Rover Maneuvering for Autonomous Vision-Based Dexterous Manipulation

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

Abstract— **Manipulators mounted on-board rovers have limited dexterity due to power and weight constraints imposed by rover designs. However, to perform science operations, it is necessary to be able to position and orient these manipulators on science targets in order to carry out in-situ measurements. This article describes how we enhance manipulator dexterity using the rover mobility system. The lack of omni-directional driving capability and the constraints imposed by the mobility mechanism requires vehicle maneuvering to supplement the manipulators' motions. Target tracking using stereo vision is integrated with rover maneuvering 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. 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 is also used to carry 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.

The objective of our project is to demonstrate autonomous manipulation on-board a rover platform subject to constraints similar to those encountered on Mars. Using the Mars rover prototype, *Rocky 7*, our goal is to autonomously pick rock samples selected from a distance of more that one meter away, and to autonomously place a science instrument on a rock from a distance of more than five meters away.

Without this level of autonomy, each objective would have taken three to five days to accomplish in a Mars mission. Tele-operation is very unlikely to succeed due to communication time-delays (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.

In the next section, we briefly present some related work that uses sensor-based manipulation algorithms. In the sections that follow, we present the mathematical formulation for rover maneuvering and demonstrate how these strategies are used on-board *Rocky 7* to achieve the above stated goals.

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 [3] [9]. Several researchers have achieved high frame-rate visual servoing [1] [2] [3] [7] [8] [10]. However, most of this work assumed a dexterous manipulator mounted on a fixed platform with the exception of [12] which was mounted on a rover prototype. 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 in the elevation map as we approach it. Our rovers operate in an outdoor environment where specular reflections, shadows, and changes in lighting conditions make the visual tracking problem quite challenging.

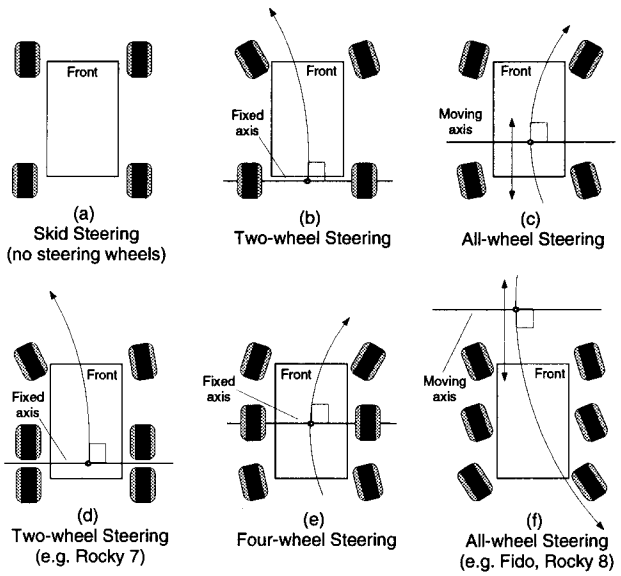

Fig. 1. Various types of rover mobility systems

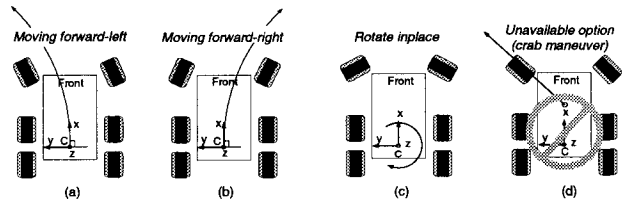

Fig. 2. The mobility of a six-wheel rover with two-wheel front steering (e.g. *Rocky 7*)

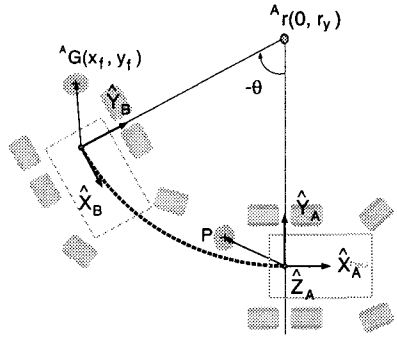

Fig. 3. Moving point P to a goal point G with no specified final heading

III. Rover Maneuvering

A. Rover Types

Because of the limited dexterity of on-board manipulators, the mobility system must be used for rover positioning and orienting to supplement the manipulators' motions. Since omni-directional driving in rough terrain is undesirable, limited, or unavailable on some rovers, we will explore the use of Ackermann steering to achieve both vehicle positioning and orienting.

Figure 1 shows the steering capability of various wheeled rovers. These rovers can be grouped into three classes: (i) skid steering vehicles (all fixed-direction wheels) Figure 1(a), (ii) Ackermann steering vehicles, where at least one pair of wheels are non-steerable Figure 1(b,d,e), and (iii) all-wheel steering vehicles Figure 1(c,f). In the latter class, the center of rotation of the vehicle for any motion is unconstrained. This is not the case for the Ackermann steering vehicles, where the center of rotation must lie along the fixed-wheel axis. In case of tandem wheels such as *Rocky 7* (Figure 1(d)), the tandem pair can be treated as one larger wheel with its axis equidistant from the two fixed-wheel rotation axes [11]. This minimizes the slippage of the vehicle which, in this case, will be considerable for sharp turns.

Rover mobility systems use some type of rocker or rocker-bogey mechanism for traversing rough terrain. Driving sideways (crab-maneuver) does not use the rocker-bogey suspension effectively, thus reducing the vehicle's stability and making such maneuvers undesirable for long traverses. Such maneuvers can only be performed with all-wheel steering vehicles.

The motion of steerable vehicles can always be described in terms of circular arcs. An arc of zero radius corresponds to a rotate-in-place motion while an arc of infinite radius corresponds to a straight line motion. The only difference between Ackermann steering and all-wheel steering vehicles is that the center of the arc for the former is constrained to the fixed-wheel axis.

Figure 2 shows the possible motions of a vehicle that has two-steerable front wheels (*e.g Rocky 7*). It is important to note that these are not accurate models of the full rover kinematics. They are merely flat terrain approximations which require minimal computational effort but can be combined with visual feedback to provide accurate vehicle positioning for manipulation.

B. Driving the Vehicle to a Goal

Given a goal point ${}^A\mathbf{G}\,(x_f, y_f)$ relative to the A frame, we would like to move the rover such that a fixed point ${}^A\mathbf{P}\,(p_x, p_y)$ in the initial rover reference frame reaches a goal point specified in the 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 $\mathbf{P}$ to point $\mathbf{G}$.

The goal point ${}^A\mathbf{G}$ which is the fixed point $\mathbf{P}$ relative to the A coordinate frame after the motion is complete can be described by the following equation:

$$ {}^A\mathbf{P} = {}^A_B\mathbf{T}\ {}^B\mathbf{P} = {}^A\mathbf{r} + {}^A_B\mathbf{R}_\theta\,[\,{}^B\mathbf{P} - {}^A\mathbf{r}\,] \qquad (1) $$

where ${}^A_B\mathbf{R}_\theta$ is the rotation matrix of frame B relative to frame A. Since the center of turning lies along the fixed-wheel axis, the vector $\mathbf{r}$ has a zero abscissa. Substituting $\mathbf{r}(0, r_y)$ into equation (1) and solving for r_y yields:

$$ \begin{aligned} r_y &= \frac{x_f - \cos\theta p_x + \sin\theta p_y}{\sin\theta} \\ &= \frac{y_f - \sin\theta p_x - \cos\theta p_y}{1 - \cos\theta} \end{aligned} \qquad (2) $$

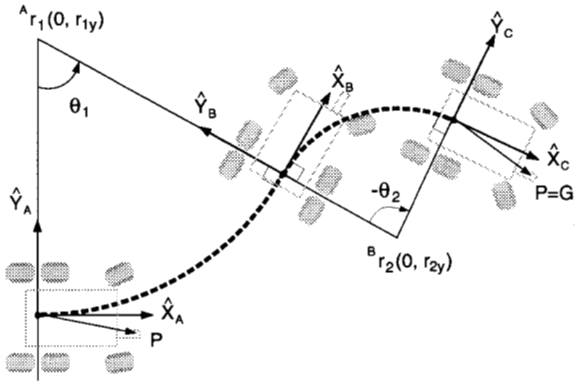

Fig. 4. Two-arc trajectory for goal point **G** and final orientation θ_f

Solving equation (2) for θ yields:

$$\theta = 2\,\text{atan2}\,(y_f - p_y,\; x_f + p_x) \tag{3}$$

Substituting equation (3) into equation (2) gives us a value for the radius of turning. Once the center and angle of turning are known, we can compute the wheel steering angles from the vehicle geometery. This derivation is used in the rock sample acquisition for *Rocky 7* since it does not require a specific rover orientation. However, orientation is constrained when placing an instrument onto a boulder. In the next section we discuss how to move the rover to a goal point with a specified final orientation.

C. Driving to a Goal with a Specified Final Orientation

Another objective is to move a specified point $^A\mathbf{P}\,(p_x, p_y)$ in the rover frame to a goal point $^A\mathbf{G}\,(x_f, y_f)$ with a specified final orientation of the rover $^A\theta_f$ relative to the A frame. For example, this maneuver is necessary when placing a mast instrument on a target, where the instrument must be oriented along the surface normal of the target. Once again we rely on the vehicle's mobility mechanism primarily for positioning and orienting the rover since the mast that carries the instruments has limited reach and degrees-of-freedom for orientation.

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 final orientation (see Figure 5). The motion of the vehicle along the two arcs can be described using equation (1) for the motion from frame A to frame B and once again from frame B to frame C (Figure 4) by :

$$\begin{aligned} {}^A\mathbf{P} &= {}^A\mathbf{r}_1 - {}^A_B\mathbf{R}_{\theta_1}\,[\,{}^A\mathbf{r}_1 - {}^B\mathbf{r}_2\,] \\ &\quad + {}^A_B\mathbf{R}_{\theta_1}\;{}^B_C\mathbf{R}_{\theta_2}\,[\,{}^C\mathbf{P} - {}^B\mathbf{r}_2\,] \end{aligned} \tag{4}$$

Once again, due to the fixed-direction wheels, the arc centers can be described by $\mathbf{r}_1 = (0, r_{1y})$ and $\mathbf{r}_2 = (0, r_{2y})$ respectively. The specified final orientation is the sum of the two rotations $^A\theta_f = {}^A\theta_1 + {}^B\theta_2$. Hence, we can write equation (4) in term of $^A\mathbf{P}\,(x_f, y_f)$ as:

$$\begin{aligned} x_f &= [\cos\theta_f p_x - \sin\theta_f p_y] \\ &\quad + r_{2y}\sin\theta_f + (r_{1y} - r_{2y})\sin\theta_1 \qquad (5) \\ y_f &= [\sin\theta_f p_x + \cos\theta_f p_y] \\ &\quad - r_{2y}\cos\theta_f - (r_{1y} - r_{2y})\cos\theta_1 + r_{1y} \qquad (6) \end{aligned}$$

By eliminating θ_1 from equations (5) and (6), we can parameterize one radius in terms of the other radius. All other parameters are known. To do so, we write an expression for $\sin\theta_1$ and $\cos\theta_1$ from equations (5) and (6) respectively, then square and sum the resultant equations. This yields:

$$r_1(r_2) = \frac{c_0 - c_2 r_2}{c_1 + c_3 r_2} \quad \text{or} \quad r_2(r_1) = \frac{c_0 - c_1 r_1}{c_2 + c_3 r_1} \tag{7}$$

where

$$\begin{aligned} c_0 &= \tfrac{1}{2}({x'_f}^2 + {y'_f}^2) && (8) \\ c_1 &= y'_f && (9) \\ c_2 &= x'_f \sin\theta_f - y'_f \cos\theta_f && (10) \\ c_3 &= \cos\theta_f - 1 && (11) \end{aligned}$$

We have dropped the y subscript since it is understood in this context. Using equations (5) and (6) we can write θ_1 in terms of the two radii as:

$$\begin{aligned} \theta_1 &= \text{atan2}\,((x'_f - r_2\sin\theta_f)\,\text{sgn}(r_1 - r_2), \\ &\quad (r_1 - y'_f - r_2\cos\theta_f)\,\text{sgn}(r_1 - r_2)) \end{aligned} \tag{12}$$

where (x'_f, y'_f) are the known location of the vehicle's origin at the goal location which are given by:

$$\begin{aligned} x'_f &= x_f - (\cos\theta_f p_x - \sin\theta_f p_y) \\ y'_f &= y_f - (\sin\theta_f p_x + \cos\theta_f p_y) \end{aligned} \tag{13}$$

Since one radius can be parameterized in terms of the other, there is an infinite number of two-arc pairs that can drive the vehicle to its goal location with the specified orientation.

D. Selecting an Optimal Path

There are many criteria that govern the selection of an optimal path. (i) One criteria is to eliminate sharp turns and rotate-inplace maneuvers which are unreliable due to slippage, require clearance around the vehicle, and are dangerous unless side cameras are mounted. (ii) Another is to keep the target in the cameras' field-of-view. (iii) A third is to avoid obstacles along the trajectory. All these are important factors to be considered when designing an optimization function for path selection. For the optimization function for *Rocky 7*, we have incorporated two of the above criteria.

$$\eta(r_{2y}) = |r_{1y}\theta_1| + |r_{2y}\theta_2| + ||r_{1y}\theta_1| - |r_{2y}\theta_2|| \tag{14}$$

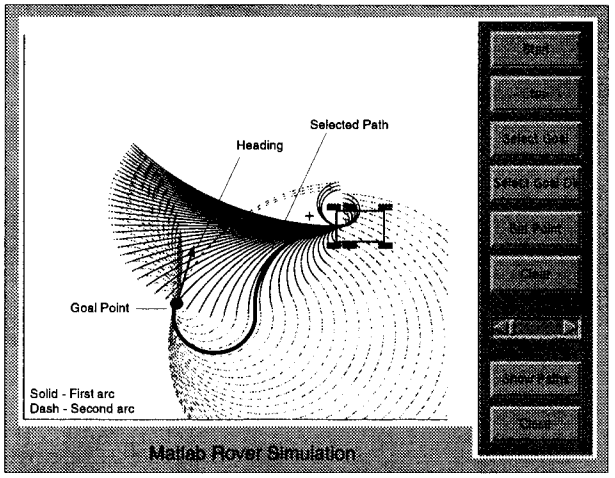

Fig. 5. Possible paths with optimal path selection with changing steering direction

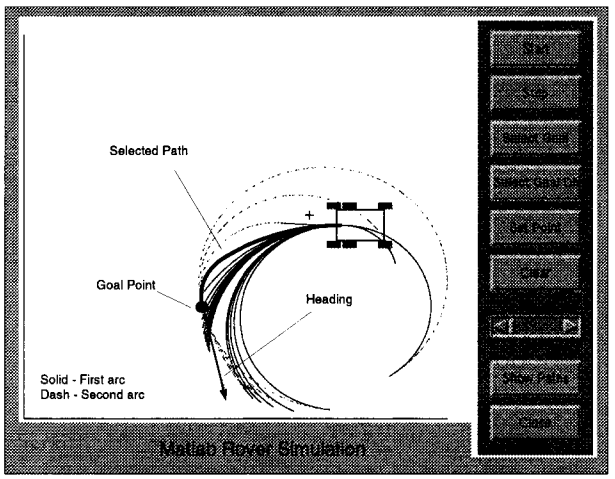

Fig. 6. Possible paths with optimal path selection with same steering direction

Our optimality function is aimed at minimizing the total traversed distance and eliminating sharp turns and rotate-in-place maneuvers. The latter is undesirable since (a) we do not have side cameras to check for clearances, (b) we do not use an absolute heading sensor to compensate for the poor odometry, and (c) the maneuver causes large wheel slippage and stress on the mobility mechanism. Hence, we minimize the sum as well as the absolute difference of the two arc lengths (see equation (14)).

E. Matlab Simulation and Visualization Tool

To validate the derived formulation, we developed a Matlab simulation and visualization tool to generate the necessary trajectories. The user selects a point P relative to the rover frame, a goal point G and an optional final heading. The program then computes and displays the possible trajectories and, if applicable, selects the optimal path based on the optimization function. The steering as well as the rover's motion are animated as the rover traverses the trajectory toward the goal (Figures 5 and 6).

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

IV. Integrating Rover Maneuvering with Visual Target Tracking

Once the scientist selects the target rock to be grasped or tested, 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 [13]. Using the x and y world coordinates of the target, a single arc is computed as described above and the rover starts its traverse toward the target. After the rover traverses 10% of the trajectory, its stops to take new stereo images. 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. Once the target is relocated, we generate a new trajectory and repeat the process until the rover is in the vicinity of (about one meter from) the goal. We continuously acquire images along the path and re-evaluate the position of the rock to compensate for the slippage in the unknown terrain, the approximate kinematics, and any other disturbances. Further details on the visual tracking algorithm can be found in [4].

A. Target Grasping

Target grasping which uses the instrument arm does not usually require a specific grasp orientation. If the instrument arm has few degrees of freedom, then the vehicle can be used to center the arm's workspace over the target using only single-arc trajectories with the visual feedback.

The *Rocky 7* manipulator arm has two degrees-of-freedom for pointing the arm and another two for orienting and opening/closing the scoops. The workspace of the 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

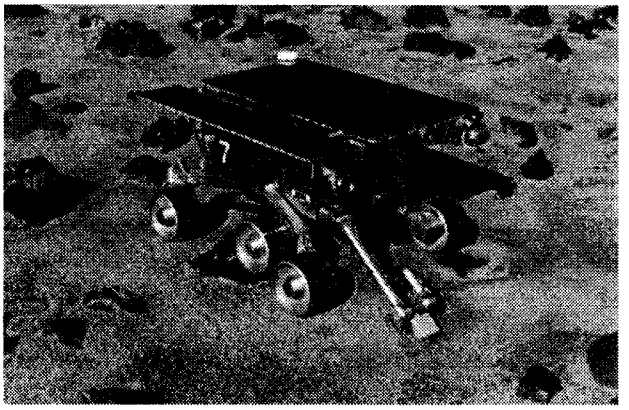

Fig. 7. The *Rocky 7* rover.

drive toward the goal and position the valid workspace region (within 1cm) over the target rock.

Once over the target, the arm is deployed. The scoops open and the arm moves downwards toward the ground sensing obstacles along its trajectory. 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.

B. Instrument Placement

The general strategy for instrument placement is similar to the rock sample acquisition, except that the rover must approach the target and place the instrument with a specific orientation determined by the target's surface normal. Using the trajectory generation and visual tracking the rover approaches the target until it is within one meter of goal. The rover stops and plans a two-arc trajectory with a final rover orientation determined by the target's surface normal. The latter is computed from the range data of the target area. The rover drives along the two-arc trajectory and stops in front of the target. 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 the surface normal 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.

V. Experimental Results

A. System & Computing Architecture of Rocky 7

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 [11] (Figure 7). *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. 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 stereo camera pair is mounted on the mast, and two wide field-of-view stereo camera pairs are mounted on the front and back sides of the vehicle 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 body 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 backplane 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). 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 subsystems such as the vision, manipulation, and 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.

D. Results

We have performed several experiments in JPL's Mars Yard and successfully demonstrated the acquisition of small rocks (3-5 cm) located over 1 meter in front of the rover. 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.

We have reported earlier results of the visual tracking in [4]. Many experiments were run, and 14 complete image/odometry datasets were collected. When run over these datasets, the visual tracker succeeded in maintaining target lock through 10 complete sequences. All but one of the failures were corrected by simply re-running the visual tracker with more appropriate parameters.

After some enhancements to the visual tracker and rock grasping strategy, 25 new trials were performed 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 the same color as background.
- 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 positional uncertainty.

VI. Future Work

We plan to introduce obstacle avoidance in the trajectory planning and optimal path selection. We will also investigate three-arc trajectories and compare them to the two-arc trajectories. We are planning to improve the robustness of the visual tracking algorithm by matching the entire shape of the terrain around the target and using visual feature tracking on the whole scene to enhance pose estimation[5]. 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 will be implementing continuous operation of the rover while images are being acquired and processed.

VII. 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 spa ce 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] 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.

[10] 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.

[11] 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.

[12] 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.

[13] Yalin Xiong and Larry Matthies. Error analysis of a real-time stereo system. In *Computer Vision and Pattern Recognition*, pages 1087–1093, 1997. http://www.cs.cmu.edu/~yx/ papers/StereoError97.pdf.