Experiments With a Long-Range Planetary Rover

D. P. Miller, T. Hunt, M. Roman, S. Swindell, L. Tan & A. Winterholler
dpmiller@ou.edu timhunt@ou.edu bellx-1@ou.edu bacchus@ou.edu litan@ou.edu alois@ou.edu
University of Oklahoma Norman, OK 73019 USA

Abstract

SR2 is a solar powered autonomous planetary rover designed and built at the University of Oklahoma in collaboration with Malin Space Science Systems. This robot was created to explore the feasibility of long autonomous traverses through Mars-like terrain – traverses in excess of a kilometer a day requiring no more than a single communications between the robot and mission control. This is a substantially different strategy than is planned for near-term missions or was used in Mars Pathfinder. This paper gives an overview of the mechanical and control design of the system, and the results from its first field tests in the Salton Sea Desert.

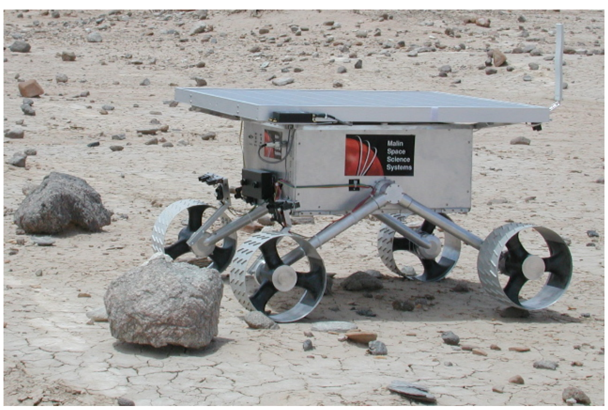

Figure 1: *The SR2 Rover Maneuvering Between Obstacles During a Field Test in the Salton Sea Desert*

1 Motivation

The Sojourner (Mars Pathfinder) Rover was incredibly successful – but it only explored an area of Mars smaller than the front yard of a typical suburban track-house. The MER mission, scheduled to launch is 2003, using two much larger "long-range" rovers, will explore a section of Mars the size of a few football fields. However, if we are to get an understanding of Mars, or even of a particular area on Mars, we need to have the capability of performing truly long traverses to get to separated sampling sites, but maintain a sense of context on how they are connected [3, 13, 7].

Sojourner and the MER rovers are mechanically capable of traversing much larger distances than have been done or are being planned [5]. But the control strategy of the rovers significantly reduces the distance the robots can travel. Effectively, the robots are never directed to travel over the horizon, or into an area that has not been carefully imaged [6]. When you are less than a meter tall and on a Mars-sized planet, the horizon is very close. More importantly, due to rocks, uneven terrain, and limited camera resolution, seldom can a significant path be imaged sufficiently to meet NASA mission managers' constraints of acceptable foreknowledge.

In addition, the kinematics and control of these rovers is fairly complex. Sojourner, like all of the planned NASA follow-on missions, is a rocker-bogie style vehicle, [1], or "rocky" for short. Rocky rovers have three wheels on each side, with the front two wheels on a passive rocker, and the two sides connected through a differential pivot onto which the chassis is affixed. Each wheel is independently driven (requiring six motors) and the four outer wheels are independently steered in Sojourner – and all six wheels can be steered in some of the more recent robots [11]. So Rocky style robots have ten to twelve motors dedicated just to the mobility system.

The Rocky rovers exhibit extraordinary mobility, especially in climbing fixed steps. However, it is not at all clear that these capabilities are needed in order to traverse significant stretches of Mars terrain.

The remainder of this paper will describe SR2 (see Figure 1). The SR2 Rover uses a simpler mechanical system than the Rocky Rovers which makes it easier to control and may also reduce the power needs of the robot. Experiments indicate that the SR2 rover has adequate mobility to easily manuever through realistic Mars terrain. The SR2 also uses a simplified control strategy. Rather than having a route planned out in great detail on the ground, the SR2 rover uses a sparse collection of waypoints to autonomously make its way through comparatively large stretches of terrain. The results from a field test carried out in June of 2002 will be reported.

2 SR2 Mechanical Approach

SR2 is a four-wheeled four-wheel-drive vehicle that uses two motors instead of twelve. It has a much simpler structure, more efficient use of energy (in most situations), and exhibits – under realistic situations – similar mobility capabilities to a Rocky style rover. While SR2 cannot climb as large a test step head on as can a rocky-style rover, such steps seldom occur in nature, and even when they do, there are a multitude of ways over them. By approaching a step at an angle, only one wheel has to climb at a time. The differential connection between the left and right sides ensures (in the same way as is done in the Rocky rovers) that all the wheels are in contact with the ground and the vehicle's weight is equitably distributed. SR2's wheels are open on both sides. The tires are wide but have almost no profile. When the robot skid steers in place, it rocks back and forth ensuring that it will not get caught against an obstacle, and that any soil that gets kicked up will pass through the tire, minimizing resistance.

These simple mechanisms, when combined with a simplified operational strategy, are able to autonomously move through hundreds of meters per hour of Mars like terrain.

3 Detailed Mechanical Design

SR2 is entirely custom machined from aluminum blocks. It is a four-wheel drive rover that uses skid steer to change direction. The motors and drive train run through a hollow tubular suspension. The main chassis is constructed of aluminum honeycomb to reduce weight while maintaining rigidity. At the center of the chassis is a geared differential that increases stability for the sensors and assures that all four wheels maintain ground contact with equal force.

3.1 Suspension

The suspension is the exoskeleton of the rover. It houses all of the bearings, shafts and gears used for motion. Three main tubes make up each side. A single large diameter (53 mm) thin walled tube houses the motor and connects the outer suspension with the internal differential. Two smaller (25 mm) thin walled tubes make up the lower external suspension that supports the chassis on the wheels. The tubes connect to each other through an upper gearbox in which the motor is also mounted coaxially with the large tube. The smaller tubes are 120^o apart and perpendicular to the large tube (see Figure 2). The end of the small tube is threaded to attach the second gearbox, which houses the wheel axle. The gearbox is designed around the internal components to keep the suspension compact. This keeps the suspension away from obstacles as it move through the terrain.

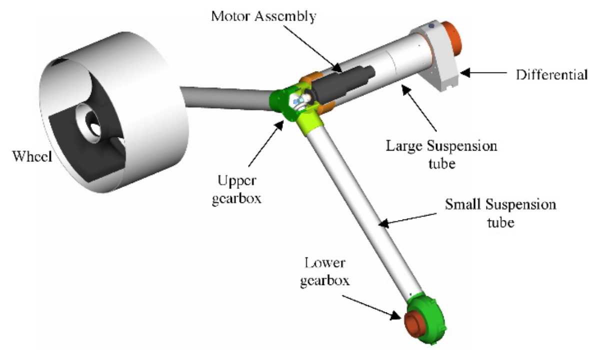

Figure 2: *Drive Train Cutaway*

3.2 Drive Train

The wheels must be driven efficiently and with enough torque to climb over obstacles of significant height. Preliminary ideas included a motor on each wheel, similar to current NASA rovers, but it was decided they would either not generate enough torque or consume too much power. However a single 10W DC motor would be able to produce enough torque to split between a pair of wheels. The driveline system must be compact enough to fit inside the suspension to keep contaminates out. A drive shaft system transfers the torque to each wheel, the efficiency is much greater than a belt drive system. Torque is increased through a 43:1 planetary gearbox connected to the motor, then passed through two beveled gear trains to the wheels. The output from a gearbox is passed to two drive shafts that are coaxial to the small suspension tubes. The other end of the shaft inputs torque into a 4:1 gear set at the wheel. The spur gear is bolted to the wheel axle in-between two bearings to decrease the effects of lateral or vertical force from the wheel. The drive train can produce 11.4Nm of torque at a wheel.

3.3 Chassis

The shape of the chassis is a simple rectangular box 45 x 35 x 20cm. The differential is in the center of the box attached to the base. The large suspension tubes protrude out of the sidewalls just above the base (Figure 3).

Weight and stiffness are major factors in the chassis design. A solid panel box would be too heavy even if pockets were milled out to decrease mass. The design used on SR2 has 2 sidewalls and a base

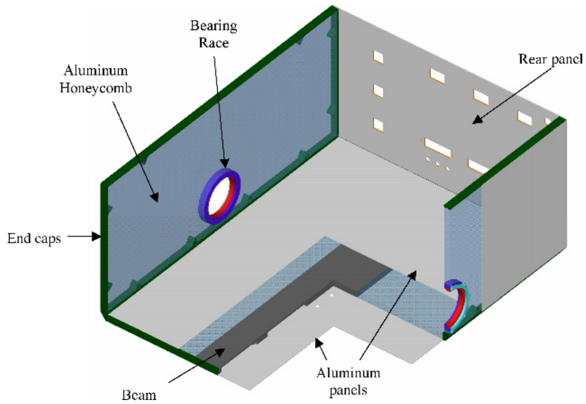

Figure 3: *Chassis Box Cutaway*

made of aluminum honeycomb. The front and back panels are thin aluminum plates with cutouts for terminals, switches, and gages. The honeycomb walls have bearing races for the suspension built into them. The differential is mounted to the base through an aluminum beam inside the panel. The beam runs down the centerline of the chassis and also serves as a place to hard mount the electronics. The top of the chassis is covered with a filter which allows air-flow while protecting the internal components from dirt. The chassis has a mass of 2.7kg.

3.4 Solar Panel Support

The mechanical interface between the solar panel (75 cm x 70.5 cm) and the main chassis is provided by a simple beam structure composed of Aluminum L-shaped beams. This structure lifts the panel above the chassis to provide adequate room for ventilation ($\approx 50mm$). In addition to the front and rear supports, side brackets extended from these main beams. These side brackets provided a stronger supporting structure that eliminated the tendency to wobble from side to side. The smaller beams (two on each side) reached outward and upward from the level of the main chassis to the outer rim of the solar panel. This supporting structure allowed for the panel to be removable.

3.5 Wheels

One of the goals in the design of SR2 was to simplify and improve upon the design and assembly of the rover wheels. Some of the shortcomings of the existing wheels used in Sojourner, FIDO and planned for use on MER are: the difficulty in assembly (each grouser in the tread is a separate welded part); the significant mass of the wheel; the closed outside wall and open inside can trap soil and rocks, further reducing performance; and the only shock absorbing structure in the drive train is the soil over which the rover travels – yielding an very hard ride.

3.6 Wheel Design and Manufacturing

The design of the wheels used on SR2 address each of these issues. Each wheel is constructed from a set of three identical carbon fiber composite spokes, a two-piece aluminum tire with an integrated formed tread, and a hollow aluminum hub which connects to the drive train. The overall dimensions of the wheels are 210mm in diameter and 110mm in width. The tread consists of three grousers per row and 26 rows spaced around the circumference. The grousers are 5mm high and run at an angle of 20 to the wheel axis. The carbon fiber composite wheel masses a total of 306g.

The wheel spokes are made of woven carbon fiber fabric and epoxy matrix material. The spokes were manufactured using the wet hand layup-method. To conduct this manufacturing process a two-piece dividable core and a four-piece outer mold had to be designed and manufactured (see Figure 4). The preshaped woven carbon fiber is laid up on the spoke core and covered with the epoxy resin. The finished spoke had a weight of 21.5g.

Figure 4: *Inner & outer wheel spoke molds*

The SR2 wheels integrate the tread into the wheel rim. No other connections like screws or rivets are needed to mount the studs on the wheel. This helps to keep the manufacturing time low. The tread patterns were specially developed for the differentially steered rover. To maximize the efficiency while steering, the tread pattern should be crosswise to the turning direction for maximum traction. Thus, the tread direction must be designed in such way that the grousers of the four wheels form an X where they contact the soil. The wheel treads were formed by punching the tread pattern into a formable material. The material used was 0.8mm thick 3003 aluminum alloy sheet metal. The punch process produces holes in the treaded rim. These holes are closed by a 0.3mm thick internal aluminum sheet metal rim. The external treaded rim and the internal rim were bonded together with epoxy resin.

In the wheel assembly process the three spokes are glued to the inner hub, and the outer rim with the tread is glued to the spokes. The parts have to be aligned carefully in this procedure. A fixture was designed to achieve this. The six components of the wheel can be seen in Figure 5. More details of the manufacture of the wheels can be found in [14].

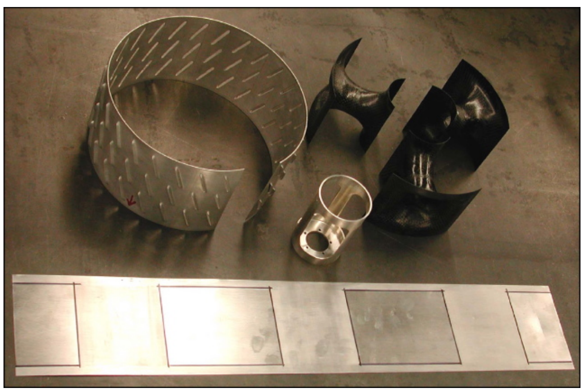

Figure 5: *The components of a wheel*

3.7 Mobility Power

The power required by the mobility system is approximately 8W for straight forward driving on level ground. The energy needed for turning in place is about 22W. Skid steering is not efficient, because the treads must slip against the ground. If a wheel is stalled, the mobility system can draw as much as 45W. The onboard control algorithms are designed to avoid stall situations and otherwise reduce the amount of turning required.

4 Navigation Hardware

The onboard computer system is designed to handle and process stereo-vision data, sensor data, commands from the operator, control motors, and transfer images and other rover data between the rover and the ground station. A PC104 Plus system was chosen to fulfill these needs. The CPU board has a PIII 400MHz processor for all computation. A Wavelan wireless PC card is used for communication between the rover and ground station. In addition, an I/O board, IEEE1394 board and quad-decoder board are used for sensor, camera and motor encoder readings.

Initially, stereo vision was intended to be the primary sensing system for obstacle and hazard detection. A Videre Design's stereo DCAM head was mounted on the front of the rover. We used SVS (Small Vision System) [4] as our stereovision process API. While the stereo system performed well during indoor testing, it did not work well in the extreme lighting conditions of the field test. Under these lighting conditions, the cameras lacked the dynamic range needed to produce images that could be correctly processed by the SVS. Because of this, stereo was not used for navigation during the primary field tests. Relieved of processing stereo, the PIII processor had many more times the computing capacity then was actually required by the robot.

The task of the sensors in the SR2 is to detect obstacles, monitor the current and voltage of the solar panel and system, and to position the rover. A combination of a magnetic compass and drive motor encoders are used to position the rover. The encoder counts between the left and right hand side are averaged to calculate distance traveled. All heading information was derived by the compass [2]. The compass and encoders were updated at 15 Hz and the compass had an accuracy of $\pm 0.5^o$. The compass and encoder readings are combined to regularly update the robot's position.

The obstacles in the environment with which the rover must deal includes big rocks, holes, and slopes over which the rover cannot climb. The characteristics of the desert environment must be considered when selecting sensors. The biggest problem we faced was high contrast and unusually bright lighting (especially in the near IR).

After extensive experimentation, we found that Sharp range sensors such as the GP2D12 [15] would not give accurate range data under the bright desert lighting conditions. However, they did prove to be reliable as threshold detectors under those conditions. The threshold range was extended by switching to a new lens package for the detector - the Sharp GP2Y0A02YK which provides longer range optics and better sun shielding. Using the sensors in a threshold mode to determine whether or not there was an obstacle in the robot's way, or whether or not the ground had fallen away proved to be much more noise tolerant.

Current sensors and voltage dividers were used to monitor system current and solar panel voltages. The system could switch to SLEEP mode if the battery voltage was too low and wake up when the battery is charged, though this capability was never tested in the field.

The sample rate of the I/O board also affects the sensor reading. If the rate is too fast, the system could not get stable readings, while the system could not respond in time if the sampling is too slow [12]. The sampling rate of Sharp sensors was set to 15Hz after all factors were taken into account.

The compass has direct heading, pitch and roll serial port outputs. The pitch and roll are used to determine if the slope is too high or the rover is rolling too much. The robot had programmed behaviors

Motor current was used to detect if a motor stalled and also as part of the power monitoring system. The power used by the sensors and computation system were also monitored.

5 Onboard Control & Rover Operations

The onboard control system has two modes of operation: manual and autonomous control. The default is manual control, in which the robot simply responds to commands issued from the remote user interface. In autonomous mode, it controls itself and avoids obstacles in order to reach a goal, or series of goals, specified by the user.

5.1 Autonomous navigation

For autonomous navigation, the user must first specify a mission. A mission consists of a goal, and optionally, a series of waypoints. The optional waypoints allow the user to set a path around large known obstacles in order to reduce the time required for the robot to reach the goal. Both waypoints and goal points are specified in rectangular coordinates, as is the current location of the rover. When starting a mission, the user will typically reset the rovers location to the origin to make specifying waypoints and goal points simpler.

When in autonomous navigation mode, the rover tries to travel in a straight line to the next waypoint or goal, but obstacles may lie in the chosen path. Not counting the stereo vision system, there are two methods the rover uses to detect obstacles: range sensors and electronic compass data. The rover's ten forward facing range sensors are clustered into three logical categories: right, middle, and left. When an obstacle is detected by the range sensors, the exact behavior chosen for avoidance varies based on the cluster that detected the obstacle. If the right cluster detects an obstacle, the rover backs up for several seconds, then rotates briefly to the left (counterclockwise) and attempts to travel forward for several seconds. This procedure is repeated until the right sensor cluster no longer detects any obstacles and the rover is able to travel forward without interruption for a specified number of seconds[8]. Once the rover has successfully gone around the obstacle, it recalculates the path to the next destination based upon its new location [9]. The same algorithm is used to avoid obstacles detected by the left sensor cluster, with the exception of rotating to the right (clockwise) rather than the left.

The middle sensor cluster is set back farther than the left and right sensors and faces slightly downward. As a result, obstacles detected by the middle cluster are typically detected when they are at a much closer distance than those detected by the left or right sensors. Additionally, since this cluster detects obstacles that are directly in front of the rover, a larger correction is needed to go around the obstacles. To compensate for this, the amount of time spent reversing away from an obstacle detected by the middle sensors was set to 2-3 times that for the left or right sensors. Aside from that, the algorithm for avoiding obstacles in the middle of the range of view was nearly identical to that for the left and right sensors. The middle cluster was actually subdivided into left and right sensors to allow easy determination of which direction to rotate when avoiding an obstacle.

In addition to range sensors, the rover would also take into account pitch and roll readings to avoid trying to navigate up or down too extreme of a slope. The rover would turn into or away from a slope to mediate the tilt between roll and pitch. If the total angle became too severe, the rover would treat that position as an obstacle, backing up and turning to avoid the severe slope area.

During field trials, this technique of obstacle avoidance proved quite effective after the sensors were properly calibrated and ideal durations of each movement phase (reverse, rotate, forward) were determined. The rover was even able to navigate out of semi-circular arrangements of rocks. By continually rotating in the same direction until no obstacle is detected it ensures, that at worst, the rover would have to backtrack the way it came to avoid complex obstacles, such as those that completely surrounded it. Ridge lines with strong slopes are large steps would be considered obstacles, and the rover would move down the ridge, periodically testing, until it found an area it could broach. The conservative settings of 15^o max slope and 8 cm max obstacle size proved to be only a minor hindrance to the rover's navigation (e.g., Figure 6). The rover was mechanically capable of going over slopes in excess of 25^o and ridges of 15 cm, however its dead reckoning accuracy would be severely degraded by the accompanying slips.

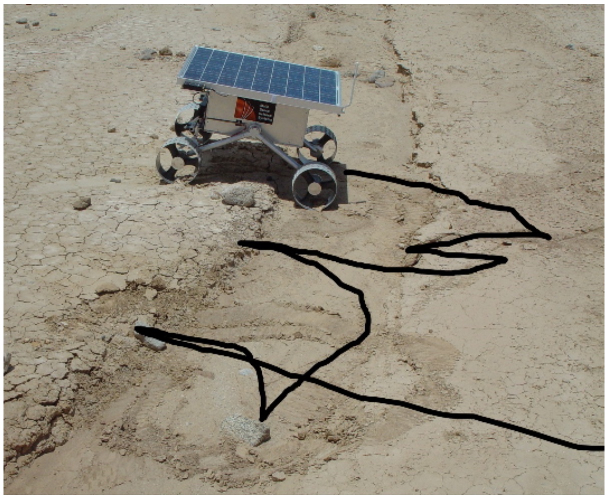

Figure 6: *SR2 moving along a ridge-line until an acceptable slope is found (pathline added)*

6 Defining Missions via the User Interface

The operator of the rover issues commands and receives feedback via a remote networked graphical interface. All Commands are human readable text and can be entered on a simple command line if desired, or more easily via graphical controls. Most of the sensor and status feedback sent from the rover is also in the form of human readable text with the exception of images, obstacle data, and mission path data. The graphical display and controls were divided into two separate panels: engineering and mission control. (see Figure 7.

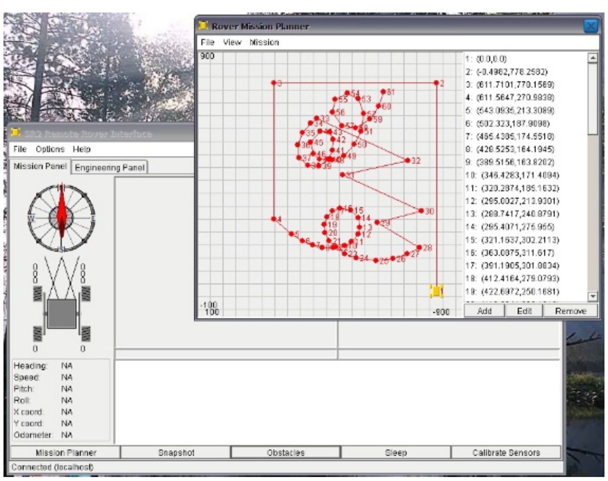

Figure 7: *The Robot Control User Interface (Note the designated path is for illustrative purposes only)*

The engineering panel is primarily intended for testing purposes and monitoring low-level system functions. It displays real-time log file updates from the rover. It also contains a simple array of nine buttons to manually control the movement of the rover in any direction, as well manual controls for each motor. In addition, this panel also displays the encoder readings of the motors to give a rough estimate of speed, as well as power meters for the computer, batteries, and solar panel.

The mission control panel offers more high-level functionality than the engineering panel. The mission panel contains a graphical compass to show the rovers current heading as well as pitch and roll readings. In addition to the compass display, the mission panel also graphically displays sensor readings arranged around a small icon of the rover to make it easier for the remote user to interpret the readings. More details about the SR2 control system may be found in [10].

7 Field Tests

The SR2 was tested in the SaltonSea desert in Southern California during the last week of June 2002. The area selected had similar rock distribution and slopes as those expected in certain areas of Mars that are of interest in the search for water [7].

During the first days of the test, it was determined that the stereo vision and range sensors did not function under desert conditions as they had indoors or in more shaded environments (see Section 4). A couple days were spent reconfiguring the sensor suite and adapting the robot's behaviors to work with the new arrangement.

The robot was then put through a series of short autonomous circular runs 100 to 300 meters in length. These runs helped to verify the robot's capability in traversing various types of terrain and also in verifying its dead reckoning capabilities. While we were not able to completely characterize the robot's ability to know its position against ground truth, in circular courses the robot was consistently able to return to its starting point with an error of less than 3% of the distance traveled. This occurred during runs of hundreds of meters over ridges and around various rocky obstacles, and was a gratifying validation of the wheel design.

On the final day of testing, the rover was given a series of way points spaced an average of ≈120m apart with the closest pair about 25m apart and the farthest just under 200m. The total length of the traverse was 1.3km. The way-points were selected in order to steer the robot around major geologic features (e.g, a steep gully). These features were all of sufficient size that they would have been easily detectable from an orbiting camera of the kind planned for future Mars missions.

The way-points were uploaded to the rover. During the next five hours, the robot autonomously traveled between way points. The robot operated completely autonomously for navigation between way-points and for obstacle avoidance. Human intervention was required at one point when one of the rover's wheels failed. The failure was due to heat softening of the epoxy that bonded the tread to the inner wheel. The air temperature was above 47^{o}C, and the ground surface temperature was believed to be far in excess of that. The wheel was repaired by adding two rivets, and the robot continued on its way.

With the exception of the wheel repair, the robot operated autonomously. All power was supplied by the solar panel (a commercial 45 watt panel). Navigation was done by the rover's onboard system following the set of eleven way-points it had been given at the start of the mission. The path taken by the rover as derived from its dead reckoning log, can be seen in Figure 8.

The mobility system was quite capable and, when going straight, very power efficient. Turning required more power and time then on a typical Rocky style robot. However, even in complex terrain, only a small portion of a traverse is spent turning. We believe that the efficiency and mechanical simplicity of this design will more than make up for its other deficiencies, when compared to Rocky style robots, for most mission scenarios.

One of the most important lessons learned was to avoid watching the robot too closely. The lack of long range sensors and precision range sensing meant that the robot's decisions on how to avoid obstacles were not always optimal. Some members of the team exhibited significant frustration when the robot went down a ridge to the right, when going to the left would have been clearly the shorter path. However, the behavior scheme used was very robust. Within a few minutes the rover was back on course. The extra couple meters traveled, while a day or more of traverses for Sojourner, were an insignificant delay for SR2. By placing the way points sufficiently far apart, there were thousands or millions of possible paths to get from here to there. The time spent taking a somewhat sub-optimal path was insignificant when compared to the time required to continually monitor and upload new seemingly optimal traverse paths to a robot in a Mars exploration situation.

We believe that SR2 is a demonstration that rover technology can be both simple enough and mature enough to be given a chance to do some serious autonomous traversing. We will not be able to remotely explore Mars until we are willing to delegate a significant amount of a robot's control to the robot.

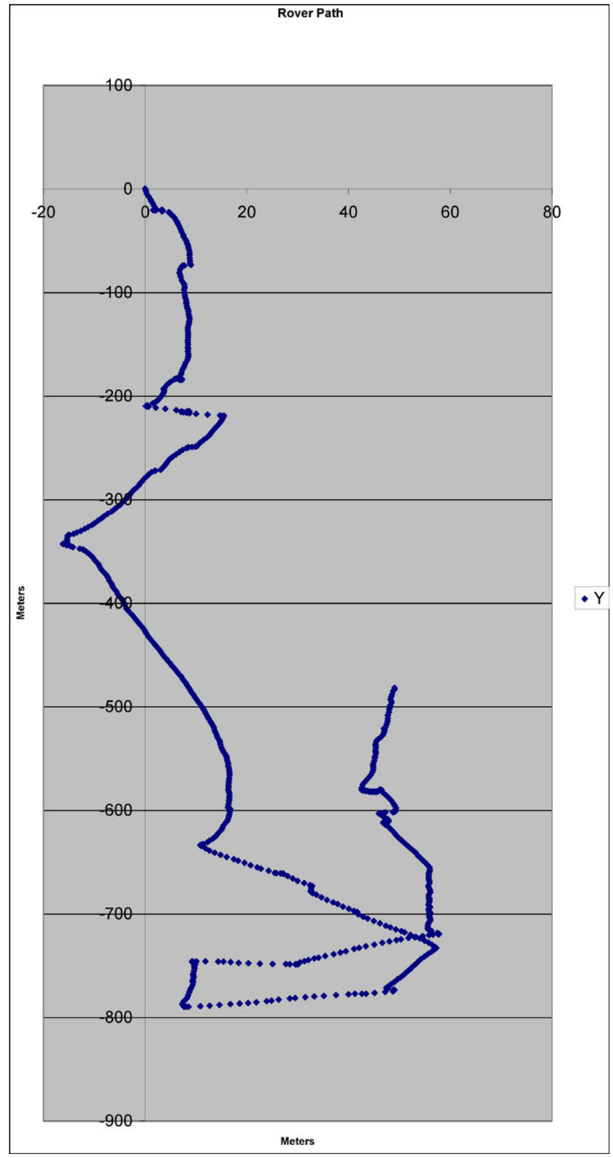

Figure 8: *SR2's 1.3km Traverse*

Acknowledgments

The authors wish to thank the SR2 science team Mike Malin & Mike Ravine. This work was funded in part by a grant from Malin Space Science Systems Inc.

References

[1] D. Bickler. The new family of jpl planetary surface vehicles. In D. Moura, editor, *Missions, Technologies and Design of Planetary Mobile Vehicles*, Toulouse France, 1993. Cepadues-Editions Publisher. IAF-97 - Q.2.07 1997.

[2] M. Bohlinger. Magnetic sensor description hmr3000. http://www.ssec.honeywell.com-/magnetic/description/desc_3000.html, 2002.

[3] N. Cabrol, E. Grin, V. Gulick, C. McKay, R. Greeley, M. Sims, and G. Briggs. Rover mobility & sampling strategy on mars: The case for gusev crater. In *Proceedings of the Lunar Planetary Science Conference #27*, 1996.

[4] K. Konolige. http://www.ai.sri.com/software/SVS, 2002.

[5] NASA Jet Propulsion Laboratory. A description of the rover sojourner. http://mars.jpl.nasa.gov/MPF/rover/descrip.html, 1996.

[6] NASA Jet Propulsion Laboratory. Rover mission overview and objectives. http://mars.jpl.nasa.gov/MPF/rover/mission.html, 1996.

[7] M. Malin. Personal communication., April 2002.

[8] D. P. Miller. Rover navigation through behavior modification. In *Proceedings of the NASA Space Operations Automation and Robotics Workshop*, Albuquerque, NM, June 1990.

[9] D. P. Miller, R. S. Desai, E. Gat, R. Ivlev, and J. Loch. Reactive navigation through rough terrain: Experimental results. In *Proceedings of the 1992 National Conference on Artificial Intelligence*, pages 823–828, San Jose, CA, 1992.

[10] D. P. Miller, L. Tan, and S. Swindell. Simplified navigation and traverse planning for a long-range planetary rover. In *Proceedings of the 2003 International Conference on Robotics and Automation*. IEEE, May 2003.

[11] P. S. Schenker, E. T. Baumgartner, R. A. Lindemann, H. Aghazarian, A. J. Ganino, G. S. Hickey, D. Q. Zhu, L. H. Matthies, B. H. Hoffman, and T. L. Huntsberger. New planetary rovers for long-range mars science and sample return. In *Intelligent Robots and Computer Vision XVI: Algorithms, Techniques, Active Vision, and Materials Handling, SPIE Proc. Vol. 3522*, page 215, Boston, MA, 1998.

[12] D. B. Stewart. How to choose a sensible sampling rate. *Embedded System*, pages 20–27, July 2002.

[13] C. Stoker. The search for life on mars: The role of rovers. *Journal of Geophysical Research*, 103(E12):28557–28575, 1998.

[14] Alois Winterholler. Development, design and verification of different wheel concepts for nasas field integrated design and operations rover. Master's thesis, Fachhochschule Augsburg, Hochschule für Technik Wirtschaft Gestaltung, University of Applied Sciences, November 2002.

[15] J. C. Zufferey. Sharp gp2d12. http://dmtwww.epfl.ch/j̃zuffere-/SharpGP2D12_E.html, April 2001.