

Autonomous Vision-Based Manipulation from a Rover Platform

Issa A.D. Nesnas, Mark W. Maimone, Hari Das
Jet Propulsion Laboratory, Pasadena, CA 91109

Abstract— Current rover designs use on-board manipulators to enhance their capabilities for planetary exploration and in-situ science. In this paper, we describe how these manipulators can be used to perform two types of operations: rock sample acquisition for return to earth and instrument placement for in-situ science measurements. We describe the computational architecture, tools, and algorithms that we developed for this task. These algorithms integrate rover odometry, stereo visual tracking, and tactile sensing for each operation. We have successfully demonstrated these operations on a self-contained Mars Rover prototype, *Rocky 7*. We have demonstrated grasping a small rock sample from a distance of more than one meter away and placing an instrument on a boulder from a distance of more than five meters away.

I. INTRODUCTION

FOLLOWING the success of the Sojourner Rover of the Mars Pathfinder mission, there has been an increased interest in adding manipulation on-board rovers to enhance their planetary exploration capabilities. Two types of manipulators have been used on several Mars rover prototypes: a mast that extends a stereo camera pair several feet above the rover's platform, and a manipulator arm that is used for sample acquisition, digging, and science experiments.

The mast's primary function is long range sensing of the surrounding terrain using narrow field-of-view cameras. To acquire a complete panorama of high resolution images, the mast must be able to pan and tilt its cameras. Taking advantage of these degrees of freedom, one or more science instruments are also mounted on the mast for placement onto a designated target. Because the mast carries sensitive sensors and instruments, it is not suitable for digging or acquiring rock samples. A second manipulator with shorter link lengths is used instead. This manipulator arm can also carry less sensitive science instruments.

Because of power consumption and mass constraints, these manipulators have limited degrees of freedom. In this work, we will demonstrate how we use these manipulators in conjunction with the vehicle's mobility system to compensate for their limited dexterity. We will also show how we use frequent visual feedback to compensate for the uncertainties and the simplified kinematic models of the system when tracking a target. Tactile sensing is used when the manipulators get close to the target.

In the next section, we briefly present some related work that uses sensor-based manipulation algorithms. In the sections that follow, we present our objectives, approach, and architecture. We also present our algorithms and provide some experimental results. We conclude with a summary and some planned enhancements to this work.

II. BACKGROUND

There have been several efforts in sensor-based manipulation especially in vision-guided manipulation. Some researchers developed algorithms to servo manipulators in the Cartesian space [1] [8], while others worked in the image plane using intensity-based approaches to vision-guided manipulation [10] [3]. Several researchers have achieved high frame-rate visual servoing [1] [8] [10] [3] [2] [7] [11]. However, most of this work assumed a dexterous manipulator mounted on a fixed platform. The relative size of the object in the images remained the same throughout the servoing process, and most of these efforts were demonstrated in an indoor environment where lighting can be controlled.

In our case, the manipulators are mounted on a moving rover platform. We rely on the mobility system to compensate for the limited dexterity of the manipulators. We compensate for uncertainties, resulting from the terrain roughness, and changes in the shape and size of the target by visually tracking the target as we approach it. Our rovers operate in an outdoor environment where specular reflections, shadows, and changes in lighting conditions make the vision problem quite challenging. Template searches on the intensity image can track well at long distances, but are less reliable at the final approach to the object [13]. To overcome these challenges, we servo on an elevation map rather than the raw image. The elevation map is generated from an on-board stereo vision system. We integrate the various sub-systems using a control architecture developed for this project to achieve intelligent autonomous behavior.

III. OBJECTIVE

The objective of our project is to demonstrate intelligent autonomous manipulation on-board a rover platform subject to constraints similar to those en-




Fig. 1. The *Rocky 7* rover.

countered on Mars. In particular, our goal is to use a Mars rover prototype, *Rocky 7*, to autonomously pick rock samples selected from a distance of more than one meter away, and to autonomously place a science instrument on a rock designated from a distance of more than five meters away.

Without this level of autonomy, each objective would have taken more than five days to accomplish in a Mars mission. Tele-operation is very unlikely to succeed due to communication time-delay (several minutes for Mars) and a restricted communication window (a few minutes twice per day for Sojourner during the 1997 Pathfinder mission). Alternatively, identifying the location of the target and then blindly driving toward it will not work either since there are many disturbances and uncertainties in modeling rover motion over the terrain.

IV. SYSTEM & COMPUTING ARCHITECTURE

A. The *Rocky 7* System

Rocky 7 is a Mars rover prototype designed and built by the Long Range Science Rover team as a testbed for autonomous and intelligent algorithms [12] (Figure 1). *Rocky 7* is a six-wheel drive vehicle with a rocker-bogey mobility mechanism. It has two steerable front wheels and four non-steerable back wheels. The mobility mechanism defines the possible maneuvers the vehicle can perform.

Mounted onto the rover platform are two manipulators: a two degree-of-freedom (DOF) arm with two independently actuated scoops (making it an effective three DOF arm), and a three degree-of-freedom mast. The arm has a shoulder roll and a shoulder pitch, while the mast has an additional elbow pitch. Three pairs of stereo cameras are mounted on the rover. A narrow field-of-view (43° measured) stereo camera pair is mounted on the mast, and two wide field-of-view (103° measured) stereo camera pairs are mounted on the front and back sides of the vehicle. The latter

pairs, also known as hazard avoidance cameras, are mounted at about 30 cm above the ground and are aimed downwards at a fixed 45° angle.

Due to the limited dexterity of the mast manipulator and the mounting of its stereo camera pair, the mast cameras cannot be used effectively for guiding the manipulator arm. So we rely on the hazard avoidance cameras when using the arm and on the mast cameras when using the mast.

B. Computing Architecture

The computing system consists of a 3U VME back-plane with a 60 MHz 68060 processor with on-board Ethernet, two frame-grabbers, a digital I/O board, and an analog I/O board. The main processor runs a VxWorks 5.3 real-time operating system. Each actuator (DC brushed) is controlled by a separate microcontroller (LM629) that is connected to the main CPU via the digital I/O board. The micro-controllers use a bi-directional shared local I/O bus connecting them to the digital I/O board. Sensors are also connected to the analog and digital boards. To optimize the performance of the system, we use the micro controller's local trajectory generation for coordinated rover motions. The on-board processor communicates with an external host via a wireless Ethernet at a maximum throughput of 1 MB/sec.

C. An Object-Oriented Software Architecture

Based on an implementation developed in [6], we developed a three-layered object-oriented system hierarchy using C++. At the lowest layer, we placed the system device drivers. The middle layer is the hardware abstraction layer which has a hierarchical structure and uses virtual mechanisms to talk to the hardware. The base classes in this layer represent the abstract and hardware independent functionality of its components. To hide the hardware dependencies, parameter passing is done using the base classes. The third layer uses similar hierarchies to represent the various sub-systems such as the vision, manipulation, mobility sub-systems. Higher level algorithms use classes from the middle and third layers to control the rover. In addition to the classes, we developed template-based hierarchies for handling data objects.

V. CONTROL STRATEGY

Since we rely on frequent sensory feedback to achieve the rock acquisition and instrument placement operations, we do not attempt to develop accurate kinematic models of the vehicle and the arms. Rather, we rely on simplified approximate solutions that require minimal




Fig. 2. The mobility of a six-wheel rover with two-wheel front steering (e.g. *Rocky 7*)

computation. Next, we present the models that were used in our approach.

A. Rover Mobility

Figure 2(a)-(c) shows the possible maneuvers that can be accomplished with a rocker-bogey mechanism having only two steerable wheels. Since the bogey wheels (back wheels) cannot change their orientation, all rover motions must occur along a circular path whose center lies along the axis of the non-steerable wheels (see the y -axis in Figure 3). Since there are two parallel axes for the non-steerable wheels, an equidistant axis is selected to minimize the slippage and stress on the bogey wheels. Unless this type of vehicle moves in a straight line (infinite radius), slippage always occurs. Circular arcs of zero radius correspond to in-place rotations. This maneuver has the most slippage and results in large heading errors (unless heading sensors are used to compensate for these errors). This type of vehicle is unable to move laterally in a crab-like maneuver Figure 2(d) which would be desirable for our manipulation task. This restriction holds for all vehicles of this class independent of what algorithm is used to move the vehicle.

The model above is accurate when the rover is traversing level terrain. We define the vehicle's motion about point C which is the centroid of the bogey wheels. However, C is not a static point on the rover. It shifts as the rocker and bogey mechanism traverses rocky terrain. Nonetheless, it does not shift much since the rocker and bogey angles have limited ranges. Because we are only interested in the initial estimates of the wheels' motion parameters, we do not need to develop the complex three-dimensional kinematics for the mechanism. Rather, we use the flat terrain approximation to provide initial input to the wheel motors. We, then, constantly reevaluate these parameters based on visual feedback.

B. Driving the Vehicle to a Goal

Using this approximate geometrical approach, we would like to move the rover such that a specified point A on the rover reaches a goal point B specified in the




Fig. 3. Generating a single arc trajectory; point A is the optimal grasp location and point B is the goal

two-dimensional terrain (see Figure 3). As long as the final rover orientation is unconstrained, a single arc motion of the rover is theoretically sufficient to drive point A to point B . This strategy forms the basis of the vehicle motions in our feedback system for the rock sample acquisition.

When using the manipulator arm which has two degrees of freedom, we have to rely on the vehicle's mobility mechanism to position the rover arm within 10 mm of the goal point. The workspace of the rover arm is merely the volume of a hollow hemisphere with a thick shell. The intersection of this workspace with the ground is an arc with a thickness of one inch. This is the reachable workspace of the arm for rock acquisition. But because the end effector on that arm does not have a wrist roll, only a small area of about 40 mm in diameter (valid workspace region) can be used reliably to pick rocks. Consequently, the rover must drive toward the goal and precisely position the optimal arm location over the target rock. Once at the goal, the rover deploys its arm and picks the rock.

The rock sample acquisition does not require a specific rover orientation. However, orientation is constrained when placing an instrument onto a boulder. We discuss this in the next section.

C. Re-orienting the Vehicle while Driving

Another objective is to move a specified point A on the rover to a goal point B with a specified final orientation of the rover. For example, when placing an instrument on a boulder, the rover must approach a boulder in a given direction and be able to place the instrument along the surface normal of the area of interest.

The problem of moving and re-orienting the vehicle is over-constrained for a single arc trajectory. A minimum of two arcs is necessary to accomplish this task. There is an infinite number of arc pairs that can drive the rover to its destination with the proper fi-

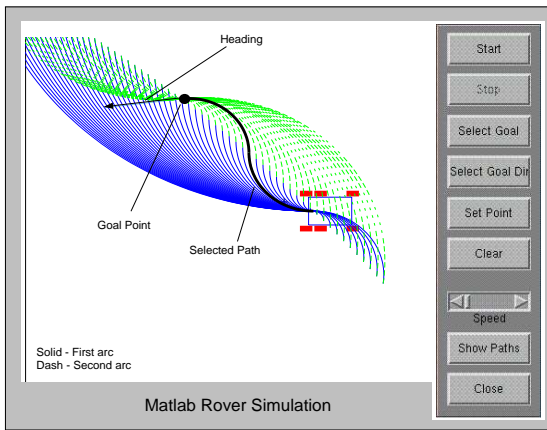


Fig. 4. Possible paths with optimal path selection (black line)

nal orientation (see Figure 4). When either arc length of the pair goes to zero, the rover does a rotate-in-place (Figure 2(c)) either at the beginning or at the end of the trajectory. However, we would like to minimize the rotate-in-place maneuvers since they have the most slippage and create the highest stresses on the bogey wheels. Besides, without a heading sensor, such as a sun sensor, a rotate-in-place yields the worst odometry causing us to lose track of the target. The criteria we use for path selection are to: (i) minimize the sum of the arc lengths if the heading direction remains the same for the two arcs, or to (ii) minimize the difference if the heading direction changes. This selects the shortest and smoothest path.

To validate the above strategies, we developed a Matlab simulation that generates the possible path and animates the rover motion (Figure 4).

In the next section, we will show how these motion strategies are updated by the visual tracker to achieve the overall task objectives.

D. Algorithm for Vision-Based Manipulation

Table I describes the algorithm that was used to acquire a small rock sample. Once the scientist selects the target rock, the selected point is transmitted back to the rover which uses stereo vision processing based on camera models to compute the three-dimensional location of the rock [14]. Using the x and y world coordinates of the rock, a single arc is computed as described above and the rover starts its traverse toward the target. The rover needs to be positioned such that the rock is inside the arm's valid workspace region. To compensate for the slippage in the unknown terrain, the approximate kinematics, and any other disturbances, we continuously acquire images along the path and re-evaluate the position of the rock.

Thus, at 10 cm intervals, the rover stops to take

new stereo images using the body navigation cameras. Using the vehicle's odometry, we compute a new estimated location of the target and a small window around that point is searched in an attempt to relocate the target. The search is done in the elevation map generated from the range image of the stereo pair. We search for the shape of the rock rather than its visual appearance. In particular, we assume that any target rock will be resting higher on the ground than its nearby surroundings, and lock in on the local elevation maximum as the new, refined 3D target point. Since we do not always have a dense elevation map, we linearly interpolate missing data from the range image before searching for the local maximum in the elevation map.

To compensate for the large errors in the odometry estimate, we also assume that our targets are visually distinct from the background sand, and use an intensity filter to focus attention in the elevation map. We have chosen this solution due to timeframe constraints. In fact, any pixel classification technique can be used instead of brightness, *e.g.*, [9]. If no range data is available, then no refinement is done, and the vehicle odometry is assumed to be correct. Further details on the visual tracking algorithm can be found in [4].

The rover tracks the target and continues to correct its motion trajectory until the arm's workspace region is centered over the target rock.

E. Target Grasping

Once the arm's workspace is centered to within 1 cm of the target, the arm is deployed. The scoops are opened and the arm moves downwards toward the ground sensing obstacles along its trajectory. Sensing is done by monitoring changes between the desired and actual trajectories of the arm's shoulder joints. The arm stops when either the target or the ground are sensed, at which point the arm goes into a grasping mode. As the scoops sense resistance, the arm is raised in small amounts while the scoops continue to close. The arm exits this mode when either a stable grasp is achieved, the scoops are completely closed, or the algorithm times out. This algorithm helps ensure that the gripper has a good hold on the target.

F. Instrument Placement

Table II describes the algorithm used for instrument placement. The general strategy is similar to the rock sample acquisition, except that the rover approaches the target and place the instrument at a specific orientation determined by the target's surface normal. The instrument placement can be started from a distance of more than five meters away. Because of this long dis-

1. Acquire stereo image pair with body navigation cameras
2. Send left image over wireless network to host
3. Scientist/Operator selects target rock on left image
4. Target location and intensity threshold sent to rover
All subsequent processing occurs on-board
5. Identify 3-D location of rock based on calibrated camera models and on-board stereo image processing
6. Compute single-arc rover trajectory to target
7. Drive rover toward target
8. Periodically (every 10 cm) poll tracking software to update target location using new stereo pair & current odometry
9. Redirect rover toward new target location using new single-arc trajectory, and repeat until target is within 1 cm of goal position
10. Deploy sampling arm, sense and pick up rock.

TABLE I
ALGORITHM FOR SMALL-ROCK ACQUISITION USING THE ARM

1. Acquire stereo image pair with mast cameras
2. Send the left image over wireless network to host
3. Scientist/Operator selects target on left image
4. Target location and intensity threshold sent to rover
All subsequent processing occurs on-board
5. Identify 3-D location of target region based on calibrated camera models and on-board stereo image processing
6. Compute single-arc rover trajectory to target
7. Drive rover toward target
8. Periodically (every 50 cm) poll the target tracking software to update target location using new stereo pair and current odometry
9. Redirect rover toward the new target location using new single-arc trajectory, and repeat until target is within 1 m from goal
10. Compute a 2-arc trajectory to a point 0.5 m from target with a final rover orientation to match the target's surface normal
11. Drive rover along the two-arc trajectory
12. Poll target tracking software for update on target location and surface normal
13. If target is out of instrument's reach, move closer to target and update location
14. Deploy mast arm instrument toward target
15. Servo mast along surface normal until mast touches the rock

TABLE II
ALGORITHM FOR INSTRUMENT PLACEMENT USING THE MAST

tance, we use the narrow field-of-view cameras of the mast to track the target. We drive the vehicle with its mast half-way up to continuously monitor the target. Every 50 cm the rover stops and acquires a new stereo pair for the target tracking algorithm. When the rover is within 1 m of its target, it stops and plans a two-arc trajectory to adjust its final approach toward the target. The final approach is determined by the surface normal which is computed from the range data of the target area. The rover drives along the two-arc trajectory and stops in front of the boulder. The mast fully deploys and approaches the boulder. It stops at about 20 cm from the target's surface. Instrument sensing is enabled and the mast moves along a straight line toward the target until the instrument touches the rock. The mast then stops and the instrument takes its measurements. The mast retracts and stows and the rover moves away from the target area.

VI. EXPERIMENTAL RESULTS

We have performed several experiments in JPL's Mars Yard and successfully demonstrated the acqui-

sition of small rocks (3-5 cm) located over 1 meter in front of the rover. Figure 5 shows a sample tracking sequence, with the target indicated in each frame by a dark square. We have also successfully placed the instrument arm onto a boulder over five meters away. However, since the visual tracking algorithm servos on the local elevation maximum, only targets on the top of rocks were specified at this time.

Out of 25 trials in the Mars Yard, eleven were completely successful in tracking and acquiring the rock. Three were marginally successful whereby the rover fails to keep hold of the rock while the arm lifts up from the ground. The remaining trials failed due to one of the following reasons:

- The visual tracker loses its target. This occurs when either the target leaves the camera FOV, no range data is available due to lighting conditions, multiple targets are visible inside the search window, poor odometry estimates move the target outside the search window, or target is same color as background. Using 14 different datasets, the visual tracker succeeded in maintaining tar-




Fig. 5. Sample tracking sequence.

get lock through 10 complete sequences. Correcting the threshold increased this to 13 successful datasets.

- The visual tracking succeeds but the rover cannot stabilize about the goal point. Since we rely on the mobility system, positioning resolution of the vehicle is less than our goal tolerances. This is mainly apparent on sandy ground where vehicle maneuvering introduces much positional uncertainty.

VII. FUTURE WORK

We are planning to improve the robustness of the visual tracking algorithm (reducing its dependency on the brightness-based filter) by matching the entire shape of the terrain around the target. We also plan to improve the position and pose estimates using visual feature tracking on the whole scene[5]. These improvements should allow tracking of targets anywhere on a rock, which would enable a more general mast placement capability. Another area that we will be addressing is the elimination of the instabilities that result from the imprecise vehicle motions on loose terrain. Improving the coordination between the vehicle and the arm trajectories will improve the overall system. We also like to introduce obstacle avoidance in the arc planning strategies and test the continuous operation of the rover while images are being acquired and processed.

VIII. ACKNOWLEDGMENTS

We wish to thank the Long Range Science Rover (LRSR) team for providing the *Rocky 7* rover and for

their assistance and support during the development of this work, especially Samad Hayati, Richard Volpe, Bob Balaram, Robert Ivlev, Sharon Laubach, Alex Martin-Alvarez, Larry Matthies, Clark Olson, Richard Petras, Robert Steele, and Yalin Xiong. The work described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract to the National Aeronautics and Space Administration.

REFERENCES

- [1] P.K. Allen. Automated tracking and grasping of a moving object with a robotic hand-eye system. *IEEE Transactions on Robotics and Automation*, 9(2):152–165, 1993.
- [2] Gregory D. Hager, Gerhard Grunwald, and Kentaro Toyama. *Intelligent Robotic Systems*, chapter Feature-Based Visual Servoing and its Application to Telerobotics. Elsevier, Amsterdam, 1995. V. Graefe, editor.
- [3] R. Horaud, F. Dornaika, and B. Espiau. Visually guided object grasping. *IEEE Transactions on Robotics and Automation*, 14(4), August 1998.
- [4] Mark Maimone, Issa Nesnas, and Hari Das. Autonomous rock tracking and acquisition from a mars rover. In *International Symposium on Artificial Intelligence for Robotic Systems in Space*, Noordwijk, Netherlands, June 1999. <http://robotics.jpl.nasa.gov/tasks/pdm/papers/isairas99/>.
- [5] Larry Matthies. *Dynamic Stereo Vision*. PhD thesis, Carnegie Mellon University Computer Science Department, October 1989. CMU-CS-89-195.
- [6] I.A. Nesnas and M.M. Stanišić. A robotic software developed using object-oriented design. In *ASME Design Automation Conference*, Minnesota, 1994.
- [7] H.K. Nishihara, H. Thomas, E. Huber, and C.A. Reid. Real-time tracking using stereo and motion: Visual perception for space robotics. In *International Symposium on Artificial Intelligence for Robotic Systems in Space*, pages 331–334, 1994.
- [8] N. Papanikolopoulos and P.K. Khosla. Adaptive robotic visual tracking: theory and experiments. *IEEE Transactions on Automatic Control*, 38:1249–1254, March 1993.
- [9] L. Pedersen, D. Apostolopoulos, W. Whittaker, T. Roush, and G. Benedix. Sensing and data classification for robotic meteorite search. In *SPIE Photonics East*, Boston, 1998.
- [10] S.B. Skaar, W.H. Brockman, and W.S. Jang. Three dimensional camera space manipulation. *International Journal of Robotics Research*, 9(4):22–39, 1990.
- [11] D.A. Theobald, W.J. Hong, A. Madhani, B. Hoffman, G. Niemeyer, L. Cadapan, J.J.-E. Slotine, and J.K. Salisbury. Autonomous rock acquisition. In *AIAA Forum on Advanced Development in Space Robotics*, Madison, WI, August 1996.
- [12] Richard Volpe. Navigation results from desert field tests of the Rocky 7 mars rover prototype. *International Journal of Robotics Research*, Accepted for Publication, Special Issue on Field and Service Robots 1999. <http://robotics.jpl.nasa.gov/people/volpe/papers/JnavMay.pdf>.
- [13] David Wettergreen, Hans Thomas, and Maria Bualat. Initial results from vision-based control of the Ames Marsokhod rover. In *IEEE International Conference on Intelligent Robots and Systems*, pages 1377–1382, Grenoble, France, September 1997. <http://img.arc.nasa.gov/papers/iros97.pdf>.
- [14] Yalin Xiong and Larry Matthies. Error analysis of a real-time stereo system. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1087–1093, 1997. <http://www.cs.cmu.edu/~yx/papers/StereoError97.pdf>.