Submitted to 2005 International Symposium of Robotics Research, October 2005, San Francisco, California, USA

Tradeoffs Between Directed and Autonomous Driving on the Mars Exploration Rovers

Jeffrey J. Biesiadecki and Chris Leger and Mark W. Maimone
Jet Propulsion Laboratory
Pasadena, CA USA
first.last@jpl.nasa.gov

May 10, 2005

Abstract

This is an extended abstract submission, not the final paper.

NASA's Mars Exploration Rovers (MER) have collected a great geological diversity of science results, thanks in large part to their surface mobility capabilities. The six wheel rocker/bogie suspension provides driving capabilities in many distinct terrain types, the onboard IMU measures actual rover attitude changes (roll, pitch and yaw, but not position) quickly and accurately, and stereo camera pairs provide accurate position knowledge and/or terrain assessment. Solar panels generally provide enough power to drive the vehicle for at most four hours each day, but drive time is often restricted by other planned activities. And driving along slopes in nonhomogeneous terrain injects unpredictable amounts of slip into each drive. These restrictions led us to create driving strategies that maximize drive speed and distance, at the cost of increased complexity in the sequences of commands built by human Rover Planners each day.

Commands to the MER vehicles are typically transmitted at most once per day, so mobility operations are encoded as event-driven sequences of individual motion commands. Motions may be commanded using quickly-executing *directed* commands which perform only *reactive* motion safety checks (e.g., real-time current limits, maximum instantaneous vehicle tilt), slowly-executing position-measuring Visual Odometry (VisOdom) commands which use images to accurately update the onboard position estimate, or slow-to-medium speed Autonomous Navigation (AutoNav) commands which use onboard image processing to perform *predictive* terrain safety checks and optional autonomous path changes.

The MER rovers have driven over more than a combined 9 kilometers of Martian terrain during their first 16 months of operation using these basic modes. In this paper we describe the strategies adopted for selecting between human-planned directed drives versus rover-adaptive AutoNav and VisOdom drives.

1 Background

NASA successfully landed two mobile robot geologists on the surface of Mars in January 2004: the Spirit and Opportunity Mars Exploration Rovers (MER). Their primary goal was to find evidence of past water at Gusev Crater and Meridiani Planum, two geologically distinct sites on opposite sides of the planet. Each rover was instrumented with a suite of tools for remote sensing (multi-filter and stereo camera pairs, and a thermal emission spectrometer) and in-situ measurement (5 DOF arm for deploying a grinding Rock Abrasion Tool, Microscopic Imager, Alpha Particle X-ray Spectrometer, and Mössbauer Spectrometer). Although the achievement of their successful landings stands out as a technological tour de force, it was their ability to traverse while on the surface of Mars that enabled both rovers to succeed in their primary goals.

The MER rovers are typically commanded once per Martian solar day (or *sol*). A sequence of commands sent in the morning specifies the sol's activities: what images and data to collect, how to position the robotic arms, and where to drive. At the end of each sol, the rovers send back the images and data human operators will use to plan the next sol's activities. The next sol's mobility commands are selected based on what is known – and what is unknown – about the terrain ahead.

1.1 Rover Mobility Configuration

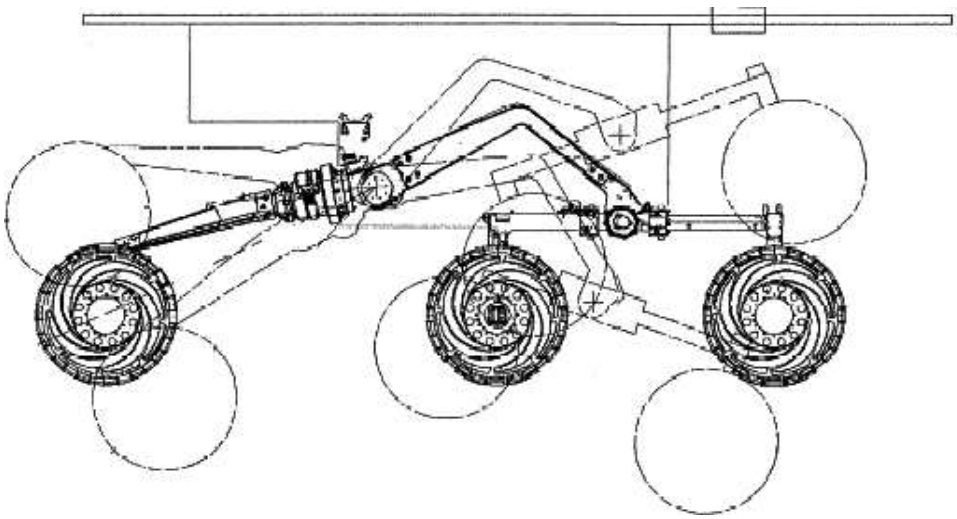

Figure 1: Rocker-bogie configuration.

MER vehicles have a very capable mobility system. Six wheels are mounted on a rocker-bogie suspension that minimizes the tilt induced by climbing over individual rocks [6]. Each wheel is 25cm in diameter and has short paddle-like cleats, and each is theoretically capable of climbing rocks over 35cm tall. However, the clearance under the body of the rover is only 29cm on a flat surface (and less when tilted), so in practise obstacles 20cm or above were avoided. All six wheels can be driven at speeds up to 5cm/sec, but only the four corner wheels are steerable. Ackermann-style arcs with turn radii as tight as 1m may be commanded, as may turns in place in which the vehicle rotates about its origin, location midway between the left and right middle wheels.

Both rovers are statically stable up to 45 degrees and have driven on hard slopes as high as 31 degrees (on rock outcrops in Endurance Crater and the Columbia Hills), but driving on slopes greater than 25 degrees requires special planning. At such high tilts, the weight reduction on the upslope wheels is enough that they can sometimes "float" off the ground. The maximum tilt on loose soil is much smaller; for example, Opportunity was failed to exit Eagle Crater because it was unable to climb straight up a slope of only 17 degrees on Sol 56, and Spirit has seen similar 100% slip on sandy soils in the Columbia Hills.

1.2 Mobility Software

The rovers are driven using three primary modes: low-level commands that specify exactly how much to turn each wheel and position steering actuators, directed driving primitives for driving along circular arcs (of which straight line driving and turn-in-place are special cases), and autonomous path selection. Low-level commands enable "non-standard" activities such as using the wheels to dig holes in Martian soil, scuff rocks, and perform mechanism health diagnostic tests. Directed drives allow human operators to specify

exactly which driving primitives the rover will perform. Autonomous path selection mode allows the rover to select which driving primitives to execute in order to reach a Cartesian goal location supplied by human operators.

Both directed and path selection modes of driving can make use of on-board Stereo Vision processing and Terrain Analysis software to determine whether the rover would encounter any geometric hazards as it drives along its chosen path. In directed driving, the rover can preemptively "veto" a specific mobility command from the ground if it appears too risky. In Autonomous Navigation (AutoNav) and other path selection modes, the rover can select its own driving primitives to steer around obstacles or recover from unplanned changes in attitude (and position, when VisOdom is enabled) and make progress toward its goal. This software provided the unique capability of enabling the vehicle to drive safely even through areas never before seen on Earth: more than 2700 meters of the rovers' combined distance was driven using autonomous hazard avoidance as of May 2005.

The rovers maintain an estimate of their local position and orientation updated at 8 Hz while driving. Position is first estimated based on wheel odometry, and orientation is estimated using an Inertial Measurement Unit that has 3-axis accelerometers and 3-axis angular rate sensors. In between driving primitives, the rover can make use of camera-based Visual Odometry to correct the errors in the initial wheel odometry-based estimate that occur when the wheels lose traction on large rocks and steep slopes. Visual Odometry software has generated over 1400 combined successful position updates on both rovers [2].

All MER surface software runs on a single 20 MHz RAD6000 computer under the VxWorks operating system. The limited computation speed, and the fact that dozens of tasks share a single address space and cache, contribute to the slow execution times of the AutoNav and VisOdom subsystems.

1.2.1 Autonomous Terrain Analysis

When information about nearby terrain is unavailable or uncertain, the rover can be commanded to evaluate terrain safety by performing stereo vision and terrain assessment autonomously. This allows the rover to *predictively* determine where a traverse hazard might be and avoid it. The procedure is summarized here; see [4, 1] for more details and [9] for the approach that inspired it. Some important parameters include max traversable obstacle size (e.g., 20cm), max terrain tilt angle (e.g., 20degrees), and average surface roughness (e.g, 7cm).

1. The rover chooses a stereo camera pair based on the goal location and its previous terrain assessment, if any. Images are downsampled in software from 256x1024 12 bit pixels down to 256x256 8 bit pixels.

2. The stereo pair of images is densely correlated, and any noisy or uncertain data are masked or filtered out.

3. Range data not filtered out above are distributed into a local map grid comprised of 20cm x 20cm cells (50x50 cells for Spirit, 60x60 for Opportunity).

4. Disc-shaped planar patches are fit to the measured range points. Each disc is 2.6 meters in diameter, slightly larger than the volume swept by the vehicle during a turn in place maneuver.

5. Parameters of the plane fit are used to look for potential geometric hazards; the surface normal is compared to the max terrain tilt, residual of the fit is compared to the max roughness, and the max cell elevation difference between cells in the patch (corrected for tilt and wheel base) is compared to the max traversable obstacle size.

6. Each of the above three quantities is linearly scaled to an 8 bit "goodness" value; the minimum goodness value is logged as the traversability of that cell. This "goodness map" feeds into a multilayered local world map, which stores knowledge of the terrain as goodness values as far as 5 meters away from the vehicle.

7. Each of dozens of potential paths (left or right, arc and/or turn in place, forward and/or backward) is evaluated by computing a weighted sum of all the cells it touches.

The rock-strewn terrain encountered by Spirit at Gusev Crater corresponded well to the exponential rock distribution models predicted using data from Viking I, II and Pathfinder missions [5]. The body-mounted 120-degree FOV HAZCAMs were designed with this terrain model in mind, and Spirit has performed all of its autonomous terrain assessment using these cameras. However, the terrain encountered by Opportunity at Meridiani Planum is vastly different. Instead of a wide variety of rocks at many scales, much of the terrain consists of very fine-grained materials; so fine, in fact, that no large scale features can be found in the wide FOV HAZCAM images at 256x256 resolution. But correlated with the lack of large scale features is a lack of large "step" obstacles. So Opportunity was reconfigured to perform terrain assessment with more narrow FOV NAVCAM images. Rock and fissure obstacles can still be detected, but the limited FOV reduces its ability to steer around them.

1.2.2 Reactive Hazard Detection

Several types of potential vehicles hazards are checked *reactively*, most of them during each Real Time Interrupt (RTI) which occurs 8 times per second.

Tilt/Pitch/Roll The accelerometers in the IMU provide an instantaneous (unfiltered) vehicle attitude estimate. If the value for any of these attitude components exceeds a threshold for some minimum number of RTIs, then motion is terminated.

Northerly Tilt A filtered version of the vehicle attitude can be compared to a minimum Northerly Tilt angle. This allows the vehicle to ensure that the rigidly-mounted solar panels stay pointed toward the sun.

Suspension Angles Three potentiometers measure the current state of the vehicle suspension; motion can be stopped if any of these lie outside a fixed range.

Motor Stalls Motion can be stopped if the current used by any motor exceeds some limit. Voltage limits are also available, but remain unused so far.

Limit Cycle If autonomous path selection is enabled, the rover can ensure that some minimum distance is covered in some number of steps. This is useful for ensuring that VisOdom drives are not slipping too much, and that AutoNav drives are able to negotiate a drive around an obstacle. Checked once per step.

Resource Contention The Activity Constraint Magager is polled at each step to ensure that no conflict exists that would prevent driving; e.g., the arm must be stowed, no communication can be taking place, no prior error may exist.

1.3 Manual Terrain Analysis

Terrain assessment is generally performed on the ground using stereo image pairs taken by any of the three types of stereo camera pairs found on MER vehicles. There are two pairs of wide field-of-view (120 degree, 10cm baseline) Hazard Cameras (HAZCAMs) rigidly mounted 53cm above the ground plane on the front and back sides, one pair of medium field-of-view (45 degree, 20cm baseline) Navigation Cameras (NAVCAMs) mounted 152cm above the ground plane on a pan/tilt head, and one pair of narrow field-of-view (18 degree, 28cm baseline) Panoramic Cameras (PANCAMs) also mounted 152cm above the ground plane on the pan/tilt head. [8] These cameras take 1024x1024 12-bit images that provide information about terrain texture throughout their images, and stereo range-derived terrain shape at different scales: 0.5m - 5m in the HAZCAMs, 2m - 20m in the NAVCAMs, 4m - 70m in the PANCAMs.

The amount of directed driving that could be commanded depended on both the terrain itself and on how much information about the terrain was available. Orbital imagery, while crucial for long-range planning, cannot resolve vehicle hazards like 20cm rocks. So after each long drive, images from each appropriate camera pair were requested. When driving long distances, typically half a dozen stereo image pairs were planned in the drive direction (some from each camera) at the end of each day's drive. Sometimes problems occurred with part of the data downlinked from Mars, resulting in fewer images than were requested and thus limiting the distance of the initial directed drive.

All stereo image pairs are processed by an automated pipeline that generates a variety of derived products including 3D range maps, texture-mapped terrain meshes, and color overlays indicating terrain properties such as slope and elevation [7]. Rover operators use image-based querying tools to measure ranges to terrain features and estimate distances and rock sizes [3]. For example, a "ruler" tool allows the operator to measure the distance between the 3D points corresponding to two pixels in an image or image mosaic. This is useful for identifying discrete obstacles such as rocks or steps. Terrain meshes give the operator a geometric understanding of the terrain and of spatial relationships between terrain features and the planned path, and allow simulation of drive sequences to predict drive safety and performance [10]. The raw images are also extremely useful in assessing traversability: operators can readily identify very sandy or very rocky areas that present hazards, though new terrain types always carry an element of uncertainty regarding vehicle performance. In some cases, there are no image cues that allow rover operators to predict the performance of a drive; patches of terrain only a few meters apart, with similar surface texture and geometry, can lead to drastically different traction or sinkage. For example, while driving uphill toward a topographic high point named "Larry's Lookout", Spirit reached 100% slip (i.e. no forward progress) on a 16 degree slope, but only a few meters further had only 20% slip on a 19 degree slope with no discernable difference in appearance.

2 Drive Techniques and Templates

Most drive sequences could be classified as either traverses (covering maximum distance) or approaches (driving to a specific position for subsequent arm operations). The techniques used for each drive type are determined based on the time allocated for driving, the amount of terrain visible in imagery, known hazards, and level of uncertainty in rover position given the terrain type. Generally, driving on level ground required a mix of blind and AutoNav driving, and driving on slopes required using visual odometry to allow the rover to compensate for slip.

The computing resources required by these different commands vary greatly. Directed driving commands execute the most quickly (achieving speeds up to 124 meters/hour), but also have greater risk since the rover can only count wheel rotations to estimate position, and never looks ahead to evaluate the terrain before driving onto it. AutoNav commands detect and avoid geometric hazards, but only achieve driving speeds from 10 meters/hour in obstacle-laden terrain up to 36 meters/hour in safe terrain, and also rely on

the accuracy of the wheel odometry to track obstacles once they leave the field of view of the cameras. VisOdom commands provide accurate position estimates (but not obstacle detection), and require close spacing between images which limits the top speed to 10 meters/hour.

2.1 Traversing the Plains

Driving on level ground typically leads to accurate and predictable mobility performance. Because of the rover's limited processing power, drives using autonomous hazard avoidance are several times slower than "blind" (manually-directed) drives. These two factors favor using the longest possible blind drives. Human operators can easily identify rocks that are large enough to be hazardous to the rover, and can plan complex paths that avoid them. The firm surfaces found on the plains of Gusev crater allowed for blind drives of up to 70m. Beyond that distance, hazards cannot be reliably resolved, and the rover operators cannot predict where the rover will travel with sufficient accuracy. Additionally, the low viewing incidence angle causes significant gaps in range coverage of the terrain; a slight rise in the terrain at 50m can lead to a several-meter occluded region in which the operator has no knowledge of the terrain. Thus, at longer ranges, the only safe course of action was to rely on AutoNav to find a path.

On the plains of Meridiani, the terrain hazards were quite different and allowed for blind drives over 100m in length. Unlike the Gusev plains, there was a near-total absence of rocks at Meridiani, and until sol 446 none of the innumerable sandy ripples posed a threat to the rover. Craters, visible in orbital imagery, and small linear depressions were the most significant hazards for Opportunity. But since the terrain was expected to be nearly flat, the rover's suspension did not articulate significantly during normal driving. This suggested that the measured suspension articulation could be used to halt driving if the rover were to encounter a depression. The April 2004 (R9.0) version of the rover's software allowed the articulation of the rover's suspension to be checked against preset limits at 8Hz, thus mitigating the hazard posed by negative terrain features that were not visible a priori. Because the reason for halting the drive (e.g. timeout, suspension check, or tilt check) is accessible to the rover sequencing language, a recovery maneuver could be performed in the event that the suspension check was tripped. This recovery consisted of backing up several meters and then continuing the drive with AutoNav, since AutoNav was able to detect and avoid negative hazards. Once Opportunity left Endurance Crater, the drive template consisted of an initial blind drive of 30-60m with relatively loose suspension checks (since the PanCam resolution allowed significant negative obstacles to be identified up to 60m away), followed by a "bonus" blind drive with stricter suspension checks, and finishing with an AutoNav drive until the allocated time was exhausted.

Both rovers used a common strategy at the end of long traverses to acquire necessary images for manipulator operations and turn to a preset heading that minimizes the multi-path interference caused by the rover's mast during communication with Earth or an orbiter. However, this presents a problem for IDD operations on the following sol: in order to determine whether the IDD can be safely deployed, a front HazCam image pair must be acquired 0.5-3m before the rover's final location. The obvious solution is to turn to the desired heading, acquire the image pair, then drive a short distance to the final location. But since the rover drives under AutoNav control until a time limit is reached, the final position of the rover cannot be predicted and thus the rover operator cannot guarantee that it is safe to travel an additional distance at the desired heading. The "guarded arc" drive primitive solves this problem by only executing the post-turn drive segment if the HazCams show that it is safe to do so.

2.2 Driving on Slopes: Mountains and Craters

While most of the distance covered by the rovers was on level ground, most of the sols and most of the approach drives occurred on slopes. The rovers invariably slip when driving on slopes, making VisOdom essential for safe and accurate driving. But VisOdom comes at a cost: the processing for one VisOdom

update takes over three minutes due to the rover's slow CPUs, and the rover can only move roughly 60cm, or turn 15 degrees, before there isn't enough overlap in successive images to reliably estimate motion. This leads to traverse rates on the order of 10m per hour when using VisOdom. Using AutoNav along with VisOdom takes roughly twice as much time, making it impractical for normal use.

This presents a challenge: the rover has the ability to know where it is, but for practical purposes is not able to detect obstacles. Additionally, in steep terrain the rover cannot identify all obstacles classes, since the rover has no means of detecting sandy, high-slip areas in advance. Even otherwise safe areas of moderate slope may represent hazards if there are steeper slopes or rocks downhill, since slippage in moderate slopes could take the rover into dangerous areas. In these cases, the rover operators can specify "keep out zones" which will cause the rover to halt driving before a hazard is encountered. The rover can thus keep track of its position using VisOdom (and can preclude driving if VisOdom fails) and can close the loop to correct for slippage, while relying on the manually-specified keep-out zones to stay safe. VisOdom also gives the rover the ability to halt driving if a high-slip region is encountered by using the "limit cycle check". This check counts the number of steps since the rover moved a set distance from a previous position, and can take a corrective action if the rover fails to move appreciably. Initial tests used a fairly low limit–40%–but this was too sensitive. We later increased the tolerance so that the rover would quickly halt driving if near-100% slip was achieved. This was crucial in driving Spirit on the steep slopes of Husband Hill, where the hollow wheels could dig in and could potentially engulf a rock and stall the drive actuator, as happened on sol 339. On sols 453 and 454, Spirit promptly halted driving after encountering slippage over 90%, and post-drive HazCam images on sol 454 showed several rocks on the verge of falling into the wheels, since the wheels had dug into the terrain by nearly one wheel radius. The combination of high slopes, sandy terrain with intermixed small rocks, and frequent obstacle-sized rocks caused us to retreat and search for a new route to the summit of Husband Hill.

2.3 Target Approach

Whereas traverse sequences focus on covering maximum distance over terrain of little interest to the science team, target approach sequences aim to place the rover at a specific target position for in-situ examination of rocks and soil with the rover's manipulator, or less frequently, high-resolution imagery of a distributed or inaccessible target region. The accuracy requirements for positioning the rover for in situ work are relatively tight, often within 10cm. This presents a challenge when driving on steep slopes with unpredictable slip behavior. The first step in sequencing a target approach is to determine the optimal rover position and heading for in situ work by simulating instrument arm placements on the target. Once the desired position and heading have been selected, a standoff location (usually 1-3m away) is chosen along the heading vector, and then the drive to the standoff is planned.

On level ground, directed drive primitives are usually sufficient for accurate target approaches from 2-10m away. On sloped terrain, visual odometry is required to close the loop on the rover's position. Conditional sequencing is often used in conjunction with visual odometry, and consists of inserting conditional tests (IF-THEN-ELSE constructs) in the sequence to allow the rover to execute different manually-specified drive primitives in response to drive performance. For example, the final leg of an approach sequence can have several 30cm conditional steps that will execute if the rover center is greater than 1.3m away from the target, and then several conditional 10cm steps that will execute if the rover is greater than 1.1m away from the target. After some or all of the steps, the rover can be commanded to execute a turn-in-place to face the target. In this way, the target will end up directly in front of the rover and within 1.0 to 1.1m away from the rover center, placing it squarely in the sweet spot of the manipulator's workspace. After each motion, visual odometry updates the rover's position knowledge, allowing it to correct for slip-induced errors. Using this technique, the rovers can accurately approach targets 5-10m away while driving on slopes in the 10 to 20 de-

gree range, with the caveat that on surfaces with sufficiently low bearing strength, the rover is mechanically incapable of making direct uphill progress.

3 Relative Merits of Directed and Autonomous Driving

3.1 Resources

There are significant differences in resource usage between manual and autonomous driving, with execution time being the most obvious. Power is also impacted by execution time, for although the power used by the mobility system is the same whether a trajectory was generated manually or autonomously, the rover's CPU and electronics draw power for the duration of the drive and thus an autonous drive will require more power than a manual drive of the same distance.

Less obvious differences in resource requirements between manual and autonomous driving also exist. The most significant is planning time: it takes a rover operator more time to identify obstacles and choose appropriate waypoints when sequencing a blind drive than when sequencing a drive using AutoNav. During the first few months of the mission when operators were still learning the basic capabilities of the rovers and were developing sequencing and imaging techniques, it often took up to 10 hours to build a drive sequence to travel 20-40m across the plains of Gusev. (This decreased dramatically later in the mission, often requiring only 2-4 hours to sequence drives over 100m in length on either rover.) Still, a directed drive places full responsibility for vehicle safety on the rover operator rather than allowing the rover to safeguard itself, thus requiring more time for manual terrain analysis and waypoint selection. This suggests an obvious trade-off between sequencing time and execution time for directed and autonomous drives.

There is an additional long-term resource trade-off: humans can rapidly adapt their techniques to deal with new terrain types or drive requirements, but updating the on-board software involves a lengthy software development, testing, and uplink process. Instead of a day-to-week turnaround in sequence development, software updates to cope with new terrain and drive techniques occur on a month-to-year cycle.

3.2 Perception and Planning

It comes as no surprise to any computer vision researcher that the human perceptual system, while qualitative and imperfect, is extremely capable. When combined with quanitative image analysis tools, humans are very good at terrain analysis for motion planning. In addition to geometric hazards such as rocks or drop-offs, humans can readily identify and classify new terrain types (e.g. sandy versus rocky slopes) on the basis of appearance alone. The MER software does not have any appearance-based terrain analysis capabilities, thus limiting hazard detection to geometric obstacles. Nevertheless, the most serious and frequent hazards (rocks, steps, and high-center hazards) can be detected by geometric analysis–assuming sufficient range data is available. At longer ranges (over 15m in NavCam images, and over 50m in PanCam images), range data becomes sparse, making it impossible to rely solely on geometric analysis. In these cases, humans manually identify rocks and, with the aid of a single range point and knowledge of camera parameters, can conservatively determine whether a rock is large enough to present a hazard to the rover. On the other hand, on-board terrain analysis is performed on data within a few meters of the rover, so dense range data is normally available when driving autonomously. The rover is better able to assess nearby hazards, but the lack of a global planner (which the human stands in for during manual drives) can cause the rover to get stuck in cul de sacs.

There is one notable circumstance in which the human has no ability to safely select paths: when the rover is driving into terrain that has not been imaged. On sol 109, Spirit was commanded to drive over local horizon 50m distant as it descended from the rim of Missoula Crater. In this case, AutoNav was the

only option available to drive further and use the available time and power, and post-drive images showed AutoNav correctly avoiding large rocks while traversing slopes up to 10 degrees. Obviously, a high degree of confidence in the hazard avoidance software is needed in situations such as this. Less severe, but more frequent, instances in which humans cannot guarantee rover safey occur when the rover drives beyond the distance at which obstacles can be resolved, or outside the available imagery, or through smaller occluded regions. In practice, even when using AutoNav the rover operator typically chooses waypoints that avoid the most hazardous areas, thus taking advantage of the perceptual strengths of both human and rover.

3.3 Execution

Directed drives have a limited ability to deal with errors or uncertainty in execution. Whereas AutoNav can close the loop on vehicle safety by imaging the terrain that the rover is about to drive through, a directed drive must make the assumption that the rover does not deviate far enough from the planned path to encounter any hazards. On firm, level ground (roughly 5 degrees or less of slope), slippage is low and execution error largely results from loss of traction while climbing over 10cm or taller rocks, which can often be avoided through manual path selction. For longer drives or in high-slip areas, the rover must be able to deal with accumulated position error, either through safeguarding itself or by using visual odometry to update its position knowledge. When using visual odometry, the rover operator is responsible for specifying the criteria for halting the drive, since it is too difficult to manually sequence reliable obstacle avoidance behavior. Typically, the halting criteria include proximity to known obstacles, the number of times VisOdom has failed to provide a position update, and a threshold on slippage.

3.4 Adaptation

Mobility performance is uncertain in any novel terrain type and can vary substantially in known terrain types, but humans can quickly learn to steer the rover clear of new hazard types. For example, after Spirit drove through a loose mixture of fine sand and rocks, a potato-sized rock became jammed in one of the wheels. When the rover encountered similar terrain over 100 sols later, rover operators knew to direct Spirit to check for slippage while driving and stop if the rover became bogged down. Post-drive images after the rover detected over 90% slip showed a similar mixture of sand and rocks, with two rocks having the potential to jam in the wheels, and we subsequently retreated to look for another route. This sort of perception and adaptation with a single training example is a key strength of manual terrain analysis.

4 Conclusion

Successful operation of the Mars Exploration Rovers has depended on both manually-directed and autonomous driving. Our experience tells us that the two methods are complementary, and careful selection of the right techniques lead to better overall performance in the face of limited time, power, imagery, and on-board computation. The rover has the advantage of being able to close the loop on execution errors, and assess terrain that is not visible in the imagery available to the human when planning a drive. Humans have enormous perceptual abilities and can adapt to new terrain types and challenges.

For both human and rover, the easiest type of driving is the one that is, in our observations, the most studied in the research community: discrete obstacles on level terrain. Both manual and autonomous driving are highly effective in this terrain, but the limited computational resources of the MER vehicles lead to a preference to start each long traverse with the longest safe directed drive and then continue autonomously until the available time is exhausted.

While most of the distance covered by both rovers has been on level ground with varying degrees of geometric hazarads, most of the time has been spent in more challenging environments coupling steep slopes with loose materials and positive obstacles. In these regimes, slippage is not always predictable and can lead to a variety of outcomes: driving can be inaccurate in the best case, or the rover can become temporarily stuck or can enter an area that it cannot escape, in the worst case. Careful terrain analysis is required in these cases, but visual odometry has also been absolutely essential for safe and accurate driving.

Future vehicles will have faster processors, allowing more advanced terrain analysis and path selection to be performed autonomously. Global and dynamic path planning, which was dropped from MER due to schedule constraints, is an obvious next step. But path planning can only be as good as the underlying obstacle avoidance methodolgoy, and if rovers are to become substantially autonomous then appearance-based terrain analysis will be required. While MER and terrestrial experience can be some guide, a truly useful terrain classification system should be capable of easy adaptation to previously unseen terrain types (possibly with substantial human involvement), since we have yet to see more than a miniscule fraction of the Martian surface.

5 Acknowledgements

This work was performed at the Jet Propulsion Lab under contract with NASA. Thousands of people were involved in the mission at JPL and throughout the world, in government, academia, and industry. We gratefully acknowledge their outstanding work, which enabled us to explore Mars through Spirit and Opportunity. We would like to thank the entire Mars Exploration Rover team, and in particular the other rover planners: Khaled Ali, Eric Baumgartner, Bob Bonitz, Brian Cooper, Frank Hartman, Scott Maxwell, Ashitey Trebi-Ollennu, John Wright, and Jeng Yen.

References

[1] Jeffrey Biesiadecki, Mark Maimone, and Jack Morrison. The Athena SDM rover: A testbed for mars rover mobility. In *International Symposium on Artificial Intelligence, Robotics, and Automation for Space (i-SAIRAS)*, Montreal, Canada, June 2001. http://robotics.jpl.nasa.gov/people/mwm/sdm-mobility/.

[2] Yang Cheng, Mark Maimone, and Larry Matthies. Visual odometry on the mars exploration rovers. In *IEEE Conference on Systems, Man and Cybernetics*, The Big Island, Hawaii, USA, October 2005.

[3] Paul Backes et al. Sequence planning for the fido mars rover prototype. In *IEEE Aerospace Conference*, Big Sky, Montana, USA, 2003.

[4] Steven B. Goldberg, Mark W. Maimone, and Larry Matthies. Stereo vision and rover navigation software for planetary exploration. In *IEEE Aerospace Conference*, volume 5, pages 2025–2036, Big Sky, Montana, USA, March 2002. http://robotics.jpl.nasa.gov/people/mwm/visnavsw/.

[5] M. Golombek and D. Rapp. Size-frequency distributions of rocks on mars and earth analog sites: Implications for future landed missions. *Journal of Geophysical Research - Planets*, 102(E2):4117–4129, February 1997.

[6] Brian Harrington and Chris Voorhees. The challenges of designing the rocker-bogie suspension for the mars exploration rover. In *37th Aerospace Mechanisms Symposium*, Galveston, Texas, USA, May 2004.

[7] C. Leger and R. Deen. Remote image analysis for mars exploration rover mobility and manipulation operations. In *IEEE Conference on Systems, Man and Cybernetics*, The Big Island, Hawaii, USA, October 2005.

[8] J. N. Maki, J. F. Bell III, K. E. Herkenhoff, S. W. Squyres, A. Kiely, M. Klimesh, M. Schwochert, T. Litwin, R. Willson, A. Johnson, M. Maimone, E. Baumgartner, A. Collins, M. Wadsworth, S. T. Elliot, A. Dingizian, D. Brown, E. C. Hagerott, L. Scherr, R. Deen, D. Alexander, and J. Lorre. Mars exploration rover engineering cameras. *Journal of Geophysical Research*, 108(E12):12–1–12–24, December 2003. http://www.agu.org/pubs/crossref/2003/2003JE002077.shtml.

[9] Reid Simmons, Lars Henriksen, Lonnie Chrisman, and Greg Whelan. Obstacle avoidance and safeguarding for a lunar rover. In *AIAA Forum on Advanced Developments in Space robotics*, Madison, WI, August 1996. http://www.cs.cmu.edu/˜reids/papers/AIAAobsAvoid.pdf.

[10] J. Yen, B. Cooper, F. Hartman, S. Maxwell, and J. Wright. Sequence rehearsal and validation on surface operations of the mars exploration rovers. In *SpaceOps*, Montreal, Canada, 2004.