Sinkage and slippage estimation for an articulated vehicle

Brian H. Wilcox*
Supervisor, Robotic Vehicles Group
Jet Propulsion Laboratory, California Institute of Technology

ABSTRACT

This paper describes an approach to estimating in real-time the degree to which an articulated robotic vehicle is undergoing wheel slip and/or sinkage in soft terrain. Robotic vehicles generally have hazard avoidance sensors which measure the shape of the sensible surface, and these can be used to predict what the articulation pose of the vehicle will be as it moves over the surface. An articulated vehicle (one with three or more wheels on each side) can directly measure the shape of the loadbearing surface by combining inclination and articulation sensing. Delays between the actual articulations and the expectations can be explained by wheel slippage. Differences between the expectation and the actual articulations can be explained by sinkage below the sensed surface. If one assumes that successive wheels on each side follow the same profile as the front wheel (sinking the same amount, if any, into the soil), then it is possible to estimate sinkage and slippage separately. A Maximum-A-Posteriori estimation procedure formalizing this heuristic approach is developed and simulated, and the results presented and discussed.

Keywords: robotic vehicles, sinkage, slippage

1. INTRODUCTION

Robotic all-terrain vehicles (such as Mars rovers landed in 1997 (Figure 1) and planned for launch in 2003 (Figure 2) and 2009) are required to safely navigate over largely unknown surfaces and to reach precise locations of interest. The dead reckoning accuracy for such vehicles is frequently 10% of distance traveled or worse[1], which is often unsatisfactory in terms of reaching targets despite the fact that such targets are clearly visible in prior imagery. One of the major sources of dead reckoning accuracy is odometry error, which is exacerbated for all-wheel-drive vehicles by the lack of any unpowered wheel giving improved estimates of over-the-ground distance traveled. The driven wheels slip in a somewhat unpredictable way, especially in soft terrain, accounting for the bulk of the dead reckoning error.

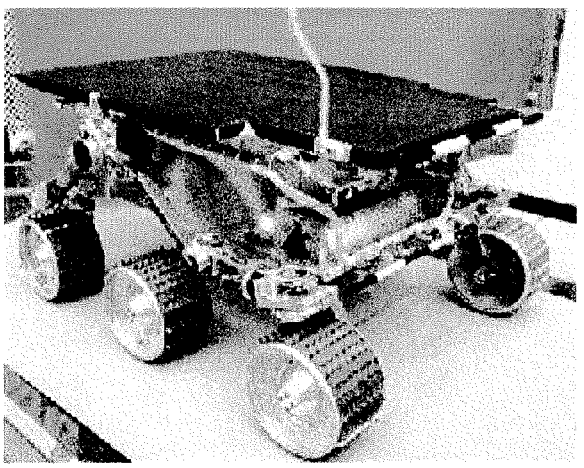

Figure 1: Sojourner Mars Rover landed on Mars in 1997

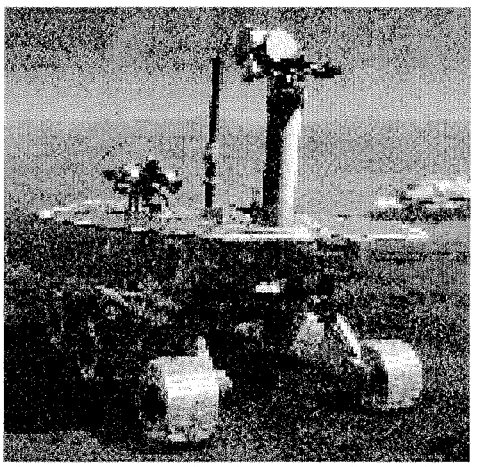

Figure 2: Mars Exploration Rover planned for launch in 2003

* Brian.H.Wilcox@jpl.nasa.gov; phone (818) 354-4625; fax (818) 354-8172; 4800 Oak Grove Dr. M/S 107-102, Pasadena, CA 91109

In addition to causing these annoying odometry errors, slippery or soft terrain can pose a threat to the safety of the vehicle. Extreme slippage can potentially move the vehicle into unsafe or unrecoverable terrain despite all efforts of the control system to correct for it. Sinkage of the wheels into soft terrain can high-center the vehicle, unloading the wheels and making further mobility impossible. Conventional ranging sensors (using lasers or stereo cameras) generally fail to warn of conditions where the loadbearing strength of the soil is too low to support the vehicle adequately. Other non-contact sensors, such as ground penetrating radar, often have ambiguous interpretations. It is desirable to find a way for a robotic vehicle to detect sinkage hazards before they become catastrophic, and also to improve upon the dead reckoning errors introduced by wheel slippage.

Such soft material is not uncommon on Mars, as on Earth. For example, at the Viking 1 landing site, about 14% at the surface was drift material[2] and one of the landing legs sank 17 cm into that material. Previous studies of non-geometric hazard detection for planetary rovers, which assumed very large (~1000 Kg) rovers, have focused on Ground Penetrating Radar to detect subsurface hazards.[3,4] However, mass and power constraints for less ambitious missions lead us to desire means to detect these hazards without requiring additional mass, power, or complexity beyond the basic vehicle configuration.

2. APPROACH

Articulated vehicles (those having more than two wheels on each side) can be easily instrumented so as to provide sufficient information to estimate the magnitude of both sinkage and slippage. Imagine an articulated six-wheeled vehicle sitting on the terrain. The needed instrumentation measures the pose of the vehicle (the height of each wheel relative to the body, and the overall body pitch and roll), and uses a ranging sensor to give a local terrain map ahead of the vehicle. (Generally such sensors are already included with most robotic vehicles since they are needed for conventional hazard detection and avoidance algorithms. Thus the proposed algorithm requires no new hardware for typical applications.) Given the local terrain map, it is possible to form an expectation of what the vehicle pose will be if it moves a short distance ahead. For example, if the vehicle moves one axle-to-axle distance "A" forward, then the rear wheel moves to where the middle wheel was, and the middle wheel moves to where the front wheel was, and the front wheel moves onto the terrain which has been mapped by the ranging sensor. We would thus expect the difference in height between the middle and back wheels to be the same as the previously-measured elevation difference between the front and middle wheels in the previous rover position. This would be true unless the rover has not moved as far as expected, from slippage, or it could be because the terrain is compressing, causing the rover to sink into the terrain.

One simple assumption that seems plausible for most situations allows us to disentangle sinkage from slippage. Specifically, we assume that successive wheels on each side of the vehicle follow in the tracks made by the front wheel on that side, and that they each compress the terrain by the same amount as the front wheel. While this may not always be true, it is plausible that in most circumstances the compression of the terrain by equally-loaded wheels will be the same. And only in the event of serious side-slip will the wheels not track along the same ruts, at least when the vehicle is nominally going straight.

Using this assumption, it is easy to see how we can estimate sinkage and slippage. If we retain a memory of the front-to center height differences on each side, and of the center-to-rear height differences, we would expect that a plot of these two curves to be identical except delayed by the axle-separation distance A. We can assume that we have moved forward by a distance A when the values of these respective curves match. We have an independent check on the validity of our assumptions by checking the right and left wheels separately. A close match in the forward motion estimate provided independently by the right and left wheels means that our assumptions that the wheels track in the same ruts and compress the terrain by a uniform amount are probably valid. A radical disagreement between the two estimates suggests that the assumption is not correct for this particular terrain. When the two values agree, we can compare them to the simple estimates based on odometry to estimate the degree of wheel slip.

Once we estimate the actual forward motion of the vehicle, we would like to estimate the sinkage as well. If we see (with the ranging sensor) that the terrain one axle-separation distance A ahead of the front wheel has a height H compared to the current front wheel contact point, then we can expect that the height difference between the front and

middle wheel when we move forward by A will be H. Since we have an independent means for estimating forward motion, we can accurately determine when we have moved forward by a distance A using the slippage estimate. Any residual in the actual height difference between the front and middle wheel and the predicted value H is due to sinkage. Thus we can estimate the magnitude of the sinkage of the front wheels in the terrain. Once again we can perform this independently for both the right and left wheels. If images of the terrain show that the texture is roughly constant, then we can guess that the terrain properties are roughly constant over the width of the vehicle. In that case, if the sinkage estimates for the right and left sides roughly agree then they are probably valid. If they radically disagree, then again the underlying assumptions are probably violated.

3.0 MAXIMUM A POSTERIORI ESTIMATION OF SINKAGE AND SLIPPAGE

We now attempt to formalize this heuristic argument mathematically[5]. We consider a planar model of one side of the vehicle as shown in Figure 3. Specifically, there are three wheels connected with passive but instrumented linkages so that they remain in contact with the surface as they roll. By combining the pitch and articulation sensor values we can compute the difference in elevation between the rear wheel and the center or front wheels (call these Z(1) and Z(2), respectively). We also assume a look-ahead ranging sensor that examines a number of discrete points on the ground ahead of the vehicle. Again, by processing the sensor data, we can compute the elevation difference between the rear wheel nominal contact point and the elevation of each sensed point on the ground ahead of the vehicle (call these Z(3)... Z(N)). Needless to say, all these measurements have noise that must be accounted for in the analysis.

We assume that undisturbed terrain in this planar model has an elevation function y(x), where y is the elevation at a point x along the horizontal axis. When the vehicle moves ahead, the front wheel sinks in the soil by an amount s(x), so that it rolls along in contact with the function y(x)-s(x). We assume that the trailing wheels do not further compress the soil (since the wheel loading of this vehicle is assumed to be roughly uniform). Thus they also track along y(x)-s(x). This is a key assumption which, if not approximately correct, will lead to a general failure of the entire approach, but which fortunately can be checked by comparing the independent estimates coming separately from the two sides of the vehicle. If the wheels all turn at the same rate (which is reasonable since they are geared so

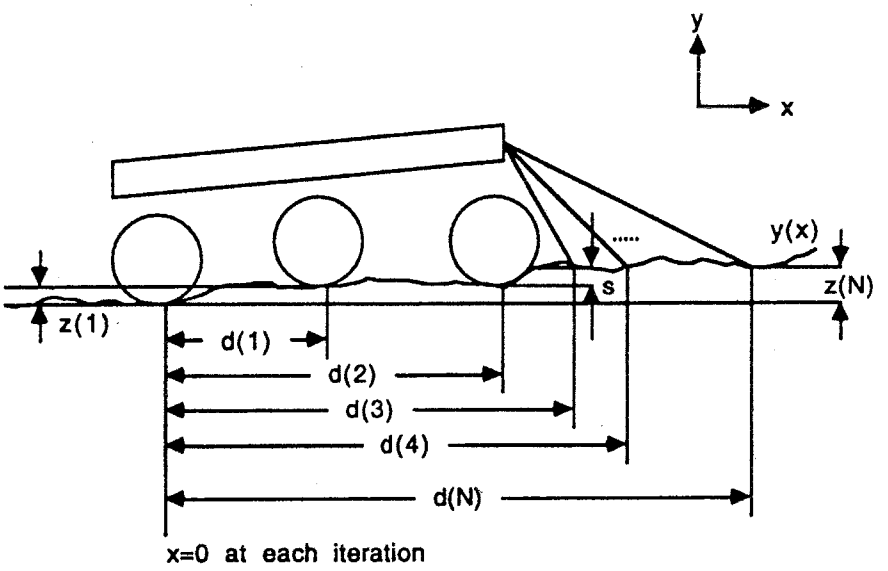

Figure 3. Planar model and symbol definitions

low that in normal terrain they run effectively at the no-load speed), then when the wheel circumference has moved a distance w the vehicle will advance some distance x in the horizontal direction, usually less than w, due to wheel slippage. This slippage will generally be a function of the type and slope of the soil.

The objective of this analysis is to estimate x and s(x) given the 'odometer' reading w, the values of z(1)... z(N) and the associated measurement noise v(1)...v(N). Intuitively this should be possible, since if y(x) and s(x) were known exactly up to the forward-most sensor (a ranging sensor for y and the front wheel for s), then for a given Δw, there would usually be a unique Δx which would allow all the sensor readings to match their predicted values. In other words one would 'slide' the rear and center wheels along the curve y(x)-s(x) until the observed elevation difference z(1) is matched between $x+\Delta x$ (the new position of the tear wheel) and $x+\Delta x+d(1)$ (the new position of the center wheel), which would fix Δx. Then one would use the measured elevation of the front wheel to compute y-s at that point (thereby extending our knowledge of s(x) forward by Δx). Similarly, we would use the measured elevation at the forwardmost range sensor to extend our knowledge of y(x) by Δx. This process would repeat so as to build an arbitrary sequence of Δx, s(x+d(2)), and y(x+d(N)) values. We would, of course, assume a $y_0(x_0)$ value as the starting elevation and position of the rear wheel. (Knowledge of the initial y(x) and s(x) functions between x and x + d(N) is trivial since the vehicle will disembark from the lander along a surface of known geometry and with negligible slip and sinkage.)

One potential problem with this approach is that values of z(1)...z(N) will not be taken densely along the vehicle trajectory. For a robotic vehicle to successfully detect and avoid obstacles, the terrain is generally mapped at least as densely as the scale of the wheel contact patch. This is adequate for our purposes, since variations finer than the size of the wheel contact patch will not be seen in the wheel excursions.

There are several issues that need to be considered with this model. First, if the vector $\mathbf{Z}(k)=col(z(N),...,z(1))$ at cycle k is measured on terrain which is very flat (compared to the measurement uncertainties v(k)) then we would still like to have a reasonable estimate of forward travel. This suggests that we should have a prior model of the distribution of the slip x(w), and that we should form a Maximum A Posteriori (MAP) estimate of the slip.[6]

Following our heuristic argument above, if we were to "slide" the vehicle along until the observed elevation difference z(i) is matched, this corresponds to generating a discrete set of values y(i), i=0,...M1 and s(i), i=1,...M2 which can be thought of as our best estimate 'histograms' (i.e. discretized piecewise constant representations) of the y(x) and s(x) functions. The horizontal density of these estimates should be sufficiently great to allow accurate models of the terrain for purposes of simulation, but not so great as to unduly burden the downlink data volume. Since the wheels mechanically average the terrain over a length equal to the tire contact patch (about a third of a wheel radius) we would tend to discretize the model at about this level. Thus we might have M2=30 or so and M1=6O.

Thus we can now outline a procedure for estimating the sinkage and slippage of the rover:

1) Measure the elevation differences z(1)...z(N).

2) use previously-estimated (described below) histograms y(i), i=0,...,M1 and s(i), i=0,...,M2, as well as a Gaussian prior distribution for Δx with mean m, and variance σ_x^2 to compute the (nonlinear) MAP estimate for Δx. We assume the distribution for measurement noise for each z(i) is also independent and Gaussian. Since the MAP estimate of independent Gaussians is a weighted least-squares estimate, we compute:

$$\min_i\Big(\sum_{j=3}^{N-1}[(1/\sigma_{z(j)}^2)(z(j)-y(d(j)+i) - y(i)+s(i))^2] + (1/\sigma_{z(1)}^2)(z(1)-y(d(1)+i) + s(d(1)+i) - y(i) + s(i))^2 + (1/\sigma_x^2)(i-m_x)^2\Big)$$

The interpretation of this expression is a follows: to maximize the posterior probability which is the product of exponentials, we need to minimize the magnitude of the exponent. If we let i be the histogram bin which we assume the rear wheel has advanced to (and changed to a elevation y(i)-s(i)), then the summation from j=3 to N-1 is of squared errors between the ranging sensor elevations and the corresponding y values in the histogram. The next term is the weighted squared error for the middle wheel incorporating the histogram data for s(1) a well as y(1). The last term is from the Bayesian prior distribution. Note that z(2) does not even appear in this expression, as the advance of the front wheel involves an unknown amount of sinkage in the soil and so there is no histogram data with which to

compare. A similar situation arises with z(N) in the summation, since y(x) unknown ahead of the forwardmost sensed point.

We implicitly assume that the forward advance is not so great as to push the next sensed point z(N-1) off the end of the histogram, although this could be accounted for if necessary. We would then perform a parabolic interpolation the weighted-sum-of-squares to get a refine estimate of Δx to a fraction of a histogram bin. While not strictly valid, interpolation of the error function should be better than taking integer bins, while not as computationally intensive as the more conceptually-correct approach of computing the minimal error function on interpolated data. Note also that we could compute m_x as a function of the data here prior to finding the minimum over i to account for the fact that our expected slip is a function of average terrain slope. For example, we could compute

$$(\sum_{j=1}^{N} (z(j)/d(j)))/N$$

as an estimate of the slope and compute some linear or non-linear function of this to compute m_x. We could also modify the estimates of m_x and σ_x^2 using prior estimates to adapt and refine our Bayesian prior,

3) Now that we have an estimate for Δx, we translate the histograms for y and s forward by Δx and up by $y(\Delta x)$. This requires interpolation, due to the non-integer nature of Δx, so we assume that linear interpolation between adjacent points is adequate (again to reduce computational complexity). We extend our knowledge of y forward by linear interpolation from the translated old y(N) value to the observed z(N) at d(N). Similarly, we extend our knowledge of s forward using linear Interpolation from the translated s(d(2)) to a new forwardmost value s(d(2))=y(d(2))-z(2).

4) We need some way to incorporate the new measurements into the histogram for y (otherwise only the forwardmost measurement y(N) will play a role in defining the function, which seems to waste a great deal of valuable information). Note that between the old d(N) and the new d(N) we have a linear approximation to y(x). When the vehicle moves forward by Δx (generally less than d(N)-d(N-1)), we will get a new value for y from z(N-l) which will, in general, not lie on the previous linear approximation to y. Since we expect that our measurement noise $\sigma_{z(N-1)}$ will be quite small compared to the grossness of the linear interpolation, we would like to force the histogram to conform to the data at this point (the new d(N-1) point). We would also expect y(x) to be a continuous function, so that nearby points should also be modified. For simplicity, we will assume that adjacent histogram bins will be updated by 'splitting the difference', i.e. they will be reassigned values halfway between the new measurements of y based on each of the z(i) measurements for i<N and the old (but translated) histogram value. This is an ad-hoc assumption made in the interests of computational simplicity which will hopefully allow a fairly accurate estimate of y(x) to be generated as all of the sensors sweep over the surface . We can perform a corresponding process for s(x) by assuming that deviations between z(1) and y(d(1))-s(d(l)) are due to errors in the measurement of s and not y, which makes some sense because by this time the histogram for y has been refined with multiple measurements while the histogram for s has been generated only by piecewise linear interpolation out to the single measurement at z(2) (i.e. the front wheel).

5) Lastly, move the vehicle forward and repeat the cycle.

This model and analysis are very simple and somewhat suspect from a theoretical point-of-view. However, as in many practical applications, real-time performance, computational complexity, and availability of sensor data are of paramount importance, with the alternative being not to do any estimation at all. Thus we would like to know what the performance of this simple estimation procedure is, and to what degree it gives improvement over use of the prior mean m_x to estimate over-the-ground distance traveled and not estimating sinkage at all (and accepting the risk of getting stuck). We would also like to evaluate the usefulness of having more ranging sensor measurements as opposed to fewer, since each additional measurement has cost.

A Monte Carlo numerical simulation of this algorithm on synthetic terrain has been performed, with the results shown in Figure 4. The synthetic terrain is developed by uniformly sampling a linear combination of sine waves, whose amplitude is random over a uniform range extending from zero to some fixed multiple of the wavelength

(thereby ensuring scale invariance), and whose phase is random over $[0,2\pi]$. Twenty different wavelengths are combined over the range from 1 cm to 1.9 meters, with each one 30% longer than the previous one. This range encompasses all scales of interest: smaller scales average to zero over the bins (and the wheel contact patch) and longer scales are virtually flat over the length of the vehicle and its look-ahead ranging sensor. (Note that the smaller scales will exhibit substantial aliasing when binned, which is an important and real effect that needs to be modeled by the analysis.) As mentioned before, a 'smooth' simulated terrain is realistic here, since presumably the normal hazard detection system will avoid rough or discontinuous terrain. Measurement noise for the articulation sensors is taken to be 0.04 mm, while the measurement noise for the ranging sensor is 2 mm, as determined for one of the Sojourner software development test vehicles.

The interpretation of Figure 4 is as follows. When the actual slip was 20% (meaning that the actual forward motion of the vehicle was 80% of the rim motion of the wheel), even relatively small surface roughness is adequate to detect both slippage and sinkage. In Figure 4, one of the curves shows the case where the standard deviation of the Bayesian prior is 0.05 cm (1% of the actual advance) when the slippage of the actual vehicle advance per cycle is 1 cm. This artificially "overweights" the Bayesian prior to show the smooth transition from 20% slippage to very small errors as the terrain gets rougher and the match between the articulation profiles of the wheels becomes overwhelming. However, the chassis articulation sensors are so accurate that we can do much better than this. When the Bayesian prior is taken to be 10% of the actual, even on relatively smooth terrain, the improvement in slippage estimation is very good whenever the surface roughness is large compared to the articulation measurement error. Improvements to the odometry to give an error of 1% or less seem quite possible with this technique. When the RMS surface roughness is above 1 cm, the measurements of the vehicle pose and articulation are so accurate that the estimate of the actual amount of slippage can be much better than 1%, even when the "goodness" of the original wheel odometry is seriously overestimated (1 sigma=10%). Given an accurate estimate of the vehicle motion, the sinkage of the front wheels can be accurately estimated as limited primarily by the noise in the ranging sensor.

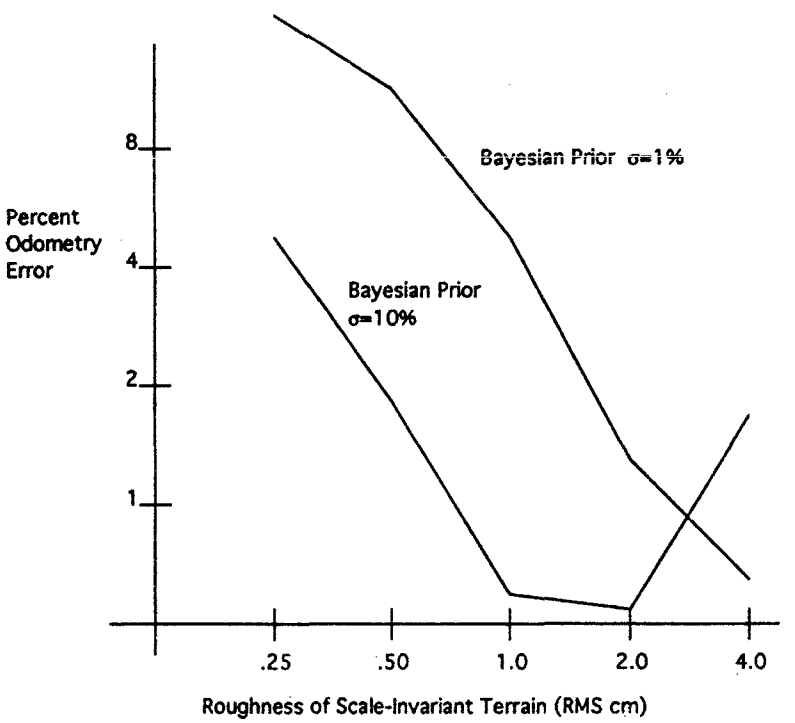

Figure 4. Odometry error as a function of terrain roughness (20% actual slip)

4.0 SUMMARY AND CONCLUSIONS

It is relatively straightforward, as was done with Sojourner, to create a hazard avoidance algorithm that deals effectively with a relatively sparse rock field. However, future robotic vehicles traveling vastly farther than the ~100 m traversed by Sojourner will occasionally face sinkage and slippage hazards that are complex and potentially catastrophic, and will in any event find such hazards make it difficult to accurately reach desired targets. We have outlined here a technique which is expected to robustly estimate the degree of sinkage and slippage using sensors which are likely to be incorporated into such vehicles anyway. This algorithm has an internal consistency check, in that it can be separately applied to the right and left side wheels of the vehicle, creating separate estimates of slippage and sinkage. If these estimates approximately agree (or account for unexpected yaw in the vehicle heading), then one can be reasonably confident that both are correct. If they disagree, then one is led to suspect that the primary underlying assumption (that the wheels track the same loadbearing surface profile) is not correct. In that case, one can always revert to the baseline assumptions concerning the forward advance expected given the odometry measurements.

Clearly it is important to validate this approach on a real vehicle, especially to determine if the key hypothesis that the wheels track the same loadbearing surface is correct on a wide variety of terrains. In the case of extreme sideslip, of course, this hypothesis will not be valid, nor will it in the case for vehicles where the wheelbase varies significantly for the different wheel pairs from front to back. However, since maximum vehicle mobility performance will be achieved when all the wheels are approximately equally loaded, it seems likely that many high-mobility vehicles (6-or-8 wheeled, or tracked) will give sufficient articulation data to feed this algorithm and give useful results.

Several extensions and improvements to this technique are possible. One is to begin to estimate the "slip versus slope" function of the terrain. This function relates the wheel slip (in percent) to slope. By estimating this function, the Bayesian prior used to initialize the MAP estimation procedure is likely to be greatly improved. This will take into account, for example, the intuitively correct fact that normal odometry tends to overestimate the distance traveled by the vehicle when going uphill, but it can underestimate the distance traveled going downhill. This slip-versus-slope function can further be estimated as a function of observed terrain visual texture. That is, one could categorize the apparent texture of the terrain (using some classification technique applied to the vector of outputs from a collection of spatial filters, for example) and then separately estimate the slip-versus-slope function for each of these terrain classes. One may quickly be able to determine that some textures imply nearly non-traversible terrain, and update the hazard avoidance algorithm so that these texture classes are treated as hazards. One may also be able to use this slip-versus-slope estimate in resource planning. For example, the energy or time needed to climb a sandy slope will be a strong function of the slip. If this slip can be accurately estimated in advance, this resource can be predicted and can be factored into the route planning algorithm or in mission planning (e.g. for reordering or pruning the activity list when resources run short).

5.0 ACKNOWLEDGEMENTS

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. The author gratefully acknowledges the assistance of Jack Morrison and Tam Nguyen in the conduct of this work.

6.0 ABOUT THE AUTHOR

Brian H. Wilcox is the Supervisor of the Robotic Vehicles Group at the Jet Propulsion Laboratory. The Robotic Vehicles Group was responsible for the control and navigation of the Sojourner rover which explored part of Mars in the summer of 1997, including design of the electronics and navigation sensors, hazard detection and avoidance, and mission operations. The group is currently involved in similar activities for the Mars Exploration Rovers to be launched in 2003. Mr. Wilcox was personally responsible for conceiving and developing the stereo camera and laser ranging system and hazard avoidance algorithm used on the Sojourner rover. Mr. Wilcox led the "Rover Performance Assessment" technology experiment on Sojourner. Mr. Wilcox holds an MS in Electrical Engineering from the University of Southern California with an emphasis in Signal and Image Processing, a BS in Physics and a BA in Mathematics from the University of California at Santa Barbara. He was awarded the NASA Exceptional

Engineering Achievement Medal in 1992 for his work on planetary rovers, and the JPL Award for Excellence in 1999 and 2001. He holds four U.S. patents, with three more pending.

7.0 REFERENCES

[1] Wilcox, B. and Nguyen, T. "Sojourner on Mars and Lessons Learned for Future Planetary Rovers", Society of Automotive Engineers publication 981695, 1997.

[2] Moore, H. J. et al, 'Surface Materials of the Viking Landing sites", *Journal Geophysical Research*, v. 81, pp 4497-4523,1977.

[3] Douglass, R. J., "Surface Property Determination for Planetary Rovers", JPL Contract 958706 Quarterly Report, Martin Marietta Corporation, Denver, CO, 16 April, 1990.

[4] Olheft, G. R. and Strangway, D. W., 'Electrical Properties of the Surface Layers of Mars, *Geophysical Research Letters*, Vol, 1, #3, July 1971.

[5] Wilcox, B. "Non-Geometric Hazard Detection for a Mars Microrover", proc. Conference on Intelligent Robotics in Field, Factory, Service, and Space (CIRFSS '94), pp 675-684, Houston, TX, 23-24 Mar '94.

[6] See, for example, Mendel, J. M. *Lessons in Digital Estimation Theory*, Prentice-Hall, Englewood Cliffs, NJ, p114, 1987.