Jeffrey J. Biesiadecki
P. Chris Leger
Mark W. Maimone
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA USA
mwm@helios.jpl.nasa.gov

Tradeoffs Between Directed and Autonomous Driving on the Mars Exploration Rovers

Abstract

NASA's Mars Exploration Rovers (MER) have collected a great diversity of geological science results, thanks in large part to their surface mobility capabilities. The six wheel rocker/bogie mobility system provides driving capabilities in a range of terrain types, while the onboard IMU measures actual rover attitude changes (roll, pitch and yaw, but not position) quickly and accurately. Four stereo camera pairs provide accurate position knowledge and/or terrain assessment. Solar panels generally provide enough energy to drive the vehicle for at most four hours each day, but drive time is often restricted by other planned activities. Driving along slopes in nonhomogeneous terrain injects unpredictable amounts of slip into each drive. These restrictions led to the creation of driving strategies that alternately use more or less onboard autonomy, to 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 limit), 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 Selection.

In total, the MER rovers have driven more than 10 kilometers over Martian terrain during their first 21 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 Autonomous Navigation, Visual Odometry and Path Selection drives.

KEY WORDS—Mars Rover, MER, space robotics, autonomy, mission planning

1. Introduction

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. Scientists and engineers on Earth successfully led the rovers to the areas where their *in situ* instruments found the data they were seeking, using a combination of explicitly directed and autonomous driving modes. Although the achievement of their successful landings stands out as a technological tour de force, it was their ability to traverse across the surface of Mars that enabled both rovers to succeed in their primary goals.

Each MER rover is 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). 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 arm, 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.

Human Rover Drivers were given the task each sol of selecting the appropriate drive mode. In Directed driving, a rover is commanded to drive along a single path without using its ability to image or adapt to the terrain. This is the fastest driv-

The International Journal of Robotics Research
Vol. 26, No. 1, January 2007, pp. 91-104
DOI: 10.1177/0278364907073777

Figures appear in color online: http://ijr.sagepub.com

ing mode, and was often used in flat areas where the chance of losing traction was minimal. In more complicated terrains, for instance on 15–30 degree slopes in craters and on hillsides, Rover Drivers relied on one or more additional autonomous capabilities to ensure the drive would complete as expected (a complete list of drive modes can be found in Biesiadecki and Maimone (2006)): Path Selection, Terrain Assessment, and Visual Odometry. Unfortunately, the vision-based autonomy required substantial time to process images (often 2–3 minutes per stereo pair), so whenever Rover Drivers concluded that the terrain was benign, the Directed mode was used to minimize the time and energy required for the drive, and hence maximize the drive speed and/or distance covered during that sol.

1.1. Related Work

The MER style of remote vehicle control is quite different from most Earth-based styles.

At one extreme, exploration robots on and near Earth might be controlled (or "joysticked") directly. With little communications delay and no substantial bandwidth constraints, this is an attractive style when especially difficult or varying terrains must be navigated. In this style a human driver watches real-time high-bandwidth visual imagery to decide where the robot should go next, sending frequent but relatively low-bandwidth commands to control the remote vehicle's operation. Undersea vehicles (Singh et al. 2005), volcanic rim explorers (Bares and Wettergreen 1999), rescue assistance vehicles (Murphy 2004), even lunar explorers have employed this mode of control. For instance, the Soviet Moon-exploring Lunokhod vehicles used this method and successfully explored over 40 km of the lunar surface in the 1970s (Vaniman et al. 1991), and researchers have proposed Safeguarded Teleoperation for future lunar missions (Krotkov et al. 1996; Wettergreen et al. 1999).

At the other extreme, the robotics community has developed vehicles that are ever more capable of driving themselves autonomously—on Earth. With highly accurate position estimation sensors (e.g., GPS, visual odometry), multiple terrain sensors (e.g., LIDAR, stereo vision, color vision, sonar, ground-penetrating radar), fast computers and dedicated computing hardware, complex mobility designs and refuelable high powered motors, such systems are capable of driving at highway speeds on open roads (Thrun et al. 2006; Urmson et al. 2006; Jochem and Pomerleau 1996), and tens of miles per hour in open unstructured terrain (Kelley et al. 2006). A nice summary of efforts prior to 1995 can be found in Gage (1995).

Other than the two Lunokhod missions mentioned above, there have been very few successful rover missions to other worlds. Not counting NASA's human-driven lunar rovers, very few others have even been attempted. Several planned rover missions (Japan's MUSES-CN nanorover, Soviet Mars 2 and Mars 3 rovers, Phobos hopper) either never launched, failed to land or failed to communicate. However, NASA's Pathfinder mission landed the Sojourner rover on Mars in 1997. Sojourner was the first spacecraft to include onboard autonomous driving capabilities (Mishkin et al. 1998), although that work mentions that the autonomy was greatly underused, partly because of the operations team's generally cautious approach to mission operations.

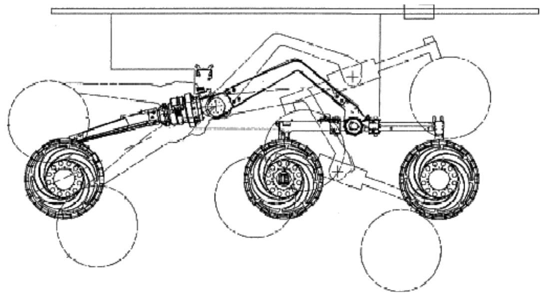

Fig. 1. Rocker-bogie configuration.

Unfortunately, the constraints of operating in a place where radio signals need from 8 to 42 minutes to make a round trip, bandwidth is limited, latency is high, and the fact that there is no way to repair hardware faults forced us to design a more conservative vehicle. Although conservative, each rover is still very capable. But more time and energy is required to use them in this more capable way.

This paper describes some of the issues involved in deciding how we operate our vehicles under complex constraints. Our focus is the topic of when it is more appropriate to use autonomous driving capabilities (spending extra minutes on Mars) versus precisely planned directed drives (requiring hours of engineering analysis on Earth). More information about mission operations in general can be found in (Mishkin et al. 2005).

In Section 2 we describe the resources available on each rover for drive planning and assessment. In Section 3 we cover the primary driving styles that have been developed during MER mission operations. Section 4 discusses the costs and benefits associated with the various drive modes. We present ideas for improvements on future rover systems in Section 5, and conclude in Section 6.

2. Background on Rover Subsystems

In this section we present the various subsystems that contribute to the MER vehicles' overall mobility. Many of the parameters described herein are summarized in Tables 1 and 2.

2.1. Rover Mobility Commands

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 (Harrington and Voorhees 2004). Each wheel is 25 cm in

Table 1. Some of the MER Vehicle Driving Parameters

System Parameter	Value	Units
Mechanical Configuration		
Drivable Wheels	6	
Steerable Wheels	4	
Wheel Diameter	25	cm
Allowed Obstacle Height	20	cm
Max Drive Speed	5	cm/s
Nominal Drive Speed	3.7	cm/s
Smallest Turn Radius	100	cm
Maximum Speed for each Drive Mode		
Directed Driving	124	m/h
Path Selection	96	m/h
AutoNav (benign terrain)	36	m/h
VisOdom	10	m/h
AutoNav Parameters		
Max Step b/w NavCams	200	cm
Max Step b/w HazCams	150	cm
Max Traversable Obstacle	20	cm
Max Traversable Roughness	7	cm
Image Resolution (all)	256×256	pixels
NavCam-based Local Map	1200×1200	cm^2
HazCam-based Local Map	1000×1000	cm^2
Rover Diameter	260	cm
VisOdom Parameters		
Max Step b/w NavCams	60	cm
Max Turn-in-place b/w NavCams	15	degrees

Table 2. Some of the MER Camera Parameters

System Parameter	Value	Units
PanCam Parameters		
Height	152	cm
Baseline	28	cm
Horiztonal Field of View	18	degrees
Max Image Size	1024×1024	pixels
4x1 Binned Image Size	256×1024	pixels
Stereo Range	4 – 70	meters
NavCam Parameters		
Height	152	cm
Baseline	20	cm
Horiztonal Field of View	45	degrees
Max Image Size	1024×1024	pixels
4x1 Binned Image Size	256×1024	pixels
Stereo Range	2 – 20	meters
HazCam Parameters		
Height	52	cm
Baseline	10	cm
Horizontal Field of View	125	degrees
Max Image Size	1024×1024	pixels
4x1 Binned Image Size	1024×256	pixels
Stereo Range	0.5 – 5	meters

diameter and has short 0.5 cm paddle-like cleats, and each is capable of climbing rocks over 35 cm tall. However, the clearance under the body of the rover is only 29 cm on a flat surface (and less when tilted), so in practice obstacles 20 cm or above were avoided. All six wheels can be driven at speeds up to 5 cm/s, but only the four corner wheels are steerable. Double Ackermann-style arcs with turn radii as tight as 1 m may be commanded, as may turns in place in which the vehicle rotates about its origin, located 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 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 greater than 100% slip on sandy soils in the Columbia Hills.

2.2. Mobility Software

Regardless of whether onboard image processing is enabled, actual motion of the rovers is controllable at three levels: low-level commands that specify exactly how much to turn each wheel and steering actuator, 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, based on its actual state (including attitude measurements made by the IMU, and optional position measurements made by Visual Odometry).

Several types of potential vehicle hazards are checked *reactively*, most of them during Real Time Interrupts (RTIs) which occur eight times per second. Available checks include Tilt/Pitch/Roll, Northerly Tilt, Rocker/Bogie Suspension Angles, Motor Stalls, Limit Cycle (no forward progress), and Resource Contention:

Tilt/Pitch/Roll The accelerometers in the IMU provide an instantaneous (unfiltered) vehicle attitude estimate. These three quantities (forward pitch, sideways roll, overall deviation from vertical tilt) are derived from the accelerometer readings and compared to preset limits. If the value for any of these derived 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. Since both rovers are exploring south of the equator, Northerly Tilt can help predict solar energy availability: the higher the Northerly Tilt, the more solar energy will accumulate. Northerly tilt is the angle between the 2D projections of the rover-body-up vector and the East/Up plane onto the North/Up plane. This check can be enabled 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 Manager 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 be raised.

In directed driving, the rover can pre-emptively "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) to make progress toward its goal. This software provides 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 (Ali et al. 2005). In between driving primitives, the rover can use camera-based Visual Odometry (VisOdom) to correct the errors in the initial wheel odometry-based estimate. VisOdom tracks terrain features in NavCam stereo images and uses the tracking information to estimate true vehicle motion during small steps; the rover can only move roughly 60cm, or turn 15 degrees, before successive NavCam images lack enough overlap to reliably estimate motion (Cheng et al. 2006).

Both directed and path selection modes of driving can make use of onboard stereo vision processing and terrain analysis software to determine whether the rover would encounter geometric hazards as it drives along its chosen path.

The computing resources required by these different commands vary greatly. Directed driving commands execute the most quickly (achieving speeds up to 124 m/h), 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 (improving safety), but only achieve driving speeds from 10 m/h in obstacle-laden terrain to 36 m/h 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 m/h. When both AutoNav and VisOdom are used, the traverse rate drops to 6 m/h. This factor of 20 difference in speed between using the least and greatest autonomy clearly points to the need to select the level of autonomy based on the specific requirements of each drive.

2.2.1. 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* locate traverse hazards and avoid them. The procedure is summarized below; see Goldberg et al. (2002) and Biesiadecki and Maimone (2006) for details and Simmons et al. (1996) for the approach that inspired it.

Some important configurable parameters that impact autonomous terrain analysis include max traversable obstacle size (e.g., 20 cm), max terrain tilt angle (e.g., 20 degrees), average surface roughness (e.g., 7 cm), and image resolution for stereo processing (256 × 256 pixels; although Earth-based development used 128 × 128 images through much of the test program, imagery from Mars demonstrated that better maps would be generated from the higher resolution stereo data).

1. The rover chooses a stereo camera pair based on the goal location and its previous terrain assessment, if any. Images are acquired from the CCD using 4-times row-binning, then downsampled in software from 256 × 1024 12-bit pixels down to 256 × 256 8-bit pixels.

2. The stereo pair of images is densely correlated, with noisy or uncertain data masked or filtered out, and the resulting range data are distributed into a local map grid comprised of 20 cm × 20 cm cells (50 × 50 cells for Spirit, 60 × 60 for Opportunity, which uses NavCams and therefore can see farther more reliably).

3. Disc-shaped planar patches are fit to the measured range points. Each disc is 2.6 m in diameter, slightly larger than the volume swept by the vehicle during a "turn in place" maneuver. Parameters of the plane fit are used to look for potential geometric hazards; the surface normal is compared to the max terrain tilt, the residual of the fit is compared to the max roughness, and the max elevation difference between cells in the patch (corrected for tilt and wheel base) is compared to the max traversable obstacle size. The most conservative assessment is logged as the "goodness" of that cell.

4. 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 (Biesiadecki and Maimone 2006), which stores knowledge of the terrain as goodness values as far as 5 or 6 m away from the vehicle.

5. 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 goodness values in all the cells it touches; see Figure 2.

6. Each potential path evaluation is weighted by additional Gaussian-shaped bias functions centered on (1) the direct path toward the current goal, and (2) the current steering direction.

7. The top 10% of all path evaluations are marked, and the best of those is chosen as the next commanded path. If its evaluation is less than some threshold, then additional images are taken in the opposite ("anti-goalward") direction and evaluated. If the best path is still below a threshold, no path is chosen and the drive terminates.

8. Motion continues until the goal is reached, the command times out, or a type of drive error occurs.

The rock-strewn terrain encountered by Spirit at Gusev Crater corresponds well to the exponential rock distribution models predicted using data from Viking I, II, and Pathfinder missions (Golombek and Rapp 1997). The body-mounted 120-degree field of view (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 4 m ahead (the nominal planning distance) in the wide FOV HazCam images at 256 × 256 resolution. Fortunately, the lack of large scale features implies a lack of large "step" obstacles. So, Opportunity was reconfigured to perform terrain assessment with narrower FOV NavCam images, which have enough spatial resolution to resolve terrain features 4 m away. Rock and fissure obstacles can still be detected, but the limited FOV means less of the terrain around any obstacles will be understood, which reduces its ability to steer around them autonomously. If Opportunity ever needs to be driven autonomously through rock-laden terrain, additional NavCam stereo images to increase the FOV can be processed autonomously, but that will slow the effective driving rate even more.

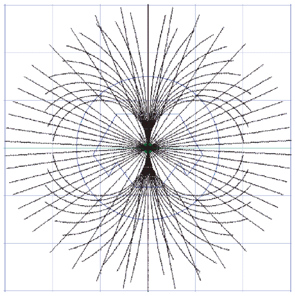

Fig. 2. Plot of the 96 paths (arcs and point turns followed by straight line drives) normally considered during Path Selection autonomous drives. The rover outline points up in this view, and the light blue grid lines indicate 1 spacing meter.

All MER surface software runs on a 20 MHz RAD6000 computer under the VxWorks operating system. The slow processor speed, and the sharing of a single address space and cache by dozens of tasks, mean Autonomous Navigation (AutoNav) and VisOdom software run slowly.

2.3. Ground-based Terrain Analysis

Directed drives will always execute more quickly than those employing autonomous terrain analysis. But directed drives are only safe when ground-based terrain assessment has determined not only that a safe path exists, but also that any deviations from that path (e.g., due to vehicle slip) can be detected and compensated for. The need for humans to assess non-geometric hazards and the desire to save tens of minutes of onboard computation combine to make human assessment of the terrain an integral part of MER operations.

Ground-based terrain assessment is generally performed 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, 10 cm baseline) Hazard Cameras (HazCams) rigidly mounted 53 cm above the ground plane on the front and back sides, one pair of medium field-of-view (45 degree, 20 cm baseline) Navigation Cameras (NavCams) mounted 152 cm above the ground plane on a pan/tilt head, and one pair of narrow field-of-view (18 degree, 28 cm baseline) Panoramic Cameras (PanCams) also mounted 152 cm above the ground plane on the pan/tilt head (Maki et al. 2003). These cameras take images at resolutions up to 1024 × 1024 12-bit pixels that provide information about terrain texture, and stereo range-derived terrain shape at different scales: around 0.5 m–5 m in the HazCams, 2 m–20 m in the NavCams, and 4 m–70 m in the PanCams.

The amount of directed driving that can be commanded depends on both the terrain itself and on how much information about the terrain is available. Orbital imagery, while crucial for long-range planning, cannot resolve vehicle hazards such as 20 cm rocks and does not provide accurate slope information. So after each drive, images from each camera pair are requested, typically eight to twelve stereo image pairs in the drive direction (some from each camera). Sometimes only portions of the images were downlinked from Mars, resulting in fewer images than were requested and thus limiting the distance of the initial directed drive.

Downlinked stereo image pairs are processed by an automated pipeline that generates derived products including 3D range maps, texture-mapped terrain meshes, and color overlays indicating terrain properties such as slope and elevation (Leger and Deen 2005). Rover operators use image-based querying tools to measure ranges to terrain features and estimate distances and rock sizes (Backes et al. 2003). 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, 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 (Yen et al. 2004). The raw images are also extremely useful in assessing traversability: operators can readily identify sandy or rocky areas that present hazards, though new terrain types always carry an element of uncertainty regarding vehicle performance. In some cases, no image cues enable 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 different amounts of traction or sinkage. For example, while driving uphill toward a topographic high point named “Larry’s Lookout” on Sol 399, 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 discernible difference in the character of the surface.

It probably comes as no surprise to a computer vision researcher that the human perceptual system, while qualitative and imperfect, is extremely capable. When combined with quantitative 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. In contrast, the MER software does not have any appearance-based terrain analysis capabilities; it only detects geometric obstacles. Nevertheless, the most serious and frequent hazards (rocks, steps, and high-center hazards) can be detected by geometric analysis as long as sufficient range data is available. At longer ranges (over 15 m in NavCam images, and over 50 m 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, onboard 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 its lack of a global planner (which the human stands in for during directed drives) can cause the rover to get stuck in cul de sacs.

3. Drive Techniques and Templates

Planned drive activities are sent to the rovers in one or more sequences of commands to be executed on a given sol. Most drive activities can be classified as either traverses (covering maximum distance) or approaches (driving to a specific position for subsequent *in situ* arm-based science operations). The techniques used for each drive type are determined based on the time and energy allocated for driving, the amount of free space remaining in flash memory, the amount of terrain visible in imagery, known hazards, and the level of uncertainty in rover position given the terrain type as discussed in Section 2.3. Generally, driving on level ground requires a mix of blind and AutoNav driving, and driving on slopes requires using VisOdom to allow the rover to compensate for unpredictable slip.

Over time these approaches evolved into patterns of commands that were able to be encapsulated into a template framework. Here we present the context that led to the development of three such templates.

3.1. Traversing the Plains

The rovers’ highest traverse rates are achieved in terrain with low obstacle density, low slopes, and a lack of high-slip hazards (e.g., deep windblown deposits of fine dust). When these conditions are met, human rover drivers can make maximum use of planning imagery to safely command fast and accurate directed drives, followed by further driving using onboard hazard avoidance. In terrain such as the plains of Gusev Crater and Meridiani, onboard hazard detection can often be traded for faster and longer directed driving – putting the burden of obstacle detection on human operators.

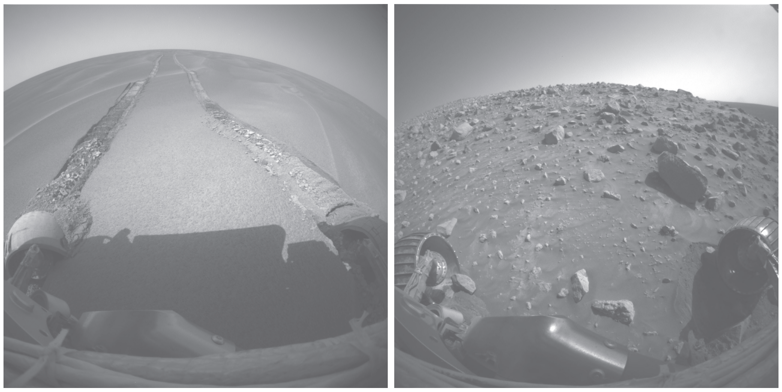

Fig. 3. Left: On Sol 446, Opportunity found its wheels more than half buried in sand. Although not a geometric hazard, the ripple of sand on which it stopped kept the human planners busy for weeks. Right: On Sol 454, Spirit terminated its drive early after detecting 90% slip. This image shows rocks that had collected next to the left front wheel.

We learned during our initial drives in each terrain that driving on level ground typically leads to accurate and predictable mobility performance; e.g., Spirit only accumulated 3% position error over 2 kilometers of driving (Li et al. 2005). Because of the rover's limited processing power, drives using autonomous hazard avoidance are several times slower than "blind" (manually-directed) drives. These two facts led us to favor long initial blind drives to achieve the longest drives in the least amount of rover execution time. 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 often allowed for blind drives of up to 70 m. 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 50 m 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 is to rely on AutoNav to find a path.

This drive template—drive blind for as long as possible given the safe terrain visible in planning images, then use AutoNav—is the fastest method of driving the rovers. The rovers can travel at 124 m/h for the first 20–70 m, then at 10–36 m/h for the remaining distance. Thus, the rover might cover 50 m in the first 25 minutes of blind driving, then another 20 m in an hour of AutoNav. In contrast, covering this same distance using AutoNav alone might take anywhere from 2 to 7 h using AutoNav the whole distance, or up to 12 hours if using AutoNav and VisOdom—both of which are unattractive given that energy and thermal limitations typically allow only 1.5 to 2 h of driving each sol. However, this drive template is limited to terrain with low slip and good coverage in planning imagery.

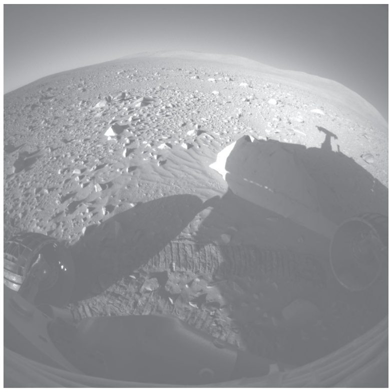

Fig. 4. On Sol 147, the onboard terrain analysis performed during a Guarded Arc prevented Spirit from driving into this rock.

On the plains of Meridiani, the terrain hazards are quite different and initially allowed for blind drives over 100 m. Unlike the Gusev plains, there has been a near-total absence of rocks at Meridiani, and until Sol 446 (see left side of Figure 3) none of the innumerable sandy ripples posed a threat to the rover. The record so far for long distance plains driving is over 390 m covered by Opportunity during Sols 383 through 385; 106 m of blind driving was followed by 284 m of autonomous driving spread over 3 sols.

Craters, visible in orbital imagery, and small linear depressions were the most significant hazards for Opportunity. While driving over flat terrain, the rover's suspension does not articulate significantly, which suggested that a measured suspension articulation change could be used to halt driving if the rover were to encounter a depression. In April 2004, the rover's software was upgraded to allow the rover's suspension angles to be checked against preset limits at 8 Hz,

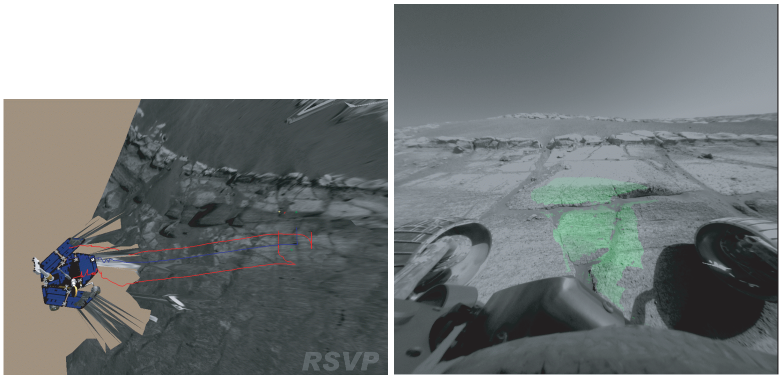

Fig. 5. Opportunity's planned 8.7 m drive along a 20–24 degree slope on Burns Cliff (see Figure 6) on Sol 304, and the front HazCam view confirming a successful single sol approach. The shaded area shows those parts of the surface reachable by the instrument arm, which includes the light bedrock that was the target of the drive. A combination of VisOdom and conditional sequencing was used to accomplish this drive.

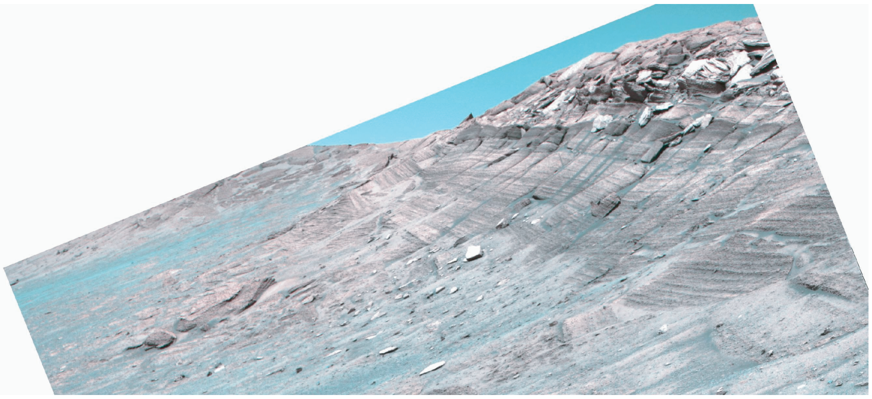

Fig. 6. View of Burns Cliff inside Endurance Crater, the area next to the Sol 304 target approach drive. Opportunity reached a tilt of 31 degrees along this slope, the highest tilt reached during its first 21 months of operation.

thus enabling the rover to stop at negative terrain features (i.e., holes) that were not visible a priori. Since the reason for halting a drive (e.g., timeout, suspension check, slip amount, or tilt check) is accessible to the rover sequencing language, a recovery maneuver could be performed whenever the suspension check tripped. The recovery consists of backing up several meters and continuing the drive with AutoNav, since AutoNav is able to detect and avoid negative hazards. Once Opportunity left Endurance Crater, the drive template became an initial blind drive of 30–60 m 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 use 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 antenna interference caused by the rover's mast during communication with Earth or an orbiter. However, this presents a problem for the next sol's Instrument Deployment Device (IDD, i.e., the instrument arm) operations: since no camera can see the part of the IDD deployment volume under the rover, a front HazCam image pair of the final terrain must be safely acquired 0.5–3 m before driving to the rover's final location to allow engineers to determine if the IDD can be safely deployed.

The obvious solution is to turn to the desired heading, acquire the image pair, then drive a short distance to the final location. But the final position of the rover cannot be predicted when driving under AutoNav control, so the rover operator cannot guarantee in advance that it will be safe to travel some additional distance at the desired heading. Instead, a turn in place followed by the "guarded arc" drive primitive forces the

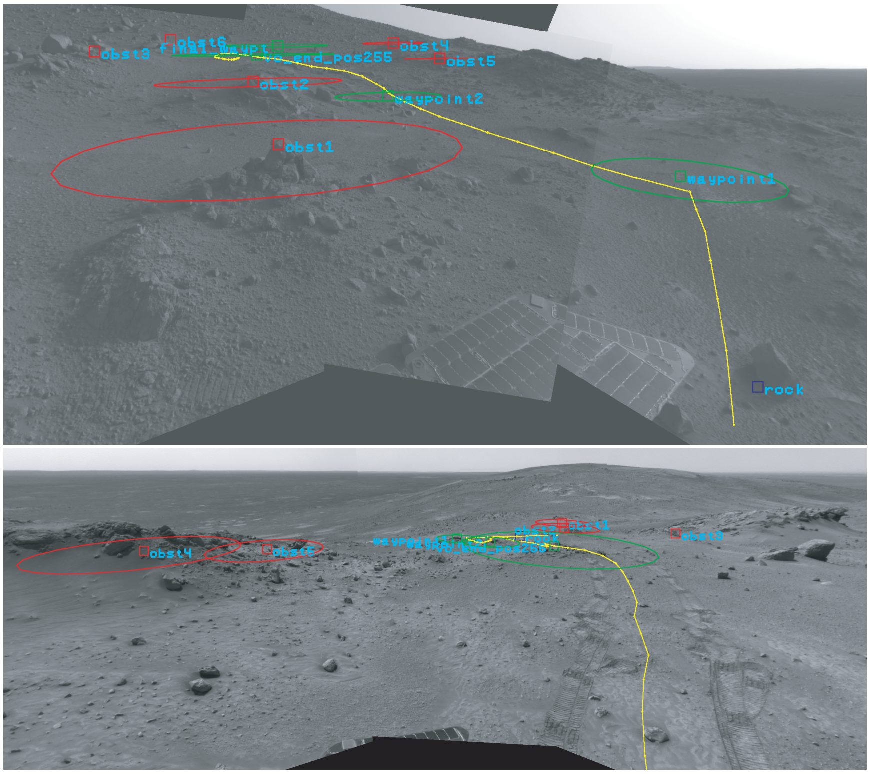

Fig. 7. Spirit's Sol 436 drive used a variety of driving modes to traverse steep terrain. Data from the actual course taken is plotted in a Sol 434 (pre-drive) NAVCAM mosaic above and a Sol 438 (post-drive) NAVCAM mosaic below. The rover course actually driven during Sol 436 is plotted in yellow, intermediate waypoints are green, obstacles are red and other features are blue. The waypoints, obstacles and keep-out zones were selected manually by human Rover Drivers studying image mosaics and stereo images, and simulating drives over a 3D mesh of the terrain (Wright et al. 2006). VisOdom was used to ensure the rover stays on track through the labeled obstacles (up to `vo_end_pos255`), then AutoNav was used to extend the drive into the area behind the ridge (only visible in the lower mosaic, taken after the drive completed). During the actual drive, 1.5 m of Directed driving, 0.5 m of Guarded motion, 19 m of VisOdom (where slip from 0 to 87%, averaging 18% was encountered), and 10.5 m of AutoNav were commanded before the sequence reached its planned timeout.

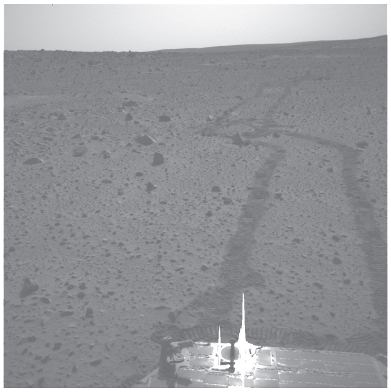

Fig. 8. On Sol 109, Spirit avoided obstacles in previously-unseen terrain.

rover to drive in a guaranteed direction, but only if the onboard terrain analysis shows that it is safe to do so. See Figure 4 for an example sol that benefited greatly from this capability.

3.2. Driving on Slopes: Mountains and Craters

While most of the distance covered by the rovers has been on level ground, the rovers have been on moderate to high slopes during most of the target approach drives and for more than 700 sols. Unlike driving on level ground, driving on slopes rarely allows a trade purely between directed drives and onboard hazard detection: the rovers invariably slip when driving on slopes, making VisOdom essential for safe and accurate driving. But using AutoNav along with VisOdom takes roughly twice as much time as using VisOdom alone, making this combination impractical for normal use.

This presents a challenge: the rover has the ability to know where it is, but in that mode cannot detect new obstacles. Additionally, in steep terrain the rover cannot identify all obstacle classes, since the rover has no means of detecting sandy, high-slip areas in advance. Even areas of only 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 specify “keep out zones” which will cause the rover to halt before a hazard is encountered (e.g., see Figure 7). The rover keeps track of its position using VisOdom and can close the loop using Path Selection 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 Sol 454, Spirit promptly halted driving after detecting slippage over 90%, and post-drive HazCam images showed several rocks on the verge of falling into the wheels, since the wheels had dug into the terrain by nearly one wheel radius (see Figure 3, right). The recurrence of high slopes, sandy terrain with intermixed small rocks, and frequent obstacle-sized rocks caused us to retreat and find a new route to the summit of Husband Hill, which was finally reached on sol 582.

The combination of using VisOdom to accurately measure rover position and detect high slip, while using Keep-Out Zones to safeguard against hazards, allowed the best daily of progress in the Columbia Hills. Due to high and unpredictable slip, blind driving—while fast—would limit the length of a sol’s drive to the distance at which a realistic amount of slip could cause a potential collision with an obstacle. In practice, this was often only 3–10 m given frequent slip up to 50% and closely-spaced obstacles. While such a traverse may take less than 10 minutes, this is less progress per sol than using VisOdom at a rate of 10 m/h when one to two hours of drive time are available. In practice, blind driving and AutoNav were sometimes feasible in the Columbia Hills for part of a drive segment: for example, VisOdom might be used for the first 10 m to negotiate a steep obstacle field, then a blind drive could quickly traverse another 10 m of level, open terrain to the edge of NavCam coverage, after which point AutoNav, with periodic VisOdom slip checks, could be used. While driving in the Columbia Hills and in Endurance Crater was characterized by the use of VisOdom, the amount of VisOdom, blind, and AutoNav driving were tuned each day to make best use of the time available while accounting for the specific challenges and opportunities presented by the terrain. The cost of this terrain-specific tuning was that significantly more human effort was required to create and verify drive plans with several hundred commands, in half a dozen nested sequences, each day.

3.3. Target Approach

Whereas traverse sequences focus on covering maximum distance over terrain of lesser interest to the science team, target approach sequences aim to place the rover at a specific target position and orientation 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 5–10 cm. Thus, target approach drives are characterized by a specific goal and relatively high accuracy, rather than a desire for covering maximum distance. Unless an approach drive is lengthy (10–15 m) and the terrain steep—where VisOdom is thus required—the amount of time available for an approach drive was rarely the limiting factor.

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 1–3 m away) is chosen along the heading vector, and then the drive to the standoff is planned.

On firm, level ground, directed drive primitives are usually sufficient for accurate target approaches from 2–10 m away. On sloped or soft terrain, VisOdom is required to close the loop on the rover’s position, 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 30 cm conditional steps that will execute if the rover center is greater than 1.3 m away from the target, and then several conditional 10cm steps that will execute if the rover is greater than 1.1 m 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 between 1.0 and 1.1 m away from the rover center, placing the target at the sweet spot of the manipulator’s workspace. After each motion, VisOdom updates the rover’s position knowledge, allowing it to correct for slip-induced errors. This conditional sequencing strategy, combined with VisOdom, allows the rover to accurately approach targets 5–10 m away while driving on slopes in the 10 to 20 degree range (e.g., see Figures 5 and 6), with the caveat that on surfaces with sufficiently low bearing strength, the rover is mechanically incapable of making direct uphill progress.

4. Relative Merits of Directed/Autonomous Driving

There are significant differences in resource usage between manual and autonomous driving, with execution time and generated data volume being the most obvious. Energy 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, IMU, and other electronics will draw power for the entire duration of the drive and thus an autonomous drive will require more energy 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 (e.g., see Figure 7). 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 h to build a drive sequence to travel 20–40 m across the plains of Gusev. This decreased dramatically later in the mission, often requiring only 2–4 h to sequence drives over 100 m 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 human-spent sequencing time and rover-spent execution time for directed and autonomous drives, though execution time was usually the limiting factor in drive length each sol.

There is an additional long-term resource trade-off: humans can rapidly adapt their sequences to deal with new terrain types or drive requirements, but changing the onboard software involves a lengthy software development, testing, and uplink process. Instead of a day-to-week turnaround in sequence development, flight software updates to cope with new terrain and drive techniques occur on a months-to-year cycle due to the rigor and complexity of developing, regression testing, validating, and uplinking new software to a spacecraft.

4.1. Driving into the Unknown

There is one notable circumstance in which the human has no ability to safely select paths: when driving into terrain that has not been imaged. On Sol 109, Spirit was commanded to drive over the local horizon 50 m 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 energy, and post-drive images showed AutoNav correctly avoiding large rocks while traversing slopes up to 9 degrees (see Figure 8). Obviously, a high degree of confidence in the hazard avoidance software is needed in situations such as this; *AutoNav has kept both vehicles safe through over 2700 m of traverse as of May 2005.* Less severe, but more frequent, instances in which humans cannot guarantee rover safety occur when the rover drives beyond the distance at which obstacles can be resolved, 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.

4.2. 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 (less than 5%) and execution error largely results from loss of traction while climbing over 10 cm or taller rocks, which can often be avoided through manual path selection. 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 VisOdom to update its position knowledge. When using VisOdom, the rover operator is responsible for specifying the criteria for halting the drive, since manually sequencing reliable obstacle avoidance is too difficult. 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.

Figure 9 summarizes the distance covered and the type of driving modes used for each rover during their first 19 months of operation.

4.3. 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 newly identified hazard types. For example, after Spirit drove through a loose mixture of fine sand and rocks on Sol 339, a potato-sized rock jammed in one of the wheels, requiring a week of cautious movements to eject the rock. 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 (see Figure 3, right). This sort of perception and adaptation with a single training example is a key strength of manual terrain analysis.

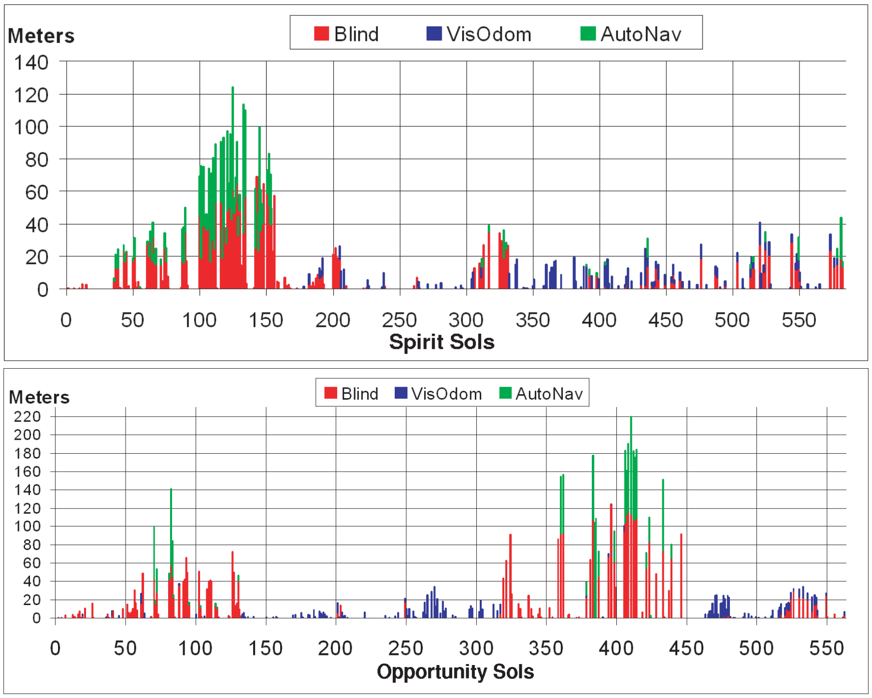

Fig. 9. Summary of distances driven by each rover (Spirit above Opportunity) per sol. AutoNav drives (in green) include any mode in which terrain assessment was done onboard (i.e., both AutoNav and Guarded motion), VisOdom drives (in blue) include both Directed and Path Selection driving modes but not AutoNav, and Blind drives (in red) include both directed arcs and rover-adapted Path Selection drives that compensated for yaw changes measured during the drive. The changing quality of the drive types suggests how human and rover driving strategies alike had to adapt to new terrains many times over the course of each mission.

5. Future Work

While Spirit and Opportunity continue to perform well beyond our original expectations, our experience operating the rovers suggests some areas for improvement. Perhaps the most obvious area for improvement is computational efficiency: driving with either VisOdom or AutoNav can slow the rovers' progress by up to an order of magnitude compared to directed drives. Aside from algorithmic and implementation optimization, some speedup can likely be obtained by accepting decreased accuracy: one use of VisOdom is to simply detect when the rover is slipping substantially, in which case a precise motion estimate is not required.

Another promising avenue for future work is terrain classification. Our current hazard avoidance software detects only geometric hazards, but areas with weak soil–particularly wind-driven drifts–have proven treacherous for both rovers. The ability to learn what high-slip terrain looks like so that it can be autonomously avoided (even dynamically updating the onboard interpretation of the terrain) would be a great benefit. One potentially useful observation is that slippage is almost always correlated with sinkage, and sinkage can be measured by observing either the wheels or the degree to which the wheels leave complete tracks in the terrain.

In terms of mobility system development, one area that seems to be underemphasized is precision mobility in natural terrain. For the types of investigation undertaken by Spirit and Opportunity, mere mobility–the ability to traverse a certain-sized obstacle, travel at a certain rate, or climb a certain slope–is not sufficient. The ability to reliably navigate the rover to within centimeters of a desired location, on slopes, near obstacles, and with external constraints on final vehicle heading, has been of the utmost importance in uncovering the water history of Mars.

Flexibility in the rovers' command language and onboard software has been critical in allowing us to encode our human and ever-changing understanding of the terrain and vehicle performance. Exploration, by definition, sends the rovers in terrain that has never been seen or driven upon before. We have responded to new terrains by adding more terrain-dependent checks—such as VisOdom slip tests—when appropriate on a sol-to-sol basis, and by reactive sequencing that can detect when a drive is not going as planned, and respond appropriately (usually by aborting the drive, but sometimes by continuing with less speed but more autonomy). While not a traditional robotics problem, it would be beneficial to introduce methods for easily formalizing and re-using new sequence idioms to reduce human errors and speed the sequence design, simulation and validation processes. Writing a sequence is writing a program, and perhaps techniques could be applied from extreme programming, automatic code generation, and other software development paradigms.

MER software development continues today. Several new robotics technologies were uplinked to the rovers in mid-2006. These include the ability to perform autonomous *in situ* instrument placement following a successful drive, something that previously required human engineering assessment prior to deploying the instrument arm; global path planning to enable intelligent backtracking using the Field D* algorithm; visual servoing to autonomously track and/or drive toward a terrain feature; and autonomous detection of dust devils and clouds in onboard imagery to optimize the science content of downlinked imagery.

Future vehicles will have faster processors, allowing more advanced terrain analysis and path selection to be performed. But path planning can only be as good as the underlying obstacle avoidance methodology, and if rovers are to become substantially autonomous then appearance-based adaptive terrain analysis will also 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 minuscule fraction of the Martian surface.

6. Conclusions

Successful operation of the MER vehicles 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 leads to better overall performance in the face of limited time, energy, imagery, and onboard 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: the case of 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 led 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 hazards, 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, and VisOdom has also been absolutely essential for safe and accurate driving.

Acknowledgements

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology under contract to the National Aeronautics and Space Administration. 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, Robert Liebersbach for his mission-wide data processing tools, and in particular the other rover planners: Khaled Ali, Eric Baumgartner, Paolo Bellutta, Bob Bonitz, Brian Cooper, Frank Hartman, Scott Maxwell, Ashley Stroupe, Ashitey Trebi-Ollennu, Eddie Tunstel, John Wright, and Jeng Yen.

References

Ali, K. S., Vanelli, C. A., Biesiadecki, J. J., Maimone, M. W., Cheng, Y., San Martin, M., and Alexander, J. W. (2005). Attitude and position esimation on the Mars Exploration Rovers. *IEEE Conference on Systems, Man and Cybernetics*, The Big Island, Hawaii, USA, October.

Bares, J. and Wettergreen, D. (1999). Dante II: Technical description, results and lessons learned. *International Journal of Robotics Research*, **18**(7): 621–649.

Biesiadecki, J. J. and Maimone, M. W. (2006). The Mars Exploration Rover surface mobility flight software: Driving ambition. *IEEE Aerospace Conference*, Big Sky, Montana, USA, March, Vol. 5.

Cheng, Y., Maimone, M., and Matthies, L. (2006). Visual Odometry on the Mars Exploration Rovers. IEEE Robotics and Automation Special Issue (MER), **13**(2): 54–62.

Backes, P. et al. (2003). Sequence planning for the Fido Mars rover prototype. *IEEE Aerospace Conference*, Big Sky, Montana, USA.

Gage, D. W. (1995). UGV history 101: A brief history of unmanned ground vehicle (UGV) development efforts. *Unmanned Systems Magazine Special Issue on Unmanned Ground Vehicles*, 13(3), Summer. http://www.nosc.mil/robots/pubs/ugvhist95-nopix.pdf.

Goldberg, S. B., Maimone, M. W., and Matthies, L. (2002). Stereo vision and rover navigation software for planetary exploration. *IEEE Aerospace Conference*, Big Sky, Montana, USA, March, Vol. 5, pp. 2025–2036.

Golombek, M. and Rapp, D. (1997). 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.

Harrington, B. and Voorhees, C. (2004). The challenges of designing the rocker-bogie suspension for the Mars Exploration Rover. *37th Aerospace Mechanisms Symposium*, Galveston, Texas, USA, May.

Jochem, T. and Pomerleau, D. (1996). Life in the fast lane: The evolution of an adaptive vehicle control system. *AI Magazine*, **17**(2): 11–50.

Kelley, A., Stentz, A., Amidi, O., Bode, M., Bradley, D., Diaz-Calderon, A., et al. (2006). Toward reliable off road autonomous vehicles operating in challenging environments. *International Journal of Robotics Research*, **25**(5-6): 449–483.

Krotkov, E., Simmons, R., Cozman, F., and Koenig, S. (1998). Safeguarded teleoperation for lunar rovers. *26th International Conference on Environmental Systems*, July.

Leger, C. and Deen, R. (2005). Remote image analysis for Mars Exploration Rover mobility and manipulation operations. *IEEE Conference on Systems, Man and Cybernetics*, Big Island, Hawaii, USA, October.

Li, R. X., Archinal, B. A., Arvidson, R. E., Bell, J., Christensen, P., Crumpler, L., et al. (2005). Spirit rover localization and topographic mapping at the landing site of Gusev Crater, Mars. *JGR-Planets, Special Issue on Spirit Rover,* to appear.

Maki, J. N., Bell III, J. F., Herkenhoff, K. E., Squyres, S. W., Kiely, A., Klimesh, M., et al. (2003). Mars Exploration Rover engineering cameras. *Journal of Geophysical Research*, **108**(E12): 12-1–12-24. http://www.agu.org/pubs/crossref/2003/2003JE002077. shtml.

Mishkin, A., Morrison, j., Nguyen, T., Stone, H., Cooper, B., and Wilcox, B. (1998). Experiences with operations and autonomy of the mars pathfinder microrover. *Proceedings of the 1998 IEEE Aerospace Conference*, Snowmass at Aspen, CO, March.

Mishkin, A. H., Limonadi, D., Laubach, S. L., and Bass, D. S. (2006). Working the martian night shift: The MER surface operations process. *IEEE Robotics and Automation Special Issue (MER)*, pp. 46–53.

Murphy, R. (2004). Human-robot interaction in rescue robotics. *IEEE Systems, Man and Cybernetics Part C: Applications and Reviews, special issue on Human-Robot Interaction*, **34**(2).

Simmons, R., Henriksen, L., Chrisman, L., and Whelan, G. (1996). Obstacle avoidance and safeguarding for a lunar rover. *AIAA Forum on Advanced Developments in Space robotics*, Madison, WI, August. http://www.cs.cmu.edu/˜reids/papers/AIAAobsAvoid.pdf.

Singh, H., Roman, C., Pizarro, O., and Eustice, R. (2005). Advances in high resolution imaging from underwater vehicles. *International Symposium of Robotics Research*, San Francisco, CA, USA, October.

Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., et al. (2006). Winning the DARPA grand challenge. *Journal of Field Robotics*, Accepted for publication, http://robots.stanford.edu/papers/thrun.stanley05.pdf.

Urmson, C., Ragusa, C., Ray, D., Anhalt, J., Bartz, D., Galatali, T., et al. (2006). A robust approach to high-speed navigation for unrehearsed desert terrain. *Journal of Field Robotics*, **23**(8): 467–508.

Vaniman, D. T., Heiken, G. H., and French, B. M. (eds). (1991). *Lunar Sourcebook: A User's Guide to the Moon*. Cambridge University Press.

Wettergreen, D., Bapna, D., Maimone, M., and Thomas, G. (1999). Developing Nomad for robotic exploration of the Atacama desert. *Robotics and Autonomous Systems*, **26**(2–3): 127–148.

Wright, J., Hartman, F., Cooper, B., Maxwell, S., Yen, J., and Morrison, J. (2006). Driving on mars with rsvp. *IEEE Robotics and Automation Special Issue (MER)*, pp. 37–45.

Yen, J., Cooper, B., Hartman, F., Maxwell, S., and Wright, J. (2004). Sequence rehearsal and validation on surface operations of the Mars Exploration Rovers. *SpaceOps*, Montreal, Canada.