Physics-Based Planning For Planetary Exploration

Shane Farritor, Hervé Hacot, Steven Dubowsky
Department of Mechanical Engineering, Massachusetts Institute of Technology
Cambridge, MA 02139 U.S.A.

Abstract

Recently a planetary rover returned important scientific information from Mars. More ambitious missions are planned. New planning methods are required that allow rovers to explore challenging areas with a high level of autonomy. This paper presents a planning methodology based on a physics-based model of the rover and environment. Plans are developed that allow a rover to perform a mission while explicitly considering constraints such as power, actuator, wheel slip, and vehicle stability limits. Results obtained from detailed rover simulations are presented.

1. Introduction

The Pathfinder spacecraft landed on Mars on July 4, 1997 with the Sojourner rover. This mission has returned important information on Martian climate and geology [1]. Additional missions are planned for 2001, 2003, and 2005. Mobile rovers will be important components of these missions. They may need to travel several kilometers over periods of months and manipulate rock and soil samples. They will need to be somewhat autonomous [2].

Areas of important scientific interest may be difficult to reach, such as ravines and craters believed to be dry river beds or dry lakes, or steep cliffs that may contain fossils in stratified rock layers. Designing and controlling a rover to explore these areas without becoming hung-up or trapped is challenging. A current approach is the planning and control scheme used for the Sojourner Mars rover [3,4]. This method is reminiscent of behavior control [5]. It requires the operator to provide closely spaced task goals. Then the action of the rover is derived from a set of achieving behaviors. The robot's actions tend to be reactive in that it quickly responds to sensor inputs, but it does not tend to plan actions far in the future. The approach does not explicitly consider the physics of the rover or the task. The scheme has been shown to be successful in low density obstacle fields [3].

Planning methods are needed to allow rovers to safely explore challenging areas with a high level of autonomy. Much work has focused on the planning problem for mobile robots [5,6,7]. The advantages of these approaches is that they do not require an extensive model of the robot. This paper presents a planning methodology that attempts to exploit a physical model of the rover and environment. Mission plans are developed that consider such constraints as power consumption, actuator saturation, wheel slip, vehicle stability and vehicle safety. In addition, in very difficult terrain, the models developed could assist a human operator. This paper describes the methodology, the physics-based rover model it uses and some results obtained on simulated rover tasks.

The LSR-1 (Lightweight and Survivable Rover), a prototype vehicle built by NASA [8], as the basis for the evaluation of the algorithm developed, see Figure 1.

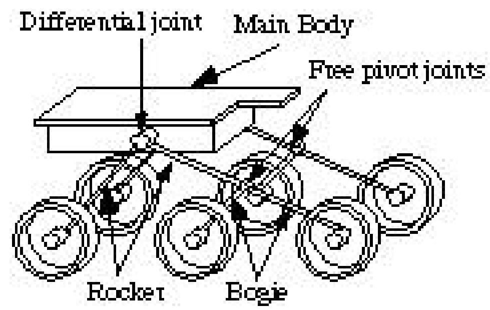

Figure 1: The LSR-1 Rover [8]

2. The Action Planning Approach

The physics-based planning methodology presented in this paper was first introduced for highly idealized systems [9,10]. It will be briefly reviewed below. In section 4 it is applied to the rover problem.

The method divides all possible activities of the robot into discrete actions, or action modules. The collection of all possible action modules is referred to as the action module inventory. The inventory used for the tasks described in this paper is shown in Table 1. The actions modules are assembled into a sequence or action plan. A successful action plan is one that allows the robot to complete the task without violating any of the physical constraints of the robot or the task. In our rover problem these constraints can include actuator saturation, static stability, energy consumption, kinematic constraints, obstacle avoidance, etc. These constraints are verified using physics-based models of the rover and environment. In this way the development of a successful plan becomes a search for the correct sequence of action modules. The size of the search space is given by equation 1.

$$D = N^m \qquad (1)$$

Where: D = number of possible action plans
N = number of possible action modules
m = number of action modules used in a plan

For an action module inventory of reasonable size, and an action plan of reasonable length, the number of possible action plans is very large making a direct application of any conventional search method difficult. For example, a task using a plan which is built from an inventory of only 9 action modules and which is only 40 action modules long leads to over 10^{38} possible action plans. Also, for challenging tasks, such as rovers in difficult terrain, the number of feasible plans is not large. This further complicates the search for a successful solution.

Table 1. Action Module Inventory

Mod. #	Action	Mod. #	Action
101	move forward 10cm	1001	move arm +x 2cm
102	move back 10cm	1002	move arm -x 2cm
103	turn left 5°	1003	move arm +y 2cm
104	turn right 5°	1004	move arm -y 2cm
105	circle (r=50cm) left 10cm	1005	move arm -z 2cm
106	circle (r=200cm) right 10cm	1006	move arm +z 2cm
107	circle (r=50cm) right 10cm	3000	turn to sample
108	circle (r=200cm) right 10cm	3001	turn from sample
901	transmit video	3002	grasp sample
902	end video	3003	release sample
910	sweep sensors	2020	retract science tray
920	sleep	2010	deploy science tray
921	communicate with ground station	4000	recharge batteries with solar array

To search this space effectively and practically, a hierarchical selection process that includes a genetic algorithm (GA) is used. First, filters using simple tests, based on the fundamental physics of the task, are used. They remove superfluous elements from the inventory before more complex evaluation procedures, such as detailed simulation, are applied. A small reduction in the inventory results in a large reduction in the number of possible action plans. For example, if a task does not require manipulation, the 10 action modules that pertain to the manipulator can be removed from the inventory shown in Table 1. Then the reduced inventory can be searched by the GA to produce a final action plan.

A genetic algorithm is a search process based on biological evolutionary heuristics such as selection, crossover, and mutation [11]. An initial population is created from random assemblies of modules from the reduced inventory. A single-point crossover operator is used, see Figure 2. It combines some attributes of two action plans to create a third and fourth unique action plan. Two modules (C1 & C2) are chosen within two action plans. Then the module and all the remaining modules (B and D) in each list are exchanged. In order to maintain solution diversity, mutation is used. From time to time a random module in a single action plan is exchanged with a random module chosen from the reduced inventory. The quality, or fitness, of an action plan is then evaluated. Here an analytical model, described in section 3, is used. The analysis determines if the action plan allows the robot to accomplish the task without violating any physical or task constraints. The analysis executes the plan until the task is completed or a constraint is violated. Then a fitness value is assigned to the action plan. Plans that complete more of a task receive a higher fitness.

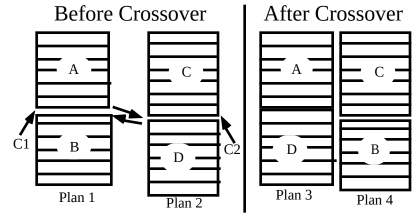

Figure 2: Genetic Crossover

The search process is formulated to find feasible solutions in spite of the large search space. The keys are the manner in which the genetic algorithm fitness function assigns a score, and how genetic crossover occurs. When a given action plan is evaluated, the plan is executed by the model (simulation) until a physical constraint is violated or the task is completed. The beginning portion of the plan, that which is successful, determines the "partial credit" of the action plan. In general a plan that can move the rover "one step closer" to completing the task receives higher fitness. The crossover is structured to ensure that the beginning or successful portion of each plan stays intact, and change occurs on the less successful portion. This is done by choosing a crossover point in a probability distribution with a mean value centered about the point where the plan fails. Therefore the planner is usually only considering how to improve the plan for the next few "steps" that move the robot along its task. The mutation aspect of the algorithm helps prevent this “short look ahead” approach from getting hung-up in dead-ends. The effect is that the planner builds on and improves action plans that are partially successful.

The effect of this approach is shown in Figure 3 for the rover problems discussed below. Here, crossover was allowed to occur at *any* point in the plans with equal probability. When crossover resulted in an improvement (increased fitness) of a plan, the change in the length of the successful portion of the plan, as compared to the length of the successful portion of the two original plans, was recorded. The results are plotted in a histogram in Figure 3. It is seen that crossovers resulting in improved fitness usually only increase the length of the successful portion of the plan fewer than 5 action modules. Hence the short look-ahead horizon.

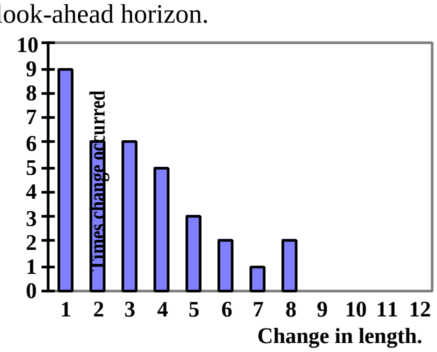

Figure 3. Histogram of Look-Ahead Horizon (5 runs).

It can be seen that the effect of limiting the search space is dramatic. For example, the search space of 10^{38} possible action plans discussed above is reduced to 10^5 when only the next 5 action modules are considered. This reduction allows the problem to be solved effectively and makes on-board computation feasible. Also, by choosing the crossover point with probability distribution, dramatic changes can occur in crossover. The helps ensure that the search is not locked into a particular region. This and other aspects of the approach such as learning and action module inventory design are discussed in [9,10].

3. Planetary Rover Physical Model

Sojourner and the LSR-1 are representative of a class of planetary rovers based on a rocker-bogie mobility design [12]. The LSR-1 vehicle has 20 cm diameter wheels and is approximately 100 cm in length, 70 cm wide, and 45 cm high. This design has six independently driven wheels mounted on an articulated frame, see Figure 1. The frame has two rocker arms connected to a main body. Each rocker has a rear wheel connected to one end and a secondary rocker, called a bogie, connected to the other. At each end of the bogie is a drive wheel and the bogie is connected to the rocker with a free pivoting joint. The rockers are connected to the main body with a differential so that the pitch angle of the body is the average of the pitch angles of the rockers. 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.

As discussed above, a physics-based analytical model of a rover and its environment is used to assign fitness to action plans. Some analysis of a rocker-bogie mechanism has been done in the design of the structure [13]. Since the analysis is used each time an action plan is evaluated, it is important to keep the model as computationally simple as possible while maintaining an acceptable level of fidelity. This is particularly important for potential on-board implementation. A key to this is to recognize that the practical constraints of space systems, such as weight and power, require rovers to travel at slow speeds, approximately 3 cm/s. Therefore, dynamic effects are small and a quasi-static model approximates its behavior.

Obviously, a rover's ability to negotiate a terrain is a function of the ground profile and surface characteristics. Rovers will have sensors such as stereo cameras and laser-based systems capable of producing maps of the terrain around the rover [2,14]. Estimates of soil characteristics will be more difficult, but are essential for any rational planning approach. On-line adaptive or identification methods, as well as pre-established estimates, may be used. This area remains an open research topic.

Inverse Kinematics. A model is required to determine the rover attitude and configuration as a function of the terrain. These are used to calculate the load distribution on the wheels, the rover's stability, actuator outputs, etc.

The rover's configuration, position and attitude can be fully defined by ten parameters: three for the position (x_B, y_B, z_B) in the fixed frame of the body, three for its attitude (ψ,Θ,φ) (yaw, roll and pitch angles respectively) and four for the configuration of the rocker-bogie mechanism (Θ_1 ,Θ_2. Θ_3 ,Θ_4), see Figure 4. Because of the differential in the rocker-bogie mechanism, φ, Θ_1 and Θ_3 are related. Therefore, nine independent parameters are sufficient to describe the system on any terrain.

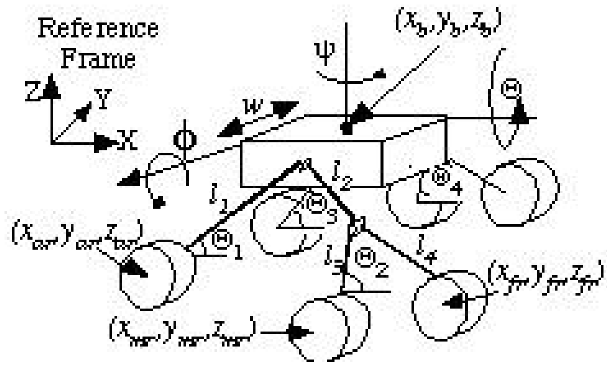

Figure 4: Kinematic parameters of the Rover

The desired heading angle (ψ) and the position of the middle right wheel (x_{mr},y_{mr}) in the ground reference frame are given as input. These variables are obtained from the action plan. With this input, six other independent parameters must be determined. We choose z_{mr}, Θ, Θ_1, Θ_2, Θ_3 and Θ_4 as unknowns. From the knowledge of all these parameters, x_b,y_b, and z_b can be easily derived. In this analysis, the ground profile is defined as a piecewise topographical map, that gives the altitude z_{center} of the center of a wheel as a function of its coordinates (x_{center}, y_{center}) in the ground reference frame and the vehicle roll, Θ. Then six loop closure equations that describe the vehicle attitude and configuration can be written as [15]:

$$z_{mr} = z_{center,mr} \quad (2)$$

$$z_{mr} - z_{center,fr} = \cos(\Theta)\ (l_3\ \sin(\Theta_2) - l_4\ \cos(\Theta_2)) \quad (3)$$

$$z_{mr} - z_{center,ar} = \cos(\Theta)\ (l_3\ \sin(\Theta_2) + l_2\ \cos(\Theta_1) - l_1\ \sin(\Theta_1)) \quad (4)$$

$$z_{mr} - z_{center,rl} = \cos(\Theta)\ (l_3\ \sin(\Theta_2) + l_2\ \cos(\Theta_1) - l_1\ \sin(\Theta_3)) + w\ \sin(\Theta) \quad (5)$$

$$z_{mr} - z_{center,ml} = \cos(\Theta)\ (l_3\ \sin(\Theta_2) + l_2\ \cos(\Theta_1) - l_2\ \cos(\Theta_3) - l_3 \sin(\Theta_4)) + w\ \sin(\Theta) \quad (6)$$

$$z_{mr} - z_{center,fl} = \cos(\Theta)\ (l_3\ \sin(\Theta_2) + l_2\ \cos(\Theta_1) - l_2\ \cos(\Theta_3) - l_4 \cos(\Theta_4)) + w\ \sin(\Theta) \quad (7)$$

where the subscripts *m*, *a* and *f* refer to the middle, aft and front wheels respectively while r and l refer to the right and left sides of the vehicle respectively. This gives six equations in six unknowns. This system of equations is highly non-linear. Numerical methods (steepest descent and Newton's method) have been applied to solving such problems. These methods often fail to converge.

The following assumptions simplify the problem. First, the model is broken into two planar problems. From the profile of the ground, Θ can be approximated. With this approximation, the planar problem for the right side can be solved (z_{mr} is given by equation (2), Θ_2 becomes the only unknown in equation (3), and Θ_1 becomes the only unknown in equation (4)). This is then repeated for the other side of the vehicle. Finally, Θ is corrected while maintaining the values for z_{mr}, Θ_1, Θ_2, Θ_3 and Θ_4. For a complete description of this approach see [15]. The average error, for the example tasks discussed below, between the wheel contact point and the ground profile is less 2% of the wheel diameter (.5 cm).

Force Analysis. A quasi-static force balance is then used to determine if the rover will tip, if the wheels will slip, if the rover will slide, the amount of energy consumed, actuator saturation, etc. This information is critical to determine the fitness of an action plan. The force analysis assumes that the ground cannot apply a large moment, in any direction, to the wheels. This assumption is valid for a wide range of designs and conditions. For example, wheels that have a crown will approximate this condition. This assumption prevents the three-dimensional force analysis from becoming highly statically indeterminate.

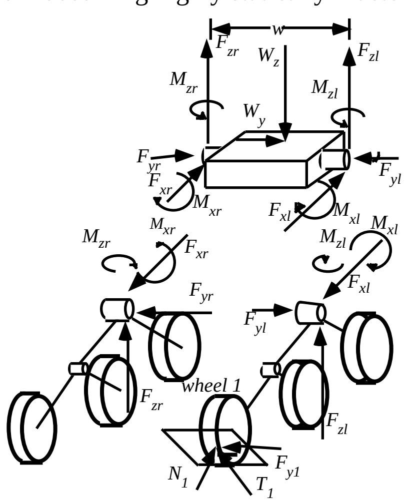

Figure 5: 3-D Force Analysis (in the Rover Frame)

The normal and the tangential forces between the i[th] wheel and ground are N_i and T_i respectively, see Figure 5. The transverse forces acting at the wheel contact point (F_{y1} to F_{y6}) will be relatively small, assuming that the vehicle roll angle is small. Further, assuming that each of these forces has the same magnitude permits the force analysis to be divided into a planar problem for each side of the vehicle (F_{yi} is known, hence so are M_{xl}, M_{xr}, M_{zl} and M_{zr}). The resulting equations of static equilibrium are:

$$F_{xr} + F_{xl} = 0 \qquad (8)$$

$$(F_{xr} - F_{xl})\frac{w}{2} + M_{zr} + M_{zl} = 0 \qquad (9)$$

$$F_{zr} + F_{zl} - W_z = 0 \qquad (10)$$

$$(F_{zl} - F_{zr})\frac{w}{2} + M_{xr} + M_{xl} = 0 \qquad (11)$$

F_{xr}, F_{yr}, M_{xr}, M_{zr} are the forces and moments acting between the rocker and the body on its right side. Again, the subscript l represents the left side and r the right and w is the width of the main body. The force balance on the body enables us to determine F_x, F_z, M_y that are applied by the body to each rocker. These forces and moments are then inputs to the planar problem, see Figure 6.

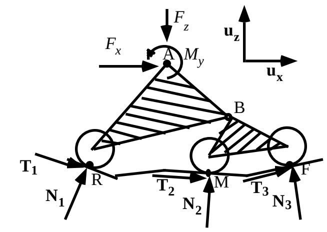

Figure 6: Planar Rover Force Analysis

The equations of equilibrium for this planar case are:

$$\sum \mathbf{F}\;\mathbf{u_x} = \sum_{i=1}^{i=3}\left(\mathbf{T_i} + \mathbf{N_i}\right)\;\mathbf{u_x} + F_x = 0 \qquad (12)$$

$$\sum \mathbf{F}\;\mathbf{u}_z = \sum_{i=1}^{i=3}\left(\mathbf{T_i} + \mathbf{N_i}\right)\;\mathbf{u}_z - F_z = 0 \qquad (13)$$

$$\sum \mathbf{M}_{\mathbf{y},body} = \mathbf{T}_1 \times \overline{\mathbf{RA}} + \mathbf{N}_1 \times \overline{\mathbf{RA}} + \mathbf{T}_2 \times \overline{\mathbf{MA}} + \mathbf{N}_2 \times \overline{\mathbf{MA}} + \mathbf{T}_3 \times \overline{\mathbf{FA}} + \mathbf{N}_3 \times \overline{\mathbf{FA}} + \mathbf{M}_y = \mathbf{0} \qquad (14)$$

$$\sum \mathbf{M}_{y,bogie} = \mathbf{T}_2 \times \overline{\mathbf{MB}} + \mathbf{N}_2 \times \overline{\mathbf{MB}} + \mathbf{T}_3 \times \overline{\mathbf{FB}} + \mathbf{N}_3 \times \overline{\mathbf{FB}} = \mathbf{0} \qquad (15)$$

These equations are solved with the following constraints on the wheel torques:

1) $\|\mathbf{T_i}\| \le \mu \|\mathbf{N_i}\|$ no slip i[th] wheel (or will just slip).
2) $N_i \ge 0$ insures i[th] wheel ground contact.
3) $\|\tau_{\mathbf{i}}\| \le \tau_{\max}$ i[th] wheel motor doesn't saturate.

Where τ_i is the i[th] wheel motor torque. τ_i is related to the wheel tangential force T_i by the wheel radius. This set of equations (12-15) is underdetermined (4 equations and 6 unknowns). Therefore two wheel torques on each side can be chosen arbitrarily (τ_1 and τ_2 for example). They are inputs to the system and the normal forces and the third wheel torque are dependent variables. A criteria is required to make the choice of input motor torques. In this analysis, τ_1 and τ_2 are chosen to minimize the vehicle's power consumption. Since the rover is driven by DC motors, power consumption is approximately proportional to the square of the motor torques, or

$$g = \sum_{i=1}^{i=3} k\|\tau_{\mathbf{i}}\|^2$$

where k is a constant that depends on the motor specifications [16]. This minimum can be computed analytically (τ_3 being a function of τ_1 and τ_2, g is a function of τ_1 and τ_2). It is found as,

$$\min(g) \quad \begin{cases} \dfrac{\partial g}{\partial \tau_1} = 0 \\ \dfrac{\partial g}{\partial \tau_2} = 0 \end{cases} \tag{16}$$

Figure 7 shows a typical solution space. Here the power consumption, an elliptical paraboloid, is plotted as a function of τ_1 and τ_2. The subspace where a solution for static equilibrium exists is shown [15]. Being able to develop an analytical solution could help make on-board implementation feasible.

Power consumption, actuator saturation, wheel slip, vehicle stability, etc. are all important factors in determining the fitness of an action plan. The analysis presented above is then a critical tool in evaluating a plan. It is also shown that the load distribution on the wheels depends on the applied wheel torques. This interesting feature could also lead to improved vehicle control. The analysis has strong simplifying assumptions, but they dramatically reduce the complexity of the problem as required for on-board implementation. Clearly, for a specific vehicle design and range of tasks, the validity of these assumptions need to be carefully evaluated prior to the use of the methodology in a flight system.

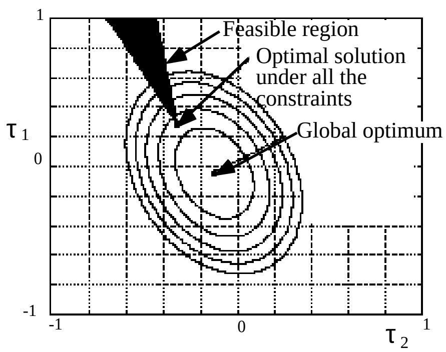

Figure 7: The Minimization of Power Consumption

4. Application of the Planning Method

The planning methodology has been tested in simulation on several representative Mars mission tasks. Two of these tasks are described here. In the first, the rover is required to travel to the top of a hill, a distance of about five meters, in one Martian day. The rover is also tasked to acquire up to five science objectives. This task and solution are shown in Figure 8.

This is a very challenging environment for the rover. A ditch exists between the rover and its goal that is approximately one wheel diameter (25 cm) in width and depth. Also, the hill is approximately 2.5 wheel diameters (65 cm) in height and has slopes ranging from 25 degrees on the left to approximately 70 degrees on the right. Finally, rocks cover the area which are up to one wheel diameter in height and width.

A fitness function is established that defines the relative importance of reaching the target, the total energy consumed, the number of science samples obtained, the stability margin maintained, the level of actuator saturation used, the task completion time, etc.

$$f = \alpha_1 D + \alpha_2 A + \alpha_3 \sum_i h(S_i) - \alpha_4 T - \alpha_5 E + \alpha_6 \delta \tag{17}$$

Where: D = distance robot has traveled toward the target
A = number of science samples acquired
S_j = stability margin for time step j
T = time required to complete the task
E = the energy consumed during execution
δ = 1 if the target is reached and 0 otherwise
α_i= weighting factor

Here vehicle stability is represented by the angle through which the vehicle must be rotated to cause it to tip. This angle is divided by the angle required to tip the vehicle on a flat surface to create a stability margin [17]. The function *h* may represent a non-linear function. For example, a stability margin above 80% may not incur a penalty, a stability margin below 20% is not allowed.

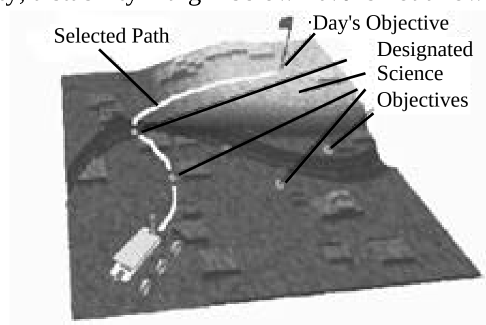

Figure 8: A Mobility Task with Science Objectives

The planning methodology was applied to search for a successful action plan using a reduced inventory of 9 modules (101, 105, 106, 107, 108, 3000, 3001, 3002, 3003). Modules giving mobility as well as manipulation were included while unnecessary modules such as deployment of the science tray were not. The GA used a population of 50 individuals and solved the problem in 13 generations then was stopped after 20 generations. A run time of approximately 2 hours was required to for this search on a 80486 type processor, which is a reasonable approximation of the computational capabilities planned for future rover missions.

The search process found an action plan that moves the rover to the target while acquiring two rock samples and maintaining a stability margin no less than 40%. The planner chose to have the rover acquire the two rock samples on the left, while forgoing the other three. This represents a trade-off between time required to complete the task and number of samples obtained. Trade-offs can be explored by altering the fitness function. For example, if a very high stability margin is required, the planner will not allow the rover to cross the ditch. Stability margin or saturation margin was used rather than using a pass/fail test based on these constraints. This was done to account for errors in the physical model and to build robustness into the plan. The robustness of the plans obtained were studied with regard to the sensitivity to the coefficient of friction. It was found that any coefficient of friction above .65 allowed the rover to complete the task. With a coefficient greater than .45 the plan moved the rover across the ditch, however the higher coefficient was required to navigate the hill. Such results can be used to determine if the task can be performed with an acceptable level of risk. Clearly, the computational efficiency of the method makes such studies possible.

In a second task, the rover is asked to cross the ditch and climb the steep portion of the hill described above. This task can be seen in Figure 9. Since the task requires no manipulation, a reasonable reduced inventory might contain 101, 105, 106, 107 and 108 (only mobility actions). This allows the rover to move forward and turn to left and right along sharp and shallow arcs. However, if a plan is developed using this inventory, a feasible solution cannot be found. The hill is simply too steep to allow the rover to maintain stability (it falls backward, as in Figure 9). However, if the additional action modules 1001, 1002, 1003 and 1004 (motion of the manipulator) are added to the reduced inventory, a successful plan is found. With these manipulator motion modules, the robot moves the manipulator forward, thereby moving its center of gravity forward. The rover is able to navigate the slope and reach the target. In this case the method is able to find a non-obvious result. More detailed descriptions of these results can be found in [18].

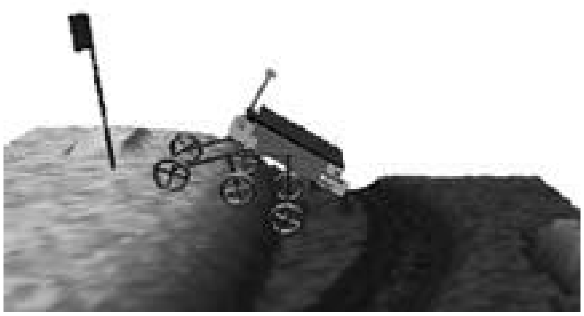

Figure 9: Instability in a Hill Climbing Task

5. Summary and Conclusions

A physics-based planning methodology is presented that has the potential to plan the action of planetary rovers. The method explicitly considers task and rover constraints such as power consumption, actuator saturation, wheel slip, vehicle stability and vehicle safety in challenging environments. The methodology is demonstrated in simulation on representative Mars mission tasks.

The methodology divides the activities of the rover into discrete, physically-realizable actions, or action modules. It searches for the proper sequence of action modules, or action plan, using a hierarchical selection process that includes a genetic algorithm. Action plans are evaluated using physical model of the rover and its environment. The method is structured to allow the search process to solve the planning problem effectively. The approach is limited by the accuracy of the model. Therefore, the model must represent the rover and its environment accurately, while maintaining computational efficiency for on-board implementation.

An experimental system is currently under construction. It will be used to determine the effect of uncertainties such as poorly characterized soil/tire interactions and to demonstrate the ability to implement the method with relatively small on-board computers.

6. Acknowledgments

This work is supported by the NASA Jet Propulsion Laboratory under contract no. 960456. The authors would like to acknowledge the support of Dr. Paul Schenker and his research team at JPL. The support and encouragement of Guillermo Rodriguez of JPL is also acknowledged.

7. References

[1] Mars Pathfinder: www.mpf.jpl.nasa.gov

[2] Hayati, S., et. al., "The Rocky 7 Rover: A Mars Sciencecraft Prototype," *Proc. 1997 IEEE International Conference on Robotics and Automation*, 1997.

[3] Matthies, L., et al, "Mars microrover navigation: performance evaluation and enhancement," *Journal of Autonomous Robots*, 2(4), 1995.

[4] Gat, E., Desai, R., Ivlev, R., Loch, J., Miller, D., "Behavior Control of Robotic Exploration of Planetary Surfaces," *IEEE Trans. Robotics and Automation*, Vol. 10, No. 4, 1994.

[5] Brooks, R., "A Robust Layered Control System for a mobile robot," *IEEE Trans. Robotics and Auto.*, 2-1, 1986.

[6] Khatib, O. "Real-time obstacle avoidance for manipulators and mobile robots," *International Journal of Robotics Research*, 5 (1), 90-98, 1986.

[7] Borenstein, J., Koren, Y., "The Vector Field Histogram-Fast Obstacle Avoidance for Mobile Robots," *IEEE Trans. Robotics and Automation*, Vol. 7, No. 3 June 1991.

[8] Schenker, P., et. al., "Lightweight Rovers for Mars Science Exploration and Sample Return," *SPIE Proc. Intelligent Robots and Computer Vision XVI*, 1997.

[9] Farritor, S., Dubowsky, S., "A Genetic Algorithm Based Navigation and Planning Methodology for Planning Methodology for Planetary Robotic Exploration," *Proc. 7th ANS Robotic and Remote Systems*, Augusta, GA, 1997.

[10] Farritor, S., and Dubowsky, S., "A Self-Planning Methodology for Planetary Robotic Explorers," *8th Intern. Conf. on Advanced Robotics*, Monterey, CA, 1997.

[11] Goldberg, D., *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addision-Wesley, Reading, MA, 1989.

[12] Bickler, B., "A New Family of JPL Planetary Surface Vehicles," *Proc. Missions, Technologies, and Design of Planetary Mobile Vehicle*, Toulouse, France, 1992.

[13] Linderman, R., Eisen, H., "Mobility Analysis, Simulation and Scale Model Testing for the Design of Wheeled Planetary Rovers," *Proc. Missions, Technologies, and Design of Planetary Vehicle*, Toulouse, France, 1992.

[14] Matthies, L., Balch, T., Wilcox, B., "Fast Optical Hazard Detection for Planetary Rovers using Multiple Spot Laser Triangulation," *IEEE Conf. on Robotics and Auto.*, 1997.

[15] Hacot, H., *The Kinematic Analysis and Motion Control of a Planetary Rover*, Masters Thesis, Department of Mechanical Engineering, MIT, Cambridge, MA, May, 1998.

[16] Dubowsky, S., et al., "On the Design and Task Planning of Power Efficient Field Robotic Systems", *Proc. 6th ANS Robotic and Remote Systems*, Monterey, CA, 1995.

[17] Papadopoulos, E.G., Rey, D.A., "A New Measure of Tipover Stability Margin for Mobile Manipulators," 1996 *IEEE Intern. Conf. on Robotics and Auto.*, 1996.

[18] Farritor, S., *On Modular Design and Planning for Field Robotic Systems*, Ph.D. Thesis, Department of Mechanical Engineering, MIT, Cambridge, MA, May, 1998.