

Autonomous Rover Traverse and Precise Arm Placement on Remotely Designated Targets

Michael Fleder, Issa A. Nesnas, Mihail Pivtoraiko, Alonzo Kelly, Richard Volpe

Abstract— Exploring planetary surfaces typically involves traversing challenging and unknown terrain and acquiring in-situ measurements at designated locations using arm-mounted instruments. We present field results for a new implementation of an autonomous capability that enables a rover to traverse and precisely place an arm-mounted instrument on remote targets. Using point-and-click mouse commands, a scientist designates targets in the initial imagery acquired from the rover’s mast cameras. The rover then autonomously traverses the rocky terrain for a distance of 10 – 15 m, tracks the target(s) of interest during the traverse, positions itself for approaching the target, and then precisely places an arm-mounted instrument within 2-3 cm from the originally designated target. The rover proceeds to acquire science measurements with the instrument. This work advances what has been previously developed and integrated on the Mars Exploration Rovers by using algorithms that are capable of traversing more rock-dense terrains, enabling tight thread-the-needle maneuvers. We integrated these algorithms on the newly refurbished Athena Mars research rover and fielded them in the JPL Mars Yard. We conducted 43 runs with targets at distances ranging from 5 m to 15 m and achieved a success rate of 93% for placement of the instrument within 2-3 cm.

I. INTRODUCTION

INTEREST in planetary rovers conducting an autonomous traverse followed by precise placement of arm-mounted instruments dates back to the first Martian rover: the Sojourner rover, which landed on Mars in 1997. This capability enables scientists to collect measurements from targets that they can designate remotely. These targets would typically fall within 10 – 20 m from the rover, primarily limited by the resolution of the imagery. The scientist would then triage these targets and revisit sites of high potential science return.

Manuscript received February 8, 2011. This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work is supported by the NASA Mars Technology Program.

M. Fleder conducted this work while at the Jet Propulsion Laboratory, Pasadena, CA (email: mfleder@mit.edu)

I. Nesnas is with the Jet Propulsion Laboratory, Pasadena, CA. (email: Issa.A.Nesnas@jpl.nasa.gov)

M. Pivtoraiko is with Carnegie Mellon University, Pittsburgh, Pennsylvania. (email: mihail@cs.cmu.edu)

A. Kelly is with Carnegie Mellon University, Pittsburgh, Pennsylvania. (email: alonzo@cmu.edu)

R. Volpe is with the Jet Propulsion Laboratory, Pasadena, CA. (email: volpe@jpl.nasa.gov)



Figure 1: The Athena Mars Research Rover in the JPL Mars Yard

Operational scenarios used on the Mars Exploration Rovers (MER) require a total of three to four sols (Martian days) for each target measurement. This autonomous capability would reduce this operational time to a single sol, thus increasing the overall science return for the mission [1]. When visiting multiple targets, the reduction in the number of sols would reach an order of magnitude.

To provide this capability, we developed and adapted a number of sensing and control algorithms and integrated them on the Athena research rover. For a robust implementation, we had to address a number of challenges including terrain variability, sensing limitations, lighting variations, and traverse challenges. Because the rover is capable of traversing *over* small rocks, analyzing terrain traversability and handling the highly-variable wheel/rock/soil traction requires careful consideration. Figure 1 shows a typical terrain that we used in our testing, where the rover would overcome rocks smaller than a wheel diameter. For MER, the traversable obstacle had to be less than 20 cm in height, the equivalent of 80% wheel diameter. Only the final placement patch was chosen to be relatively free of obstacles to minimize rover slippage as the center-of-mass shifts during instrument placement. By using a rover prototype in the JPL Mars Yard, we tried to mimic the subtle conditions that would arise in an environment similar to that encountered on Mars.

A. Related Work

Motion planning and control of mobile manipulator systems have received significant attention over the past several

decades. Even though our work shares some of the motivations with the general topic of mobile manipulation, the need for determinism coupled with the limited available sensing and computational resources on-board planetary rovers preclude the use of many sampling-based approaches that are commonly used in mobile manipulation. A variety of field robotics applications, such as planetary exploration considered here, admit representations of low enough dimensionality that deterministic approaches can be applied directly. The present work can be viewed as a result in leveraging this property to design an efficient fielded system.

The autonomous capabilities of planetary rovers have continued to increase with each rover deployment on the Red planet. Back in 1997, the Sojourner rover achieved the first autonomous rover traverse on another planet. However, this autonomous capability was limited. The hazard avoidance system used laser stripes with a camera system to detect rocks and determine contour lines [2]. By repeating this process at small, three-inch increments, the rover was able to build sparse terrain maps and avoid obstacles. Using the above rock detection algorithm, engineers were able to command the rover to position itself in front of designated rocks. These capabilities were exercised over distances of only a few meters.

In 2004, the Jet Propulsion Laboratory landed two more capable rovers on the opposite side of Mars. Both Spirit and Opportunity enjoyed a greater level of sensing and compute capabilities compared to their Sojourner predecessor. Each rover had a suite of stereoscopic cameras: front and rear camera pairs with wide field-of-view (FOV) lenses for hazard avoidance (“hazcams”). Each rover also carried an articulated mast head with two stereo camera pairs with both wide and narrow FOV camera pairs (“navcams” and “pancams”). These rovers were designed to traverse longer distances than their predecessor. To date, the Spirit and Opportunity rovers have logged a combined 30.5 km on the Martian surface [3]. A sixth of this traverse distance was accomplished with some level of autonomy for hazard detection and avoidance. The rovers would either use active obstacle detection and avoidance or employ hazard detection only to affirm the safe traversal of a predefined rover path. Unlike the Sojourner rover that used laser stripes for generating terrain maps, the Mars Exploration Rovers generated three-dimensional maps using dense stereo at quarter resolution from their hazcams. Then, they used goodness maps to assess terrain traversability. Using their autonomous navigation capability, the MER rovers demonstrated, one time, a 6m autonomous traverse and precise placement [1]. While this marks another major milestone in autonomous capabilities for planetary rovers, the execution of this capability was done in a relatively benign

environment and without any obstacles in the path of the rover¹.

In addition to these developments on flight missions, active research in the autonomous traverse and instrument placement for planetary rovers was on-going at several institutions over the past decade. Early work focused on instrument placement for single and multiple rock targets from a distance of 3-5 m [4]. This work was demonstrated on the Rocky 7 research rover [5] on fairly benign terrain and had a final instrument-placement precision on the order of 5-10 cm. Planning and execution for such tasks has been investigated at LAAS-CNRS [6]. Work by Pedersen *et al.* [7] demonstrated multiple-target single cycle instrument placement in terrains with only a few obstacles.

B. State-of-the-Art

To acquire measurements with the Spirit and Opportunity rovers at designated targets, MER scientists and operators spend a significant amount of time carefully planning and preparing a sequence of rover steps to (i) traverse and position the rover relative to the target (ii) verify a collision-free path for the arm and (iii) deploy and orient the instrument on the target to acquire measurements. When a rover is within 10 – 20 m from the designated target, it typically spends one or two sols navigating to a nearby location and positioning itself for the final approach to the target. Then it approaches the target such that it is within the arm’s workspace with a high manipulability index [8]. After completing the final approach, the third sol will deploy the arm and acquire a measurement. Each sol requires significant human oversight and control. Were the rovers able to navigate to targets and take measurements autonomously—human input only for target(s) selection—the speedup for taking certain kinds of measurements would increase by at least three fold for a single measurement and an order of magnitude for multiple targets in a single sol.

C. Challenges Addressed

This work builds upon and extends prior work done by members of this team, other researchers at JPL[1][4][9], and researchers at Ames Research Center [7]. Our work focuses on advances to motion planning and terrain analysis for addressing the challenges of environments with denser rock distributions. For instance, in very rocky environments the rover often needs to execute tight maneuvers with small clearances between rocks such as in “threading the needle” between multiple obstacles. Such a capability is not available on the Mars Exploration Rovers: they always maintain a safe clearance for an in-place turn since their path planners require more obstacle clearance than [10]. Furthermore, we address situations that require the rover to traverse *over* small to

¹ Hardware limitations on the rovers have limited further utilization of this capability as of this writing.

medium-sized rocks (up to a wheel diameter in height) where such traversals result in rover undulations and tilts of the mast of about 15°-30° from the horizontal. Maintaining 2-3 cm precision on the final instrument placement – after autonomously traversing terrain with highly-variable wheel/rock traction – requires constant tracking of the target during the traverse. Such precision also becomes difficult as the target’s size and appearance changes drastically as the rover closes in on the target. We address all these challenges within the computational constraints of radiation-hardened processors and image acquisition systems of planetary rovers. To the best of our knowledge, prior work did not address such challenging rover traverses followed by precise instrument placement.

To address the traverse challenge, we adapted and integrated a new continuous-curvature path-planning algorithm [10] with a modified version of our traversability analysis algorithm. These algorithms control all twelve rover drive and steering actuators simultaneously. The rover drives along clothoid trajectories as opposed to the traditional straight line and fixed-arc traverses of prior work. The continuous curvature traverses are the first step towards continuous driving for planetary rovers without the need to stop the rover at every navigation step. Currently, planetary rovers—including the Mars Science Laboratory —use the think-steer-drive-stop cycle; they cannot drive, steer and think concurrently due to (i) the limited number of available controllers and (ii) limited available power. However, future rovers, similar to today’s terrestrial systems, are currently planning to overcome this limitation to allow continuous driving for faster traverses.

D. Operational Constraints

Depending on the instrument, the time savings and increase in science return may be limited. Certain instruments take a significant amount of time to acquire a measurement. For example, spectroscopic measurements can take hours of integration time; so for this instrument, the “traverse and precise placement” capability would be used on a single target at a time. However, for instruments such as microscopic imagers, the time savings could reach an order of magnitude when investigating multiple targets in a single sol. A triage on several designated targets can arm scientists with information that would later help in deciding on which targets to deploy resource consuming instruments. This can become particularly important for future missions, such as the proposed Mars Sample Return, where the rover would likely have consumable resources in its sample collection and handling. With an ability to efficiently carry out a close-up assessment of multiple targets within a single site, scientists could improve the quality of samples collected for the mission.

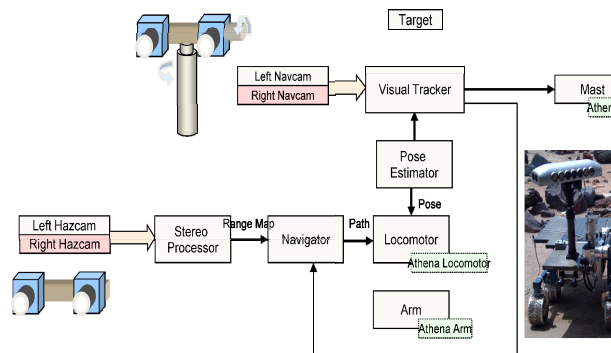


Figure 2: Functional diagram for rover navigation during “go and touch.”

II. SYSTEM OVERVIEW

The “traverse and precise placement” rover capability is a highly autonomous operation where the user interacts with the system only to designate the target(s). The rest is accomplished without user intervention.

Figure 2 shows a system-level functional diagram of the key subsystems. Our current system uses two sets of stereo cameras: the mast mounted 45° FOV navcams to track the target, and the 110° FOV hazcams to navigate and avoid obstacles. The rover visually tracks the designated target from its navcams while autonomously avoiding obstacles using the hazcams. It also uses the hazcams to plan its traverse to the goal. When the rover reaches the vicinity of the goal to within 2 meters, it re-orient itself toward the goal and then approaches the target along the vertical plane of the target’s surface normal. Once the target is within the arm’s reach, the rover deploys the instrument to acquire the measurements. The key subcomponents are:

Stereo processor – generates three-dimensional point clouds from stereoscopic images at different image pyramid levels.



Figure 3: Scientist selects a target (white circle centered on target pixel) from a rover image

For both the hazcams and navcams, we used pyramid level 2 (256×192) on the original (1024×768) images. Since the navcams have narrower field-of-view, we acquired multiple wedges for terrain analysis.

Traversability Analyzer – bins the above three-dimensional point cloud into a grid and then computes the traversability map using a least-squares slope fit evaluation of a rover-sized patch on that grid. The slope, roughness and height measures are then combined with the rover’s mobility model to generate a goodness value for each cell of the map[9].

Motion Planner – generates continuous curvature paths and evaluates them across the goodness map.

Locomotor – uses the best-candidate path generated from the motion planner to drive the rover. The locomotor coordinates both drive and steering wheel trajectories to control the six-wheeled rocker-bogie rover mechanism.

Pose estimator – provides an estimate of the rover pose with associated uncertainty. The pose estimator fuses multiple sensing modalities using an Extended Kalman Filter [11]. These include wheel odometry, inertial sensing, and - depending on the terrain slip conditions - visual odometry.

Visual Tracker – uses a monocular camera for the image-based feature tracking of the designated target; the visual tracker also uses the corresponding three-dimensional target information from stereo imaging for pointing the mast.

Mast/arm – provides control of the articulated pan/tilt mast and the five degree-of-freedom arm.

All these sub-systems were designed using the CLARAty[12] reusable software infrastructure to enable the use of different implementations for its subsystems. It also enables the deployment of this capability on rovers with different hardware configurations and sensing capabilities.

III. TARGET SELECTION

We applied the following constraints for selecting targets: (1) the target must be visible in both stereo navcam images (so we can compute its 3-D location) and (2) the user should avoid targets near depth discontinuities (edges of rocks). Once selected, the rover re-points its mast to center the target in the next image. Once the target is centered in the image, visual tracking locates the target pixel and updates its template with the newest target image. At this point, the rover has not yet started to move. Centering the target in the image produces a target template with reduced lens distortion, which will be more robust to large changes in appearance and view angles for the tracker as the rover weaves between obstacles. Figure 3 shows a typical terrain and the initial




Figure 4: Edge effects in binning algorithms. All cells are flat except for the area marked “ROCK,” which is not traversable. A 3x3 rover-sized patch is highlighted (not to scale). Cell 7 is being evaluated for traversability. The rock straddles four connected cells. The rock contributes a small number of stereo points to the rover-sized patch; consequently, the rock only lightly affects the plane fit. Rover placement in the 3x3 patch could result in a collision.

target selection. The camera’s resolution limits the accuracy of target specification (determined by a single pixel). We set a goal of 2-3 cm placement accuracy.

IV. TRAVERSABILITY ANALYSIS

The terrain analyzer, which is based on the CLARAty [12] implementation of the Morphin algorithm [9] processes the point cloud data in two steps: (1) the terrain analyzer bins the points into a grid map. We typically use 100×100 cells in the map with each cell representation a 20 cm × 20 cm area. After binning the points, the analyzer evaluates the traversability of rover-sized patches centered at each cell. For each patch, the analyzer fits a plane to the raw point-cloud data that falls within that patch. Each patch-sized plane fit produces three results: the two-angle tilt of that plane, the height of the plane, and the roughness of the fit. (2) In the second step, we compute goodness and certainty scores for each cell. The goodness and certainty values derive from the plane-fitting statistics computed in the first step. For a particular cell, we consider the rover-sized patch centered at that cell and then use the associated plane fit to determine goodness and certainty for that cell. Thus, the goodness of an adjacent cell derives from translating the rover-patch one cell (to re-center) and re-computing the plane-fit. Prior work used a circular patch with a radius corresponding to the largest rover dimension. We used a rover-sized rectangle with a few centimeters buffer in each direction.

Our experience, as well as that of others, has repeatedly shown that tuning navigation parameters can be quite challenging. Several thresholds can be very sensitive and dramatically affect the success of the navigation, especially in challenging terrains. As we investigated more challenging rover maneuvers in dense obstacle fields, we identified the following as potential sources of errors that can result in navigation failure. These considerations are likely to apply to other field systems that utilize terrain analysis via similar

variants of occupancy grids. Below is a brief description of each:

- **Edge effects:** there are certain challenges that arise in binning algorithms that are related to the boundary effects of binning points and statistically analyzing bins. One such effect occurs when a small part of an untraversable obstacle spills into the rover patch being evaluated (see Figure 4). The patch in Figure 4 would be dangerous for the rover, but determining that from plane-fit statistics can be difficult. In practice, the uncertainties in the pose estimate alleviate this problem as maps are merged: the grid slides around slightly as pose-estimation varies.
- **Filtering artifacts:** a plane fit approach on rover-sized patches reduces the effect of some obstacles: e.g. tall and narrow “tent pole” obstacles. The plane-fit height will get averaged out over the patch size and the residual histogram will be bimodal with lots of low-magnitude residuals for the flat part and some high ones representing the tall obstacle. Because this histogram will look very close to an all flat patch, the thresholds must be carefully set to separate out obstacles. Here we carefully tuned that threshold.
- **Merging of goodness maps:** because of errors in pose estimation, the most recent data is generally weighed more heavily when merged with older data. However, it is important to factor in the number of stereo points that are recorded in the cell at every instance. This number of points is captured in the certainty (of goodness) measure used in our algorithm.

To address these sensitivities, we added hysteresis to our goodness map calculation. More specifically, we bound the rate at which the goodness measure of a cell can increase to 5% but did not bound the rate at which it can decrease. This causes the rover to be cautious about terrain that previously had a low traversability in recent steps but suddenly appears more benign.

V. MOTION PLANNING

While dual local and global planners have been fielded extensively in robot navigation, we opted for a single, multi-resolution motion planner. Dual planner methods are currently in use by the Spirit and Opportunity rovers on Mars. Even though such methods are computationally fast and perform very well in benign to medium-difficulty terrains, they have been noted to struggle in complex natural environments with very rocky terrain. The reasons for these difficulties stem from the dual-nature of the planning component. Involving two separate planners to accomplish the global and local planning task requires explicit methods of integrating the planners and getting them to agree on compatible notions of costs of motions through the




Figure 5: Goodness map and linear-approximation to rover path (splines not shown). The path terminates at the target.

environment. Moreover, there are often representational differences between the two planners, as it is typical for the local planner to satisfy the model of robot motion, e.g., differential constraints, and for the global planner to disregard them entirely for the sake of efficiency. These differences may result in conflicted behavior when the constraints of robot motion are most pronounced: e.g., aggressive maneuvering in difficult terrain. Furthermore, local planners in this setting typically draw from a small set of arc motions, for the sake of efficiency. This simplified design operates in a reactive manner by picking a motion that is the best fit. Unfortunately, a small, limited set of possible motions is often a poor representation of the overall vehicle mobility. Furthermore, many motion planners, including those used by the MER rovers, utilize configuration space expansions that result in excessive obstacle clearance that is related to the size of the vehicle footprint as it were to perform a turn in place. While this is reasonable for robots of nearly circular shape, this can be restricting for robots with elongated footprints, especially as the terrain gets denser with obstacles. In order to address the above difficulties, we explored alternative methods of planning. We investigated the use of a state lattice motion planner[10]. The planner generates motion trajectories with continuous curvature that maneuver through dense rock distributions. This requires the rover to drive and steer simultaneously. While this capability is not available in today’s flight rovers due to power limitations, it is available on the research rover prototypes.

The planner is configured to plan aggressive paths through regions that would be too difficult for the current MER planner. Generated rover motions are represented as cubic splines. It is important for the motion planner to be able to choose among most, if not all, *feasible* paths (the paths the rover is capable of executing). Paths featuring continuous




Figure 6: Tracking the target. Red dot: tracked target. Green dot: predicted target location (pre-tracking). Green frame: search window. Red frame: reference

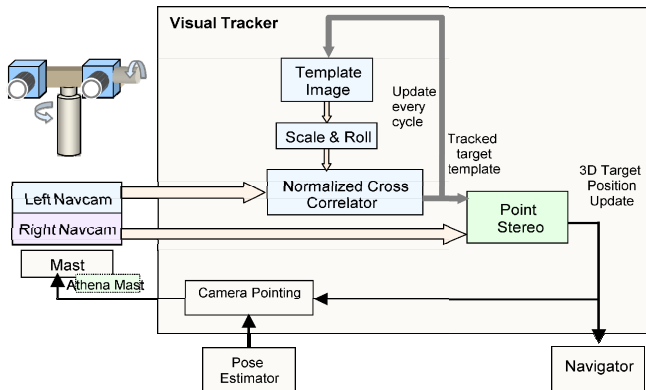


Figure 7: Functional Diagram for the visual target tracker (VTT)

steering have a greater expressive capacity and can approximate a larger collection of the feasible paths than paths that are constrained to discrete steering, e.g., consisting of constant-curvature arcs. A rendering of a linear-approximation to a path generated by the planner is shown in Figure 5. Because the lattice planner supports paths across different fidelity of representation, it is scalable over long distances and can simultaneously reason at local and global scales. Thereby, we obtain the computational benefits of the dual local-global planners, while avoiding the challenges of their different representations and separate processes [10].

A. Search space and algorithm

Our motion planner consists of two primary components: the search space that represents the feasible motions of the rover and the search algorithm that evaluates these motions and selects the one for the rover to follow. The goal of the planner is to select the motion that is optimal with respect to the relevant measures of motion quality. In the presented implementation, the search algorithm was D* Lite[13] that

was modified to allow graduated fidelity planning, as described further in Section V-B. The algorithm sought to minimize path length. The search space was a *state lattice*, a directed graph that consisted of (i) vertices, pre-defined regular state samples, and (ii) edges, pre-computed motions connecting the above state samples. We sampled a three-dimensional state space, consisting of 2D position and heading. The motions were steering functions computed by the boundary value problem solver in [14]. The motions were represented as cubic polynomial curvature functions of path length.

B. Graduated Fidelity

Because the terrain information is processed from the rover’s on-board sensors, we use a high-fidelity representation in the immediate vicinity of the rover (within its sensor range), and a lower fidelity representation in the areas that are either less known or less relevant for the planning problem. Lower fidelity of representation is designed to increase search speed, while higher fidelity provides better quality solutions. Since, traditionally, grids have been utilized in D*-like replanning, the notion of varying the quality of problem representation has been identified with varying the *resolution of the grid*. However, the proposed motion planner relies on the discretization of both the state and motions. The term “resolution” is typically used in 2D scenarios; so, to avoid confusing vocabulary use, we refer to managing the fidelity of state lattice representation as *graduated fidelity*.

In designing the connectivity of regions of different fidelities, care must be taken to ensure that all fidelities consist of motions that are feasible with respect to the robot’s mobility model. Robot navigation quality degrades if motion constraints are violated. Requiring that all levels of fidelity include feasible motions avoids such difficulties. For example, suppose the high fidelity regions consist of feasible motions; then the above requirement is satisfied when the connectivity of low fidelity regions is a strict subset of that of the high fidelity regions.

VI. VISUAL TARGET TRACKING

Tracking a designated target across a 10 – 20 m traverse can be challenging for a number of reasons: (1) the view of the target changes significantly as the rover approaches the target (Figure 6), (2) the vantage point of the tracking cameras can change significantly due to the motion of the rover as it tries to avoid local obstacles, (3) the lighting conditions can change on a slow moving rover, and (4) the rover can experience significant tilt as it traverses rocky terrain.

We have integrated the visual target tracker (VTT) designed by Kim, Nesnas *et al.* [1]. This algorithm uses a normalized cross-correlation (NCC) matcher with a scale and roll-corrected template. NCC was adopted because it performed quite well using limited computation compared to potentially more capable but computationally more expensive



Figure 8: Expanded view of tracking

algorithms such as SIFT (Scale Invariant Feature Transform). NCC smoothly handled the (expected) large target-pixel displacements between frames. We did not use iterative search based on image gradients since such methods typically require small target displacements.

We run the tracker at every rover step, which is typically every 0.75 m. To compensate for rover movement, the tracker rolls and scales the template image based on its odometry and inertial pose estimates. It then uses NCC within a search window. Figure 7 shows a functional diagram for the tracker described below:

1. Setup: after target selection (and after the cameras are centered on the target), we initialize a fixed-sized image template— 21×21 pixels—centered on the target pixel. This template represents the target.
2. For each tracking iteration, we predict the new 3-D position of the target based on (i) last known target location and (ii) estimated change in rover pose.
3. We point the mast cameras at the predicted 3-D target location (after rover movement). Based on the estimated position change and the left camera model, we compute (i) a magnification / shrinking factor and (ii) a roll angle for the target template. The target template is then scaled and rolled accordingly.
4. We then run the NCC search within a search window centered on the predicted 2-D image coordinates of the target.
5. The algorithm replaces the old target template image with a new one centered on the found target pixel.
6. The algorithm uses stereo vision to compute the 3-D coordinates of the target.

We added a tracker recovery mode. In the event of tracking failure, the rover (i) retracts part of the last step before loss of target (ii) reacquires the image and (iii) re-attempts tracking. The tracker can lose the target for up to two consecutive

steps. A third tracking failure results in the rover stopping, declaring a fault, and calling home for help from the operator to confirm the target location.

VII. EXPERIMENTAL RESULTS

We adapted and integrated all these algorithms on the Athena rover and tested the system in the JPL Mars Yard. The Athena rover measures 1 m length by 0.8 m in width and stands 1 m off the ground. We conducted 43 runs of the “traverse and precisely place instrument on target” under different lighting conditions (early morning, noon and afternoon) and terrain topographies (different rock distributions). A typical, challenging terrain is shown in Figures 1 and 8, which had rocks and boulders varying in size from 7 cm to 0.5 m in radius. Twenty eight tests were conducted with the target selection being 3-5m away. Fifteen of the tests put the target 10 m away. The 15 long-range tests all involved at least one major obstacle in the straight-line path to the target (see Figure 5). The rover had a 93.33% success rate. Failures were all attributed to tracking failure. Failures were graceful however. For example, in a typical failed test run, the rover failed to track the target (despite the recovery mode). The tracker declared a fault and called “home” for help from the operator to confirm the target location. Once we reselected the target, the rover continued successfully. The terrain conditions were such that the optimal path required the rover to incur the maximum change in azimuth (and therefore stress the tracker significantly). We show a typical iteration of target tracking in Figure 6.

In several test cases, it was not apparent from the initial images whether or not a feasible path to the goal existed. In those cases, the rover determined a feasible path’s existence after several move-sense iterations. We focused our test cases on targets that can be approached from the visible side of the rock, because more complex scenarios that would require the rover to approach the target from behind the rock are very unlikely to be considered in an actual mission due to its higher risk.

The locations of (i) the predicted target pixel and (ii) the tracked target pixel differed on average by 22.48 pixels with a standard deviation of 30.4 pixels. This average target-pixel-prediction error can be used as a rough measurement of the accuracy of several of subsystems. The predicted target pixel is based on (i) last rover pose (ii) predicted rover pose after moving (iii) accurate pointing of the cameras (to point at the predicted 3-D location of the target before acquiring an image for tracking). Since 96% of the tracked-targets pixels were within 174 pixels of the predicted-target pixels, we safely used a tracking search window of 200×200 pixels.

The placement of rocks that were obstacles (as opposed to traversable rocks) was challenging enough to require the rover to come within 8–12 cm of significant boulders. The rover did so without collision as a result of the aggressive

paths of the planner. Using this path planner, we demonstrated that (a) we can move precisely through tough terrain and (b) we can compute precise paths quickly; replanning time took less than 1 second on all runs. The rover's movement through rocky terrain resulted in zero collisions over the course of the test runs. The combination of the terrain analyzer and path planner has resulted in aggressive but safe traverses. The terrain analyzer is precise enough that we plan paths (i) without a configuration space expansion and (ii) with only a minor buffer space around obstacles. The final placement accuracy was measured to be within 3 – 5 cm of the initially selected target.

VIII. CONCLUSION

In this paper, we have presented a fully autonomous "traverse then precisely place an instrument on target" capability for planetary rovers. We have successfully demonstrated this capability on the Athena research rover in the rocky outdoor terrain of the JPL Mars Yard. To enable the rover to traverse rock-dense terrains, we modified the navigation system to handle tight rover maneuvers around rocks. We extended the traversability analyzer and integrated it with the Lattice motion planner, which generates continuous curvature paths maneuvering within 8-12 cm of obstacles without collision.

Our preliminary results showed good promise of the potential of this level of autonomy for future planetary exploration. While we made several enhancements to the overall reliability, additional validation under different terrain and lighting conditions would still be warranted. A more detailed error budget on final placement would need to be assessed; and the traverse of challenging terrain would need further validation for future insertion into flight missions.

This level of autonomy would have significant impact on the science return when multiple targets could be assessed in a single sol. Compared to state-of-the-art planetary operations, the saving could amount to an order of magnitude reduction in the number of sols. However, this technology would need to be assessed relative to available on-board instruments and the time that they would need to acquire and process their measurements.

Future work would include demonstrating autonomous instrument placement on multiple targets on rocky and sloped terrains. Several components of this capability are being considered for the Mars Science Laboratory mission and the entire system could be integrated onto the proposed joint NASA/ESA ExoMars mission currently planned for 2018.

IX. REFERENCES

- [1] W Kim et al., "Targeted Driving Using Visual Tracking on Mars: from Research to Flight," *Journal of Field Robotics*, 2008.
- [2] B Wilcox, "Experiences with Waypoint Navigation in Planetary Exploration and DOD Applications," in *Association of Unmanned Vehicles, Winter '02*, 2002.
- [3] (2010) Mars Rovers. http://marsrovers.nasa.gov/mission/traverse_maps.html
- [4] I Nesnas, M Maimone, and H Das, "Autonomous Vision-Based Manipulation from a Rover Platform," in *IEEE Symposium on Computational Intelligence in Robotics & Automation*, CA, 1999.
- [5] R Volpe, "Navigation Results from Desert Field Tests of the Rocky 7 Mars Rover Prototype," *International Journal of Robotics Research, Special Issue on Field and Service Robots 18(7)*, 1999.
- [6] Félix Ingrand, Simon Lacroix, Solange Lemai-Chenevier, and Frédéric Py, "Decisional Autonomy of Planetary Rovers," *Journal of Field Robotics, Volume 24, Issue 7*, pp. 559 - 580, 2007.
- [7] L Pedersen, M Deans, D Lees, S Rajagoplan, and D Smith, "Multiple-Target Single Cycle Instrument Placement," *International Symposium on AI, Robotics, and Automation in Space*, 2005.
- [8] T Yoshikawa, "Manipulability of robotic mechanisms," *Int. Journal of Robotics Research 4(2)*, pp. 3-9, 1985.
- [9] C Urmson, R Simmons, and I Nesnas, "A Generic Framework for Robotic Navigation," *IEEE Aerospace Conference, Montana*, 2003.
- [10] M Pivtoraiko and A Kelly, "Differentially constrained motion replanning using state lattices with graduated fidelity," *IEEE/RSJ Int'l Conf on Intelligent Robots and Systems, (IROS)*, 2008.
- [11] A Mourikis, S Roumeliotis, and J Burdick, "SC-KF Mobile Robot Localization: A Stochastic Cloning-Kalman Filter for Processing Relative-State Measurements," *IEEE Trans on Robotic*, 2007.
- [12] I Nesnas et al., "CLARATy: An architecture for reusable robotic software," in *SPIE Aerosense Conference*, Orlando, 2003.
- [13] S Koenig and M Likhachev, "D* Lite," in *Proceedings of the AAAI Conference of Artificial Intelligence (AAAI)*, 2002.
- [14] A Kelly and B Nagy, "Reactive nonholonomic trajectory generation via parametric optimal control," *Int'l J. of Robotics Research*, 2002.
- [15] John Vannoy and Jing Xiao, "Real-time adaptive motionplanning (ramp) of mobile manipulators in dynamic environments with unforeseen changes," *IEEE Trans on Robotics*, 2008.
- [16] L Pedersen et al., "Mission Planning and Target Tracking for Autonomous Instrument Placement," in *IEEE Aerospace Conf*, 2005.
- [17] M Maimone, P Leger, and J Biesiadecki, "Overview of the Mars Exploration Rovers' autonomous mobility and vision capabilities," in *IEEE Int'l Conf. on Robotics and Automation (ICRA)*, Rome, 2007.
- [18] Devendra P. Garg and Manish Kumar, "Optimization techniques applied to multiple manipulators for path planning and torque minimization," *Engineering App of Artificial Intelligence*, 2002.
- [19] Jaydev P. Desai and Vijay Kumar, "Nonholonomic motion planning for multiple mobile manipulators," in *ICRA*, 1997.
- [20] Maria Bualat, et.al., "Operating Nomad during the Atacama Desert Trek," in *Int'l Conf on Field and Service Robotics*, 1997.
- [21] Y Cheng, M Maimone, and L Matthies, "Visual odometry on the Mars Exploration Rovers," in *IEEE Robotics and Automation Magazine Special Issue (MER)*, 13(2), 2006, pp. 54 - 62.
- [22] Mingwu Chen and Ali M.S. Zalzal, "A genetic approach to motion planning of redundant mobile manipulator systems considering safety and configuration," *Journal of Robotics Systems*, 1997.
- [23] A Bowling and O Khatib, "The dynamic capability equations: a new tool for analyzing robotic manipulator performance," *IEEE Transactions on Robotics 21(1)*, pp. 115-123, 2005.
- [24] P Backes, A Diaz-Calderon, M Robinson, M Bajracharya, and D Helmick, "Automated Rover Positioning and Instrument Placement," in *IEEE Aerospace Conference*, 2005.
- [25] K Ali et al., "Attitude and Position Estimation on the Mars Exploration Rovers," *IEEE Systems, Man and Cybernetics Conf*, 2005.