Small Robots, Big Missions

David P. Miller

May 23, 1992

To move about and explore the surfaces of distant planets, humans are the obvious mobility agent of choice. They are capable, have wonderful sensing and control, and are quite capable of deploying any instruments and sensors for observations that they are unable to make themselves. Unfortunately, humans usually travel with a lot of excess baggage: food, water, and air are just a few examples. For exploring the surfaces of Mars or the Galilean satellites, toting all those niceties in addition to the astronauts and instruments make such missions prohibitively expensive and logistically almost impossible, at least for the near future.

Robots have provided a promising alternative. Robot explorers have flown by and observed every planet in the Solar System with the exception of Pluto. Robot landers have set down instrument packages and relayed observations from the surfaces of Venus, Mars, and our Moon. Also on the Moon, five rovers have driven along the surface, greatly expanding our knowledge of the makeup of the Lunar surface and underlying structures. These rovers performed admirably in their job of exploring the Moon, but there are definite problems in using similar designs elsewhere in the Solar System.

The problems with three of the rovers are quite apparent. These three rovers were driven by the Apollo astronauts, and were not instrumented for autonomous control, or even extended remote operation. The other two rovers, the Soviet Lunakods, are much closer to a Mars or Ganemede rover. However, these vehicles were designed to be remotely driven with no more than a few seconds between the time the images were viewed by their remote driver and the time the steering and speed commands were received and executed by the robot. Depending on Mars' position relative to the Earth, it would take between seven and forty minutes for a signal to get from Mars to Earth, and back to Mars. A remote driver could move the rover only a meter or two and then wait half an hour for the next set of images, or drive blind and hope for the best. The situation for Jupiter's Moons and elsewhere in the Solar System is much worse.

To get around this problem, one either needs faster than light communications, or a rover that can do some of the driving by itself. During the late '70s and all through the '80s quite a bit of research was performed on getting a robot to find its own way through rough terrain. This work was usually performed by researchers working the area of artificial intelligence. While a variety of approaches were tried, they pretty much all followed the same basic pattern, and it was an attempt to automate the process that the Lunakod, and its human operators used.

The robot had a video camera (or pair of cameras, or scanning laser rangefinder) which was used to image the terrain in front of the robot. A computer program would then examine those images in order to *perceive* the shape of the terrain. This program created an internal topographic map of the terrain inside the robot's computer memory. Then another program would examine this map and *plan* a path for the robot take. In order for the system to be able to *monitor* its own progress, the robot would simulate executing its own plan through its own internal model of the terrain. This simulation would allow the robot to generate expectations of what it should sense during different portions of the traverse. Finally the steps of the plan would be *executed* by the robot and the real sensor values compared to the expectations. If the expected values matched reality, then everything was going according to plan; if not then either the robot had drifted off course, had a faulty internal model, had a bad sensor, a bad actuator, or some other unforseen problem. In any of these cases the traverse was aborted and an updated model was created.

After several years of work, a variety of autonomous vehicles have been created. These include DARPA's Autonomous Land Vehicle (ALV) a four ton eight wheeled single-bodied robot which navigated autonomously offroad over rolling hills, avoiding rocks, trees and cliffs; Ambler, a three ton six-legged walker developed at Carnegie-Mellon University; and the Jet Propulsion Laboratory's Robby, a two ton six-wheeled three-bodied roller that navigated its way through the boulder fields and dry washes outside of JPL.

These and some other robots with similar designs are able to autonomously perceive the terrain, plan a path, execute that plan, and monitor their progress. To do these tasks they use computers capable of operating at hundreds of millions of instructions per second, and programs that are hundreds of thousand to millions of bytes in size. These robots are also capable of negotiating some terrain that is utterly impassable by the family car or even your favorite 4x4. The only major problem with using any of these designs for planetary exploration is their mass. A two ton rover requires a similarly sized lander. The transfer vehicle to boost the rover/lander from Earth orbit into a Mars injection is larger still. The booster to get the rover, lander, transfer vehicle, and their associated accoutrements would be of a size not seen since the Apollo landings. It would also be accompanied by an Olympian sized bill. The mass of these "intelligent" robots makes them too expensive to launch.

The obvious solution is to make the rovers lighter, and the most straightforward way to do that is to make them smaller. The motors, gears, stucture, most of the electronics, and even the computers can be shrunk in size so that the robot weighs in at twenty kilograms rather than two thousand. But there are a few catches. If the wheel size of the rover is reduced from one-hundred centimeters to ten, you will also reduce the size of the obstacle that the rover can traverse by an order of magnitude. Fortunately, this is not as much of a problem as it might at first appear. While a small robot will not be able to go over as many things as a large vehicle can, it will have an easier time going around things. Since the vehicle is smaller, it will be able to find paths through crowded terrain that would be difficult or impassable for a larger vehicle. All in all, it might take longer for a small vehicle to get there, but in most situations it will get there.

The other catch is a little more complicated. Electronics in general, and computers in particular have been getting rapidly smaller over the years. It is now possible to put very powerful computers (such as those used on Ambler and Robby) into a small package weighing just a kilo or two. Unfortunately, while the size has been greatly reduced, the power requirements to flip a bit or store a number in computer memory has stayed pretty much the same since the advent of CMOS a few years ago. The computers on Robby require several hundred watts to perform their computations. The batteries, solar panels and/or radio-isotope thermal generators to supply that much power on a continuous basis would mass dozens to hundreds of kilograms. So even though the computers can be made small and light, the power to run their software is very heavy.

Several approaches have been tried to get around this difficulty. One approach was to go back to the Lunakod method of remote driving, and suffer with only having the robot move a couple meters an hour. But even this performance would be optimistic. To perform remote driving, the robot has to send a set of images back to the Earth every time the rover is to move. Another spacecraft subsystem that does not shrink well is the communications system. To send detailed images back from Mars requires either a large antenna and tens of watts of power or a small antenna and hundreds of watts, or hours of transmitting time. Since a small rover can not carry a large antenna, and since if it had hundreds of watts of power it could compute its own path, a small rover would probably take many hours to broadcast back the information operators on Earth would need to direct the vehicle. Using this method, the rover would only be able to move a few meters a day.

A more promising approach is to reduce the software on the robot without reducing its capability to autonoumously make its way to the various points of interest it is to explore. Most artificial intelligence researchers study how people do things, and try and duplicate either the form or functionality in a computer program. However, some researchers have been studying insects, and other simpler lifeforms. These researchers have developed a style of robot programming that has become know as behavior control. Insects seem to be dominated by collections of reflex behaviors. These reflexes keep bugs out of trouble (for the most part) and allow them to be very successful without requiring any deep cognitive reasoning on their part. While each of these reflexes can be as simple as the human knee-jerk reflex (or jumping backwards when a shadow crosses their field of view, in the case of a housefly) when combined with other reflexes the combined effect is one of seeming intention.

Researchers at the Massachusetts Institute of Technology, Teleos Research, IS Robotics, the Jet Propulsion Laboratory, and elsewhere have created many examples of behavior controlled robots. These robots have a variety of capabilities ranging from mapping an office building and delivering bagels, to navigating though Mars-like terrain, collecting soil samples and returning them to an analysis box. Perhaps the most impressive thing about these robots is that most of them use computers that execute less than a million instructions per second and have programs that are only a few thousand bytes in size.

To a large extent, behavior control relies much more heavily on the intelligence of the robot system designer then it does on the encoded intelligence in the robot. The designer must carefully examine what classes of situations the robot will probably get into, and what are simple schemes for getting out of or negotiating those situations. Rather than decomposing the robot's activities into functions (e.g., perceive, plan, execute, and monitor) behavior control breaks the robot's activities and programs into sets of simple skills (e.g., move around obstacles, maintain a given heading, deploy the arm). The robot may not have any awareness that it is negotiating hazardous terrain to get a sample from a designated spot, but the interaction of all of these behaviors with the environment will result in that being accomplished.

Small, behavior-controlled rovers (often referred to as "micro-rovers") are now being given serious consideration as part of several upcoming planetary missions. A rover massing less than five kilograms is a possible addition to parts of NASA's upcoming MESUR mission. MESUR will place several seismic and meterologic stations on the surface of Mars. The micro-rover would be used to deploy the seismometer away from the lander, and to bring minerological instruments directly to the rocks they are to observe. Another, slightly larger rover (approximately fifteen kilograms) has been studied for use as part of a Mars sample return mission to take place early in the nect century. The Russians are planning to launch a moderately sized (sevety-five kilogram) rover as part of their Mars-96 mission. They are also considering having a five kilogram micro-rover ride piggy-back on the larger rover, and then get off to gather samples, or explore especially hazardous or crowded terrain.

In the research laboratories, work on small rovers is continuing. New techniques are being developed to automatically relink sets of behaviors so that behavior-controlled robots can handle different sorts of tasks while using the same program. Researchers have also made great strides in micro-machining and the creation of micro-motors. These results have led to electrical motors smaller than a millimeter and etched out of the same materials (silicon wafers) as the computer chips used to control them. As these techniques progress, it may be possible to make entire planetary rovers all etched out of a single piece of silicon. The resulting robot would weigh less than a gram.

Future planetary missions may involve an orbiter that releases thousands of ant sized robots to float down to the surface. Once on the ground these robots would comb the surface testing for a particular set of useful materials (such as water). If they find what they are looking for, they signal the orbiter (perhaps by using a corner reflector to return laser light emitted by the orbiter). The refections could be imaged by the orbiter, and the results could tell scientists back on the Earth, where and howm much of the material they were looking for, is located on the distant planet.