Proceedings of the 1999 IEEE
International Conference on Robotics & Automation
Detroit, Michigan • May 1999

Rapid Physics-Based Rough-Terrain Rover Planning with Sensor and Control Uncertainty

Karl Iagnemma *Frank Genot Steven Dubowsky

Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA, 02215 U.S.A.

* Institut National de Recherche en
Informatique et Automatique
INRIA, France

Abstract

In future planetary exploration missions, rovers will be required to autonomously traverse challenging environments. Much of the previous work in robot motion planning cannot be successfully applied to the rough-terrain planning problem. A model-based planning method is presented in this paper that is computationally efficient and takes into account uncertainty in the robot model, terrain model, range sensor data, and rover path-following errors. It is based on rapid path planning through the visible terrain map with a simple graph-search algorithm, followed by a physics-based evaluation of the path with a rover model. Simulation results are presented which demonstrate the method's effectiveness.

1 Introduction and Related Work

Current paradigms for planetary exploration are strongly dependent on the use of small, capable mobile robots ("rovers") to seek out and retrieve scientific information [7]. The recent Pathfinder mission is a successful example of the capabilities of such systems [11]. However, the scope of the Pathfinder mission was limited to short traverses under relatively close human supervision. In future missions rovers will be required to negotiate more challenging terrains, with limited human supervision [7]. In a representative future mission command cycle, a rover might be expected to obtain mid-range (i.e. 5-10 rover length) terrain data via a mast-mounted ranging system, identify a scientific "goal" or traversal waypoint, and autonomously plan and negotiate a path through challenging terrain. In this paper, the problem of autonomous rover rough-terrain path planning is addressed. A rapid path-planning technique which considers uncertainty sources relevant to the rough-terrain planning problem is presented.

Many of the "traditional" robot motion planning methods cannot be successfully applied to the rough-terrain planning problem due to their inherent characteristics. Many planning methods ignore vehicle kinematics and dynamics, assume perfect knowledge of the environment, and represent obstacles and free space in a binary format [10]. Additionally, many traditional planning methods are relatively computationally inefficient. These factors are critical to rough-terrain planning for several reasons. First, in rough terrain the concept of an obstacle is not clearly defined, as it depends on an understanding of the terrain and the mobility characteristics of the rover. Second, terrain data cannot be assumed to be perfectly known, due to errors in range sensing techniques [8, 13]. Third, the planned path may not be faithfully followed by the rover due to path-following errors [20]. Finally, planetary exploration systems will generally have limited computational resources to devote to path planning.

Some researchers have begun addressing the rough-terrain planning problem. First works were dedicated to the problem of computing time-optimal paths through rough terrain [18]. Researchers have developed dynamic vehicle models and utilized them to examine the feasibility of proposed paths [2, 9, 15]. Employing a rover model to evaluate "traversability" at a large number of points in the configuration space of the rover's position and heading has also been proposed in [4, 19]. A search algorithm is then used to find the optimal path through the terrain. Related work has focused on the uncertainty present in range-sensor data and rover navigation error [8, 13, 20].

With the exception of [6], most proposed planning methods do not consider uncertainty. In rough terrain, failure to account for uncertainty can lead to mission failure, an unacceptable result. A strong argument can be made that for rough-terrain rover planning, it is better to plan a safe path than an "optimal" one (i.e. one that optimizes a criteria, such as path length, but causes the rover undue risk).

The planning method presented in this paper is unique in that it considers terrain, modeling and path-following uncertainty, and also utilizes both kinematic and force analyses of rover-terrain interaction. The method is composed of two steps. The first step uses the A* algorithm and an estimate of the local terrain roughness to rapidly plan a path through the range map from the rover start position to the goal position [14]. Terrain roughness is defined with respect to rover kinematic parameters.

0-7803-5180-0-5/99 $10.00 © 1999 IEEE

The second step is a rigorous evaluation of the proposed path using a physical model of the rover. Uncertainty in the terrain data, soil/tire interaction model, rover model, and rover path following are considered. If the model-based evaluation determines that the rover could potentially be subject to undue risk along the proposed path, the A* cost function is increased at the high-risk location, and the path is replanned. Simulation results are presented using a rover model based on the NASA LSR-1 rover and Mars-analog terrain, and it is shown that the method is able to effectively plan safe routes through rough terrain [17].

2 Rough Terrain Planning

2.1 The Rapid Search

The goal of the rapid search is to quickly find a direct, reasonable path from the rover start position to the goal. The rapid search utilizes range map data, such as would be obtainable from a stereo-vision or laser-based range sensor. The terrain is represented as a map of elevations z associated with a regularly-spaced grid in (x,y).

The A* algorithm is used to find a path through the terrain grid from the start position to the goal [14]. A* is a graph-search technique, and is attractive due to its optimality and high speed for relatively small graphs (i.e. approximately 10^4 cells in this work). An optimal path is computed based on a cost function, f.

The cost function is formed to consider three factors. Path length, l, is considered since it is desirable for the rover to travel as short a distance as reasonably possible. Terrain "unevenness", u, is considered since unevenness is related to traversability and rover safety. Finally, rover turning action, t, is considered, since in rough terrain excessive turning may not always be desirable or possible. Since the cost function must be evaluated a large number of times, it is desirable that it be mathematically simple to speed computation.

To define "unevenness," the spatial dimensions of the rover must be considered. If the rover is moving from a point (x,y) it will have soil/tire contact at some points $(x\pm\alpha, y\pm\beta)$. To account for this, first let R be the set of points included in a rectangle of dimensions w' x l' centered at (x,y) on the path and oriented in the direction of rover motion along the path, where w' and l' are defined by the convex hull of the soil/tire contact points of the rover on flat ground. The terrain roughness S at (x,y) is defined as the variance of z of all points in R:

$$S(x,y) = \mathrm{var}\sqrt{(z(R))} \tag{1}$$

Although soil/tire contact occurs only at known discrete terrain points, it is important to include the "interior" terrain (i.e. within the rover footprint), as this provides an estimate of potential rover hang-up failure.

Knowledge of the physics of the system is then used to define unevenness, u. For the class of rovers considered here (i.e. rovers with rocker-bogie suspensions), it is known based on prior kinematic studies that it is generally possible to traverse obstacles that are approximately one wheel diameter in size, see Figure 1 [5]. Larger obstacles could lead to hang-up or instability. The function u is therefore defined to increasingly penalize arcs which violate this heuristic rule, as:

$$u(x,y) = \left(\frac{S}{d}\right)^{\alpha} \tag{2}$$

where d is the rover wheel diameter, and α is positive. In this work α=3. Terrain data mesh size is not considered in this function, but is assumed to be well scaled to the rover system.

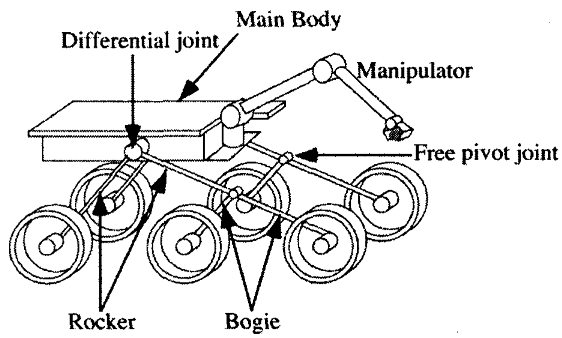

Figure 1. A rocker-bogie class rover

The cost t associated with rover turning is defined assuming that skid steering is utilized as the turning mode. On flat terrain, skid steering allows point turns. On rough terrain, however, it may not be possible to skid-steer due to high terrain-induced transversal forces. Thus, a metric based on terrain unevenness is again required. The function t is defined as:

$$t(x,y) = \left(\frac{\mathrm{var}(z(R'))}{d/2}\right)^{\alpha} \tag{3}$$

where var(R') is the variance of the points included in the set R'. R' is defined as the set of points included in a circle of radius $\sqrt{w'^2 + l'^2}$, and can be viewed as a superset of R augmented to include the area swept during turning. The factor $d/2$ is used as a scaling factor and is chosen based on *a priori* system knowledge [5]. Note that t is invoked only when the path involves a turn at (x,y).

The final cost function f is formed as a weighted sum of the functions considering terrain unevenness, turning, and path length. Weighting factors are chosen to normalize l, u, and t for a nominal rover motion and can be adjusted to address various mission constraints.

A least-cost path is found from the start point to the goal using the A* algorithm. The algorithm returns a path which combines short distance, terrain smoothness, and minimum turning requirement. The path is not guaranteed to be free of hazards, due to the heuristic nature of the cost function. Again, the cost function was selected due to its ease of evaluation and intuitive elationship to rover mobility. It is intended to rapidly lead to a reasonable path through the terrain.

2.2 Model-Based Evaluation

Since the path generated by the rapid planner is not guaranteed to be free from hazard, a more rigorous evaluation of the proposed path is applied in order to ensure rover safety. It uses a physical model of the rover and terrain to evaluate the traversability of the path. Rover stability, wheel slip, and actuator saturation are evaluated at discrete configurations along the path. Since the rover moves slowly, dynamic effects are neglected.

If failure (or unacceptable risk) is detected at any point along the path, the cost at the failure point is increased, and the rapid planner is called to re-plan the path.

Rover Model and Stability Analysis

To evaluate stability along the proposed path, the rover position, attitude, and configuration must be computed at closely-spaced discrete locations. The rover state can be defined by ten parameters: three for the position of its center of mass, $(x,y,z)_b$, three for its attitude, (Θ, Φ, Ψ), and four for the configuration of the right and left side rocker-bogie mechanisms, $(\theta_1, \theta_2, \theta_3, \theta_4)$. Because of the differential in the rocker-bogie mechanism, body pitch Φ is related to θ_1 and θ_3 (see Figure 2). Therefore, nine independent parameters are sufficient to describe the system on any terrain. Taking the heading of the rover Ψ and the location of the right middle wheel $(x,y)_{mr}$ as inputs from the rapid planner, the six remaining unknown quantities must be computed as a function of the terrain. Taking as unknown the rover wheel elevation z_{mr}, body roll Θ, and configuration $(\theta_1, \theta_2, \theta_3, \theta_4)$, six kinematic loop closure equations for a rocker-bogie rover can be written to solve for these six unknowns, as [5]:

$$z_{mr} = z_{mr}(x_{mr}, y_{mr}, \Theta) \tag{4}$$

$$z_{mr} - z_{fr} = \cos(\Theta) \cdot (l_3 \cdot \sin(\theta_2) - l_4 \cdot \cos(\theta_2)) \tag{5}$$

$$z_{mr} - z_{rr} = \cos(\Theta) \cdot (l_3 \cdot \sin(\theta_2) + l_2 \cdot \cos(\theta_1) + l_1 \cdot \sin(\theta_1)) \tag{6}$$

$$z_{mr} - z_{rl} = \cos(\Theta) \cdot (l_3 \cdot \sin(\theta_2) + l_2 \cdot \cos(\theta_1) - l_1 \cdot \sin(\theta_2) + w \cdot \sin(\Theta)) \tag{7}$$

$$z_{mr} - z_{ml} = \cos(\Theta) \cdot (l_3 \cdot \sin(\theta_2) + l_2 \cdot \cos(\theta_1) - l_2 \cdot \cos(\theta_3) - l_3 \sin(\theta_4)) + w \cdot \sin(\Theta) \tag{8}$$

$$z_{mr} - z_{fl} = \cos(\Theta) \cdot (l_3 \cdot \sin(\theta_2) + l_2 \cdot \cos(\theta_1) \cdot l_2 \cdot \cos(\theta_3) - l_4 \cos(\theta_4)) + w \cdot \sin(\Theta) \tag{9}$$

where the subscripts *f*, *m*, and *r* refer to the front, middle, and rear wheels, *r* and *l* refer to the right and left sides of the vehicle, respectively, and Θ is the body roll.

These equations are highly nonlinear, and convergence of a numerical solution is not guaranteed. The problem can be simplified and solved by estimating the rover roll based on the ground profile, and treating each side of the rover as a planar problem [5]. The roll estimate is then corrected. With this formulation, the rover state can be computed with minimal error.

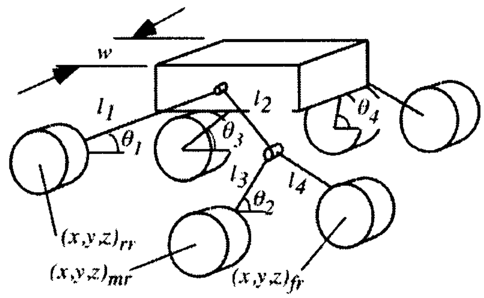

Figure 2. Rover kinematic notation

The rover state is computed at discrete locations along the proposed path. To verify that rover configuration constraints are not violated, at each discrete location the rocker and bogie angles are required to remain in pre-specified joint limits:

$$\theta_{i\,\min} < \theta_i < \theta_{i\,\max} \tag{10}$$

If these configuration constraints are violated, failure is declared, and the cost at the failure point is increased.

Defining stability statically via the force-angle representation, rover stability is considered along the path [16]. As shown in Figure 3, rover wheel/ground contact points are denoted as p_i, and the line which connects points p_i and p_{i+1} is defined as a_i. Letting η_i be the angle between the vertical direction and the line orthogonal to a_i containing the center of mass of the system, the system will be unstable when η_i goes to zero. The stability margin α is then defined as:

$$\alpha = \min\left(\eta_i / \eta_{i\,nom} \right) \tag{11}$$

where $\eta_{i\,nom}$ is the angle when the rover is on flat ground. The value of α is used to determine plan failure. At every point along the path, the condition:

$$\alpha > \alpha_{cr} \tag{12}$$

must hold, where α_{cr} is a user-specified minimum stability ratio. If this stability condition is violated, failure is declared, and the cost at the failure point is increased.

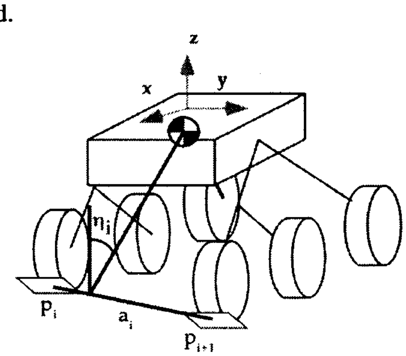

Figure 3. Rover stability criterion

Force Analysis

A quasi-static force analysis is performed at each configuration along the proposed path to determine if the rover is able to move in its desired direction subject to wheel slip and actuator saturation constraints.

A three-dimensional force analysis yields the forces and moments F_x, F_y, and M_y acting on the rocker joint. The force analysis can then be divided into two planar problems (see Figure 4). The resulting equations of static equilibrium are [4,5]:

$$\sum \mathbf{F} \cdot \mathbf{u}_x = \sum_{i=1}^{i=3} (\mathbf{T}_i + \mathbf{N}_i) \cdot \mathbf{u}_x + F_x = 0 \quad (13)$$

$$\sum \mathbf{F} \cdot \mathbf{u}_z = \sum_{i=1}^{i=3} (\mathbf{T}_i + \mathbf{N}_i) \cdot \mathbf{u}_z - F_z = 0 \quad (14)$$

$$\sum \mathbf{M}_A = \mathbf{T}_1 \times \overline{\mathbf{RA}} + \mathbf{N}_1 \times \overline{\mathbf{RA}} + \mathbf{T}_2 \times \overline{\mathbf{MA}} + \mathbf{N}_2 \times \overline{\mathbf{MA}} + \mathbf{T}_3 \times \overline{\mathbf{FA}} + \mathbf{N}_3 \times \overline{\mathbf{FA}} + \mathbf{M}_y = \mathbf{0} \quad (15)$$

$$\sum \mathbf{M}_B = \mathbf{T}_2 \times \overline{\mathbf{MB}} + \mathbf{N}_2 \times \overline{\mathbf{MB}} + \mathbf{T}_3 \times \overline{\mathbf{FB}} + \mathbf{N}_3 \times \overline{\mathbf{FB}} = \mathbf{0} \quad (16)$$

These underdetermined set of equations can be solved by taking two wheel input torques as free variables and optimizing with respect to an appropriate criteria (e.g. power consumption) [5].

If the results of this analysis predict that the rover cannot move in the desired direction, or that the rover will slide in an unintended direction, failure is declared and the cost at the failure point is increased.

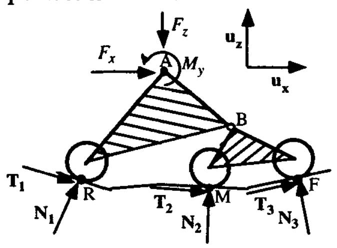

Figure 4. Planar rover force analysis

If none of the stability, configuration, or force requirements are violated, the proposed path is deemed safe, and the planning problem is considered solved.

3 Incorporating Uncertainty in the Planning Method

There are several important sources of uncertainty which should be considered in a rough-terrain path planning algorithm. These include uncertainties in terrain measurement, the soil/tire model, the vehicle model, and rover path following. In this section methods for incorporating uncertainty estimates into the planning method are described.

3.1 Terrain and Rover Model Uncertainty

The stability and force analyses presented in Section 2.2 assumed perfect knowledge of the terrain geometry, soil/tire model, and rover model. In practice, substantial uncertainty will be present. Range sensor systems include error due to range distance, noise, and systematic errors such as miscalibration [8, 13]. Terrain characteristics may not be well known, which will lead to erroneous soil/tire interaction models. Rover physical modeling error can be caused by idealizations or simplifications. Additionally, the acquisition of science samples may strongly influence the rover mass distribution [17].

Rover Model and Stability Analysis

The rover configuration analysis described in Section 2.2 includes a requirement that the rover joint angles stay within a prescribed range of motion. This requirement is defined in Inequality (10), and can be influenced by terrain uncertainty due to range sensor error.

Assuming that the actual location of the ground contact point at each wheel lies inside a sphere of radius R_s about the assumed contact point, the worst-case kinematic uncertainty, θ_{uncer}, based on the worst-case error in the ground contact location can be computed:

$$\theta_{i\ uncer} = \sin^{-1}(R_s / l_i) \quad (17)$$

where R_s is a function of the assumed range sensor uncertainty, and can be a function of distance. This can be conservatively incorporated into Inequality (10):

$$(\theta_{i\ \min} + \theta_{i\ uncer}) < \theta_i < (\theta_{i\ \max} - \theta_{i\ uncer}) \quad (18)$$

The stability metric described in Section 2.2 depends on accurate knowledge of the location of the soil/tire contact points and the rover center of mass, which is a function of vehicle model uncertainty. The stability analysis with uncertainty is simplified by assuming that the true center of mass lies within a sphere of radius R_m about the assumed location. This radius is a function of nominal parameter modeling uncertainty and, if the rover will acquire samples, the ratio of the expected sample mass to the rover mass. Again it is assumed that the location of the ground contact point at each wheel lies in a sphere of radius R_s about the assumed point.

The worst-case stability uncertainty, η_{uncer}, can be accurately estimated based on the worst-case error in center of mass and ground contact location, as:

$$\eta_{uncer} \approx \max\left(\sin^{-1}\left((R_m + R_s) / \left\| (I - \hat{a}_i \hat{a}_i^T) P_{i+1} \right\| \right)\right) \quad (19)$$

with I as the 3x3 identity matrix and $\hat{a}_i = a_i / \|a_i\|$.

This uncertainty can then be conservatively incorporated into Inequality (12) by modifying Equation (11) as:

$$\alpha = \min\left(\frac{\eta_i - \eta_{uncer}}{\eta_{i\ nom}}\right) \quad (20)$$

The configuration and stability uncertainties in Equations (18) and (20) are worst-case estimates and are therefore highly conservative. Worst-case estimates are employed due to the extreme cost associated with mission failure.

Note that the uncertainties can be pre-computed. Thus, there is negligible additional computational cost.

Force Analysis

The force analysis described in Section 2.2 relies on accurate knowledge of the location of the center of mass of the rover, and more importantly on knowledge of the soil/tire interaction characteristics. Various models have been developed for soil/tire interaction, all of which agree that soil/tire interaction forces can vary strongly with terrain properties and wheel slip [1]. It has also been found that terrain on the Martian surface varies significantly [12].

It is difficult to accurately measure or estimate soil/tire interaction properties. However, given knowledge of the type of terrain the rover might encounter (i.e. rock, loose silt, dense soil) we can define a set of tractive force coefficients for a nominal wheel slip ratio on each expected terrain type. For example:

$$M_{dense\ soil} = \{\mu_{\min}, \ldots, \mu_{\max}\} \quad (21)$$

In this work, the model-based evaluation selects the force coefficient for an assumed terrain type conservatively, as:

$$\mu_{model} = \min(M) \quad (22)$$

Current work to improve this estimate includes investigation of on-line estimation of soil/tire interaction properties, and terrain property learning based on correlation of visual cues and rover behavior.

3.2 Rover Path Following Uncertainty

The rapid planner generates a least-cost path from the rover start location to the goal location. Implicit in this approach is the assumption that the rover will follow the proposed path. In practice, rover path following error is present [20]. Failure to account for this uncertainty source in the planning method could lead to proposed paths which are dangerously close to obstacles.

Path following error is considered in both the rapid planning step and the model-based evaluation. In the rapid planning step, a pre-processing routine replaces each terrain point by the output of a gaussian filter, centered at the point with size proportional to the assumed spatial path following error P_f. This formulation effectively "blurs" the terrain (and, therefore, potential obstacles) in a region proportional to the estimated path-following error. Thus, proposed motion plans will include a safety margin proportional to P_f.

Rover path-following error is also accounted for in the model-based evaluation. In the model-based evaluation, the stability and force analyses described in Section 2.2 are applied at discrete configurations along the proposed path. Using the uncertainty estimates P_f of path following error and P_h of heading error, the set of configurations that are examined with the model-based evaluation is expanded. For a configuration (x_{mr}, y_{mr}, Ψ) on the proposed path, all configurations included inside the convex hull $(x_{mr} \pm P_f, y_{mr} \pm P_f, \Psi \pm P_h)$ are also examined. The model-based evaluation is then applied at equally-spaced discrete points in this space. The resolution of evaluation is chosen equal to the spatial resolution of the terrain image data. The angular resolution is chosen similarly.

As in Section 2.2, if any configuration, stability, or force conditions are violated, failure is declared. The cost at the failure point is increased, and the path is re-planned with the rapid planner.

4 Application of the Planning Method

4.1 Rover Model

The rover system considered in this application is a scale model of the LSR-1 Rover [17]. This rover has been modeled extensively and a prototype has recently been built in the Field and Space Robotics Laboratory [5]. The rover has 9 cm diameter wheels and is approximately 35 cm in length, 19 cm wide, and 34 cm high. This design has six independently driven wheels mounted on a rocker-bogie articulated frame (see Figure 1). With this design each wheel tends to remain in contact with the ground while on uneven terrain and the weight of the rover is well distributed over the six wheels. This allows each wheel to develop good ground traction and results in a highly mobile rover.

4.2 Terrain Model

The planning method was examined on a variety of challenging terrains, including such features as hills and gullies several wheel diameters in height (see Figure 5 for a contour map). The spatial terrain data resolution for all trials was 1.25 cm. The visible range of the rover was assumed to be approximately 8 rover lengths. This resulted in a range-data map of 250x200 points.

The terrain was modeled with a force coefficient that differed depending on the type of terrain. Force coefficients for sandy and rocky areas were assumed as 0.5 and 0.6, respectively.

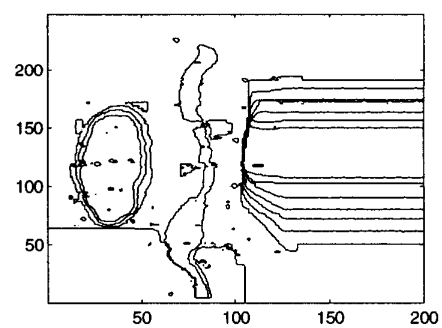

Figure 5. Representative terrain

4.3 Simulation Results

The planning method was simulated on a IBM Pentium 300 MHz PC. All algorithms were written in C. Numerous trials were conducted. For a representative set of 20 trials, it was found that the rapid search generated a safe path on the first attempt 12 times. This result implies that the cost function is a reasonable indication of terrain difficulty, but is clearly insufficient as a stand-alone path planner. The model-based evaluation must be utilized to verify path safety.

It was also observed that the cost function is sensitive to parameter changes. Trials were run with α varied 25% and the cost fucntion weighting factors varied 15%. With these variations, it was found that the rapid search generated a safe path on the first attempt 9 out of 20 times. This lack of robustness again highlights the necessity of the model-based evaluation.

The rough-terrain planning method was compared to a purely heuristic planning method that did not utilize a rover physical model and treated potential obstacles in a binary manner. In the heuristic method, any neighboring *(x,y)* points which had elevation changes greater than one wheel diameter were deemed impassable obstacles. Neighboring *(x,y)* points which had elevation changes less than one wheel diameter were deemed free.

The heuristic planning method returned paths to 12 of the 20 goal locations. Eight goal locations were deemed unreachable, and the planner returned no solution. The 12 returned paths were then examined with the model-based evaluation. Two of these "safe" paths were determined to place the rover in undue risk.

In contrast, the rough-terrain planner was able to successfully plan and verify a path to all 20 goal locations. Eighteen of the 20 paths required two or fewer planning iterations. All paths were successfully planned within 3 iterations. Thus, it can be claimed that for the simulated terrain, the rough-terrain planning method allowed access to 66% more goal locations than the heuristic method.

Computation time results for the 20 trials including path re-planning are presented in Table I. The planning method is rapidly able to generate and evaluate motion plans through difficult terrain.

Table I. Average computation time for 20 trials

	Time (sec)
Rapid Search	91.3
Model-Based Evaluation	83.1
Total:	174.4

5 Conclusions and Future Work

A rapid physics-based planning method for rough-terrain rovers was presented. The method utilizes range sensor data to plan a route through the visible terrain. This route is then evaluated rigorously with a rover physical model. Uncertainty in terrain range data, terrain modeling, rover modeling, and rover path-following are accounted for in the method. Robustness is ensured by conservatively accounting for unknown or difficult-to-model parameters. A simulation has been developed which demonstrates the potential effectiveness of the method.

We are currently implementing this approach on the prototype rover described in Section 4.1. The rover will operate in a Mars-analog environment in the MIT Field and Space Robotics Laboratory.

Future work will focus on incorporation of methods for dealing with unknown or partially-known terrain data. In these situations, it may be possible to "infer" terrain data based on gross terrain knowledge or visual cues.

We are also looking into incorporating a local re-planning algorithm into the method, for cases where the model-based evaluation detects possible failure. This may lead to more efficient plan generation.

References

[1] Bekker, M., *Introduction to Terrain-Vehicle Systems*, University of Michigan Press, 1969

[2] Ben Amar, F., and Bidaud, P., "Dynamic Analysis of Off-Road Vehicles," *Proceedings of Fourth International Symposium on Experimental Robotics*, pp. 363-71, 1995

[3] Cherif, M., and Laugier, C., "Dealing with Vehicle/Terrain Interactions when Planning the Motions of a Rover," *IEEE International Conference on Robotics and Automation*, pp. 579-586, 1994

[4] Farritor, S., Hacot, H., and Dubowsky, S., "Physics-Based Planning for Planetary Exploration," *IEEE Internatl. Conf. on Robotics and Automation*, pp. 278-83, 1998

[5] Hacot, H., *Analysis and Traction Control of a Rocker-Bogie Planetary Rover*, M.S. Thesis, Mass. Institute of Technology, Cambridge, MA, 1998

[6] Hait, A., and Simeon, T., "Motion Planning on Rough Terrain for an Articulated Vehicle in Presence of Uncertanties," *IEEE/RSJ International Symposium on Intelligent Robots and Systems*, pp. 1126-1133, 1996

[7] Hayati, S., Volpe, R., Backes, P., Balaram, J., and Welch, W., "Microrover Research for Exploration of Mars," *AIAA Forum on Advc. Developments in Space Robotics*, 1996

[8] Hebert, M., and Krotkov, E., "3D Measurements from Imaging Laser Radars," *Image and Vision Computing*, **10**, 3, pp. 170-178, 1992

[9] Kelly, A., and Stentz, A., "Rough Terrain Autonomous Mobility—Part 2: An Active Vision, Predictive Control Approach," *Autonomous Robots*, **5**, pp. 163-198, 1998

[10] Latombe, J., *Robot Motion Planning*, Kluwer, 1991

[11] Mars Pathfinder Website, http://mars.jpl.nasa.gov/ index1

[12] Matijevic, J., *et. al.*, "Characterization of Martian Surface Deposits by the Mars Pathfinder Rover, Sojourner," *Science*, **278**, 5, pp. 1765-1768, 1997

[13] Matthies, L., and Grandjean, P., "Stochastic Performance Modeling and Evaluation of Obstacle Detectability with Imaging Range Sensors," *IEEE Trans. on Robotics & Automation*, **10**, 6, pp. 783-92, 1994

[14] Nilsson, N., *Principles of Artificial Intelligence*, Tioga Publishing Company, 1980

[15] Olin, K., and Tseng, D., "Autonomous Cross-Country Navigation," *IEEE Expert*, **6**, 4, pp. 16-30, 1991

[16] Papadopolous, E., and Rey, D., "A New Measure of Tipover Stability Margin for Mobile Manipulators," *IEEE Internatl. Conf. on Robotics and Auto.*, pp. 3111-16, 1996

[17] Schenker, P., Sword, L., Ganino, A., Bickler, D., Hickey, G., Brown, D., Baumgartner, E., Matthies, L., Wilcox, B., Balch, T., Aghazarian, H., and Garrett, M., "Lightweight Rovers for Mars Science Exploration and Sample Return," *Proceedings of SPIE XVI Intelligent Robots and Computer Vision Conference*, **3208**, pp. 24-36, 1997

[18] Schiller, Z., and Chen, J., "Optimal Motion Planning of Autonomous Vehicles in 3-Dimensional Terrains," *IEEE Internatl. Conf. on Robotics and Auto.*, pp. 198-203, 1990

[19] Siméon, T., and Dacre-Wright, B., "A Practical Motion Planner for All-terrain Mobile Robots," *IEEE Internatl. Conf. on Intel. Robots and Systems*, pp. 1357-1363, 1993

[20] Volpe, R., "Navigation Results from Desert Field Tests of the Rocky 7 Mars Rover Prototype," submitted to *International Journal of Robotics Research*, 1998