A Cooperative Multi-Robot Approach to the Mapping and Exploration of Mars

Paul Rybski, Sascha Stoeter, Chris Wyman and **Maria Gini**
Department of Computer Science, University of Minnesota
200 Union St SE, Minneapolis, MN 55455
{rybski, stoeter, gini}@cs.umn.edu wyman@itlabs.umn.edu

Abstract

In the AAAI [1] "Life on Mars" competition this year, we intend to employ a multi-robot team which combines traditional methods of autonomous navigation with experimental group arbitration strategies to explore and map a simulated extra-terrestrial environment. By communicating with each other to optimize the search, the robots will be able to explore the environment faster than could a single robot. Group strategies will also be used to coordinate the robot's actions to optimize the retrieval of objects in the environment.

Introduction

Space travel is an inherently hazardous undertaking which requires redundant components in hardware and software systems to prevent catastrophic mission failure. This is especially true when a robot is used for extra-terrestrial exploration. A malfunction or accident which damages the robot can place the entire mission into jeopardy. To help ensure a successful mission, we suggest the use of several robots, each capable of accomplishing the task individually, but fully capable of cooperating and sharing the task load. These robots should be capable of dynamically delegating responsibility to solve their tasks. Our interest in this problem is how to maximize the efforts of a group of robots to make the best use of their capabilities. For the "Life on Mars" competition, we will have our team of two robots dynamically prioritize and segment the mapping and object retrieval tasks to make the best use of their abilities in the environment.

Local Navigation

Our robots are equipped with an array of ultrasonic sensors which is used in concert with shaft encoders to accomplish navigation. In order to overcome the noise inherent with the sonars, each robot makes use of a certainty grid (Moravec 1988) localized around itself. The certainty grid uses a probabilistic model of the sonars to filter out erroneous readings caused by specular reflections. Because of the positional errors that accumulate in shaft encoders, an object's position encoded in the certainty grid becomes less precise the further the robot moves away from it. This limits the effective range of the certainty grid to use with local path planning only.

We are experimenting with two methods in order to come up with a robust sensor model for the certainty grid. The first approach is to hard-code the certainty grid parameters through experimentally determined values. This is done by aiming a stationary ultrasonic sensor at a fixed object and taking a number of readings at known distances and angles in order to accuratey model the sensor's behavior. The second approach is to tune the certainty grid parameters using a genetic algorithm. This approach would be done off-line before the run in order to produce an optimal sensor model.

The robots use potential field navigation to maneuver around obstacles (Arkin 1989). Cells in the certainty grid that have a high probability of containing an obstacle will emit a repulsive vector field, the magnitude of which is proportional to that probability. In environments where the obstacle density is sparse, the problem of running into local minima is mostly avoided. Where the concentration of obstacles is quite dense, global planning methods are used to free a robot should it become trapped in a potential well.

Exploration and Map Building

A different approach must be used to maintain a global map of the environment, since the positions of objects in certainty grids become less accurate as the robot moves away from them. To compensate for this, we use a method of hierarchical certainty grids (Moravec 1988). As a robot explores and creates its local certainty grid, it updates a global certainty grid that has

---

a much coarser granularity. A rough description of the objects discovered the local certainty grid is stored in each elemtn of the global grid. The size of the grid cells in this global map is determined by the amount of positional error that creeps into the robot's shaft encoders.

While it explores its environment, each robot updates the global map with its current position, as well as with references to interesting objects it has discovered. This way, the robots can share information with each other and determine the best way to partition the tasks at hand. Initially, the robots will start in the known grid space containing the lander and the pen. Once released, they will begin to explore the unknown spaces and build the map. Once the global map contains more information, the robots can use it to plan the best path between two grid cells.

Path Planning

For path planning, the robots combine the information from the global and local certainty grids (Ratering & Gini 1995). Each cell of the global certainty map is assigned a value representing the density of the objects discovered inside it. If the cell contains danger spots, it will be assigned a value corresponding to a high density of objects as well. Path planning between two cells in the global map is done to minimize the amount of clutter through which the robot has to travel. A tradeoff is made between the distance a robot must travel and the amount of clutter it has to go through, since the former wastes battery life, and the latter risks trapping the robot. Once the path has been determined, a potential attraction field will be assigned to each grid cell that points from one cell to the next along The chosen path. This field will direct the robot toward the best path in the direction of its intended goal. The global attraction fields are used in concert with the local heuristic potential field algorithms to fine-tune the navigation of the robot.

Task Division for Multiple Robots

Deciding how to effectively share the responsibility of exploring the environment between the two robots is our primary research interest. Initially, the global map that represents the environment will be filled with unexplored grid spaces. The robots must decide among themselves how best to segment the area such that they can conduct a thorough search while minimizing the chance of interference, collision, or duplication of effort. When deciding what actions to take next (gather objects or continue to explore), each robot takes the other's capabilities into account in order to most effectively share the tasks.

We are currently experimenting with several different kinds of arbitration strategies (Goldberg & Mataric 1997). One strategy is to define a pack structure where one robot is dominant over the other. The dominant robot decides initially where it wants to go and the other robot chooses its actions in such a way as to not interfere with the first. Should a robot become incapacitated in some way such that it cannot continue to complete some goal (i.e. a disabled camera or gripper), the fully functional robot can change its priorities to accomplish more of the tasks that the damaged one is no longer capable of doing. Another strategy we are considering is to give the task of exploring the environment to one robot and to have the other robot follow in the first robot's footsteps, gathering interesting items. This way, any interesting discoveries made by the first robot could be capitalized on by the second, leaving the first robot free to continue exploring. In the event that a robot becomes incapacitated in some way (a gripper breaks down, for instance), the robots could swap tasks in order to continue the mission.

Hardware

Our robotics exploration team consists of two RWI Pioneer robots. Both are outfitted with ActiveMedia grippers for gathering objects from the environment. The robots are also outfitted with Newton Labs Cognachrome Vision Systems for recognizing and discriminating between the various colored martians and recognizing danger areas. Computational power is provided with Dell Latitude LM Pentium 133 Mhz laptops riding on the backs of the robots and communicating with each other over BreezeCOM wireless ethernet

Acknowledgments

We wish to thank all the people who have helped in this project. We would like to acknowledge the support of NSF under grant NSF/DUE-9351513, the Undergraduate Research Opportunities Project of the University of Minnesota, Honeywell Military Avionics, and Rosemount Inc.

References

Arkin, R. C. 1989. Motor schema-based robot navigation. *Int'l Journal of Robotics Research* 8(4):92–112.

Goldberg, D., and Mataric, M. J. 1997. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings AAAI-97*.

Moravec, H. P. 1988. Sensor fusion in certainty grids for mobile robots. *AI Magazine* 9(2):61–74.

Ratering, S., and Gini, M. 1995. Robot navigation in a known environment with unknown moving obstacles. *Autonomous Robots* 1(2):149–165.