AIAA-94-1275-CP

NON-GEOMETRIC HAZARD DETECTION FOR A MARS MICROROVER

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

Abstract

Non-geometric hazards (i.e., those which cannot be characterized soley by their shape, but instead are related to mechanical properties such as strength and friction) may pose a significant risk to planetary rovers. This paper describes a means for an articulated vehicle to detect sinkage and slippage in such material so as to prevent entrapment and to correct for dead-reckoning errors. Simulation results and preliminary indications of test data are described.

Introduction

For an exploring vehicle to move safely over the surface of another planet, it is potentially important to know if the vehicle is sinking into very soft surface material or is experiencing high levels of wheel slip. For example, at the Viking 1 landing site, about 14% of the surface is drift material, and one of the landing legs sank 17 cm into that material.[1] Previous studies of non-geometric hazard detection for planetary rovers, which assumed very large (~1000 Kg) rovers, have focussed on Ground Penetrating Radar to detect subsurface hazards.[2,3] However, mass and power constraints for microrover missions lead us to desire means to detect these hazards without requiring additional mass, power, or complexity beyond the basic vehicle configuration.

For the purposes of this discussion, we consider the mission model of NASA's Mars Environmental Survey (MESUR) Pathfinder project, scheduled for launch to Mars in November, 1996. In this mission, a microrover with a mass of under 10 Kg will traverse over the terrain within a few tens or hundreds of meters from it's lander to goal points selected by ground-based analysis of images taken by lander stereo cameras. These goal points are selected for their scientific interest, and it is important that they be approached quite accurately (for example to take a spectrum of a particular rock). Thus safety and improvement in the accuracy of dead-reckoning navigation are important reasons to develop a reliable means for estimating the sinkage and slippage of the rover wheels. Means for detection and avoiding geometric hazards are described elsewhere.[4]

The Pathfinder rover is a six-wheeled "rocker-bogie" articulated vehicle. It will be functionally equivalent and the same size as our research vehicle "Rocky 3.2", shown in Figure 1.

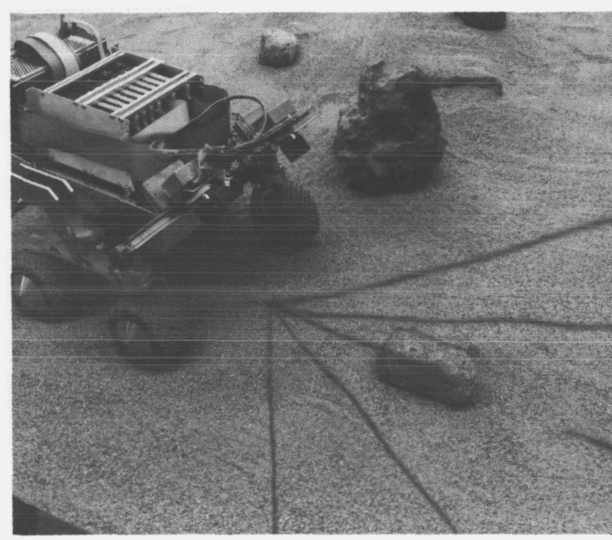

Figure 1. Rocky 3.2 vehicle with laser stripe sensors

Rocky 3.2 (one of a long line of Rockys) has sensors for wheel speed (all wheels are driven) and for determining the articulation angles of the chassis. (The articulations are passive so that each wheel follows the terrain contours independently.) It also has a look-ahead ranging sensor based on detecting, in a CCD image, the position of laser stripes projected ahead of the vehicle (Figure 1 is reproduced from a color original with a blue filter so the red laser stripes show as dark lines). Because the computation on-board the rover is very limited (an 8085 CPU, about 20 times less powerful than a typical personal computer), it is important that only a small amount of sensor data be taken and processed. Thus it is important to formulate simple algorithms for estimation of slippage and sinkage, and to do performance evaluation based on the concept that only the wheel, chassis, and a minimum number of discrete measurements from the look-ahead sensor can be used as input to the system. If a simple algorithm gives good performance, in terms of improvement of dead reckoning vs basic odometry and in detection of hazardous sinkage conditions, then the increased computational load will be justified.

The Sinkage and Slippage Model

We consider a planar model as shown in Figure 2. Specifically, there are three wheels connected with passive but instrumented linkages so that they remain in contact with the soil as they roll. By processing 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 have a look-ahead ranging sensor which 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 which must be accounted for in the analysis.

Assumptions

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 roughly uniform). Thus they also track y(x)-s(x). This is a key assumption which, if not approximately correct, will lead to a general failure of the entire approach. If the wheels all turn at the same rate (which is reasonable since they are geared so 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.

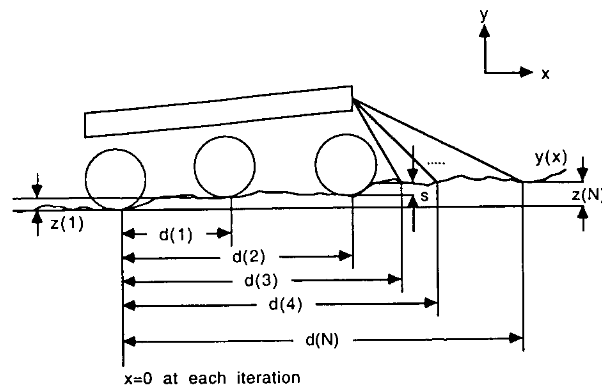

Figure 2. Planar model and symbol definitions

Objective

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 rear 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 ramp 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. Actually, the processor on the vehicle is sufficiently slow and burdened with other activities that the navigation and mobility sensors are only monitored roughly every wheel radius of forward traverse. This is often enough to ensure that rocks, craters and other geometric hazards can be detected and avoided (an issue not addressed in this paper).

There are several issues which need to be considered with this model. First, if **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.[5]

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 processor. 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=60 or so (the actual Rocky 3.2 vehicle has 13 cm dia. wheels and an overall length of 60 cm, with the look-ahead sensor reaching about one vehicle length).

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

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_x 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:

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

The interpretation of this expression is as 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 an 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) as 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) is 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 of the weighted-sum-of-squares to get a refined 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

$$\big({}_{j=1}\Sigma^{N}\,(z(j)/d(j))\big)/N$$

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-1) 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(1)) 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 and computational complexity 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 travelled and not estimating sinkage at all (and accepting the risk of getting stuck). We would also like the evaluate the usefulness of having more ranging sensor measurements as opposed to fewer, since each additional measurement has cost and may only be needed for this purpose (as rocks and craters may be detectable with as few as two look-ahead range points). If possible, we would like to also have a way of choosing the distances d(3)...d(N).

Thus what remains to be done is 1) perform an evaluation of the performance of the system by estimating the variance in the slippage and sinkage estimates by Monte Carlo numerical simulation (since the nonlinear MAP formulation is intrinsically iterative and because we want to explicitly incorporate the effects of quantization into the histogram bins, the effects of resampling and interpolation, etc.). This simulation will evaluate the effects when the data are not drawn from a Gaussian distribution, such as a uniform distribution of equal or different mean. Lastly, we would like to evaluate the effect on performance of varying the number of sensed values N, of modifying the mean m_x of the prior slip distribution based on experience, and of changes in the sensor noise $\sigma_{z(i)}$, which we might adjust in an ad-hoc way to account for the aliasing which the point-range measurements will have in estimating the average elevation over the histogram bins, where the spectrum of y(x) might grossly violate the Nyquist sampling theorem when binned in this manner.

The assumed model for y(x) in the simulation needs to be chosen with some care. A scale-invariant (fractal) model is attractive, but we need to recognize that the hazard-detection and avoidance system will effectively clip the distribution of terrain features at some particular scale. Similarly, a model for s(x) needs to be formulated, which will be slowly varying and of low amplitude. It would be good to assess the performance of the system when the slippage and sinkage are correlated, as one would intuitively expect, even though the model does not incorporate that effect (although it easily could). Another interesting correlation which would be good to model in the simulation is the fact that the mechanical linkages in the vehicle chassis cause the noise in the measurements z(i) to be highly correlated (since wheel pairs are at opposite ends of links), even though they may be Gaussian (from digitizing analog potentiometer values or peak detection in analog CCD scan lines).

The simulation model for y(x) and s(x)

As mentioned above, we desire to test the sinkage and slippage estimation algorithm on terrain which is "scale invariant". Specifically, we wish to create a sample random terrain in the form of a histogram (i.e. sequence) at the same resolution as that maintained by the estimation algorithm. This is accomplished 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 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 modelled by the analysis.) As mentioned before, a "smooth" simulated terrain is realistic here, since the geometric hazard detection system will avoid rough or discontinuous terrain.

The terrain we construct here is characterized by a single parameter: the maximum slope of each sine wave component. We call this parameter the "roughness" of the terrain. Both y(x) and s(x) are created by this technique, but s(x) is clipped at zero so that only positive values of sinkage are allowed. The "estimated" histograms of y and s are initialized with the "actual" values from this simulation; from that point on the estimation procedure extends them. This is reasonable since, as mentioned above, the first meter or so of traverse will be on the lander exit ramp and therefore known. We arbitrarily set the roughness of the sinkage function s(x) to be 20% that of y(x), based on the philosophy that the terrain mechanical characteristics are more slowly-varying than the surface topography.

It is perhaps worth mentioning that the approach of combining sine waves over a large number of different scales is computationally intensive, but need only be done once to simulate a large number of different terrain types, since to change the "roughness" only requires rescaling the vertical coordinate of a "standard" terrain, i.e. with unity roughness. Another approach to generating scale-invariant terrain, the use of Gauss-Markov random sequences, needs to be fairly high-order to get the needed range of scales and thus becomes extremely complex to analyze.

Specifics parameters for initializing the model are drawn from the actual design of the Rocky 3.2 microrover. Thus, for example, the distances from the sensed points to the rear wheel contact point are 25, 50, 60, and 80 cm for the middle wheel, front wheel, downlooking range sensor, and outlooking range sensor, respectively. The sensor noise (standard deviations) associated with these elevation differences are 0.04 mm for the wheel sensors, and 2 mm for the look-ahead sensors. We normally expect the vehicle to advance about 5 cm in each sensing cycle.

Simulation Trials

For each trial run, we evaluate the odometry error and sinkage error as a function of bin size and terrain roughness for different input assumptions. We evaluate bin sizes from 0.2 cm to 8 cm, which spans the range from very fine to very coarse compared the the expected forward advance per cycle. We evaluate terrain roughness ranging from a maximum slope at each scale of 0.25% to 8%, which spans terrain from very smooth (with typical elevation differences of 2 mm over the length of the vehicle) to very rough (with 8 cm typical elevation differences over the length of the vehicle, about the limit which the hazard avoidance system would permit). The simulation covers 62.5 meters of simulated terrain (25,000 bins at the finest bin size), which is created once and then resampled for the different simulations so that the effects of aliasing can be evaluated on identical terrain.

One important issue not addressed in the previous description of the algorithm is the choice of the search range for the weighted-sum-of- squares (WSS). Initially, the search was extended to $4\sigma + 4$ bins beyond the Bayesian prior mean. However, it was found that the simulation would occasionally get "stuck" and fail to advance the rover by the proper amount for several cycles, whereupon the simulation lost track of the terrain (i.e. presumably the internal histogram for y(x) had no relation to the actual y(x)). This was caused by the global minima of the WSS function not corresponding to the actual forward advance. A simple fix for this problem was to compute the secondary minima, and if it was beyond the global minima and nearly as good (within a factor of 3), then the search range on the next cycle was extended to include that minima. Note that, in all cases, the global minima is chosen for the simulation, and only that the search range is extended if another minima shows promise, so that it can be selected as the global minima on the next cycle. This effectively cured the problem, and subsequently the simulation was not observed to lose track of terrain.

Since we expect the look-ahead ranging sensor to have much worse measurement accuracy than the chassis articulation sensors, we will characterize the slippage estimation with the articulation-based elevation sensing noise, while the sinkage is based on the look-ahead sensor noise.

Thus we represent the results of this analysis by plotting the sinkage or slippage error against the terrain roughness value. Typically we would expect to have little or no cumulative error when the terrain is very rough, and if the terrain is smooth the algorithm will just return the Bayesian prior mean value as the result, so the error that accumulates is just the difference between the Bayesian prior mean and the actual mean value. Thus for slippage, for example, if the Bayesian prior is in error by 20% (that is, the actual expected distance advanced per sensing cycle is different from the Bayesian prior mean by 20%), then we would expect the algorithm to smoothly transition from small error to 20% error in estimating traverse distance as the roughness is increased from zero to a large value. We wish to establish the nature of this curve for both slippage and sinkage. Furthermore, to reduce computational complexity, we wish to determine how coarse the histogram bin size can be without excessive degradation of these results.

Table 1 shows the program output for the first test case, where the Bayesian prior overestimates the forward advance by 20%. Each entry in the table is the percentage odometry error over the 62.5 meter course, followed by the RMS sinkage error in parenthesis (in cm). Note that, indeed, the odometry error more-or-less smoothly falls from 20% for smooth terrain to near zero for rougher terrain. Furthermore, note that the performance improves as the bins get larger up to a point, and then declines for larger bin sizes, especially on rough terrain.

The two effects which seem to be occurring are severe aliasing for large bins (when the bins are larger than the advance of the vehicle), and poor terrain modelling for small bins. The former effect is compounded by the fact that we cannot fit a parabola to the WSS function if the minima is at zero bins of advance, since we do not compute the function for negative advance and so cannot bound the integral minima with values on each side, as needed for a parabolic interpolation. In this case we merely set the forward advance estimate to exactly zero. For large bins (e.g. 8 cm when the expected forward advance is 5 cm) this occurs commonly, and is only sometimes compensated for in later cycles. This produces a strong tendency to underestimate the distance travelled.

For very small bins, on the other hand, the algorithm we have selected for modelling the sparsely-sampled terrain is inadequate. Remember that we incorporate new z[i] measurements into the histogram by forcing the value at bin i to be consistent, and then "split the difference" on the i-1 and i+1 bins. When the bins are very fine this will produce narrow

SLIPPAGE ERROR AS A FUNCTION OF ROUGHNESS AND BIN SIZE
(each entry percent odometry error, RMS sinkage error (cm))

bin (cm)	roughness -- max slope at each scale: 0.0025	0.0050	0.0100	0.0200	0.0400	0.0800
0.25	20.87 (0.2)	15.24 (0.3)	10.36 (0.6)	6.62 (0.9)	3.57 (1.5)	-0.83 (6.3)
0.50	20.58 (0.2)	15.73 (0.3)	9.02 (0.5)	5.16 (0.8)	2.58 (1.4)	-1.20 (4.5)
0.75	18.94 (0.2)	13.57 (0.3)	7.86 (0.7)	3.69 (0.8)	1.84 (1.4)	-0.46 (3.4)
1.00	20.98 (0.2)	14.37 (0.3)	8.45 (0.4)	3.94 (0.7)	1.99 (1.2)	2.27 (2.8)
1.25	21.56 (0.2)	15.63 (0.3)	7.33 (0.4)	2.91 (0.6)	2.05 (1.3)	-2.17 (4.6)
1.50	19.30 (0.2)	11.15 (0.3)	5.56 (0.4)	1.42 (0.6)	0.10 (1.0)	-0.63 (2.9)
1.75	18.16 (0.2)	9.80 (0.3)	3.78 (0.4)	0.35 (0.5)	-14.61 (6.2)	-1.07 (2.6)
2.00	18.74 (0.2)	11.66 (0.2)	5.51 (0.4)	1.39 (0.4)	-0.32 (1.0)	-1.04 (2.9)
2.25	16.30 (0.2)	11.07 (0.2)	5.10 (0.3)	1.32 (0.5)	-0.01 (0.9)	-0.04 (2.2)
2.50	17.52 (0.2)	11.44 (0.2)	6.75 (0.4)	3.79 (0.6)	1.67 (1.2)	1.67 (2.9)
2.75	16.29 (0.2)	11.02 (0.2)	6.09 (0.3)	2.29 (0.6)	-2.21 (2.0)	2.01 (3.4)
3.00	12.97 (0.2)	7.60 (0.3)	1.97 (0.4)	-2.03 (0.6)	-4.16 (1.1)	-3.48 (3.0)
3.25	15.20 (0.2)	8.33 (0.2)	1.01 (0.3)	-3.58 (0.6)	-7.01 (1.4)	-4.99 (3.5)
3.50	15.98 (0.1)	10.44 (0.2)	5.64 (0.3)	1.14 (0.5)	-0.88 (1.2)	0.63 (3.3)
3.75	14.47 (0.2)	7.45 (0.2)	0.41 (0.3)	-5.57 (0.6)	-8.25 (1.5)	-4.91 (3.6)
4.00	14.82 (0.1)	8.97 (0.2)	2.74 (0.3)	-2.24 (0.6)	-4.85 (1.3)	-3.43 (3.5)
4.25	12.83 (0.2)	4.81 (0.2)	-5.49 (0.4)	-12.58 (0.8)	-14.81 (1.9)	-14.12 (4.0)
4.50	15.14 (0.2)	7.66 (0.2)	-0.33 (0.3)	-8.00 (0.7)	-8.91 (1.7)	-7.28 (4.0)
4.75	14.61 (0.2)	8.82 (0.2)	1.41 (0.3)	-5.23 (0.6)	-9.01 (1.4)	-8.88 (3.2)
5.00	15.04 (0.2)	9.75 (0.3)	4.66 (0.4)	-0.94 (0.8)	-2.37 (1.8)	1.74 (4.8)
5.25	12.31 (0.2)	3.32 (0.2)	-7.19 (0.4)	-15.85 (0.8)	-18.43 (1.8)	-13.87 (4.6)
5.50	12.66 (0.2)	6.01 (0.3)	-4.00 (0.4)	-12.23 (0.8)	-15.58 (1.6)	-14.52 (3.5)
5.75	14.54 (0.2)	6.20 (0.2)	-3.20 (0.3)	-13.75 (0.7)	-15.86 (1.6)	-13.18 (3.5)
6.00	16.26 (0.2)	9.72 (0.2)	1.25 (0.3)	-6.40 (0.8)	-8.48 (1.7)	-4.22 (4.1)
6.25	14.44 (0.2)	10.00 (0.3)	1.46 (0.4)	-4.67 (0.9)	-5.12 (1.8)	-0.16 (4.7)
6.50	10.56 (0.2)	-2.75 (0.2)	-18.84 (0.5)	-27.70 (1.1)	-29.74 (2.4)	-51.35 (5.8)
6.75	11.23 (0.2)	-2.80 (0.3)	-19.46 (0.4)	-29.11 (0.8)	-28.83 (1.7)	-54.61 (3.5)
7.00	11.60 (0.2)	0.36 (0.3)	-16.44 (0.4)	-25.05 (0.8)	-26.24 (1.6)	-23.42 (3.8)
7.25	14.54 (0.2)	1.71 (0.2)	-15.08 (0.5)	-25.49 (1.1)	-28.44 (2.6)	-27.28 (5.8)
7.50	15.21 (0.2)	3.99 (0.2)	-12.62 (0.5)	-23.38 (1.2)	-25.50 (2.7)	-22.45 (5.3)
7.75	13.48 (0.2)	4.52 (0.3)	-13.74 (0.4)	-24.38 (0.7)	-25.00 (1.8)	-30.19 (4.1)
8.00	17.31 (0.2)	6.79 (0.2)	-12.49 (0.4)	-21.23 (0.8)	-17.86 (1.8)	-15.76 (4.7)

Simulation Parameters
Actual mean advance per cycle: 5.0 cm, Sigma: 1.00 cm
Bayesian prior mean advance per cycle: 6.0 cm, Sigma: 0.05 cm
Unit-Roughness Terrain RMS Amplitude 1.11 meters over 62.5 meters

Statistical Attributes of Unit-Roughness Simulated Terrain by Bin Size
(each entry RMS bin-to-bin elevation change (cm),
RMS error in bin-to-bin linear projection (cm))

Bin Size (cm) X	X.00		X.25		X.50		X.75	
0.--	0.00	0.00	2.40	1.29	4.62	4.07	6.58	6.38
1.--	8.30	7.47	9.90	7.85	11.50	8.12	13.14	8.74
2.--	14.81	9.40	16.50	10.14	18.18	10.82	19.86	11.50
3.--	21.52	12.51	23.16	13.81	24.77	14.78	26.40	15.74
4.--	27.87	16.84	29.73	18.87	31.38	20.58	33.06	22.26
5.--	34.70	24.27	36.33	26.42	38.02	28.56	39.60	30.45
6.--	41.24	32.49	42.69	33.57	44.22	35.31	45.69	37.22
7.--	47.24	39.24	48.90	41.42	50.43	43.60	51.91	45.52

Table 1

"spikes" in the histograms, and not at all correspond to realistic terrain. The proper fix for this would be to "remember" when and where each prior measurement was taken, and try to perform a statistically-valid terrain reconstruction (based on some assumed terrain Fourier spectrum), incorporating all prior measurements and their uncertainties. However, this would be computationally demanding, and the procedure we have adopted seems to work quite well for intermediate-sized bins, about 2 cm long.

Note that the sinkage estimates in Table 1 are all about the same for a given roughness, and increase more-or-less proportionally to roughness. This is intuitively pleasing, since the high accuracy of the wheel sensors compared to the look-ahead sensors makes the estimate of the forward advance of the vehicle (i.e. the minima of the WSS function) almost entirely a function of the wheel sensors. Thus, the primary function being estimated accurately is the loadbearing surface y(x)-s(x), with both y and s being much more uncertain than their difference. Then sinkage is estimated using the look-ahead sensor(s), with their large attendent noise. This suggests that a more appropriate implimentation for the actual vehicle is to use the wheel sensors alone to estimate travel along the loadbearing surface, and to use only one look-ahead sensed value to estimate sinkage. Thus it is irrelevant to examine the case of additional look-ahead sensing values so long as their noise is very large compared to the chassis articulation sensing. The "roughness" scale used corresponds approximately to the RMS elevation differences in meters over the scale of the vehicle, i.e. a roughness of 0.08 gives 8 cm of typical elevation difference across the vehicle. At the bottom of Table 1 is a chart showing some of the statistical properties of the unit-roughness simulated terrain: the RMS bin-to-bin elevation change and the RMS error in a bin-to-bin linear projection to the next bin, each as a function of bin size. This table has cm of bin size along the left, with fractions of a cm along the top. Note that the values for zero bin size, which in fact don't exist, are set to zero for printing purposes.

There is one striking fact represented in Table 1: we have selected the standard deviation of the Bayesian prior to be 0.05 cm (1% of the actual advance), when the sigma of the actual vehicle advance per cycle is 1 cm. This artificially "overweights" the Bayesian prior to show the smooth transition from 20% error to small error as the terrain gets rougher. However, the chassis articulation sensors are so accurate (σ=0.04 mm) that we can do much better than this. Figure 3 shows the results for different values of the Bayesian prior (1% and 10% of the actual). As one can see, with lower confidence in the prior, even on smooth terrain, the results are very good for bin sizes of about 2 cm (ranging from 5% error on very smooth or rough terrain to under 1% error on moderate terrain). This, again, is to be expected, since even the smooth terrain has large excursions compared to the sensor noise. If we reduce the prior variance further, however, the estimator performance degrades rapidly. This presumably occurs because much more error exists in the terrain histogram reconstruction than would be apparent from the sensor noise alone. Thus, if the simulation is not "driven" strongly by the Bayesian prior, it is "free" to choose any match to the sensor data, weighted artificially heavily due to the low sensor noise. Thus, even though the sensors are good, the terrain estimates which result from our sparse sampling and crude interpolation are not nearly so good. Thus weighting the perceived errors from this function by one over the sensor variance is unrealistic; we compensate by making the Bayesian prior very tight. Thus there is no particular value to be gained in evaluating somewhat different levels of sensor noise.

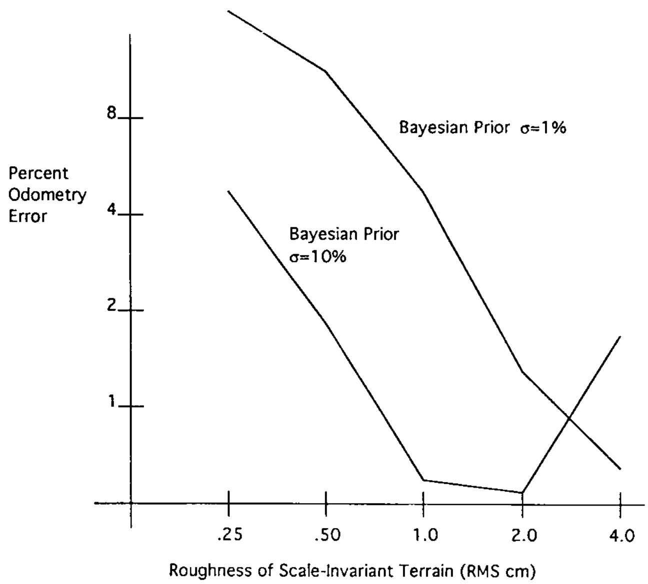

Figure 3. Odometry error as a function of terrain roughness (20% actual slip)

Numerous additional runs analogous to Table 1 were performed using different simulated terrain (using different seeds for the random number generator), and the results were virtually identical. Note that there are occasional anomalies where the performance is poor (such as in Table 1 at roughness 0.04 and bin size 1.75 cm). These anomalies presumably result when the terrain and binning processes conspire to give ambiguous terrain for matching purposes. This is to be expected, but so long as it is rare and does not give worse estimates than doing nothing (i.e. using just the prior mean estimate), then no harm is done. This is another reason to overweight the prior distribution.

Additional runs explored several issues. For example, when the prior distribution underestimates the forward advance, the performance is generally good for bin sizes between 2 and 3 cm, but that very bad performance is not uncommon. Another issue considered was the estimation performance when the actual forward advance is not Gaussian. Once again, the performance was excellent. Lastly, we considered the system performance when the actual slip is a function of sinkage and slope, as one would expect. The results were evaluated for the case when the mean of the actual advance per cycle drops linearly with increasing sinkage and/or slope, (and continuing with the non-Gaussian uniform actual distribution). Since the very rough terrain will probably have slopes and sinkages which would literally stop the vehicle under such an assumption, we clipped the left end of the uniform distribution at zero advance per cycle, so that the simulation doesn't get in an infinite loop (as would the actual vehicle). Here we have assumed that the linear coefficients are such that a 60% grade will stop the vehicle, as would sinkage of 5 cm. The performance on smooth terrain was poor, as the Bayesian prior of 6 cm/cycle was much larger than the actual average, which is 5 at best and 1 at worst, depending on terrain conditions. However, as soon as the roughness increases to 1 cm or so over the length of the vehicle, the odometry performance improves to within 10% and at 2-4 cm roughness. Only a few percent of odometry error is observed for bin sizes between 2 and 3 cm. This performance is very encouraging considering the simplicity of the model and the gross deviations which "reality" makes with the assumptions underlying the model.

Computational Requirements

As described above, the optimal bin size is in the neigborhood of 2.5 cm, which means that there are only 20 bins of data over the length of the vehicle to be accumulated and maintained, so the compuation and storage requirements are small. Good performance can be anticipated with only 2 terms in the WSS function-- one for the Bayesian prior (which can be precomputed and stored in a table) and one for the center wheel, since the look-ahead sensor is so noisy as not to affect the forward-advance estimate. (The front wheel moves onto unknown terrain, and so is not used in the matching process.) Since squaring can also be accomplished as a table look-up, the computation is of the order of 1 add and 2 table look-ups per bin, with typically 5 bins searched. Finding the global and secondary minima requires roughly 2 comparisons per bin. Maintenance of the histogram requires a relatively few operations also, since the histogram data can be in a ring buffer with a pointer, to avoid actually shifting the data in an array. Thus only the linear interpolation and "split the difference" operations are needed, which are simple. This implies that, with of the order of 100 operations per cycle, the odometry estimates of the vehicle can be markedly improved, and sinkage estimates provided.

Preliminary Test Results

The algorithm described above has been implemented on Rocky 3.2 and, as of this writing, a few test runs have been conducted. The preliminary indication is that the articulation sensor noise is substantially larger than anticipated, leading to odometry results which are somewhat degraded compared to the simulations. However, it appears that, even with the noisy data, the algorithm will give a very reasonable hazard alarm for sinkage and slippage. (In this case, we set the confidence in the Bayesian prior to be very high, and then threshold the WSS function to trigger a slip alarm.) Work is continuing on reducing the noise in the analog-to-digital portion of the system. Extensive tests for this system are planned for 1994.

Conclusions

The MAP estimation procedure developed here, based on a simple weighted-sum-of-squares computation, seems to give sinkage and slippage estimates which will allow planetary microrovers to detect and avoid a wide range of non-geometric hazards. Simulations suggest that it may also improve odometry markedly over simple wheel revolution counting, and thereby lead to a significant improvement in dead-reckoning accuracy for this class of vehicle.

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 assistance of Jack Morrison and Tam Nguyen in implementing and testing the algorithm on the research vehicle is gratefully acknowledged.

References

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

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

[3] Olheft, G. R. and Strangway, D. W., "Electrical Properties of the Surface Layers of Mars," *Geophysical Research Letters*, Vol. 1, #3, July 1974.

[4] Wilcox, B. H., "Robotic Vehicles for Planetary Exploration", *Journal of Applied Intelligence* 2, p181-193, 1992.

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