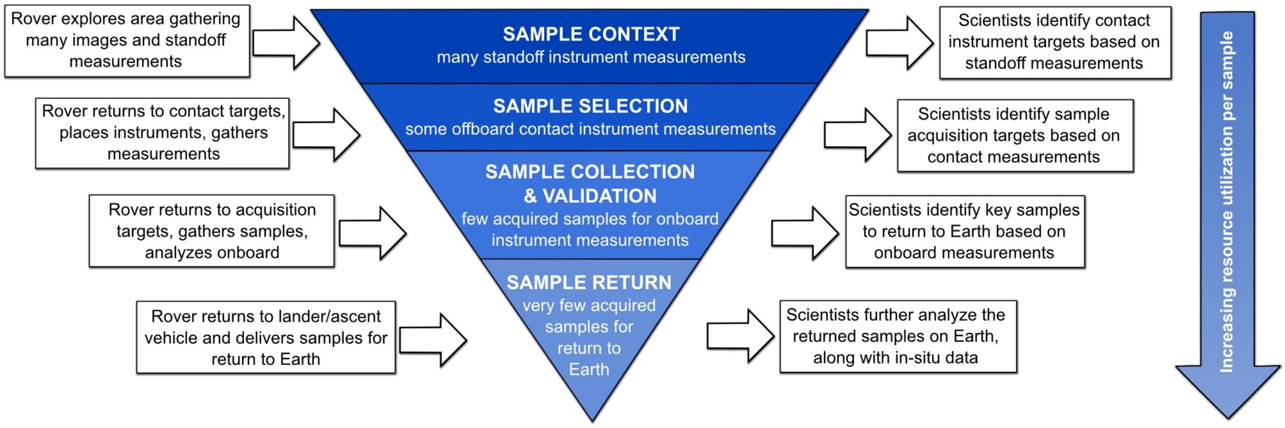

Figure 5. A methodical down-selection process is enabled by a network of reusable paths. There are a decreasing number of samples at lower levels to accommodate the higher resource usage per sample.

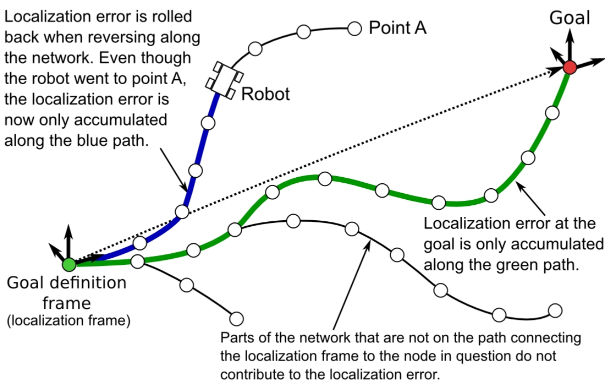

Figure 6. The localization error is only due to the errors in the transformations at the edges that connect the node used as the localization base frame, and the node against which the rover is localizing. This means the localization error at the goal is only accumulated on the final path to the goal from the goal definition node, and that error is rolled back when reversing along a previous route.

The relative approach also extends to defining waypoints (goals, target positions), in situ analysis tasks, or sampling tasks. These can be defined relative to any node on the network, rather than some global reference frame or the current frame of the robot. In Fig. 4 the goal location, G, is designated as a point in one of the reference frames at a node, called the goal definition node, x_{gd}. This is not necessarily the initial robot pose, x_0, or the current robot pose, x_r. Similarly, a sampling location near Site B (see Fig. 3) can be defined in a nearby reference frame, even when the robot is not physically present at that site.

Using local reference frames to define waypoints and tasks leads to more accurate results. In NRP, the robot uses relative localization when adding to the network, but when reusing a previous path the system localizes against the network and does not accumulate additional pose uncertainty. As in Fig. 6, the localization error at a node is only due to the errors in the transformations along the path through the network that connects the current node with the node in which the localization is expressed.

This means the localization error at the goal is only accumulated on the final path to the goal from the goal definition node, and that localization error is rolled back when reversing along a previous route, i.e., if the robot encounters a dead end, that dead end does not influence the accuracy of the goal acquisition. Systems that use only dead-reckoning will accumulate localization uncertainty along the entire length of the traverse, including when the robot retreats from a dead end to return to a previous position.

2.2. A NRP Approach to Exploration

In the sample-return scenario, NRP allows for a methodical down-selection process as shown in Fig. 5. This process is possible because the robot can return to any previous position, and therefore tasks and waypoints can be defined relative to any previous position. In a sense, this allows the instructions to be parallel, in that in a sequence of complex steps, each step is not defined relative to the predicted end-point of the previous step. Instead, NRP encourages setting short-range, parallel objectives that can be reliably completed in a single command cycle and then built upon in later command cycles once mission controllers have reviewed the resulting telemetry.

Consider, again, the example network in Fig. 3. Here, there are three sites of interest that are being investigated in parallel. While operators on Earth discuss a decision on where to sample at site A, they can send the robot to site B and then C to collect imagery before returning to site A. Then, while the robot is sampling at site A the mission team can use the data from Site B and C to select another potential sampling site. The rover does not need to loiter at a particular site of interest until all the work there is done. It is able to leave and return.

In this way the mission team can use a methodical approach to selecting the most promising samples to return to Earth. As in Fig. 5, a great deal of data about the area are collected, at many sites in parallel, using imagery and standoff measurements. Scientists then use these data to identify targets for contact measurements. The robot returns to the selected sites and carries out the contact measurement tasks. Scientists use the results of the contact measurements to select the best candidates for sampling. Again, the robot returns to previous points, this time to collect samples. Once the samples are collected the scientists can determine the key samples to be returned to Earth. The rover can return the selected samples to the lander/ascent vehicle on the network with one command.

Many variations on these ideas are possible, and more of these options are discussed in Section 4. In the above example the benefit is that mission operators have the flexibility to efficiently delay sampling decisions pending a more thorough investigation of the data already on Earth. This can dramatically improve the efficiency of the system. In practice we found it to be as though there were multiple rovers with offset command cycles.

3. A MOCK LUNAR SAMPLE-RETURN MISSION USING NRP

The mock Lunar sample-return mission was conducted near the impact crater in Sudbury, Ontario, Canada. This was one of three missions [3] we conducted that were funded by the Canadian Space Agency. In this section, we give a brief overview of the mission, discuss the robot configurations that were used, and then present details of the mission time line, with an emphasis on the parallel exploration made possible by NRP.

3.1. Mission Overview

The field test was a robotic Lunar analogue mission in support of future sample-return missions to the Moon (and Mars), with a target of the South Pole Aitken basin. The primary objectives of the mission were to perform: (i) an in situ investigation of geology in a Lunar analogue environment, and (ii) an investigation of the formation processes and resource potential of impact crater(s).

The mission scenario lasted two weeks. For an overview of the mission operations, see Moores et al. [6]. Command cycles were nominally two hours in length and communication was only available during a window at the beginning and end of the cycle. Instructions were sent to the robot at the beginning of the cycle and telemetry was sent back at the end. This meant that there was very little time to review the results from the previous command cycle before the instructions for the next were sent.

In the first week, 24 command cycles were carried out, creating a network with 0.23 km of paths while driving a total of 1.0 km. The second week had 19 command cycles, a 0.44 km network and 2.92 km of total driving (3.9 km in total in the two weeks). Of the 17 samples that were collected and returned to the lander, ten were selected as the sample retention set (i.e., the samples that would have been returned to Earth for analysis). An overview at the end of the second week is shown in Fig. 8.

3.2. Rover Configurations

The mock mission used two different robot configurations, as shown in Fig. 7. The week-one configuration is on the left, and the week-two configuration is on the right. In both, the robot used a stereo camera as the primary sensor for adding to, and repeating paths on, the network of reusable paths. The vehicle, a robuROC6 made by Robosoft SA, had three passively articulated body segments.

Figure 7. The robot as it was in the week-one configuration (left) and the week-two configuration (right).

All the primary guidance, navigation & control (GN&C) sensors and software were onboard the vehicle, but there were scientific sensors and tools that were not integrated with the vehicle.

In week one, the configuration favoured onboard scientific capabilities over mobility. The vehicle had only simple GN&C capabilities. The stereo camera used for NRP was at the back of the vehicle, facing backwards. None of the other onboard sensors were integrated into the rover GN&C. The robot would turn and drive directly toward the current waypoint with only simple safety monitoring. If an obstacle was detected the robot would stop and attempt to reach the next waypoint in the sequence of closely-spaced waypoints. The CBRN Crime Scene Modeler (C2SM), made by MDA, was on the front of the robot. A ground penetrating radar (GPR) was pulled behind the robot (and removed when the vehicle was repeating a path in reverse). Panoramic imagery was obtained using a DSLR camera on a GigaPan pan-tilt unit. Also available, but not onboard the robot, were a handheld Raman spectrometer, an XRF spectrometer, and a drill used to obtain core samples. We had an Optech ILRIS-3_6D lidar that was located at the lander (see the first-week configuration in Fig. 8). A GPS antenna was located near the stereo camera; however, GPS was only used to measure the ground-truth localization, and it was not available during the scenario.

In week two, the configuration favoured mobility over sensor integration. The GPR and C2SM were not used, instead, the lidar was mounted on the front of the vehicle, and on top of that, a forward-facing stereo camera that was used for NRP. The onboard lidar allowed for many local scans to be taken. The rover GN&C was more sophisticated. The robot used the stereo camera, inclinometers in each segment of the vehicle, and wheel odometry to identify hazards. Onboard planning let the robot plan to avoid detected hazards. The planner was able to reuse the existing network, as well as plan paths into previously untraveled terrain [2]. In these tests we used a differential GPS for the ground-truth localization, and again, the resulting data were not available to the robot or mission controllers during the tests.

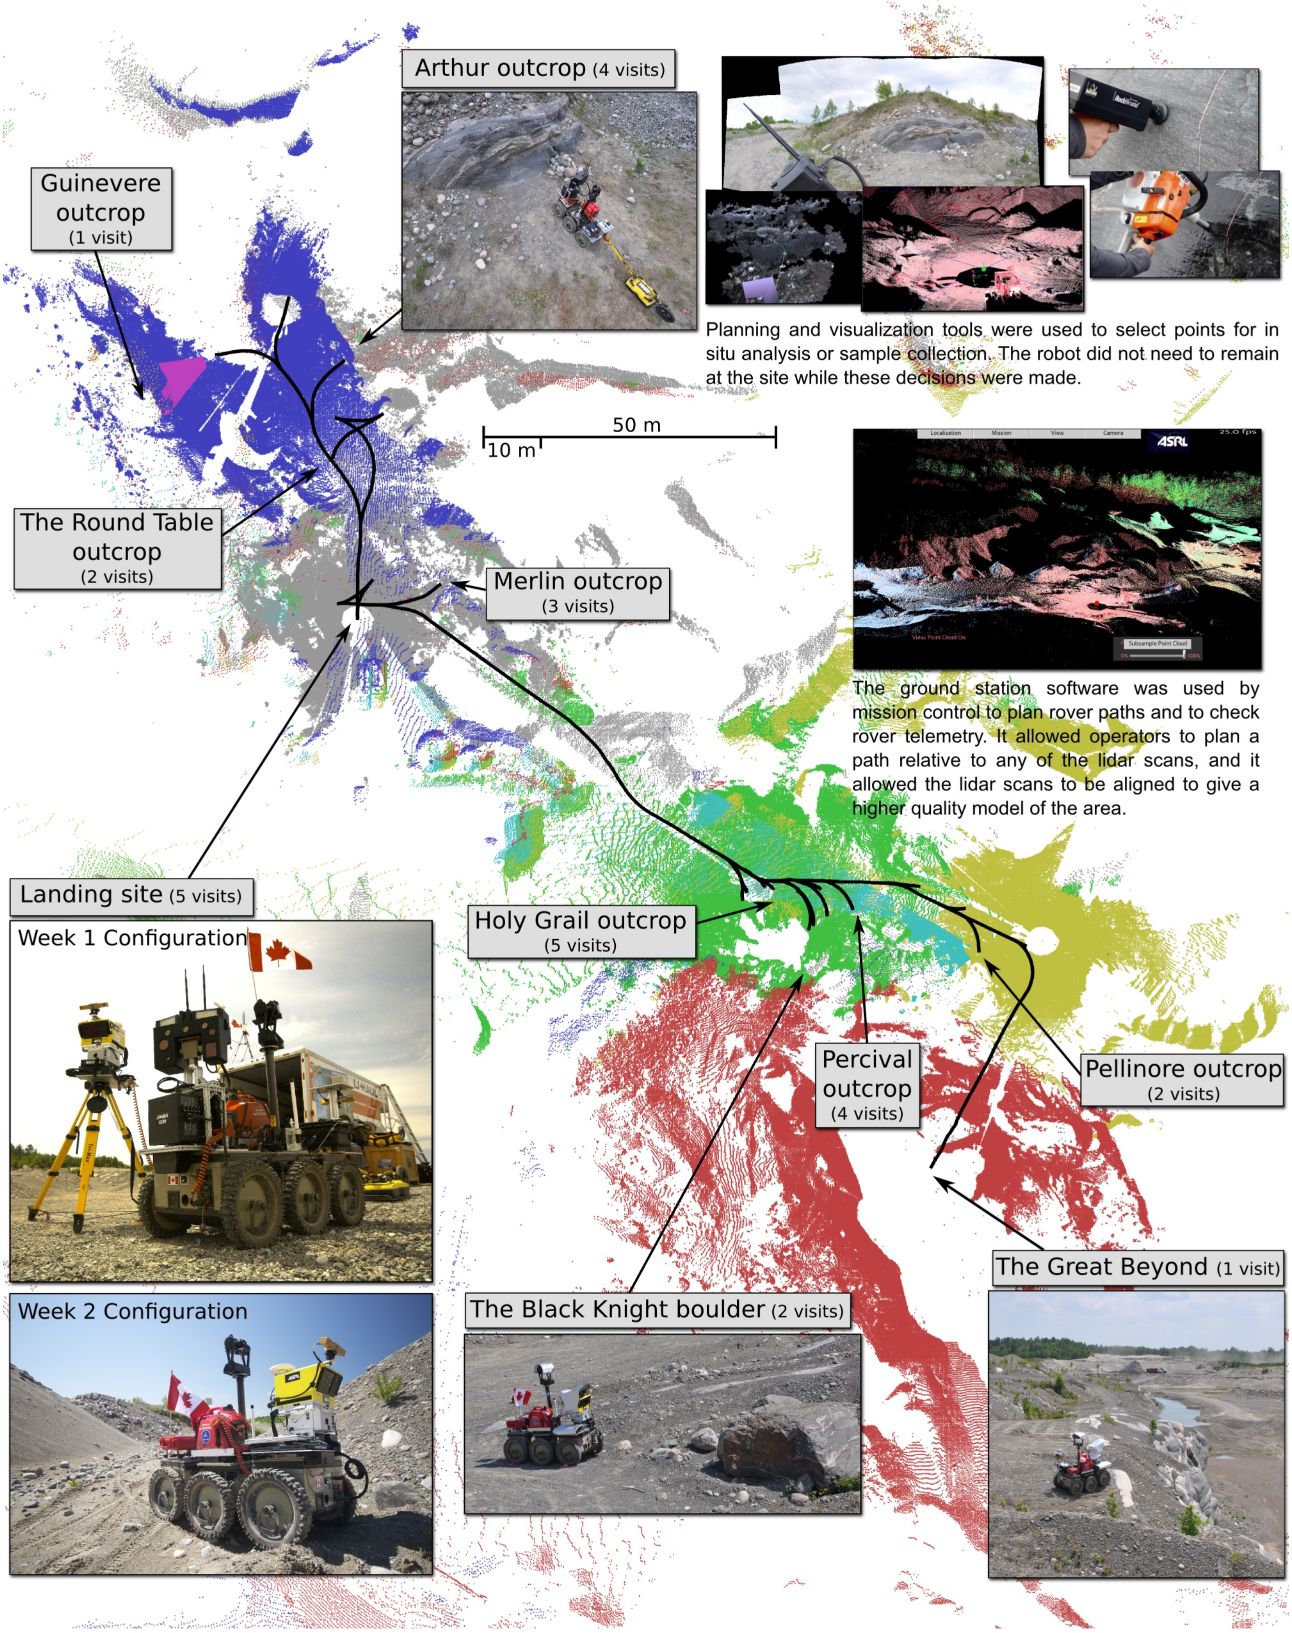

Figure 8. An overview of the network of reusable paths (the black line) at the end of the mission. NRP allowed the robot to return to any position that was previously visited. This meant that mission control could delay analysis or sampling decisions at one site, and still continue to carry out operations at other sites. At the end of week two, the total length of the network of reusable paths was 440 m, and by then the robot had driven a total of 3 920 m.

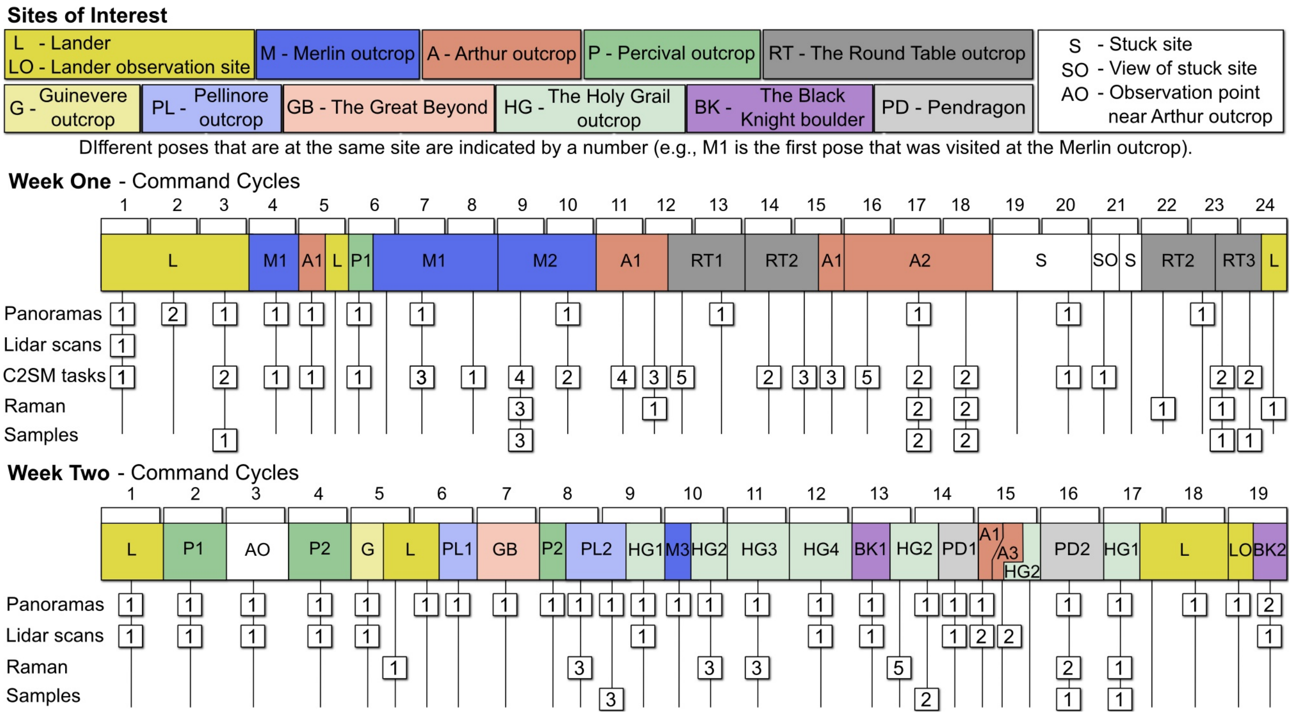

Figure 9. An overview of the command cycles carried out in the first and second weeks of the mission. The type and quantity of the tasks that were done at each site are shown in the squares below the command cycle. A total of 24 command cycles were carried out in the first week, and 19 command cycles were carried out in the second week.

In both weeks, mission control created traverse plans and reviewed the rover telemetry by using the ground station software (see Fig. 8). In the second week, when the lidar was onboard the rover, the ground station could also be used to manually tie together multiple lidar scans rather than relying on the dead-reckoning localization from the visual odometry onboard the rover.

3.3. Results from the Mock Mission

As this paper is about the use of NRP, rather than this specific mock mission, we omit further details about the scientific sensors, sampling methods, and the resulting findings. Instead we limit the presentation of results to those that pertinent to this discussion, namely, traverses to and between sites of interest, and broadly, the tasks that were carried out at those sites, as those tasks fit into the down-selection process.

An overview of all the command cycles in the mission is shown in Fig. 9. The main sites of interest are colour-coded at the top. Many of these sites had more than one distinct pose that was visited, these poses are distinguished by a unique number (i.e., M1 is a unique pose at the Merlin outcrop). In the overview we can see that in the first week, the rover spent the first three command cycles at the landing site. In the first command cycle, it took a panoramic image, a lidar scan, and observed the terrain using C2SM. The second cycle was used to collect two detailed panoramic images. The third cycle was to collect another panoramic image, take two measurements using C2SM, and collect a sample of the material at the landing site. In the fourth cycle, the robot approached the Merlin outcrop and collected panoramic imagery and standoff measurements. At this point we begin to see the pattern of parallel exploration. In the fifth and sixth cycle the robot visited other sites and collected imagery and standoff measurements, before returning to the Merlin outcrop. While the robot was there, mission control reviewed the data from Arthur and Percival and planned future tasks for when the robot returned to these sites.

In command cycle 19 of week one, the robot attempted to reach an observation point near the Arthur outcrop (later visited in cycle 3 of week two), but it became briefly stuck in soft soil. In cycle 20, the rover collected imagery to aid in determining what had happened. The next command cycle was to back up and observe the point where the rover became stuck, and then attempt a similar traverse, which again, the robot could not complete due to soft soil.

At the end of the first week we marked the rover wheel positions at the sites of interest and manually drove the rover to teach it a new network that reached those same physical locations. This was necessary as the significantly different camera location (now pointed forward rather than reverse) changed the appearance of the scene and the paths could not have been recognized.

Fig. 8 gives an overview of the mock mission at the end of week two. Multiple long-range lidar scans are displayed in different colours. The network of reusable paths is shown in black. Fig. 10 takes a more detailed look at command cycles 2 through 5 in the second week. This is an example of parallel exploration. In the even-numbered cycles, the rover was exploring in the bottom right of the map. In the odd-numbered cycles the rover was exploring in the upper right. Mission controllers reviewed the even-cycle data during the odd cycles, and had instructions ready for the robot at the beginning of the next command cycle. In this case, it was as if there were two robots exploring two different areas. With a serial approach, it would take three cycles to explore each individual branch (one where the rover did not move while mission controllers reviewed the telemetry and waited for the next communication window), and in order to begin exploration of the second branch, the rover would need another two or three command cycles to return to nearby the lander. Thus, in a serial approach, the same exploration might be expected to take 8 or 9 command cycles, rather than 4 (100% to 125% more cycles to explore the same two areas).

We have attempted to determine how many command cycles the entire mission would have needed had a serial exploration strategy been used. We began by estimating how many command cycles the robot would take at each site. For example, we estimated that the Merlin outcrop would have taken $8-9$ command cycles to investigate to the same degree. This range comes from assuming the rover sits idle for one cycle while mission control reviews data and waits for the next communication window. So, one cycle as in week one, cycle 4, then an additional cycle idle, then $4.5-5$ cycles as in week one, cycles $6-10$, then another cycle idle, then $0.5-1$ cycle as in week two, cycle 10. Following this example, and including getting stuck and returning to the lander, we get $63-75$ command cycles to do the same work as done in 43 (47% to 77% more cycles). Recall that without NRP it is unlikely that the robot can return to the lander in a single cycle; if we discount the cycles the robot spent at the lander we get 35 cycles for NRP, and $56-64$ cycles for the serial approach (60% to 83% more cycles).

Considering that many of the parallel techniques were being developed and refined during the mock mission, we expect that further operator experience will only increase the improvements made possible by NRP.

4. OTHER APPLICATIONS OF NRP

NRP can be used for other applications as well. For example, a second mission scenario was done using a robot as an astronaut assistant [3]. The robot was at times manually operated by an astronaut on site, and at other times it was operated in the same way as described earlier in this paper. This allowed the system to efficiently leverage the expertise of the astronaut in order to quickly reach sites of interest, and then let the astronaut leave while mission control remotely operated the rover. A work-site mapping scenario [19] has also been carried out. In this, NRP was used onboard the robot as part of the GN&C system.

NRP can also be used to extend the window for many types of opportunistic investigations. For example, if the robot were to drive past an interesting site, but the site was not identified as interesting until much later, the robot can at a later time, with a single command, return precisely to that previous position. In essence, it provides an insurance policy against leaving a site before all the useful science data has been acquired.

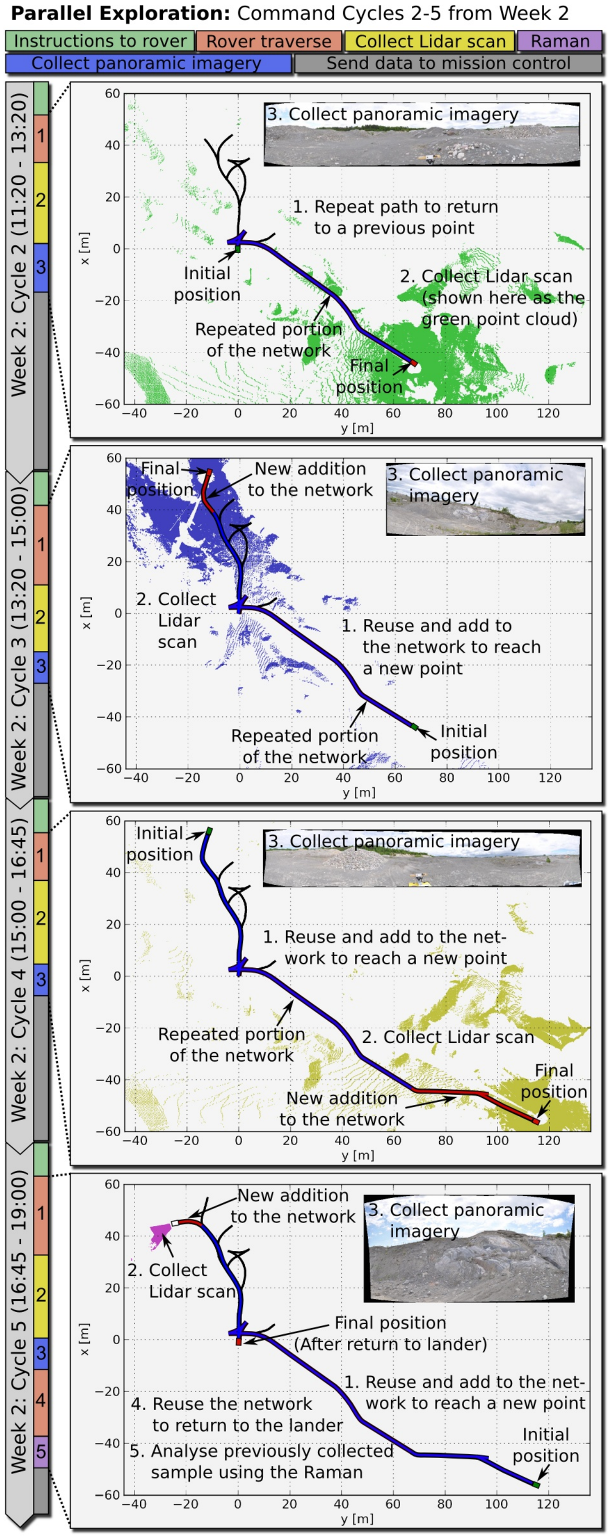

Figure 10. Parallel exploration was done between the odd and the even command cycles. In cycle 2 of week two, the robot collected a lidar scan and panoramic imagery. In cycle 3 the robot was sent to another area to collect more data while mission control reviewed the data from cycle 2. In cycle 4 the robot reused the network to return to the end point of cycle 2 and continue exploring while the data from cycle 3 were reviewed.

5. CHALLENGES AND FUTURE WORKS

One of the byproducts of using NRP is that the robot tends to travel a greater distance in a shorter time span. This raises two questions: (i) is there sufficient power available to do this additional driving, and (ii) is the additional wear due to more driving compatible with the mission lifetime? Both of these concerns are addressed in rover design. The question of power largely disappears when sources other than solar power are used (for example, an onboard reactor). In cases of operation in permanently shadowed regions, such as at the Lunar South Pole, the use of solar power is likely not viable anyway. Even in cases of limited power, it is likely that parallel exploration can be useful by scaling the problem to something that fits within the power budget. Outcrop characterization is one such scenario (see site A in Fig. 3). In this case the rover might move along the base of an outcrop and collect imagery at many points. It would then return to specific points and continue the down-selection process. In this scenario the rover may only move a few metres between sites of interest, but these different sites would give different vantage points and allow the rover to reach a larger area in less time.

The additional wear could be factored into the design of the rover, or the parallelism could, as above, be scaled to fit within the desired reliability requirements. This scaling is a trade off available to mission controllers.

It should also be noted that NRP offers benefits beyond parallel exploration, such as more accurate goal acquisition and more robust navigation [2]. These benefits remain even when parallel exploration is not carried out.

We also need to consider the question of what happens when the vehicle is not able to repeat a path. In all of our testing, this has happened only rarely, and by using a lighting-invariant sensor such as a high-framerate lidar, we can eliminate the inability to localize due to lighting changes or lack of illumination [8, 10]. However, there will still be cases where the appearance of the scene changes, or the traversability of a previous path changes and this leads to the desire to be able to repair paths; this is an area of ongoing future work.

6. CONCLUSIONS

A network of reusable paths offers a new approach to planetary surface exploration using a mobile robot. NRP has many benefits in the context of robust autonomous navigation [2]. It also allows mission-level improvements by allowing parallel exploration of multiple scientific targets, and it inherently includes sample return. During the analogue mission, this capability enabled nearly twice as many sites to be visited within the mission time frame. Such a capability would be extremely useful for sample-return missions to the Moon or Mars.

ACKNOWLEDGMENTS

We extend our thanks to Ethier Sand and Gravel for providing support and an excellent test site. We also wish to thank NSERC, the Canada Foundation for Innovation, DRDC Suffield, the Canadian Space Agency, and MDA Space Missions for providing us with the financial and in-kind support necessary to conduct this research.

REFERENCES

[1] C. Leger et al. Mars Exploration Rover Surface Operations: Driving Spirit at Gusev Crater. In *In Proceedings of the 2005 IEEE Conference on Systems, Man, and Cybernetics*, 2005.

[2] B. Stenning and T. D. Barfoot. Path Planning on a Network of Paths. In *Proceedings of the IEEE Aerospace Conference*, Big Sky, MT, 5-12 March 2011.

[3] C.L. Marion et al. A Series of Robotic and Human Analogue Missions in Support of Lunar Sample Return. In *Proc. of the 43rd Lunar and Planetary Science Conference (LPSC)*, #2333, 2012.

[4] B. Stenning, G.R. Osinski, T.D. Barfoot et al. Planetary Surface Exploration using a Network of Reusable Paths. In *Proc. of the 43rd LPSC*, #2360, 2012.

[5] R. Francis, G.R. Osinski, J. Moores, T.D. Barfoot, and ILSR Team. Co-operative Human-Robotic Exploration of Lunar Analogue Sites. In *Proc. of the 43rd LPSC*, #1996, 2012.

[6] J. Moores, R. Francis, M. Mader, G.R. Osinski, T.D. Barfoot et al. A Mission Control Architecture for Lunar Sample Return as Field Tested in an Analogue Deployment to the Sudbury Impact Structure. *Advances in Space Research,* special issue on "Lunar Exploration", 2012. Accepted on May 9, 2012.

[7] P.T. Furgale, T.D. Barfoot. Visual Teach and Repeat for Long-Range Rover Autonomy. *Journal of Field Robotics*, 27(5):534560, 2010.

[8] C. McManus, P.T. Furgale, B. Stenning, T.D. Barfoot. Visual Teach and Repeat using Appearance-Based Lidar. In *Proc. of the International Conference on Robotics and Automation*, 2012.

[9] D. Nistr, O. Naroditsky, J. Bergen. Visual Odometry for Ground Vehicle Applications. *J. of Field Robotics*, 23(1):3–20, 2006.

[10] T.D. Barfoot, B. Stenning, P.T. Furgale, and C. McManus. Exploiting Reusable Paths in Mobile Robotics: Benefits and Challenges for Long-Term Autonomy. In *9th Canadian Conference on Computer and Robot Vision (CRV)*, pages 388–395, 2012.

[11] J. Marshall, T.D. Barfoot, and J. Larsson. Autonomous Underground Tramming for Center-Articulated Vehicles. *Journal of Field Robotics*, 25(6-7):400–421, June-July 2008.

[12] A.M. Zhang and L. Kleeman. Robust Appearance Based Visual Route Following for Navigation in Large-Scale Outdoor Environments. *The International Journal of Robotics Research*, 28-3:331–356, 2009.

[13] T. Ohno, A. Ohya, and S. Yuta. Autonomous Navigation for Mobile Robots Referring Pre-recorded Image Sequence. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, 1996.

[14] S. Jones, C. Andresen, and J. Crowley. Appearance Based Processes for Visual Navigation. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, 1997.

[15] Y. Matsumoto, K. Sakai, M. Inaba, and H. Inoue. View-Based Approach to Robot Navigation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 1702–1708, Takamatsu, Japan, 2000.

[16] L. Tang and S. Yuta. Vision Based Navigation for Mobile Robots in Indoor Environment by Teaching and Playing-Back Scheme. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, pages 3072–3077, 2001.

[17] E. Royer, M. Lhuillier, M. Dhome, and J. Lavest. Monocular Vision for Mobile Robot Localization and Autonomous Navigation. *International Journal of Computer Vision*, 74(3), 2007.

[18] G. Sibley, C. Mei, I. Reid, and P. Newman. Vast-Scale Outdoor Navigation using Adaptive Relative Bundle Adjustment. *The International Journal of Robotics Research*, 29-8:958–980, 2010.

[19] R.S. Merali, C.H. Tong, J. Gammell, J. Bakambu, E. Dupuis, and T.D. Barfoot. 3D Surface Mapping using a Semi-Autonomous Rover: A Planetary Analog Field Experiment. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, Turin, Italy, September 2012.