Autonomous Navigation and Control of a Mars Rover

David P. Miller
David J. Atkinson
Brian H. Wilcox
Andrew H. Mishkin

Jet Propulsion Laboratory/California Institute of Technology
M.S. 301-440
4800 Oak Grove Drive
Pasadena, California 91109
USA

Abstract

A Mars Rover will need to be able to navigate autonomously kilometers at a time. This paper outlines the sensing, perception, planning, and execution monitoring systems that are currently being designed for the rover. The sensing system is based around stereo vision. The interpretation of the images use a registration of the depth map with a global height map provided by an orbiting spacecraft. Safe, low energy paths are then planned through the map, and expectations of what the rover's articulation sensors should sense are generated. These expectations are then used to ensure that the planned path is correctly being executed.

Keywords: stereo vision, navigation, path planning, mobile robots.

1. Introduction

In the next decade, a mission to the planet Mars is planned that will involve a robot moving across the surface of the planet and collecting samples. The robot will have to travel dozens of kilometers through rough terrain. It will have to be able to find its way to sites of scientific interest and to the return vehicle. The Rover will have to be able navigate through its local environment autonomously, due to round trip light-time delays. This implies that the vehicle must be able to sense its environment, plan a course through that environment, and react appropriately to unexpected situations as they appear. All this must be done while guiding the vehicle towards the goals that have been given to it from its human operators on the Earth. Our research concentrates on the sensing, perception, planning and execution monitoring that must be carried out by the rover to ensure that a safe and efficient path is found, and that that path is traversed correctly.

The overall approach we are taking tries to minimize the interaction between the planning system and the properties of any particular sensor. This is done in two ways. First, the perception system builds up a world model from the sensors' outputs. Only this world model, including the uncertainties of the data, is given to the path planning system. Once a suitable path has been determined then sensory commands must be generated that will help the rover to monitor its progress through the traversal. To maintain the separation between planning and specific sensory hardware, this process is divided into two steps: expectation generation and sensor scheduling. The generation of expectations can be accomplished using a model of the vehicle and the terrain. The sensor scheduling planner can then assign the relevant sensors to monitor the rover's state.

Section 2 of this paper describes the semi-autonomous navigation scenario. Section 3 presents the sensor and perception systems of the rover. Section 4 describes the planning systems that are used to accomplish path planning and execution monitoring. Architectures for both systems are given. Implementation and testing testbeds are presented in section 5, and section 6 presents some conclusions and results we have achieved so far.

2. Semiautonomous Navigation Scenario

In the semiautonomous navigation (SAN) approach we are currently investigating, local paths are planned autonomously using sensor data obtained by the vehicle; however, local path planning is guided by global routes planned less frequently by human beings using a topographic map, which is obtained from images produced by a vehicle orbiting Mars. The orbiter could be a precursor mission which would map a large area of Mars in advance, or it could be part of the same mission and map areas only as they are needed.

The sequence of operations in the portion of SAN involving Earth is as follows. As commanded from Earth, the orbiter takes a stereo pair of pictures (by taking the two pictures at different points in the orbit) of an area to be traversed (if this area is not already mapped). These pictures might have a resolution of about one meter, although poorer resolution could be used. The pictures are sent to Earth, where they are used by a human operator (perhaps with computer assistance) to designate an approximate path for the vehicle to follow, designed to avoid large obstacles, dangerous areas, and dead-end paths. This path and a topographic map for the surrounding area are sent from Earth to the rover. This process repeats as needed, perhaps once for each traverse between sites where experiments are to be done, or perhaps once per day or so on long traverses.

The sequence of operations in the portion of SAN taking place on Mars is as follows. The rover views the local scene and, by means discussed below, computes a local topographic map. This map is matched to the local portion of the global map sent from Earth, as constrained by knowledge of the rover's current position from other navigation devices or previous positions, in order to determine the accurate rover position and to register the local map to the global map. The local map (from the rover's sensors) and the global map (from the Earth) are then combined to form a revised map that has high resolution in the vicinity of the rover. This map is then provided to the path planning system. A new path then is computed, revising the approximate path sent from Earth, since with the local high resolution map small obstacles can be seen which might have been missed in the low-resolution pictures used on Earth. Using the revised path, the rover then drives ahead a short distance (perhaps ten meters), and the process repeats.

3. Sensing and Perception

The rover's sensing and perception system must be able to sense the surrounding terrain, convert the sensed data into a representation of the local environment, and correlate this local representation with the global map, thereby determining the rover's position; the local and global maps must be integrated, and slope and roughness estimates determined for nearby terrain, in order to provide the appropriate world model for use by the planning system [Wilcox88]. Figure 1 details the elements of the sensing and perception system. The major operations of this system are briefly described below.

3.1 Stereo Correlation

We have implemented a two camera stereo correlation approach [Gennery 77, 80], which operates on a pair of images

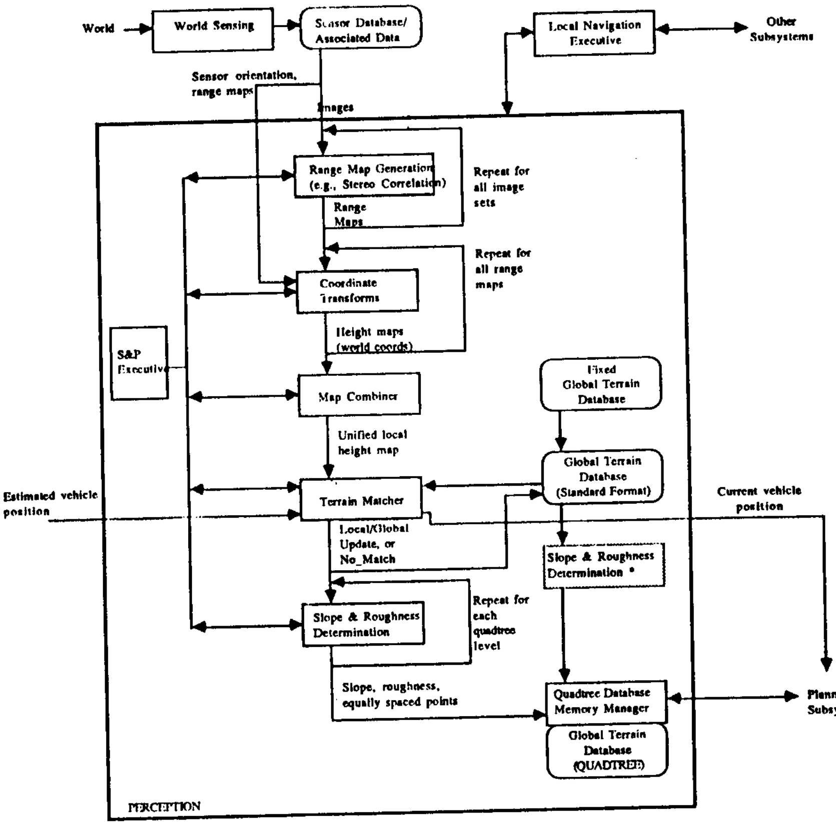

Figure 1. Sensing and perception system

from the left and right cameras on the rover. Area correlation is used, with same-size windows into the two images being moved in one-pixel increments relative to each other. Area correlation was chosen, rather than edge correlation, because a Mars rover (as well as the current experimental rover) will operate in a natural terrain environment. Scenes in natural environments tend to contain significant texture, but few well-defined edges.

The result of stereo correlation is a disparity map, which specifies the difference in position between images of pixels representing the same point in space. Triangulation is then used to build a range map from the disparity map. Each point in the range map has an associated covariance matrix representing the uncertainty in its location. Finally, a coordinate transformation is applied to the range map to put it into a coordinate frame with the same orientation as the global database coordinate frame. The coordinate transformation is based on the known heading of the cameras on the rover at the time images were captured.

3.2 Terrain Matching

The terrain matcher merges the elevation map generated by stereo correlation with the elevation map in the database by a process of correlation and averaging, which also produces the best estimate of vehicle position as that which produces the best correlation. However, this computation is more complicated than ordinary correlation because of the following: the points are not equally spaced, there may be significant uncertainties in their horizontal positions, and there may be occasional mistakes in the stereo data.

The terrain matcher performs a variable resolution search, beginning at the lowest resolution to produce a rough match, then refining the match as the resolution is increased. The uncertainty estimates associated with each point in the height maps are incorporated into the representations used for matching. In the matching process, points in the local database that disagree significantly with the global database are eliminated. Note that a feature in the local database must be large enough to appear in the lower resolution global database to be thrown out. Features in the local database which are too small to show up in the global database are retained. The process for rejecting points is similar to that in the ground surface finder approach previously described in [Gennery 77, 80].

The output of the terrain matcher is: 1) the position of the rover in the global coordinate frame, together with a goodness of fit measure for the determined location, and 2) an edited local height map with apparently erroneous points removed.

3.3 Slope and Roughness Determination

Traversability can be determined by analyzing the database to determine the slope and roughness of the ground at each horizontal position. This can be done by local least-square fits of planar or other surfaces and analysis of the residuals.

The slope and roughness determination module accepts as input the current rover position, the corresponding local height map, the global database, and any existing local height maps generated previously from rover data collected at earlier stopping points along the route. The rover position provides the offsets necessary to properly position the local database in the global database. Only a rectangular subset of the global database in the vicinity of the rover location is read in for integration with the local database. A Gaussian function is defined around each (equally spaced) output point, and a least squares fit is performed of all local database points that fall close enough to the output point, with weights derived from the Gaussian function and the covariance matrices of the input points. A plane is fit in three dimensions to the data; the z value of the plane at the (x, y) coordinates of the output point becomes the output height. The scatter of the points used for the least squares fit is used to determine the roughness measure.

The output of the slope and roughness computation is a height map of equally spaced points. Associated with each point is the following data: the slope at the location, a 3 x 3 covariance matrix indicating the uncertainty in the height and slope, a roughness value, and an uncertainty measure associated with the roughness value. All of this information is provided to the planning system.

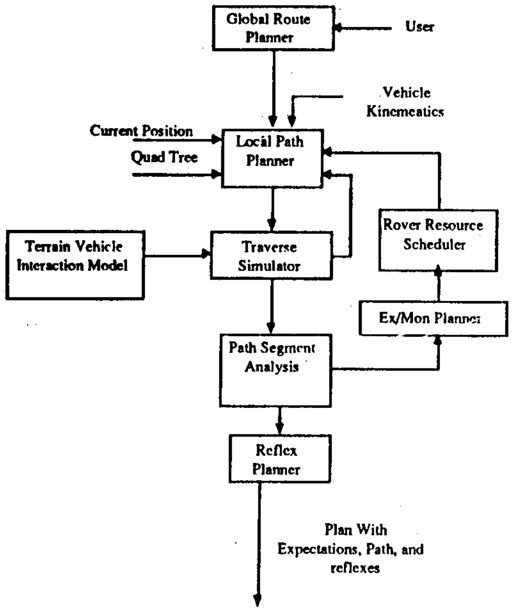

Figure 2. Planning systems

4. Path Planning & Execution Monitoring

The rover must be able to autonomously plan a path and monitor its execution of that path so that it may take corrective action as needed. These needs can be met by using a set of three integrated planning systems, each operating at a different level of detail. The first is a spatial path planner which plans gross rover movements through the terrain. The second is a resource planner; this system allocates and schedules the rover's resources such as power, pointable sensors, and the rover's exact position. The final planner analyzes the path and activities produced by the first two systems and then plans out contingencies, i.e., reactions the rover will take if certain sensor alarms fire, indicating that the rover has deviated from the original plan. The functional modules of such a planning system are outlined in the figure. Details of the tree major planning modules used to accomplish these tasks are given below.

4.1 Path Planning

Much work on path planning has been done for mobile robots. However, with few exceptions (e.g., [Linden87]) all of that work has been done for vehicles designed to travel in indoor environments (for examples, see [Brooks82, Chatila85, Chubb87, Crowley85, Elfes87, McDermott84, Nguyen84, Slack87, and Thorpe84]). In those situations a piece of space is either traversable or not. For a planetary rover, traversing rough natural terrain, such a black and white analysis is not usually possible. A path plan must take into account what kind of terrain the rover is currently on, what the transitions are like to the adjoining types of terrain, the size of the rocks, whether they are expected to be loose or solid, what the battery level of the rover is charged at, and what the possibilities are of getting into a trapped situation if something turns out to be harder to traverse than expected. Additionally the route planner must model non-geometric hazards that cannot be directly detected by the sensing and perception system, but can only be inferred by the geologic context, and/or the behavior of the rover as it makes its traverse.

As described in the previous section, the rover first takes a stereo image of the surrounding terrain The resulting depth map is integrated into an Earth provided orbiter-view map of the general area. The orbiter data can be first assembled on the Earth, and then, as the rover needs the information, the data can be sent to the rover a few kilometers at a time. The combined map of orbiter and stereo data gives the rover a high resolution map of its immediate surroundings with lower resolution in the outlying areas. This information is encoded statistically into a quad-tree representation. A Bayesian model of the sensor uncertainty is used to combine information from multiple sources and represent features of the terrain in the quad-tree.

The path planner is then passed the quad-tree. Path planning is done using a hierarchical message-passing algorithm that traverses through the high resolution area of the quad-tree in the general direction of the goal [Miller87]. Most of the time the high resolution area does not contain the goal (the goal is hundreds or thousands of meters away and the high resolution area is only tens of meters long). For this reason the path planning can only be in the general direction of the goal. Additionally, the path should be in a direction that, according to the low resolution data, will probably not lead to a blocked area that would require backtracking. A variety of heuristics are used, in addition to the hierarchical representation, in order to assure that the computation remains tractable. This is done by having the path evaluation function incorporate a crude simulation of the rover. Path messages will not be sent through a portion of space that is untraversable. Similarly, messages will be held back that accrue a high traversal cost. The cost function is dominated by power consumption and safety considerations (i.e., the probability of the space being traversable, and the probability of the rover getting trapped if it is not).

So the path planner plots a low cost path between the rover's current position and the edge of the high resolution area of the quad-tree that lies in the general direction of the goal. The path contains more than just directional information. As a result of the simulation process used in path planning, the path also has annotations about expected power draw, and expected stability of the rover for each segment of the path. Additionally, expected rover orientation, tilt, and acceleration may be calculated from the path information.

Once a path has been computed, the navigation system performs a terrain/vehicle interaction simulation of the rover traversing the path. This simulation is performed in order to locate sensing landmarks for execution monitoring. The simulation also serves to verify the traversability and safety of the selected path.

4.2 Resource Scheduling

The resource planner uses all the information in the path plan to allocate the rover's resources. This planner uses models of the rover's sensors to decide what sensor will need to be used to verify the rover's successful traverse of each segment of the path. In particular, the planner needs to know the sensor's range, visibility requirements, processing time, power draw, setup time, and pointing accuracy to calculate when a particular sensor is to be used. Additionally, the planner examines the path from the path planner to make sure that there are adequate landmarks visible to verify the rover's successful traverse. This latter is not usually a problem. In rough terrain there are usually enough landmarks, simply by virtue of the terrain being rough. If there are no landmarks, then the terrain is usually smooth and safe, so an exact following of the plan is both simple, and unnecessary (if the rover is in a flat open desert, then it does not matter if it strays a couple meters here and there). If absolutely necessary, the resource planner can request a specific waypoint to be added to the path; this info is passed back to the path planner which replans that section of the path.

The resource planner uses a directed beam-search through state/time search space [Miller86]. This is in effect a directed multiple worlds discrete event simulation. The best performing world is preferentially selected until a simulation with a successful outcome is found. The schedule followed in that simulation is used as the resource plan for the rover.

4.3 Execution Monitoring

The expectations generated by the path planner for rover position and orientation are combined with the resource allocations to derive a set of sensor expectations. These consist of a period of time/space where a specific sensor should be reading within a certain range. A reading outside that range at that time would indicate the the traverse is not proceeding as planned [Munson73]. A sensor can either be a physical sensor on the vehicle, or a virtual sensor – a function of several of the sensors' data.

There are too many possible virtual sensors for the rover's computation system to successfully monitor them all. Therefore, one job of the execution monitor planner is to select which are the important sensors for the particular context that the rover is in. This is a form of predictive monitoring [Doyle89].

The final planner on the rover uses the expectations and the path plan to derive a set of reflexes that the rover should have at the ready for each segment of the traverse. Different contingency reflexes would be triggered by different sets of sensor alarms. The sensor alarms occur whenever the sensor reads outside the expected range.

In most cases, the reflexes might just be to have the rover halt, take a new panorama, and replan the remainder of the traverse. Some situations may demand more complex reflexes. For example, if the front wheels suddenly start spinning free, and the suspension encoders indicate that those wheels have suddenly dropped, then the front of the rover has probably broken through the surface. If this was not expected, then the rover should stop and backup immediately, to avoid having the entire vehicle getting mired. In other situations, (e.g., if the sensor system was initially unable to tell if an area of the surface was dust covered or bedrock, but thought it was probably dust-covered) if the rover suddenly encounters less surface drag then had been expected, the rover might be able to modify its model of the surface on the fly, and push through the new power projections into the rest of the plan (perhaps canceling a recharging stop).

External landmarks are also used to monitor the rover's progress along its traverse. These landmarks, and their associated sensor tasks, can be of several forms including tracking a particular boulder in the distance, or checking the time and odometer reading when a bump of a particular size should be encountered. Particularly tight monitoring will be scheduled on those parts of the path that might approach any of the vehicle's safety margins. Under these conditions (e.g., passing between two narrowly spaced boulders) the execution monitoring system will also select and load the appropriate reaction software (e.g., if both left and right proximity sensors turn on, perform an emergency stop and slight back-up — to ensure that the rover does not become wedged between the two boulders). This way, in an emergency, the rover will never have to stop and think before making an appropriate response.

5. Implementation and Testing

The navigation system described in the previous sections is being implemented and tested on two testbed robots at the Jet Propulsion Laboratory. For ease of prototyping and testing, parts of the navigation system are being implemented on a modified Cybermation K2A industrial mobile robot. The K2A is basically an indoor robot, but is capable of surmounting small obstacles. This allows us to test the algorithms in terrain that has finer gradations in quality than simply traversable or non-traversable.

The second testbed is a custom designed six-wheeled fully articulated outdoor robot that should have most of the mobility characteristics of a real Mars Rover. It is capable of surmounting obstacles (using significant power) that are approximately a full wheel diameter (one meter) in size [Miller89]. The two vehicles are software compatible with one another.

6 Conclusions and Results

Due to safety, high level mission goals, and resource conditions it is necessary that the rover perform some strategic planning. Interaction between a Mars Rover and the Earth-based mission control is very costly. In addition to the light-time delay, there are also limited communication opportunities, and long decision-making times. It is therefore a very important goal to ensure that the rover has two-way contact with the Earth only when necessary. An onboard semiautonomous navigation system is a necessary ingredient for achieving that goal. To ensure that the plan does not leave the rover stranded in awkward spot, or that the plan be executed faultily, resource allocation is also needed. However, due to the possibility of a time critical, life threatening situation, the rover will also need reactive abilities. Since the requirements on these abilities change with the rover's situation, they should be dynamically installed in the rover as part of the planning process to keep them consistent with the task and environment context in which the rover finds itself. To meet the power and reliability requirements of an actual Mars mission, a robust and passive sensory system is needed. The architectures outlined above allow long range planning, yet incorporate real-time performance.

Acknowledgements: The research described in this paper was carried out by the Jet Propulsion Laboratory — California Institute of Technology under a contract with the National Aeronautics and Space Administration.

References

[Brooks82] Brooks, R. A., Solving the find path problem by a good representation of free space, in *Proceedings of AAAI 82*, AAAI, pp. 381-386, 1982.

[Chatila85] Chatila, R., Position referencing and consistent world modeling for mobile robots, in *Proceedings of the International Conference on Robotics and Automation*, IEEE, pp. 138-145, 1985.

[Chubb87] Chubb, D.W., An introduction and analysis of a straight line path algorithm for use in binary domains, in the *Proceedings of Workshop on Spatial Representation and Multi-Sensor Fusion*, St. Charles Illinois, pp. 220-229, October 1987.

[Crowley85] Crowley, J.L. "Dynamic world modeling for an intelligent mobile robot using rotating ultra-sonic ranging device." *Proceedings of the 1985 International Conference on Robotics and Automation,* pp 128-135, IEEE, St. Louis, MO. (March 1985).

[Doyle89] Doyle, R., S.Sellers and D.Atkinson, Enhancing Aerospace Systems Autonomy Through Predictive Monitoring, *27th Aerospace Sciences Meeting*, AIAA, paper #89-0107, Reno, NV, January 1989.

[Elfes87] Elfes, A., "Sonar based real world mapping and navigation", *IEEE Journal of Robotics & Automation*, pp249-265, RA-3, #3, 1987.

[Gennery80] Gennery, D.B., "Modelling the Environment of an Exploring Vehicle by Means of Stereo Vision," AIM-339 (Computer Science Dept. Report STAN-CS-80-805), Stanford University, (Ph.D. dissertation) (1980).

[Gennery77] Gennery, D.B., "A Stereo Vision System for an Autonomous Vehicle," Fifth International Joint Conference on Artificial Intelligence, MIT, 576-582 (1977).

[Linden87] Linden, T. & Glicksman, J., "Contingency planning for an autonomous land vehicle," in *Proceeding of the 10th IJCAI, AAAI, IJCAI*, Milano, Italy, pp 1047-1054, 1987.

[McDermott84] McDermott, D. V., Davis, E., Planning routes through uncertain territory, *Artificial intelligence*, v22, pp. 107-156, 1984.

[Miller86] Miller, D.P. "Scheduling robot sensors for multisensory tasks," *Proceedings of the 1986 Robots West Conference*, SME, Long Beach, CA. (September 1986).

[Miller87] Miller, D., P., Slack, M., G., Efficient Navigation Through Dynamic Domains, in the *Proceedings of Workshop on Spatial Representation and Multi-Sensor Fusion*, St. Charles Illinois, October 1987.

[Miller89] Miller,D.P., A.Mishkin, K.Lambert, D.Bickler, and D.Bernard, Autonomous Navigation and Mobility for a Planetary Rover, *27th Aerospace Sciences Meeting*, AIAA, paper #89-0859, Reno, NV, January 1989.

[Munson72] Munson, J., H., Robot Planning, Execution and Monitoring in an Uncertain Environment, in the *Proceedings of the International Joint Conference on Artificial Intelligence*, Stanford California, pp. 338-349, September 1972.

[Nguyen84] Nguyen, V., *The Find-Path Problem in the Plane*, MIT AI Technical Report 760, 1984.

[Slack87] Slack, M. G., Miller D. P., Path Planning Through Time and Space in Dynamic Domains, in *Proceeding of the 10th IJCAI, AAAI, IJCAI*, Milano, Italy, pp 1067-1070, 1987.

[Thorpe84] Thorpe, C. E., Path Relaxation: Path Planning for a Mobile Robot, in *Proceedings of AAAI*, AAAI, pp. 318-321, 1984.

[Wilcox87] Wilcox, W.H., D.B. Gennery, A.H. Mishkin, B.C. Cooper, T.B. Lawton, N.K. Lay, and S.P. Katzmann, "A Vision System for a Mars Rover," Proc. SPIE Mobile Robots II, vol. 852 (5-6 Nov 1987).