Multiple Behavior-Controlled Micro-Robots for Planetary Surface Missions

David P. Miller

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

Abstract[1]

Robotic planetary surface missions cover a wide range of scenarios from grabbing a piece of rock near the spacecraft landing site to building and maintaining a large radio observatory on the Lunar farside. This paper discusses how teams of small autonomous mobile robots can be applied to many of the missions now on the drawing board. Teams of small autonomous robots potentially have many advantages over using a small number of larger robots. These include savings in fabrication costs, launch mass, landing mass, and an increased mission reliability.

1. Introduction

It is not yet clear of what type, and how much, "intelligence" is needed for a planetary rover to function semi-autonomously on a planetary surface. Current designs assume an advanced AI system that maintains a detailed map of its journeys and the surroundings, and that carefully calculates and tests every move in advance [7, 3]. To achieve these abilities, and because of the limitations of space-qualified electronics [6], the supporting rover is quite sizable, massing a large fraction of a ton, and requiring technology advances in everything from power to ground operations.

An alternative approach is to use a behavior driven control scheme. Recent research has shown that many complex tasks may achieved by programming a robot with a set of behaviors and activating or deactivating a subset of those behaviors as required by the specific situation in which the robot finds itself, such as going around a rock [1, 2, 9, 4]. Behavior control requires much less computation than is required by traditional AI planning techniques. The reduced computation requirements allows the entire rover to be scaled down to the micro-rover (1-5 kg) level.

A micro-rover is a small self-contained robot that is capable of autonomous local-navigation. This means that the robot is capable of avoiding obstacles and traps while heading in the desired direction. The chosen direction is in reference to a constantly readable feature such as: a compass heading or a reference direction from a radio beacon. Indoor prototype rovers have been developed [2, 9] and more capable vehicles are now under construction at MIT, JPL and elsewhere. These robots are programmed using any of the several behavior languages now developed. All of these programming paradigms require relatively low computation rates allowing the robot to move continuously and reacting dynamically to its environment.

The next logical question is: how can micro-rovers be used? The next section will briefly go over some of the logistical problems of micro-rovers; their overall capabilities, delivery mechanisms, communications, etc will be outlined. The third section will cover some basics on how to get teams of micro-rovers to act in concert. The sections that follow will give examples of how teams of rovers can be used to accomplish useful missions of science and space operations. Finally some thoughts on the overall utility of micro-rovers will be presented.

2. Micro-rover Logistics

Micro-rovers are small high-mobility autonomous mobile robots. Between the design of their mobility systems, and their onboard intelligence, micro-rovers should be able to go almost anywhere a large vehicle can go, and several places where a large vehicle cannot. A micro-rover cannot go over some of the large rocks that may be traversable by a larger vehicle, but its autonomous obstacle avoidance will allow it to go around those rocks. Its high mobility will allow it to handle obstacles of the characteristic size of the vehicle. Very large obstacles can either be gone around, or are so large that, at the scale of the micro-rover, they have a sufficient fractal dimension to be traversable by the small robot. One should remember that small animals can go most places large animals can go and are normally only blocked by smooth artificial structures. Any artificial structure that a micro-rover needs to traverse can be made traversable when the object is designed. Our preliminary experiments in natural terrain with one of our robots leads us to believe that a micro-rover can go anywhere it needs to for the purposes of planetary exploration.

Getting micro-rovers to a planetary surface is much easier that landing their larger alternatives. A smaller size and lower mass can make a micro-rover much stronger and more shock resistant than a vehicle massing the greater part of a ton. On planets such as Mars, it should be possible to land a micro-rover using a combination aeroshell and parachute. The rover itself will be encased in a shock absorbing package, which opens up upon landing. The *Capsule System Advanced Development* system, developed in 1967 to survivably land a science package on the surface of Mars, used a mass ratio of approximately 1.6 between the support hardware (e.g., aeroshell, parachute, and shock absorbers, etc) and the payload that was landed in working condition. CSAD was designed for landing payloads in the 30-50kg range.

The CSAD capsule was derived from the Ranger landing capsule [11] which had a slightly higher mass ratio (approximately 2:1). The system developed for Ranger is applicable for getting micro-rovers down on surfaces without an atmosphere. The Ranger system used a rocket motor to slow the capsule to a few hundred miles an hour. The shock absorbers in the capsule (originally made out of radially cut balsa wood) were designed to absorb the remainder of the impact. Both the Ranger landing capsule and the CSAD capsule are capable of safely landing several micro-rovers and a beacon system on the appropriate planetary surfaces. With approximately a 2:1 mass ratio, they are more efficient at delivering payload to a planetary surface than any propulsive lander so far developed.

---

[1] This work was sponsored by, and performed for, IS Robotics, 501 Marin St, Suite 214, Thousand Oaks, CA 91360

90CH2930-6/90/0000-0289$01.00 © 1990 IEEE

Micro-rovers operate autonomously. This means that little communication between the robots and the Earth is necessary for the safe operation of the robots. Micro-rover programs so far developed have been on the order of ten thousand bytes of program and data. Such a small program size means that critical portions can easily be stored in ROM and that the robot could receive a totally new program in less than an hour with only a few bytes per second of uplink speed. The larger communications needs come from relaying the data that the rover acquires back to Earth. From the Moon, this is still a trivial matter, because of the short distance. From Mars and beyond, the communications subsystem is still simple if it is assumed that a relay orbiter is available. Since communications is one of the few things that does not scale with the size of the robot, care needs to be taken to reduce communications overhead and usage.

While maintaining minimal communications, there is no good way to teleoperate these robots. Time-delay makes most forms of teleoperation infeasible anyway. The way to instruct these robots to accomplish a task is to give them an evaluation function to satisfy. For instance, a useful task for a Mars rover to accomplish would be to check for carbonates near the edges of dried-up riverbeds. The rover can be directed to the riverbed by giving it a position relative to beacons landed with it, and at other rover landing sites. Once the rover is near the riverbed, behaviors can be activated that search around for a ridge going down, and have the rover try to maximize its depth. Once the rover is in the riverbed, it should just follow the ridge-wall, sampling for carbonates. It should then follow any carbonate gradient to its maximal concentration. By changing the evaluation function, and by controlling which behaviors are active and available at a given time, it is possible to direct a behavior-controlled micro-rover without requiring the communications exchange needed for teleoperation.

3. Micro-Rover Teams

A single micro-rover can orient itself based on signals received from fixed radio beacons. A team of micro-rovers can orient itself internally by having each rover position itself relative to beacons on its companion rovers. To get micro-rovers to swarm together, it is only necessary that they all have some sort of sensible beacon mounted on their structure. Each rover then activates behaviors to move it towards the greatest concentration of beacons, while keeping it some small distance away from any other vehicle. This will cause the rovers to clump up, and the clumps to eventually converge. The minimum distance behaviors will keep the rovers from actually running into one another, and will keep the swarm at a constant density.

Once you have a swarm of rovers, it is easy to get them moving as a coordinated pack. For this to happen the rovers need to select a leader. It is not assumed that all rovers will always function; therefore it is is not a good idea to hardwire in a single lead rover. However, it is assumed that all functional rovers can sense the beacons of all rovers in the pack. If each rover has a serial number encoded in its beacon, and each rover has a leader-behavior that is suppressed when sensing a higher serial number than it has, then only one rover will have an active leader behavior. The leader-rover disables its swarm behavior, and starts moving on the appropriate course. The followers bias their swarm behavior with a follow-the leader behavior. This combination will cause the pack to move in formation after the lead-rover. Should the leader ever be disabled, the next highest numbered rover will take over.

The leader-based serial number scheme can also be used to have the rovers take turns doing things. For example, in a team of three, if rover-three finds an interesting sample, it can apply its instruments to the sample. When finished, rover-three signals "next", if rover -two is functioning, it signals its response (suppressing rover-one's response) and applies its instruments to the sample. Finally, rover-one responds to rover-two's next signal. When rover-one sends out a "next", there is no response within the allotted time causing rover-three to take command of the group once again.

Similar behaviors can also be used to have the rovers form particular formations. For example, suppose there is a group of rovers, and directionally distinct radio-beacon such as an airport VOR. It is easy to get all the rovers to form a line along a specified radial starting at a specified distance. Each rover has a behavior to get it to the radial, if there is no rover in front of it, it should move along the radial until it is at the correct distance from the beacon. If there is a rover in front of it, it should cut in front of that rover if that rover has a higher serial number, otherwise it should maintain a specified distance behind that rover. When they settle down, all the rovers will be equally spaced, in order, along the specified radial, at the specified distance.

These basic techniques of moving as a group, forming formations, and taking turns can be used to accomplish many useful missions.

4. Mission Possibilities

4.1 Planetology Science

While great progress has been made in the miniaturization of various science instruments [12, 13], it is still true that science instruments do not scale with the size of the vehicle, while the vehicle's capability to carry them does. Therefore, one of the important reasons for using teams of micro-rovers is to ensure sufficient payload capability for science instrumentation [8]. Using the group movement and turn-taking behaviors described above, it should be possible to have several small rovers duplicate virtually all the capabilities of a larger rover. One small rover might just have multi-spectral imaging, and processing facilities. This rover would then normally be the leader since its main purpose would be to locate the appropriate sample. Another rover might use all its payload capability in a rock crusher and rock drill. A small rotor saw may be the most effective way of getting a fresh rock sample from a low-mass, low-power sample acquisition system. Other rovers might have chemical analysis facilities. gas-chromatographs, and other analysis tools can now be packaged in a few tens of grams. Finally, it might be possible to have very specialized rovers; a mole rover that could dig down into the Martian sand and sniff for volatiles, would be potentially very valuable.

4.2 Planetary Site Surveys

One important task to be done by robotic precursor missions is to certify a site for a human follow-on mission. A major part of the certification process is a sounding of the proposed area. The purposes of the sounding are to locate the depth at which bedrock begins, soil density, rock sizes, locations of accessible and useful deposits, etc. Much of this information can be discovered through performing an electromagnetic sounding and a seismic sounding. The em-sounding is performed by broadcasting a radio signal at the ground and constructing a differential radar-image from the reflection. The equipment to do this is a pair of several meter long antennas. The seismic sounding equipment is similar in appearance. A string of seismaphones are attached on a long cable; a known charge is then set off at a known distance from the seismaphones.

The antennas for both types of sounders can be pulled by a micro-rover (the antennas would mass less than a kilogram) however, care must be taken to keep them from getting tangled. One way to accomplish this would be to have a team of micro-rovers following at different positions along the antennas, with special manipulators to lift the wires over any obstacles that might otherwise ensnare them. Other rovers could be used to emplace the charges for the seismic soundings.

Several teams of rovers could be used to form a pack of teams. In other words, each team would take positions relative to its leader, and the team leaders would form a formation. In this way, a thorough site survey could quickly be accomplished of a large area.

4.3 Instrument Emplacement

Several network missions are currently under study. A network mission involves the emplacement of several similar instrument packages at various points on the planet. Most of the network missions being explored use approximately one dozen penetrators which embed themselves a meter or two into the surface [5, 10]. Some of the problems with penetrators is that they put strong shock resistance requirements on the instrument packages, and that there is no fine tuning of the instrument position. The latter makes it difficult or impossible to do a detailed network mission over a small area. It also makes it impossible to use penetrators for emplacing an array of instruments that need to be in a certain pattern, such as the receiver elements of very-low frequency array (VLFA) radio-observatory.

Teams of micro-rovers provide a new way of approaching these network missions. The simplest method is to build the instrument package into the micro-rover as an integral part. The rovers then distribute themselves into the correct pattern, and start returning the data. This is an especially useful configuration if it might make sense to shift the entire network to a new location at a later time.

Another approach is to have a central cache of instrument packages and a team of rovers to distribute them. If the rovers can sense where the emplaced packages are as well as the position of the cache, then it is not necessary for any complex plan or map to be used for the emplacement. By choosing appropriate values for the maximum distance an instrument pack can be from the cache and the proper distance between instrument packages, the rovers will quickly fill the designated area with an even distribution of instruments. By making the distance between instruments a function of the distance from the cache, a varying density of instrument packages can be achieved that would be suitable for the VLFA. All that is necessary is that each rover can track the position of the central cache, and of the deployed instrument packages nearest the rover.

4.4 Construction and Mining

Before humans can reside permanently on the Moon or Mars a variety of construction activities will have to have occurred. These include digging out large areas, leveling, covering habitats with regolith, etc. Some of these activities are probably outside the purview of micro-rovers. However, activities such as digging and covering structures should be able to be efficiently accomplished by teams of micro-rovers.

For example, covering a habitat will require only sufficient teamwork so as the rovers do not interfere with one another. The rovers could start off surrounding the structure to be buried. There is a non-directional radio beacon on the habitat. Each rover scoops up some dirt that is at least some preset distance from the habitat. It then carries the dirt and deposits it at the nearest point of the structure that is not yet covered. If a certain minimum depth of coverage is necessary, each rover should have a small sounding device that can gauge the depth of coverage. When the rover drops its load, it then goes in search for some more dirt. The preset distance requirement ensures that some rovers will not uncover one part of the habitat to bury another. By maintaining a minimum separation from the other rovers, interference between rovers will be kept to a minimum.

5. Conclusions

Behavior-controlled micro-rovers have been constructed in the laboratory that have a high degree of mobility, and are capable of maneuvering relative to beacons or other sensible stimuli. Space capable versions of these rovers should be constructible at relatively low cost. Because of their small size, mass and computational needs, these robots are also relatively easy to launch, land, and power. The behavior-control paradigm also allows a high-degree of autonomous operation, reducing the cost and complexity of ground-based operations. Many sensors and instruments necessary for geologic, seismic, meteorologic, and biologic surveys of a planets surface can now be made in a form that masses a fraction of kilogram, and requires only a few milliwatts of power.

Given all of the above, it would appear that micro-rovers can be used to great advantage for most aspects of the scientific exploration of planetary surfaces. There are very few places that a large rover can go where a micro-rover cannot. Further, because of their relative low-cost (both in the robot and getting it to its destination) micro-rover missions can afford a higher degree of risk, and afford to drop some rovers at sites of high-scientific interest, that are perhaps also contain very difficult terrain.

One objection that has been posed to micro-rovers is that they would be unable to carry sufficient payload to accomplish a useful mission. Since virtually all instruments are now available in a size that can be carried by a single rover, the team rover concepts that have been outlined above should put this objection to rest. The original Mars Rover Sample Return Mission (MRSR) proposed a rover massing near one ton, and carrying 40kg of science instruments. A team of forty micro-rovers each carrying one kilogram of science instruments would mass no more than two-hundred kilograms, and would have a much greater component redundancy (and likely mission reliability) than would the one ton rover.

The key to micro-rover team activities is the behavior control. This low-computation technique allows several competing stimuli to result in coherent action by the robot. Beacon following and obstacle avoiding behaviors have been well developed. By putting coded beacons and beacon readers on each rover, other rovers can be considered obstacles, or attractors as appropriate. These team activities fit very well into the behavior-control architecture, and require no directed communication between micro-rovers. The only communication between rovers is for them to broadcast their ID and an occasional "ready" and "next" signal. The individual rover's behaviors, and a careful system design will take care of the rest.

Further work is needed to determine what is the best distribution and organization of capabilities among rovers. Experimentation is also necessary to determine if micro-rovers can compete with large rovers on a kilogram of rover verses kilogram moved material basis. But it seems definite that when it comes to exploration, micro-rovers have an important role to play.

Acknowledgements

The author would like to thank Rod Brooks, John Loch, Jim Firby, Erann Gat, Brian Muirhead, and Bruce Bullock for discussions leading to some of the ideas in this paper.

References

[1] R.A.Brooks, "A Robust Layered Control System for a Mobile Robot", *IEEE Journal of Robotics and Automation*, vol RA-2#1, March 1986.

[2] Brooks, R.A., Flynn, A.M., Fast, Cheap, and out of Control: A Robot Invasion of the Solar System, *Journal of the British Interplanetary Society*, vol 42, #10, pp478-485, October 1989.

[3] Gat, E., Slack, M.G., Miller, D.P., Firby, R.J., Path Planning and Execution Monitoring for a Planetary Rover, in *Proceeding of the IEEE International Conference on Robotics and Automation*, Cincinnati, OH, May 1990.

[4] Gat, E., Miller, D.P., BDL: A Language for Programming Reactive Robotic Control Systems, JPL Working paper, 1990.

[5] L.A.Manning "Mars Surface Penetrator - System Description", NASA TM-73243, June 1977.

[6] Miller, D.P., Mishkin, A., Lambert, K., Bickler, D., & Bernard, D., Autonomous Navigation & Mobility for a Planetary Rover, in the *Proceedings of the 1989 AIAA Meeting on Aerospace Sciences*, paper #89-0859, Feburary 1989, Reno NV.

[7] Miller, D.P., Execution Monitoring for a Mobile Robot System, in the *Proceedings of the 1989 SPIE Conference on Intelligent Control and Adaptive Systems, vol 1196*, pp. 36-43, Philadelphia, PA, November 1989.

[8] Miller, D.P., Mini-rovers for Mars Exploration, appears in the *Proceedings of the Vision 21 Workshop*, Cleveland, OH, April 1990.

[9] Miller, D.P., Rover Navigation Through Behavior Modification, in *The Proceedings of the Space Operations Automation and Robotics Workshop*, NASA, Albuquerque, NM, June 1990.

[10] J.P.Murphy, et.al. "Surface Penetrators for Planetary Exploration:Science Rationale and Development Program", NASA TM 81251, March 1981.

[11] "Final Technical Report, Lunar Rough Landing Capsule Development Program," Aeronutronic Division Publication Number U-2007, February 1963.

[12] M.A.Ravine & T.A.Soulanille, "Cameras for Microspacecraft" to appear in *Journal of the British Interplanetary Society*, vol 42, #10, pp460-467, October 1989.

[13] Waltman, S.B., Kaiser, W.J., Electron Tunnel Sensor Technology, Journal of the British Interplanetary Society, vol 42, #10, pp474-477, October 1989.