Chapter 3

SURFACE NAVIGATION AND MOBILITY INTELLIGENCE ON THE MARS EXPLORATION ROVERS

Mark Maimone, Jeffrey Biesiadecki, Edward Tunstel, Yang Cheng, Chris Leger
NASA Jet Propulsion Laboratory, USA

1. INTRODUCTION

The NASA Mars Exploration Rovers (MER) mission represents one of the longest deployments of robotic intelligence on remote planetary surfaces thus far. In addition to establishing a landmark in planetary *in situ* scientific exploration, MER represents a new benchmark and the current state-of-the-art for planetary surface robotics in the first decade of the new millennium. The twin rovers, *Spirit* and *Opportunity*, were required to use mobile autonomy to enable success of their distinct surface missions on roughly opposite sides of planet Mars (each of which has been underway since January 2004).

Surface mobility software designs for these rovers include robotic autonomy capabilities of varying complexity for navigation, science instrument placement onto surface samples, onboard resource management, and science data gathering. This chapter focuses on the autonomous navigation capabilities of the rovers. Requirements on the autonomous mobility subsystem can be summarized as: the rover(s) must be able to safely traverse some substantial minimum distance per day for operations in terrain of some reference complexity, while maintaining estimated position knowledge within some small percentage of distance traversed. This requires a level of onboard intelligence sufficient to achieve mobility and navigation objectives on Mars, given only daily guidance from human mission operators on Earth. Autonomy software for surface mobility comprises algorithms for vision-based autonomous navigation and, to some extent, lower-level sensor-based motion control. We refer to autonomy software for surface mobility as onboard software designed for data processing associated with sensing, perception, reasoning and decision-making, for the ultimate purpose of directing servo-level motion execution. This robotic intelligence software is at a higher level than servomotor control and at a lower level than symbolic planner-schedulers; the latter was not implemented onboard the *Spirit* and *Opportunity* rovers (the equivalent planning function is performed by mission controllers on Earth using planning and scheduling software tools).

The MER mission achieves planetary surface exploration using a human-robot system that operates in a semi-autonomous fashion [1-3]. That is, it incorporates remote planning, command sequencing and visualization of rover activity sequences and related data products by an Earth-based science-engineering team, all under extreme time delay with limited, intermittent communication afforded by daily uplink and downlink cycles of deep space networks. As the principal robotic part of this semi-autonomous system, the rovers perform autonomous execution of a daily sequence of commands. Autonomous navigation is commanded by specifying explicit surface coordinates to be reached using onboard sensing, perception, motion planning and execution. Mission operators provide a global path plan in the form of a series of waypoints that is executed and evaluated onboard each rover. Often, the goal waypoint at the furthest extent of a long traverse is beyond the reliable field of view of the rovers' mast-mounted stereo navigation cameras, in which case the rover is commanded to drive itself into areas never before seen in images sent to Earth (at least not at sufficient resolution to locate potential obstacles in advance).

The surface navigation system performs terrain hazard detection via geometric analysis of the terrain near the rover, combining multiple range data snapshots generated by the stereo vision system into a map that represents the traversability of local terrain by the rover mechanical mobility system. This local traversability map is then used by the onboard software to automatically select the best incremental path towards a waypoint or goal coordinate while avoiding terrain hazards/obstacles. The set of algorithms and software developed for interpreting these data is called the Grid-based Estimation of Surface Traversability Applied to Local Terrain (GESTALT) system [4]. Execution is enabled by onboard wheel actuation and control facilitated by proprioceptive sensor-based safeguarding and dead-reckoned position estimation. Stereo images acquired during a traverse may also be used to perform onboard visual odometry to estimate position changes based on displacement of many image features tracked between successive image captures of overlapping scenes. This is typically done when wheel odometry is expected to be highly unreliable due to non-deterministic wheel-terrain interactions and in low-traction mobility regimes.

At the time of this writing, the MER robotic vehicles had been exploring the Martian surface for approximately 1.75 Earth years. During that time period, the onboard surface mobility intelligence was used successfully to achieve a substantial portion of nearly 11 km of combined distance traversed by the rovers, each of which had traversed over 5 km. This chapter covers aspects of the autonomous mobility flight software design, algorithms, operation, and improvements associated with the rovers' exploration experiences on Mars.

2. PROBLEM STATEMENT

The MER surface operations began in January 2004 when *Spirit* and *Opportunity* drove off their spacecraft landing systems (Fig. 1) and set all wheels on Martian soil at their respective landing sites. The planned duration of their *Prime Missions*, i.e., the baseline mission for which they were designed, was 90 Martian Solar days (or *sols*). In order to fulfill the mobility related objectives of remote field geology for the MER science mission, the rovers required the capability to traverse the Martian surface safely and reliably. They had to drive to a variety of rock and soil science targets selected daily by mission scientists over many terrain types and with various densities of mobility hazards. Terrain types included hard-packed soil, soft sand, rough rock fields, and combinations of these all on slopes ranging from flat to something less than the magnitude of the rovers' maximum tilt angle of static stability, 45°. Terrain features that are considered to be navigation or mobility hazards include isolated rises/depressions of certain height/depth, extreme slopes, deep ditches, crevasses, cliffs, sand/dust-covered pits and otherwise unstable surfaces of insufficient bearing strength to support the rovers' weight.

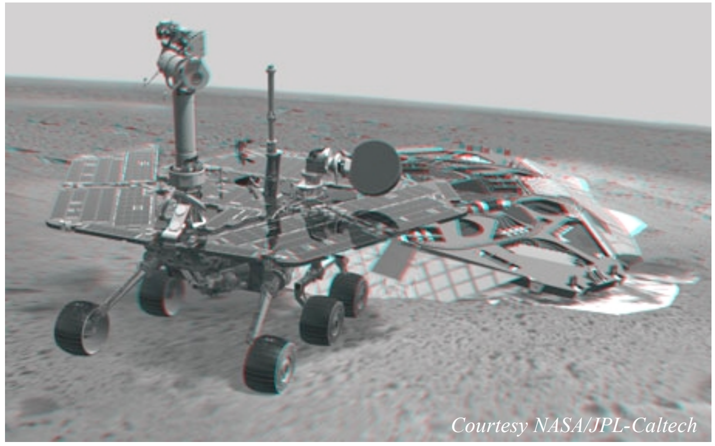

Figure 1. Spirit *and lander (computer models of hardware combined with actual Mars 3-D terrain surface data acquired by* Spirit's *cameras).*

Most of the navigation and mobility hazards are related to the rovers' kinematic constraints (e.g., rock climbing and ditch crossing capabilities, ground clearance at which body high-centering could occur, tilt stability, etc.) and can be detected by onboard proprioceptive or look-ahead sensors and software-based sensor data processing; but non-geometric hazards like insufficient terrain bearing strength require human detection and remediation.

In general, surface mobility and navigation software capabilities were required that would enable autonomous mobility, including 3-D stereo image range mapping, mobility hazard detection, local hazard avoidance path selection, and position estimation while satisfying certain performance related requirements. Specifically, *Spirit* and *Opportunity* had to be able to safely navigate at an average velocity of 35 m/hour in autonomous mode in potentially rocky terrain to designated positions on the surface while maintaining estimated position knowledge within an accuracy of 10% of integrated distance traversed (relative to starting points for traverses of <= 100 m). They were required to achieve such performance reliably in terrain types of traversability complexity on the order of that observed at the site where the NASA Viking-1 Lander spacecraft landed in 1976 (~7% rock abundance) [5]. In addition, the rovers were required to traverse a total accumulated path length of at least 600 meters (with a goal of reaching 1000 meters) over the course of their Prime Missions.

Direct teleoperation of the rovers was not a feasible strategy. Science goals demanding hundreds of meters of traverse, logistical restrictions on use of the Deep Space Network (the ground antennae on Earth used to communicate to spacecrafts), and a round-trip communication delay averaging 20 minutes make direct teleoperation impossible. As a result, the MER rovers are typically commanded only once per sol [6]. A sequence of commands sent in the morning specifies the activities for that sol: what images and data to collect, which robotic arm-mounted instruments to use and where to place them, and where to drive. Then, at the end of each sol, the rovers send back the images and data that human operators will use to plan the next sol's activities. The next mobility activities are determined based on what is known – and what is unknown – about the terrain ahead.

The capability of the MER vehicles to autonomously execute robotic activities was essential for effective command and control from Earth. They were required to navigate to surface locations on their own given specified goal coordinates, and sometimes additional guidance from human traverse planners on Earth in the form of waypoints selected with the intention to avoid certain areas of terrain. The rovers autonomously execute all such traverses commanded in a given sol without receiving further instructions from Earth until the next sol. Mobility traverses (and many other commanded activities) are typically executed within a four-hour period around local noon at the rovers' Mars locations. The need to operate unsupervised for hours at a time and manage complexities of wheeled mobility while negotiating natural terrain, combined with the infeasibility of remote teleoperation, make autonomy and intelligent motion control absolute necessities for *Spirit* and *Opportunity*. In this regard, their missions required sufficient autonomy in the form of vision-based perception and terrain modeling with supporting onboard computational intelligence to make safe navigation decisions, invoke the necessary mobility maneuvers, as well as estimate and maintain pose knowledge. This autonomous surface navigation functionality is supported by basic mobility software that performs low-level reactive fault protection based on proprioceptive sensing to achieve safeguarded mobility. All rover software, including navigation software as a subset, runs on a single 20 MHz RAD6000 computer under the VxWorks real-time operating system.

The following section provides an overview of the rovers' mobility and navigation hardware subsystems including mechanics, sensing, and associated control modes. This is followed by descriptions of autonomous navigation software functionality in ensuing sections.

3. MER MOBILITY AND NAVIGATION SUBSYSTEM

The MER vehicles can be commanded directly, or given autonomous control over multiple aspects of mobility and navigation, e.g., which mobility motions to execute, measurement of actual motion, even the selection of targets of interest (although this latter mode remained largely under-used during most of their mission). This full suite of capabilities for navigation and mobility intelligence is built on a foundation of robust mechanical capability, a range of sensing capabilities from simple analog measurements to stereo image processing, and a range of motion control capabilities from individual wheel motor control to fully autonomous navigation. An overview of this underlying foundation follows.

3.1 Mechanical Mobility Subsystem

Spirit and *Opportunity* have a very capable mechanical mobility subsystem. Six aluminum wheels are mounted on a rocker-bogie suspension (Fig. 2) that minimizes the overall chassis tilt induced by climbing over individual rocks [7]. Each wheel is 25 cm in diameter with a tread of short paddle-like cleats (aligned perpendicular to the wheel driving direction). The mechanical mobility subsystem consists of the wheel and rocker-bogie suspension assembly and is theoretically capable of climbing mobility terrain hazards (e.g., rocks, mounds, depressions, or other obstacles) that are over 35 cm tall. However, the ground clearance under the body of the rover is only 29 cm on a flat surface (and less when the rover body is tilted); so, in practice, mobility software is designed to avoid obstacles at or above a certain height. The minimum height that defines an obstacle is set by a software parameter whose nominal value is 20 cm, after taking potential vehicle tilt, rocker/bogie articulation limits, high centering, and stereo range error into account. All six wheels can be driven at rotational speeds resulting in vehicle translational speeds up to 5 cm/sec, but only the four corner wheels are steerable. The mobility system provides forward and backward straight-driving capability. Turns are accomplished using double-Ackermann steering, which enables execution of driving arcs with curvature radii as tight as 1 m, as well as turns-in-place in which the vehicle rotates about its prescribed origin (located midway between the left and right middle wheels).

Both rovers are statically stable up to 45º of tilt magnitude (pitch and roll) and have driven on hard slopes as high as 31º on rock outcrops in Endurance Crater [2] (Fig. 3) and on the Columbia Hills [1], but driving on slopes greater than 25º requires special approval by mission controllers as a means of managing mission risk. At such tilts, the weight reduction on the upslope wheels is enough that they can sometimes lose contact with and seemingly "float" above the terrain surface. The maximum tilt on loose soil is much lower. For example, *Opportunity* failed on her first attempt to exit Eagle Crater (in which she landed) because she was unable to climb straight up a soil slope of only 17º on Sol 56 after encountering 100% wheel slip for the first time, and *Spirit* encountered similar 100% slip on sandy soil slopes in the Columbia Hills.

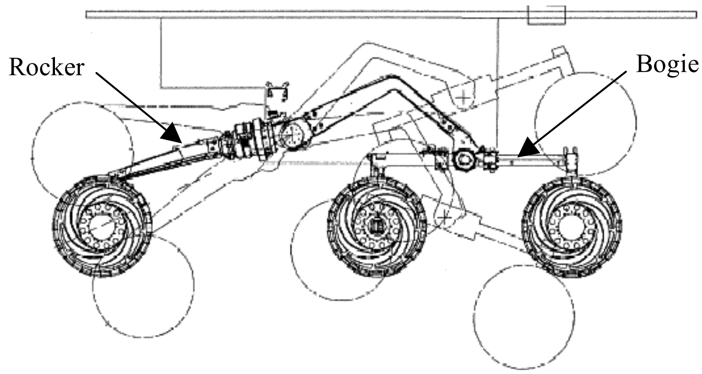

Figure 2. MER Rocker-Bogie suspension and ranges of motion.

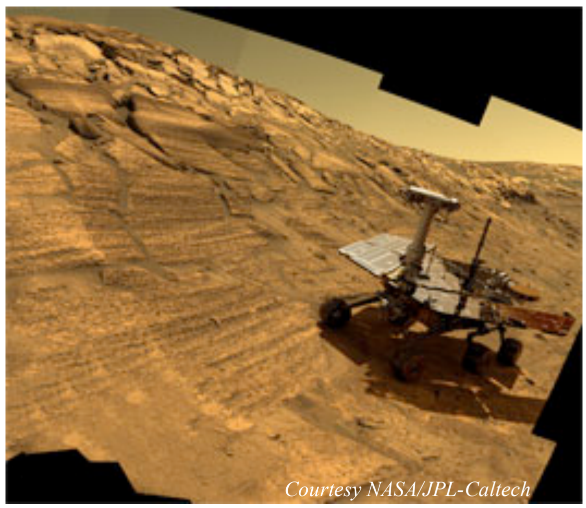

Figure 3. Opportunity *(simulated) on "Burns Cliff" in "Endurance Crater" (Image created using photorealistic rover model and approximately full-color mosaic of images acquired by* Opportunity. *Rover model size approximated based on size of rover tracks in actual mosaic).*

3.2 Sensing and Safeguarding for Mobility

Sensing for control of basic mobility functions includes wheel encoders (for dead-reckoned odometry), potentiometers for articulated suspension kinematic state, inertial attitude sensing, and celestial (sun) sensing for absolute heading determination. The rovers maintain an estimate of their local position and orientation updated at 8 Hz while driving. Position is first estimated based on how much the wheels have turned (wheel odometry). An Inertial Measurement Unit (IMU) that has 3-axis accelerometers and 3-axis angular rate sensors uses the accelerometers to measure raw, unfiltered instantaneous tilt, and the rate sensors to update vehicle attitude with high accuracy (less than 3º of drift per hour of integration).

The rover software performs multiple layers of safety checks to ensure safeguarded vehicle performance during basic mobility. These include command limits, reactive safety checks, and predictive safety checks. Examples of command limits include prescribed time limitations on the duration of motion command execution, including latest time of day by which execution of all mobility commands should terminate, as well as limits on local, repetitive motions attempted during navigation (e.g., to avoid motion limit cycles in which backups or high slip impede progress while negotiating rough terrain). Such monitoring enables the software to halt execution (and/or perform other tasks like recovery maneuvers) in the event that motion commands take longer to execute than expected due to various factors such as low-level sensor faults, impeded progress over terrain obstructions that bog the wheels down, encounters with impassable obstacles, etc. Other command limits include the use of activity constraints, which manage onboard resource use and conflicts so that certain activities or combinations of activities can only be executed when other software modules are not using common system resources, such as pointable cameras. Reactive checks include sensing of exceeded thresholds prescribed for motor currents, vehicle tilt, suspension articulation, actuator ranges of motion, etc. Reactive safeguarding enables the software to halt execution during the triggering of an unexpected event such as low-level hardware safety faults (e.g., motor stalls) and software execution errors. It also enables the software to trigger intelligent control actions in response to anticipated proprioceptive sensor state changes. The primary predictive safety checks derive from look-ahead vision-based perception of terrain hazards to the mobility system such as step hazards, tilt hazards, and roughness hazards (see Section 4). A complete list of mobility faults can be found in [4].

Sensing for navigation includes several stereo camera pairs [8]. Each rover has body-mounted, front and rear stereo camera pairs (Hazcams), each with a 120º horizontal field of view (FOV), used for local terrain hazard detection and avoidance during autonomous navigation, as well as stereo cameras for global path planning (Navcams) that are mounted on a pan/tilt platform above a fixed mast at a height of about 1.3 meters above the ground plane, with a 45º FOV. In between driving primitives during a traverse, the rover can make use of images from these cameras to perform visual odometry (see Section 5) to correct the errors in the initial wheel odometry-based estimate that are introduced when wheels lose traction on large rocks and steep slopes. This sensing of the vehicle's actual position enables many modes of vehicle safeguarding during navigation, including explicit checks for wheel slippage (based on differences between wheel and visual odometry estimates for short drive segments) and enforcement of position-referenced “keep-out zones” around predetermined obstacles or terrain hazards. In addition to the Hazcams and Navcams, the rovers have stereo Panoramic Cameras (Pancams) on their masts, which are science cameras that facilitate long-range navigation planning; they also serve as the rover sun sensor for absolute heading determination (relative to true north). The Pancams have a much narrower 16º FOV with higher resolution than Navcams, making them very useful for scanning the horizon and resolving far-field terrain features that can be accounted for in long traverse plans.

3.3 Basic Mobility and Navigation Control Modes

Vehicle motion on terrain surfaces can be commanded using multiple layers of control. Three primary basic mobility and autonomous navigation command modes are used to drive the rovers:

1. Low-level motor control commands that specify exactly how much to rotate each drive wheel and turn each steering wheel actuator
2. Directed driving primitives for driving along circular arcs (of which straight line driving and turn-in-place are special cases)
3. Autonomous path selection for reaching a goal while reacting to unexpected changes along the way.

Low-level commands enable ”non-standard” soil properties experiments such as scuffing of rocks and digging shallow trenches using a single wheel (with all other wheels stationary) as well as performing mobility mechanism health diagnostic tests. Directed drives allow human operators to specify exactly which driving primitives the rover will perform. The rovers executed directed drives with vision-based hazard avoidance *disabled*, thus directed drives are also referred to as “blind” drives. In contrast, autonomous path selection mode allows the rover to select which driving primitives to execute in order to reach Cartesian goal locations supplied by human operators in commands of a traverse sequence. Both directed and path selection modes of driving can make use of onboard stereo vision processing and terrain traversability assessment software to predict whether the rover would encounter any geometric hazards as it drives. During “guarded” directed driving, the rover’s software can preemptively “veto” a specific mobility command supplied by Earth operators if the onboard traversability assessment indicates that the perceived terrain is too risky or non-traversable. During Autonomous Navigation (AutoNav) and other path selection modes, the rover can select its own driving primitives to steer around obstacles and make progress toward its goals over low-risk and traversable terrain. This waypoint-based mode can compensate for yaw changes during blind drives, compensate for position changes during visual odometry-enabled drives, and ensure safe, obstacle-free navigation during hazard-avoiding autonomous drives.

Controlled motion is achieved via commands to onboard software functions that exploit the rover sensing and kinematics. The primary functions include three basic driving commands for translation and rotation in the plane (while the suspension system conforms to the 3-D terrain topography) as illustrated in Fig. 4. The figure illustrates capabilities of forward and reverse motion and arcing turns with a range of radii of curvature, including turns-in-place (yaw) about the vehicle center of rotation. Turns-in-place can be commanded using absolute or relative reference headings or specific Cartesian coordinates of a location toward which to face.

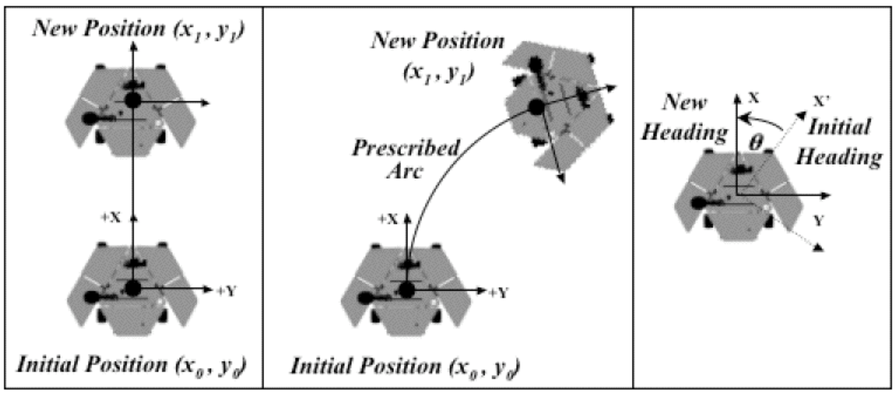

Figure 4. Basic mobility for motion along linear/curved paths and turns in place.

4. AUTONOMOUS NAVIGATION

The baseline requirements for the MER mission included the capability to drive tens of meters per Martian day *safely*. The basic mobility safeguarding capabilities described in Section 3.2 could safely stop a vehicle once it had already entered a risky situation, but to ensure the most progress a *predictive* system that could avoid even entering the risky areas was needed as well. The requirement was for a system that could conservatively evaluate the geometry of nearby terrain to avoid mobility hazards such as rocks, ditches, and areas of high slope. However, terrain stability prediction and identification of hazards not related to the observable surface geometry of the terrain were tasks that would be left to the human mission operators.

The MER onboard-software-based intelligence for autonomous surface navigation performs stereo vision-based perception, local terrain hazard mapping, traversability assessment, incremental goal-directed path selection, and vision-based pose estimation. Autonomous navigation is commanded in terms of desired destinations or explicit surface coordinates to be reached using onboard sensing, perception, and motion planning to perform hazard detection and avoidance. The following three orthogonal capabilities for autonomous driving that embody these functions are available in any combination, and are discussed in turn below: Terrain Assessment, Path Selection, and Pose Update.

4.1 Terrain Assessment

The first technology needed to implement the required navigation capability was the ability to sense the shape of the nearby terrain. MER rovers use passive stereo image processing to measure geometric information about nearby terrain. This is done by automatically matching and triangulating pixels from a pair of stereo-rectified images to generate a "cloud" of 3-D points representing the imaged terrain that will be critical for subsequent processing (see Fig. 5). Various algorithms are available for stereo image processing and are covered in numerous textbooks (see [9] for example). Stereo vision is an attractive technology for planetary exploration because it has low power requirements and nominally requires no moving parts. The stereo vision software developed at the Jet Propulsion Laboratory (JPL) has a long history [10, 11], especially as implemented on real robotic systems (Mars rover research [12-16], Unmanned Ground Vehicles [17], RedZone Robotics' Pioneer [18], Urbie [19], and PerceptOR [20, 21]).

The Mars Pathfinder rover, *Sojourner*, demonstrated the first use of autonomous stereo triangulation on a planetary rover [22]; but the *Sojourner* system relied on active projection of 5 laser stripes, and only found at most 20 XYZ (3-D) points from each pair of stereo images. In contrast, the low-power passive stereo vision used for MER relies on sunlight to illuminate the terrain. By virtue of their relatively faster processor (20 MHz compared to 0.1 MHz on *Sojourner*) and new software, the MER vehicles compute many more point measurements: *Opportunity* measured an average of 48,000 XYZ points in each of 69 Navcam image pairs as of its 322nd sol, and *Spirit* measured an average of 15,000 XYZ points in each of 1687 Hazcam image pairs as of its 342nd sol.

Simply sensing the shape of the terrain is not enough, however. The micro-properties measured by stereo vision (X,Y,Z locations of small patches of terrain) need to be processed into rover-sized macro-properties, to determine where vehicle safety might be compromised. So another technology used by the MER vehicles is the aforementioned GESTALT system for terrain assessment [4].

Originally demonstrated on the JPL Athena Rover prototype [16], GESTALT uses stereo-generated 3-D geometric data to build and maintain a grid-based local traversability map in rover memory. The map is centered on the rover position and travels with the rover while being updated with new 3-D information as new images are acquired. That is, one or more stereo pairs of images are processed into the traversability map, and merged with an existing map (generated in the same manner from previously acquired stereo images). The size of the map is configurable and the grid cell resolution is nominally 20 cm (about the size of the rover's wheel).

Inspired by the Morphin algorithm developed at Carnegie Mellon University [23, 24], GESTALT uses simple plane fits to the stereo-generated 3-D geometric information to estimate how safe the rover would be at each point in the traversability map. In this case, the plane is a disc of diameter large enough to encircle the rover shape or footprint as seen from above. This rover disc is a gross simplification of the rover body and its use is motivated by computational efficiency. For the MER vehicles, the rover disc is 2.6 meters in diameter, which is large enough to circumscribe the shape of the solar arrays (the largest dimension of the vehicles). A less conservative disc based on the smaller six-wheel footprint is also used internally by the step hazard detector, described below.

At every 20 cm map position, a set of 3-D points representing (most of) a rover disc is sampled from the stereo point cloud. The 3-D data are analyzed for tilt, roughness, and step hazards (Fig. 6). That is, each rover disc that has very high tilt, too high an overall residual (indicating that the underlying terrain is either very rough or not planar), or deviations from the best fit plane (greater than a pre-set obstacle height) cause the corresponding 20 cm by 20 cm grid cell to be marked as impassable. Less extreme deviations from a flat disc result in a continuum of terrain "goodness" evaluations of varying desirability. For example, in Fig. 7, tall rocks are assigned red "impassable" evaluations, medium rocks are assigned yellow "moderate" evaluations, and flat areas are assigned "perfect" green evaluations.

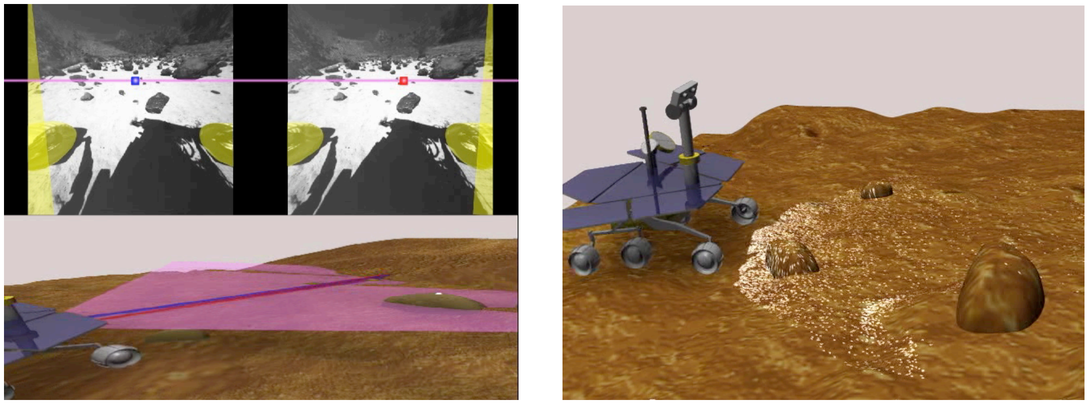

Figure 5. Stereo image processing (left) and resulting 3-D point cloud (right) representing the imaged terrain.

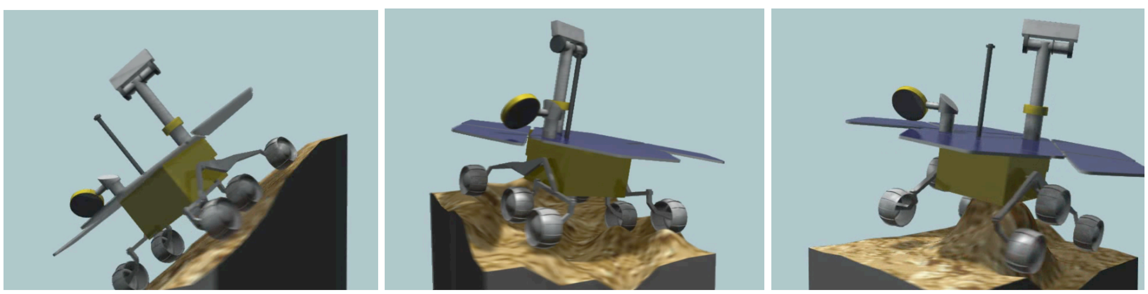

Figure 6. Mobility hazards (left to right): excessive tilt, excessive roughness, step obstacle.

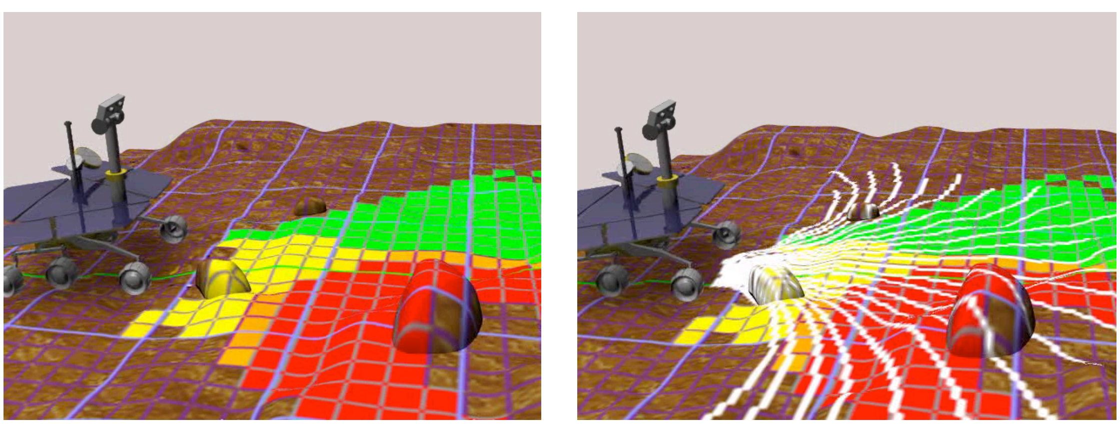

Figure 7. Traversability map: grid cell goodness evaluations (left) and possible paths (right).

For terrain assessment, *Spirit* uses the body-mounted 120º FOV Hazcams, but *Opportunity* uses the mast-mounted 45º FOV Navcams because the terrain at her landing site (Meridiani Planum) has very fine-grained soil that is not resolvable in the Hazcam images at the 256x256 pixel resolution used for Hazcam image processing by AutoNav.

4.2 Path Selection

The path selection capability gives the rover autonomous decision and control authority to select its next drive direction (in contrast to directed driving, in which it follows a single pre-commanded path). With Path Selection enabled, candidate motion paths are projected onto the traversability map (Fig. 7, right), and a weighted evaluation of the constituent grid cells is assigned to each path. This results in a set of *Obstacle* path evaluations for which low values indicate a less traversable path and high values indicate a more traversable path.

To select from among multiple paths, *Waypoint* path evaluations are assigned to all possible candidate paths according to how effectively each path would drive the rover toward its goal point. The path that would lead directly toward the goal is given the highest evaluation; other paths are assigned lesser values according to a Gaussian distribution. The variance of each distribution is configurable, nominally 3.2 curvature units for arcs, 97º for point turns. Goals that lie inside the tightest turning circle (an arc with 1 meter radius) cannot be reached using arcs, so in those situations a backup arc is selected as the preferred heading. When backing up, the rover executes a mobility behavior referred to as a "K-turn," driving in its tightest arc away from the goal location and then forward along an arc toward the goal location. As of flight software uplinked to the rovers after their Prime Missions (software releases are discussed in Section 7), when a backup behavior is chosen the highest path evaluation is halved, which in practice allows a point turn to be selected instead of an arc. This processing results in a set of *Waypoint* evaluations, which are merged with the *Obstacle* path evaluations when Terrain Assessment is enabled. Thus, using the resulting traversability map, many possible paths through the grid cell evaluations are evaluated for safety. These path-based safety evaluations, together with the goal location and current steering direction, are used to make the final selection of the next navigation step to execute. The navigation step size is configurable, but nominally 50 cm in arc length. The distance between image acquisitions is configurable to longer distances limited only by the quality and extent of stereo range data, and conservatism of flight controllers responsible for rover mission operations.

In order to navigate from one surface location to another, the rover is provided with a global path plan in the form of a series of waypoints leading to the goal location. Each waypoint is reached through incremental execution of navigation steps decided by cycles of Terrain Assessment and Path Selection. Waypoint navigation is repeated until the designated goal location is reached. This process is depicted in Fig. 8 where waypoints (and goals alike) have associated position tolerances, i.e., radial distances within which the waypoint or goal is considered reached. This capability is encapsulated in a single rover command called *Goto Waypoint*, which directs the rover to drive until the estimated position of its coordinate system origin is within a specified tolerance of its commanded goal location, or until a specified timeout period expires. Note that the rovers can execute *Goto Waypoint* equally well in either the forward direction (using front Hazcams) or reverse direction (using rear Hazcams). Fig. 9 is a portion of a mosaic of post-drive Navcam images looking backward at *Spirit's* wheel tracks and showing one of many instances wherein AutoNav had correctly avoided several large rocks while executing *Goto Waypoint*.

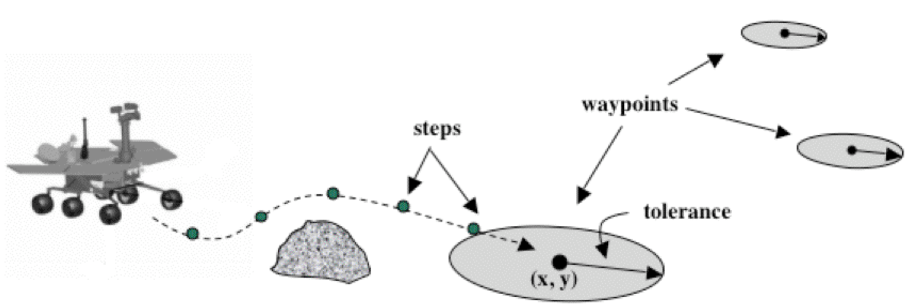

Figure 8. Autonomous waypoint navigation.

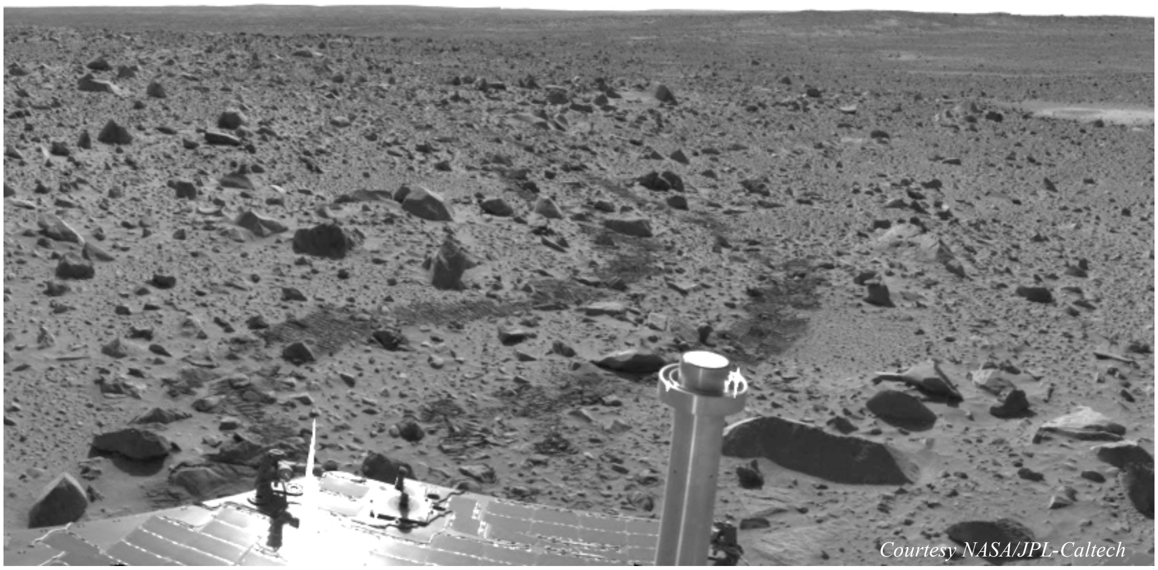

Figure 9. Spirit *wheel tracks made while avoiding obstacles on a 10º slope on sol 107.*

4.3 Pose Update

Estimates of rover position and orientation (a 6 degree-of-freedom pose) must be made by the onboard system and updated as the rover traverses in order to successfully perform autonomous navigation to specific surface locations. As mentioned in Section 3.2, dead-reckoning via wheel odometry is often used for this purpose but has severe limitations in rough and natural terrain. In relatively benign terrain, the accumulated error of wheel odometry with distance traveled may be sufficiently bounded that waypoint/goal tolerances alone are effective as a means to compensate for dead-reckoning errors. In more challenging terrain or mobility regimes, higher levels of sensing, perception, and autonomy are required to compensate for such errors.

The MER rovers make use of stereo camera images to correct and compensate for the errors in wheel odometry-based position estimates introduced, for example, when wheels lose traction on large rocks and steep slopes. The same approach can be used for orientation/attitude estimation, but so far has not been necessary due to the high quality of the IMU-derived attitude estimate.

The Pose Update capability enables the rover to update its current position and/or attitude by comparing locations of features found in stereo image pairs taken before and after a small motion step. Both rovers use Navcams for this processing, since the scale changes induced by even small motions in the wider FOV Hazcam images make autonomous tracking of features difficult. This processing only converges successfully if the imaged terrain has a sufficient number of detectable features; so the human rover driver (who plans and sequences the traverse) must ensure that Navcams are pointed toward useful features anytime the technique is used.

With the initial flight software version used in the MER surface mission, rover drivers also had to explicitly command each step of the drive to ensure that there would be sufficient overlap between successive images (nominally 60%). As of later versions, the *Goto Waypoint* command can be configured to restrict autonomously-commanded motions to ensure sufficient overlap, assuming the actual rover heading change after the motion step is no larger than what was commanded. This image processing approach and algorithm for updating rover pose are discussed in more detail in the next section.

5. VISUAL ODOMETRY

Spirit and *Opportunity* are typically commanded using precise metric specifications, for example: ”drive forward 2.34 meters, turn in place 0.3567 radians to the right, drive to location X,Y, take color pictures of the terrain at location X,Y,Z” [1, 2]. So the maintenance of the rover position estimate is of critical importance. As mentioned above, the MER vehicles can process visual cues in pre- and post-motion images to correct data derived from wheel odometry; as such, the technique is referred to as Visual Odometry. The visual odometry algorithm, or ego-motion estimation, used on MER was originally developed by Matthies [25]. Following his work, some minor variations and modifications were suggested for improving its robustness and accuracy [26, 27]. Related work is described in Chapters 4 and 7 of this volume.

Our Visual Odometry system (or VisOdom) computes an update to the 6 degree-of-freedom rover pose (x, y, z, roll, pitch, yaw) by tracking the motion of ”interesting” terrain features between two pairs of stereo images in both 2-D pixel coordinates and 3-D world coordinates. A maximum likelihood estimator is applied to the computed 3-D offsets between features in successive images to produce the final motion estimate. This motion estimate can be used to facilitate precision pointing of mast-mounted instruments or ensure accurate driving, even when the wheels slip on sloped and/or loose terrain. However, if any of the many internal consistency checks of the algorithm fails, too few feature points are detected, or the estimation procedure fails to converge, then no motion estimate update will be produced. If VisOdom does not produce an estimated motion update, the initial estimate (nominally based on wheel odometry and the IMU) will be maintained, and the operator who plans the traverse must always consider that possibility when constructing sequences.

5.1 Visual Odometry Algorithm

The basic VisOdom algorithm can be broken down into functional steps of feature detection, stereo matching of features, feature tracking, and motion estimation. The underlying mathematical formulation and additional details can be found in [28].

Feature Detection: First, features that can be easily matched between a stereo pair of images and tracked between image steps are selected. An interest operator (e.g., Forstner, Harris) is applied to one of the image pairs. Pixels with high interest values are then selected. In order to reduce the computational cost, a grid whose cell size is smaller than the minimum distance between features, is superimposed on the left image. In each grid cell, only one feature with the strongest corner response is selected as a feature candidate. Then all candidate features are sorted into descending order and features are selected in the order of the feature list. A constraint of minimum distance between features is enforced to ensure features are selected evenly across the image scene.

Feature Stereo Matching: The 3-D positions of selected features are estimated by stereo matching. Because the stereo cameras are calibrated, the stereo matching is only done along the epipolar line with a few pixels of offset buffer above and below it. Pseudo-normalized correlation is used to determine the best match. In order to obtain sub-pixel accuracy, a biquadratic polynomial is fitted to a 3x3 neighborhood of correlation scores and the peak of this biquadratic polynomial is chosen as the correlation peak.

Under perfect conditions, the rays of the same feature from left and right image should intersect in space. However, due to image noise and matching error, they do not always intersect. The gap (shortest distance between the two rays) indicates the quality of the stereo matching: features with large gaps are eliminated from further processing. In addition, the gap is a factor of the error model, which is incorporated into the covariance matrix computation. Details of this mathematical formulation can be found in [28].

Feature Tracking: After the rover moves a certain distance, a second pair of stereo images is acquired. The features selected from the previous image can be projected into the second pair using the prior knowledge of the approximate motion provided by the onboard wheel odometry. Then a correlation-based search determines 2-D positions precisely in the second image pair.

Stereo matching is then performed in these tracked features on the second pair to determine their new 3-D positions. Because the 3-D positions of those tracked features are already known from previous step, their stereo search range can be greatly reduced.

Robust Motion Estimation: If the initial motion is accurate, the difference between two estimated 3-D positions should be within a prescribed error ellipse. However, when the initial motion is off, the difference between the two estimated positions reflects the error of the initial motion and it can be used to determine the change of rover position. The motion estimation is done in two steps. First, a less accurate motion is estimated by least-squares estimation, for which a closed-form solution exists [25]. The advantage of this least-squares method is that it is simple, fast and robust. The disadvantage of it is that it is less accurate because it only takes the quality (the volume of the error ellipsoid) of the observations as a weight factor. Because it is an inexpensive operation, we use it and a random sampling method (RANSAC) to do outlier removal as follows:

1. A small set of features (e.g., 6) is randomly selected and the motion is then estimated.
2. All features from the previous step are projected to the current image frame by the newly estimated motion. If the gap between a re-projected feature and its correspondent is less than a threshold (e.g., 0.5), the score of this iteration will be incremented.

We repeat steps 1 and 2 for a maximum number of iterations or until the residual error is small enough. Finally, we pick the motion with the highest score and all features that passed this iteration will be used in the following, more accurate estimation — the maximum likelihood motion estimation.

The maximum likelihood motion estimation takes account of the 3-D position difference and associated error models to estimate position. Two nice properties of the maximum-likelihood estimation make the algorithm powerful. First, it estimates the 3 axis rotations directly so that it eliminates the error caused by rotation matrix estimation (made, for example, by the least-squares estimation). Secondly, it fully incorporates error models (the shape of the ellipsoid) in the estimation, which greatly improves the accuracy.

5.2 Benefits and Examples

For the MER mission, several benefits were realized with Visual Odometry. Vehicle safety was maintained by having the rover terminate a planned drive early if safety would have been compromised by continuing the drive, for example, if it realized via Visual Odometry that it was making insufficient progress toward its goal (Slip Check), or was nearing the pre-specified location of an obstacle (Keep Out Zone). Improvements in accuracy of traverses planned in new or mixed-soil terrains were facilitated by re-pointing to the drive goal or re-computing the distance remaining to the goal after each step along the way. The improved drive accuracy also yielded a greater number of science observations by reducing the number of sols needed to accurately position the rover within instrument-arm's reach of science targets. In addition, use of remote sensing instruments on the rover's mast (Pancam and Miniature Thermal Emission Spectrometer) required precision pointing of the mast (via pan and tilt actuators) to make particular science observations, which were often scheduled in the middle of a traverse. At these times, use of Visual Odometry obviated the need for human confirmation of the pointing angle to ensure accurate pointing at the intended science target. Since VisOdom maintained accurate pose estimates as the rover traversed, any pointing at absolute coordinate locations specified before the drive remained equally accurate.

To provide a sense for the positive effect of using Visual Odometry relative to wheel odometry alone, consider the following example. Over the course of short drives from sol 188 to sol 191, *Opportunity* drove up-slope and cross-slope backward on a 15-19º slope for 8.1 meters, followed by a 9.4 meter forward drive down-slope. Fig. 10 shows a comparative view of the end positions on sol 191 according to the wheel odometry estimate and the VisOdom estimate for which the position offset (wheel odometry error) is 5 meters (over 28% of the distance traversed). Fig. 10 is a screenshot from the Rover Sequencing and Visualization Program (RSVP), which is the primary MER ground data software tool for creating, visualizing, and validating rover command sequences before uplink to Mars [29].

RSVP is also used to visualize playbacks of rover kinematic state histories, returned to Earth in rover telemetry, thereby revealing the actual (estimated) motions executed on Mars [30]. Shown in the top left of the screenshot is the animated rover and a trailing line representing the path of its body-fixed coordinate frame based on wheel odometry; the same is shown in the bottom of the screenshot based on vastly improved VisOdom position estimates. The rover graphical icons and paths are superimposed atop a synthetic terrain mesh generated from a mosaic of actual stereo images acquired by *Opportunity* on prior sols.

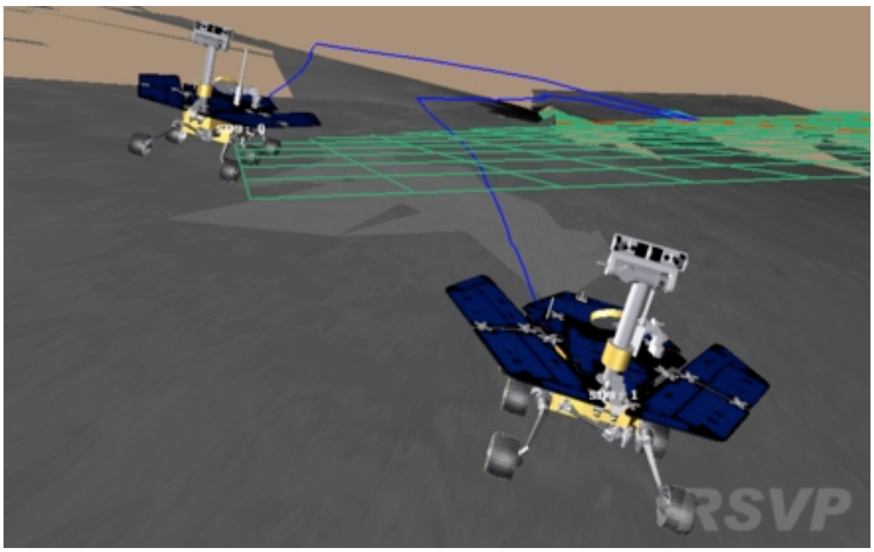

Figure 10. Opportunity *rover position estimate comparison after 17.5 m traverse on slopes (sol 188-191): wheel odometry (top left); visual odometry (bottom right).*

On higher slopes and in a variety of terrain conditions, VisOdom enabled precise approaches to designated science targets despite reduced wheel traction or low soil trafficability. The following is one example of a precise approach to a difficult science target executed by *Opportunity* on slopes slightly higher than the slopes in the previous example. On *Opportunity's* Sol 304, a drive of over 8 meters was planned on an outcrop whose slope varied from 20º to 24º. Because the drive plan took a wide range of potential wheel slip percentages into account, *Opportunity* was able to drive just far enough across slope, then turn and drive just far enough up-slope, to perfectly position the desired science target within its instrument arm work volume in a single sol.

The left image of Fig. 11 is an RSVP screenshot depicting the planned drive, with the rover graphical icon shown at *Opportunity's* initial position. The planned drive was a backward traverse (toward the right in the figure) followed by a turn-in-place to face the science target and a very short drive forward. The planned path is depicted by the lines superimposed on the synthetic terrain, which represent the rover position and wheel tracks.

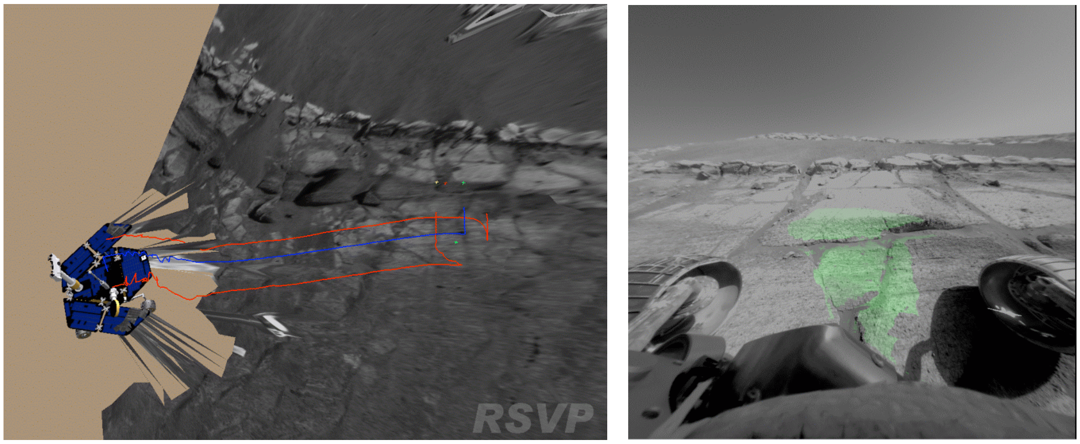

Figure 11. Precise science target approach using visual odometry: initial position and planned path (left); front Hazcam image and target reachability at final position (right).

On the right of Fig. 11 is an actual image acquired by the body-mounted front Hazcams from the position reached at the end of the drive. It shows the science target area perfectly located between *Opportunity's* front wheels. The faint green areas superimposed on the image represent the many 3-D points that are reachable by science instruments on the rover's instrument arm (still stowed in the image).

Visual Odometry has been a highly effective tool for maintaining vehicle safety while driving near obstacles on slopes, achieving difficult drive approaches to science targets in fewer sols, and ensuring accurate science imaging. Although it requires active pointing planned by human drivers in feature-poor terrain, the improved position knowledge enables more autonomous capability and better science return during planetary operations.

6. TRAVERSE PLANNING AND EXECUTION

The MER rovers are typically commanded once per Martian day. A sequence of commands sent in the morning (at the rover's location on Mars) specifies the day's activities. Surface mobility plans designed to reach areas of scientific interest or specific science targets can include one or more long traverses (tens of meters), short approaches (less than 10 meters), and fine-positioning maneuvers (within 2 meters). These three traverse types are central to the MER traverse planning approach.

Command sequences for traverses on the Martian surface by *Spirit* and *Opportunity* use the following types of mobility commands: (a) basic mobility commands alone (*manual* driving without hazard avoidance enabled); (b) basic mobility commands for directed driving with *guarded* execution (hazard detection enabled to build local traversability map, but hazard avoidance disabled); (c) fully *autonomous* navigation (hazard detection and avoidance enabled with autonomous path selection and execution); (d) visual odometry to provide the best position estimates. The selected mobility and navigation approach for a given traverse plan is determined based on what is deemed most appropriate given combined human and rover perception of the terrain and risks perceived by engineers and mission managers.

6.1 Semi-Autonomous Operations: Humans in the loop

To achieve effective semi-autonomous missions at remote sites on Mars, the onboard robotics software functionality is complemented by human functionality at a local operations and command facility on Earth. Mobility Engineers, Rover Planners, and the rovers form a closed-loop human-robot control system (notwithstanding the NASA Deep Space Network of ground antennae and supporting teams and systems beyond the scope of this chapter). Humans collaborate with the rovers to achieve best performance of onboard mobility and robotic arm software as it affects actual robotic motions and execution of mobility and instrument placement command sequences [31]. Mobility Engineers effectively function in the feedback loop of the human-robot system (Fig. 12) as human observers of mobility and robotic arm kinematic state as well as maintainers of best-known state knowledge for delivery to the uplink planning team. Rover Planner functions are manifested in the feed-forward loop and can be thought of as providing reference inputs and serving as compensators for the rover system given input from Mobility Engineers.

Within this closed-loop human-robot system the science instrument, image, and engineering data telemetered to Earth on a given sol, from either rover, determine its exploration plan for the next sol. Typically, the next sol's planning cannot begin without certain critical data and necessary knowledge representing the last known state of the rover at the termination of the previous sol's activity.

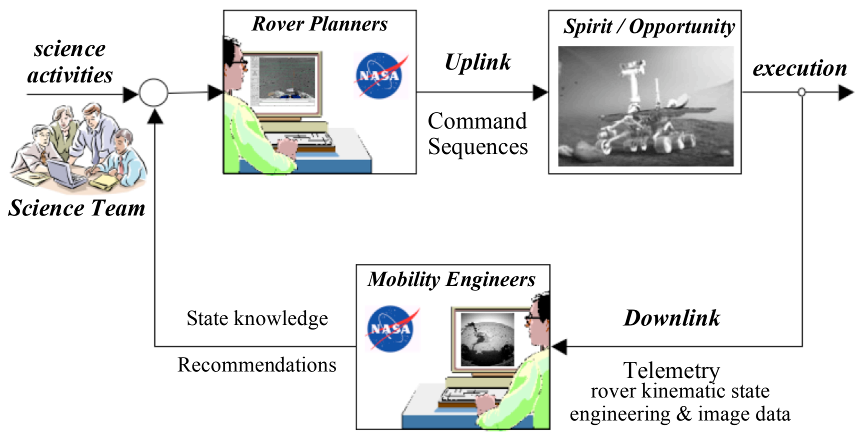

Figure 12. Simplified human-robot control system for remote surface mobility operations.

Mobility Engineers determine the best-known state of the rover and deliver that knowledge to Rover Planners on the uplink team. With significant direction from the MER body of scientists, the uplink team plans and sequences the agreed upon activities for the next sol. The mobility and robotic arm planning process proceeds with generation of rover motion command sequences that will carry out the intended activity. High-level (autonomy) and low-level motion commands are refined by Rover Planners using their perception of the rover surroundings and knowledge of rover behavior [1, 2]. This is facilitated by analyses performed by Mobility Engineers that result in engineering recommendations for making the best use of the rover functionality. This collaborative loop of human and rover functionality serves to facilitate proper autonomous execution of the sol's command load on Mars. Nominally, each rover is sent a command load once per day and executes uplinked sequences throughout a period of 3-6 hours around local noon (with occasional nighttime communications or science activity). In this manner, human-guided robotic execution leads to exploration progress, which generates new data and images that feedback into the cyclic process, ultimately leading to scientific discovery.

6.2 Adjustable Traverse Autonomy

Adjustable autonomy refers to properties of an autonomous system that enable different levels of system operation between manual and autonomous. The level of operational autonomy may be adjusted by human operators, other systems, or the system itself [32]. The MER surface operations approach achieves a form of adjustable traverse autonomy via mixed modes of sequencing the rover manual and autonomous drive commands as follows.

MER mobility operations are encoded as event-driven sequences of individual motion commands. The various drive command types were discussed earlier and include directed (blind) driving, guarded motion, blind Goto Waypoint, visual odometry and fully autonomous navigation. Combinations of these command types avail the human Rover Planners with a variety of *driving modes* with which to plan different traverses executed using different degrees of manual and autonomous driving.

The planning techniques used for each long traverse, short approach, or fine positioning drive are determined based on the time allocated for driving, the amount of terrain visible to operators in imagery received on Earth, known hazards, and level of uncertainty in rover position given the terrain type. Depending on the situation and feasible tradeoffs, one of several strategies is followed for selecting between human-planned directed drives versus rover-adaptive AutoNav and VisOdom [3]. Generally, driving on level terrain required a mix of blind and AutoNav commands that achieved many tens of meters of traverse, and driving on slopes required using VisOdom to allow the rover to compensate for expected wheel slippage while driving only a few tens of meters each sol.

There are significant differences in resource usage between manual and autonomous driving, with execution time being the most obvious. In general, the time allocated for driving (versus other exploration activities), is limited by combined thermal, power, communications, and science activity constraints. Power is also impacted by execution time, for although the power used by the mobility system is the same whether a path was generated manually or autonomously, the rover's CPU and electronics draw power for the duration of the drive and thus an autonomous drive requires more power than a manual drive of the same distance.

The computing resources required by the different command types also 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 performs only *reactive* motion safety checks, never looking ahead to evaluate the terrain before driving onto it. AutoNav commands use onboard image processing to 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; AutoNav commands also rely on the accuracy of the wheel odometry to track obstacles once they leave the FOV of the cameras. VisOdom commands execute slowly while processing images more frequently to provide the best position estimates (but not obstacle detection), and require close spacing between images, which limits the top speed to 10 meters/hour. Contributing factors to the relatively slow execution times of the AutoNav and VisOdom subsystems are the limited computation speed of the 20 MHz RAD6000 and the fact that dozens of tasks in the real-time system share a single address space and cache.

The marked difference between regional terrain types at the Gusev Crater and Meridiani Planum landing sites resulted in different experiences for *Spirit* and *Opportunity* Rover Planners, respectively [1, 2]; but even within each regional site the human and rover driving strategies alike had to adapt to new local terrains many times over the course of each mission. This is evident in Fig. 13, which summarizes the distances covered by each rover in a single sol using the various driving modes during their first 19 months of operation. In Fig. 13, 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 adaptive driving modes but not AutoNav, and Blind drives (in red) include both directed arcs and rover-adapted drives that compensated for yaw changes measured during the drive.

The data displayed in Fig. 13 is summarized numerically in Table 1 where informal names given to each driving mode are listed in the left column. All modes except Directed Driving indicate some degree of autonomous driving and indicate whether onboard terrain assessment, path selection, visual odometry, or all of these autonomy capabilities were enabled. The final columns indicate how much each driving mode was used, and traverse distances achieved, on each vehicle as of *Spirit's* sol 573 and *Opportunity's* sol 555.

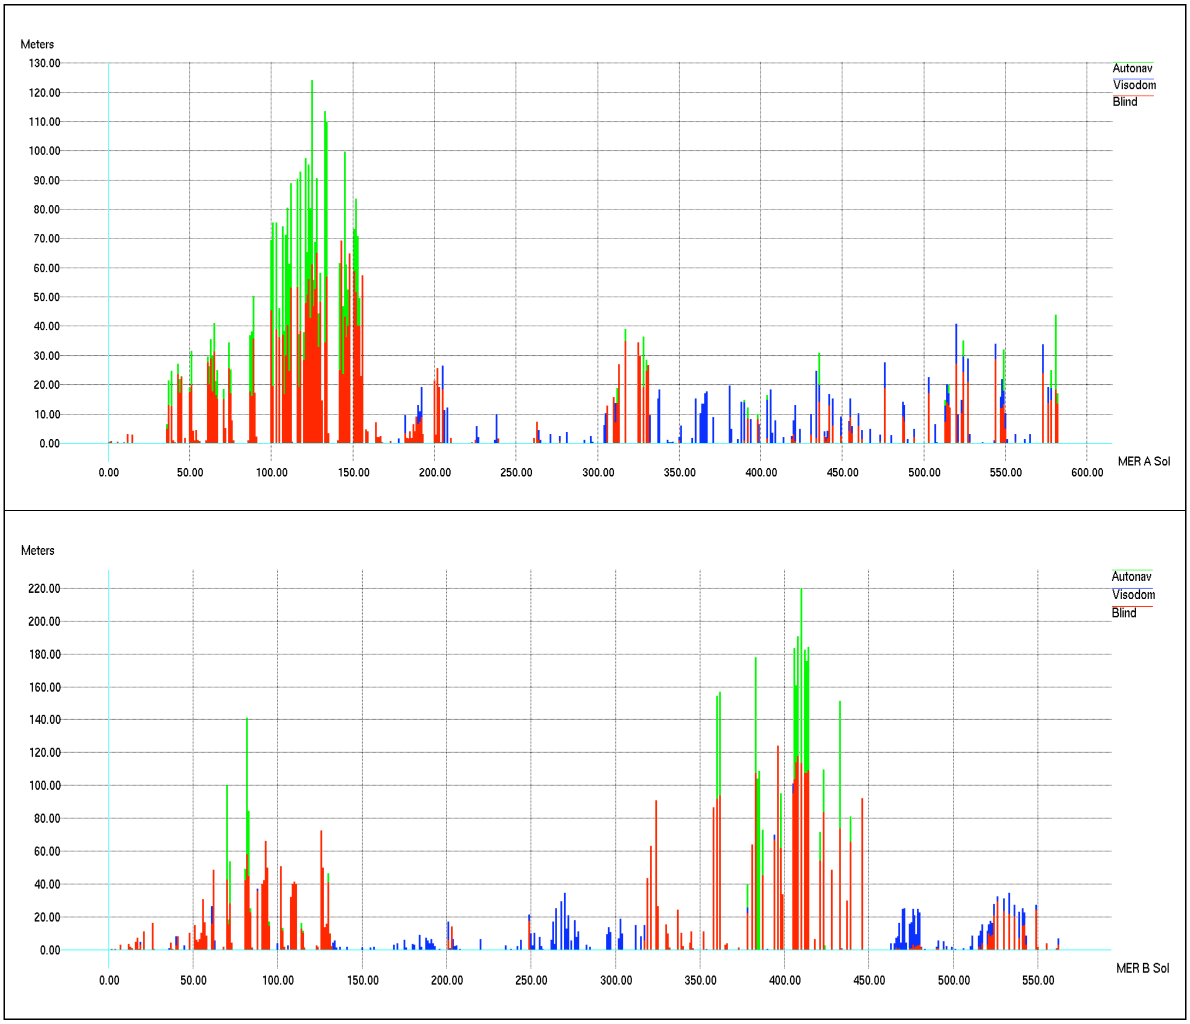

Figure 13. Summary of distances driven by Spirit *(top) and* Opportunity *(bottom) per sol.*

Driving Mode	Terrain Assessment	Path Selection	Visual Odometry	*Spirit*		*Opportunity*	
Directed Driving	no	no	no	451 m	9%	1973 m	33%
VisOdom	no	no	YES	410 m	8%	561 m	9%
Blind GotoWaypoint	no	YES	no	2196 m	46%	1911 m	32%
VisOdom GotoWaypoint	no	YES	YES	379 m	7%	121 m	2%
Guarded Motion	YES	no	no	36 m	1%	117 m	1%
Guarded VisOdom	YES	no	YES	0 m	0%	0 m	0%
AutoNav	YES	YES	no	1315 m	27%	1262 m	21%
AutoNav with VisOdom	YES	YES	YES	3 m	0%	0 m	0%
				4798 m	100%	5947 m	100%

Table 1. MER driving mode usage as of 15 August 2005, counting 573 sols for *Spirit* and 555 sols for *Opportunity*. (Distances are as measured onboard, and can be overestimates of actual distance traveled if wheels slip during non-VisOdom drives. Only rover translation distances are reflected here; turn-in-place and single-wheel trench digging motions are not reported.)

6.3 Shared Navigation Support

Thorough planning and safe execution of traverses using any of the driving modes listed in Table 1 is facilitated by collaborative navigation support provided by human operators and the rovers. In particular, aspects of terrain assessment and traversability analysis are shared, with tasks allocated appropriately based on their respective capabilities for performing these navigation support functions as follows. 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 to the human rover operators. Imagery acquired by Mars-orbiting spacecraft, while crucial for long-range traverse planning, currently cannot resolve mobility hazards like 20 cm rocks. So after each long drive, images from each appropriate camera pair are requested. The post-drive imagery provides operators with situational awareness and a visual sense of the new terrain about to be visited by the rover.

All stereo image pairs received on Earth are processed by an automated pipeline that generates a variety of derived products including 3-D range maps, texture-mapped terrain meshes, and color overlays indicating terrain properties such as slope and elevation [33]. Rover operators use image-based querying tools to measure ranges to terrain features and estimate distances and rock sizes [34]. For example, a "ruler" tool allows an operator to measure the distance between the 3-D 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 operators 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. 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 wheel traction or sinkage. For example, while driving uphill toward a topographic high point dubbed "Larry's Lookout," *Spirit* reached 100% slip (i.e., achieved no progress in the drive direction while spinning her wheels) on a patch of terrain with 16º slope, but encountered only 20% slip on terrain of 19º slope that was only a few meters away and had no discernible difference in appearance from the previous terrain patch.

When provided with quantitative 2-D and 3-D image analysis tools, humans are very good at terrain analysis for motion planning. In addition to geometric hazards such as rocks or sudden elevation drop-offs, humans can readily identify and classify new terrain types (e.g., sandy versus rocky slopes) on the basis of appearance in images alone. This, along with consideration of rover mobility system capabilities and/or operational constraints, allows operators to plan a safe series of waypoints. The MER software does not have any appearance-based terrain analysis capabilities, thus onboard hazard detection is limited to geometric obstacles. Research is underway to develop appearance-based techniques that may be useful for future missions [35-38]. Nevertheless, the most serious and frequent obstacles (rocks, steps, and high-center hazards) can be detected by geometric analysis — assuming 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 normally becomes available when driving autonomously. As such, the rover operators use a combination of their appearance-based and geometric hazard identification skills to supplement the rovers' capabilities for geometric hazard detection. While the rovers are better able to assess nearby hazards, the lack of a global planner (which the human stands in for during manual drives) can cause the rovers to get stuck in cul de sacs and dead-ends.

The autonomous navigation software enables *Spirit* and *Opportunity* to drive safely even through areas never before seen in images sent to Earth, that is, to surface locations beyond the extent of surface imagery discernable by human rover operators. As such, rover operators could, and occasionally did, command autonomous traverses to goal waypoints (in images acquired prior to the traverse) beyond the reliable field of view of the rovers' mast-mounted stereo cameras. In such cases, the human has no ability to select safe waypoints and the rover must drive into terrain that has not been imaged prior to planning the traverse. One notable instance of this was on *Spirit's* sol 109 when she was commanded to drive over the local horizon (50 m distant) as she descended from the rim of a crater dubbed "Missoula Crater." In this case, AutoNav was the only option available to drive further and use the available time and power. Post-drive images of *Spirit's* wheel tracks revealed that AutoNav correctly avoided large rocks while traversing slopes up to 10º. 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 safety occur when the rover drives beyond the distance at which obstacles can be resolved, outside the imagery available during traverse planning, or through smaller occluded regions. In practice, and 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.

In the approach described above, humans apply their perception and navigation skill to supplement the rovers' capabilities and limited onboard intelligence to accomplish mobility objectives of the surface mission. This collaboration is enabled through the use of sophisticated ground-based software tools that help to bridge the gap between rover and human perceptions.

7. DISCUSSION

Of the nearly 11 km of combined distance traversed on Mars by *Spirit* and *Opportunity* by the time of this writing, more than 2500 meters was driven while autonomously detecting and/or avoiding obstacles, and over 2500 images were processed for Visual Odometry position updates. In this Section, we highlight and briefly discuss some of the key challenges, notable practicalities, and particular difficulties encountered throughout the experience. We also point out some of the software enhancements or improvements that were motivated by our experiences and uplinked as software upgrades to the rovers at different times over the course of the mission.

In the first images returned by *Opportunity* shortly after landing in the small crater dubbed Eagle Crater, scientists were thrilled to see bedrock outcrops a mere 9 meters from the rover. Bedrock outcrops are of very high scientific value and this was the first time they had been imaged by a landed spacecraft on Mars. However, mobility engineers were horrified to see very little 3-D information recovered from those first stereo images. The fine-grained soil particles covering the majority of Eagle crater were unresolvable by stereo image processing of the one-bit-per-pixel compressed 1024x1024, 45º FOV Navcam images. Fortunately, Navcam images with more bits per pixel were acquired quickly and proved easy to process. But at the nominal size used by onboard autonomy software (256x256 squared pixels), even eight bits per pixel were insufficient to resolve the soil particles in the 120º FOV Hazcam images. This meant that the autonomous navigation system, tested almost entirely using rigidly mounted Hazcams, had to be reconfigured to use the pointable mast-mounted Navcams instead. The Navcams were successfully used to better process images of the smooth and nearly featureless terrain at *Opportunity's* landing site. Even though their FOV is significantly narrower than that of the Hazcams, the navigation algorithm was still able to perform its intended functions. This was a testament to the flexibility of the navigation software design.

While *Spirit* and *Opportunity* were designed to traverse in Viking Lander-I terrain types (mostly flat terrain, with many small non-obstacle rocks and occasional obstacles), they both encountered and successfully traversed terrain of significantly higher variability. During actual operations, the MER vehicles were driven over sloped terrain at tilts as high as 30º, and over soil textures comprised of slippery sandy material, hard-packed rocky material, and mixtures of the two. Although models of wheel slip in sandy terrain, derived from Earth-based testing of MER rover engineering models [39], correlated remarkably well with some Meridiani Planum terrains, slip behavior was extremely difficult to predict in areas where the rover was driven over non-homogeneous terrains (e.g., when climbing over rock for one part of the drive and loose soil for another). The uncertainty in the amount of slip resulting from drives on high slopes or loose soils initially forced the operations team to spend several days at a time driving toward some questionable targets, even some that were nearby.

As the rovers continued to explore through extended mission durations well beyond their design lifetime, only few signs of mechanical wear and tear of the mobility subsystem became apparent. Most notable are problems that arose with each rover's right front wheel actuators (coincidentally). By *Spirit's* sol 154, the drive motor for her right front wheel had begun to draw over twice as much current as the other wheels [1]. A period of diagnostic test activities led to the tentative conclusion that the actuator was near the end of its usable life due to its high current draw. This led to the use of a 5-wheel driving technique devised by the operations team and designed to minimize the use of the right front wheel. The technique required *Spirit* to drive backward with the right front wheel disabled and dragging (since the high gear reduction of its actuator prevented the wheel from rolling freely). After months of continued progress in this condition, attempts of short 6-wheel drives indicated that the right front wheel drive motor current draw had returned to normal. The best explanation continues to be that infrequent driving and diurnal temperature cycles allowed viscous lubricant to redistribute itself throughout the wheel's harmonic drive. On *Opportunity's* sol 433, her right front steering actuator stalled and after a period of diagnostic tests was declared failed. The actuator failed with the right front wheel oriented at a nearly forward steering angle (left of straight ahead by less than 10º). *Opportunity* continued to explore with operations restrictions on use of her right front steering actuator.

Although visual odometry processing could have been beneficial during all rover motion, each step required nearly three minutes of processing time on the 20 MHz RAD6000 CPU, and thus it was only commanded during relatively short drives that occurred either on slopes (typically more than 10º), or in situations where a wheel was being dragged, (e.g., when digging a trench, or when conserving drive motor lifetime on *Spirit's* right front wheel during a period when it drew significantly higher currents than the other five wheel drive motors). The onboard IMU exhibited a very small drift rate (usually less than 3º per hour of operation) and therefore maintained attitude knowledge very well: so from January through September 2005, for example, visual odometry was typically used to update rover position only. There were some instances in which visual odometry did not converge to a solution. These are primarily attributable to either too large a motion (e.g., commanding a 40º turn-in-place which resulted in too little image overlap between successive images) or lack of features in the imaged terrain. Wheel slips as high as 125% on *Spirit's* sol 206, when she tried to drive up a more than 25º slope, were successfully measured by visual odometry. Overall, visual odometry software generated over 1400 combined successful position updates on both rovers as of March 2005, or 15 months of surface operations.

There have been three versions of the mobility flight software used on Mars during MER surface operations. The version used during the rovers' Prime Missions had autonomy software that was overly conservative. It required that terrain processing occur after each autonomously-commanded motion, and also checked for 20 cm step differences everywhere within each 2.6m-diameter rover disc used for onboard terrain assessment. In practice, this meant that the rovers would perceive small mounds (similar to baseball pitcher's mounds) as non-traversable obstacles and would not attempt to autonomously drive over them; this occurred several times on *Spirit*.

In early April 2004, a new version of software was uplinked to both rovers, initiating their first Extended Missions. This version incorporated several robustness enhancements made to the autonomy software during pre-landing outdoor field tests on Earth, as well as some lessons learned during the Prime Mission. The most dramatic changes included the ability to skip terrain assessment when the existing data were sufficient to ensure that all paths are safe, and the ability to perceive small mounds as traversable and autonomously traverse them.

Another software upgrade was made after December 2004 to help streamline commanding and data analysis, improve position estimation in highly sloped terrain, and enable terrain assessment using more than only two mast-mounted Navcam images. This upgrade also included the first changes made to the low-level mobility software since 2002: use of look-up tables to generalize the steering angles used by primitive commands, use of look-up tables to estimate slip based on rover attitude, and the ability to alternate between wheel dragging and non-dragging modes in the event of a disabled drive motor. These capabilities were added in response to problems encountered during the Prime and Extended missions, but as of September 2005 remained unused because alternative solutions or workarounds were being used effectively. As examples: the effect of *Opportunity's* stuck right front wheel steering motor was minimized simply by avoiding point turns; VisOdom was used for precise (not approximate) slip measurements; and since *Spirit's* right front wheel drive motor current draw had returned to normal, it no longer needs to be dragged. Of course these capabilities may well prove useful as the rovers enter their second Martian year of operations in 2006.

8. CONCLUSIONS

The Mars Exploration Rovers autonomous navigation system running the GESTALT algorithm has kept both *Spirit* and *Opportunity* safe through over 2500 meters of autonomous driving in the first 21 months of surface mission operations. Successful operation of the 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 leads to better overall performance in the face of limited time, power, imagery, and onboard computation. Humans have enormous perceptual abilities and experiences that are useful for navigation and can adapt their application to new terrain types and challenges. The rovers have the advantage of being able to close the loop on execution errors, and can assess terrain that is not visible in the imagery available to the human when planning a traverse. This semi-autonomous human-robot system ensures safe navigation in near-field and far-field terrain visible in planning imagery as well as over the local horizon and beyond. The nature of the algorithm (e.g., modeling the rover as a disc) makes it an extremely conservative system, however. Very often, human traverse planners are able to safely command the vehicle through areas that the onboard system declares unsafe. Future missions would benefit greatly from terrain assessment techniques that could be used not only during long drives on relatively flat terrain, but also during approaches to science targets in areas with a greater density of obstacles and to targets on slopes.

Visual odometry has been a highly effective tool for maintaining vehicle safety while driving near obstacles on slopes, achieving difficult drive approaches in fewer sols, and ensuring accurate science imaging. Although it requires active pointing by human traverse planners in feature-poor terrain, the improved position knowledge enables more autonomous capability and better science return during planetary operations. For the MER mission, the slow execution time of visual odometry limited its feasible use to relatively short drives (less than 30 meters) in high-risk areas, but future missions would benefit from running it all the time. Position uncertainties resulting from drives in slippery terrain present difficulties for human planners attempting to schedule activities using improperly-aligned images from many locations; running visual odometry all the time could ease ground operations significantly in this respect.

Throughout the first 21 months of operation, MER vehicles were typically driven relatively short distances between consecutive uplinks of command sequences (from an average of tens of meters to a maximum of 370 m). If future missions require the ability to navigate further (well beyond the edge of what is visible from a starting point), a path planner should be used that is better able to deal with extended obstacles and large-scale terrain hazards like cliffs, long ridges, and cul de sacs. As part of its next planned flight software update in 2006, MER is considering incorporating the Field D* algorithm [40, 41], which can address these issues. But a planner that could take even more of the overall system constraints into account would be needed to enable truly autonomous long-term activities. The key overall system constraints include available power, optimal communication attitude and position, opportunities for additional science measurements, and onboard data volume.

Additional plans for the 2006 flight software update include capabilities to: autonomously update a drive goal position onboard by visually tracking it throughout the approach; autonomously evaluate the safety of deploying the rover instrument-arm at the end of a drive, without requiring human confirmation, and proceed to deploy the arm when safe before any subsequent Earth command cycle; and prioritize transmission of data based on its content (e.g., only log those images known to have captured clouds or dust devils).

The MER experience and capabilities represent the state-of-the-art for planetary surface navigation. The autonomous navigation system is being evaluated for potential use as a baseline capability for the planned NASA Mars Science Laboratory mission as well [42]. The long lasting exploration missions performed with *Spirit* and *Opportunity* have enabled us to learn a great deal about remote robotic surface operations for long duration missions. The experience and knowledge gained will benefit Mars Science Laboratory, the ExoMars mission planned by ESA, as well as lunar and planetary robotics missions beyond. The enormous amount of engineering data generated by the missions is unprecedented and is as relevant to Earth-based field robotics as it is to planetary robotics. With the definition of suitable functional performance metrics applicable to performance data for MER and other rovers, systematic assessments can be made of future flight rover technologies relative to MER as a state-of-the-art baseline [43].

ACKNOWLEDGEMENTS

The research described in this chapter was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Many thanks to the entire MER development and operations teams, especially the Mobility and Navigation test team, for without their tireless efforts the amount of onboard autonomous capability would have been much more limited.

REFERENCES

[1] C. Leger, A. Trebi-Ollennu, J. Wright, et al, "Mars Exploration Rover surface operations: Driving Spirit at Gusev Crater," Proc. IEEE Intl. Conference on Systems, Man, and Cybernetics, Waikoloa, HI, October 2005, pp. 1815-1822.

[2] J. Biesiadecki, E. Baumgartner, R. Bonitz, et al, "Mars Exploration Rover surface operations: Driving Opportunity at Meridiani Planum," Proc. IEEE Intl. Conference on Systems, Man, and Cybernetics, Waikoloa, HI, October 2005, pp. 1823-1830.

[3] J.J. Biesiadecki, C. Leger and M.W. Maimone, "Tradeoffs between directed and autonomous driving on the Mars Exploration Rovers," Proc. 12th Intl. Symp. on Robotics Research, San Francisco, CA, Oct. 2005.

[4] J.J. Biesiadecki and M.W. Maimone, "The Mars Exploration Rover surface mobility flight software: Driving ambition," Proc. IEEE Aerospace Conf., Big Sky, MT, 2006.

[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*, Vol. 102, No. E2, 25 Feb. 1997, pp. 4117-4129.

[6] A.H. Mishkin, D. Limonadi, S.L. Laubach, D.S. Bass, "Working the Martian night shift: The MER surface operations process," *IEEE Robotics and Automation Magazine*, in press.

[7] B. Harrington and C. Voorhees, "The challenges of designing the rocker-bogie suspension for the Mars exploration rover," Proc. 37th Aerospace Mechanisms Symp., Galveston, Texas, USA, May 2004.

[8] J.N. Maki, J.F. Bell III, K.E. Herkenhoff, S.W. Squyres, et al, "The Mars Exploration Rover engineering cameras," *Journal of Geophysical Research*, Vol. 108, No. E12, 2003, pp. 12-1–12-24.

[9] D.A. Forsyth and J. Ponce, Computer Vision: A modern approach, Prentice Hall, 2003.
[10] L.H. Matthies, "Stereo vision for planetary rovers: Stochastic modeling to near real-time implementation," *International Journal of Computer Vision*, Vol. 8, No. 1, July 1992, pp. 71-91.
[11] Y. Xiong and L.H. Matthies, "Error analysis of a real-time stereo system," Proc. IEEE Conf. on Computer Vision and Pattern Recognition, IEEE Computer Society, San Juan, Puerto Rico, 1997, pp. 1087-1093.
[12] L. Matthies, E. Gat, R. Harrison, B. Wilcox, R. Volpe, and T. Litwin, "Mars microrover navigation: Performance evaluation and enhancement," *Autonomous Robots*, Vol. 2, No. 4, Kluwer Academic Publishers, Dordrecht, Netherlands, 1995, pp. 291-312.
[13] R. Volpe, "Navigation results from desert field tests of the Rocky 7 Mars rover prototype," *International Journal of Robotics Research*, Special Issue on Field and Service Robots, Vol. 18, No. 7, July 1999.
[14] S. Singh, K. Schwehr, R. Simmons, et al, "Recent progress in local and global traversability for planetary rovers," Proc. IEEE Intl. Conf. on Robotics and Automation, 2000.
[15] P.S. Schenker, T. Huntsberger, P. Pirjanian, E. Baumgartner, et al, "Robotic automation for space: Planetary surface exploration, terrain-adaptive mobility, and multi-robot cooperative tasks," Proc. SPIE Intelligent Robots and Computer Vision XX: Algorithms, Techniques, and Active Vision Conference, Vol. #4572, October 2001.
[16] J. Biesiadecki, M. Maimone, and J. Morrison, "The Athena SDM rover: A testbed for Mars rover mobility," Proc. Intl. Symp. on Artificial Intelligence, Robotics and Automation in Space, Montreal, Canada, Paper No. AM026, June 2001.
[17] L. Matthies, A. Kelly, T. Litwin and G. Tharp, "Obstacle detection for unmanned ground vehicles: A progress report," Proc. 7th Intl. Symp. on Robotics Research, Springer-Verlag, 1996.
[18] M. Maimone, L. Matthies, J. Osborn, E. Rollins, J. Teza and S. Thayer, "A photo-realistic 3-D mapping system for extreme nuclear environments: Chernobyl," Proc. Intl. Conf. on Intelligent Robots and Systems, Victoria B.C., Canada, 1998, pp. 1521-1527.
[19] L. Matthies, Y. Xiong, R. Hogg, D. Zhu, et al, "A portable, autonomous urban reconnaissance robot," Proc. Intelligent Autonomous Systems, Venice, Italy, July 2000.
[20] P. Bellutta, R. Manduchi, L. Matthies, K. Owens and A. Rankin, "Terrain perception for Demo III," Proc. Intelligent Vehicles Symposium, Dearborn, MI, Oct. 2000.
[21] A. Rankin, C. Bergh, S. Goldberg and L. Matthies, "Passive perception system for day/night autonomous off-road navigation," SPIE UGV Symposium, Orlando, FL, April 2005.
[22] A. Mishkin, J.C. Morrison, T.T. Nguyen, H.W. Stone, et al., "Experiences with operations and autonomy of the Mars Pathfinder microrover," Proc. IEEE Aerospace Conference, Aspen, CO, March 1998.
[23] R. Simmons, L. Henriksen, L. Chrisman and G. Whelan, "Obstacle avoidance and safeguarding for a lunar rover," Proc. AIAA Forum on Advanced Developments in Space Robotics, Madison, WI, Aug. 1996.
[24] D. Wettergreen, D. Bapna, M. Maimone and G. Thomas, "Developing Nomad for robotic exploration of the Atacama Desert," *Robotics and Autonomous Systems*, Vol. 26, No. 2-3, Elsevier, Feb. 1999, pp. 127-148.
[25] L.H. Matthies, *Dynamic Stereo Vision*, Ph.D. Thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA, October 1989.
[26] C.F. Olson, L.H. Matthies, M. Schoppers, M.W. Maimone, "Robust stereo ego-motion for long distance navigation," Proc. IEEE International Conference on Computer Vision and Pattern Recognition, Hilton Head, SC, 2000, pp. 453-458.
[27] T. Huntsberger, H. Aghazarian, Y. Cheng, et al, "Rover autonomy for long range navigation and science data acquisition on planetary surfaces", Proc. IEEE International Conference on Robotics and Automation, Washington, DC, May 2002, pp. 3161-3168.
[28] Y. Cheng, M. Maimone and L. Matthies, "Visual odometry on the Mars exploration rovers," Proc. IEEE International Conference on Systems, Man, and Cybernetics, Waikoloa, HI,

October 2005.

[29] S. Maxwell, B. Cooper, F. Hartman, J. Wright, J. Yen, "The design and architecture of the rover sequencing and visualization program (RSVP)," Proc. 8th Intl. Conf. on Space Operations (SpaceOps), Montreal, Canada, 2004.

[30] B. Cooper, F. Hartman, S. Maxwell, J. Wright, J. Yen, "Using RSVP for analyzing states and previous activities for the Mars Exploration Rovers," Proc. 8th Intl. Conf. on Space Operations (SpaceOps), Montreal, Canada, 2004.

[31] E. Tunstel, M. Maimone, A. Trebi-Ollennu, J. Yen, R. Petras and R. Willson, "Mars Exploration Rover mobility and robotic arm operational performance," Proc. IEEE Intl. Conf. on Systems, Man, and Cybernetics, Waikoloa, HI, October 2005, pp. 1807-1814.

[32] G.A. Dorais, R.P. Bonasso, D. Kortenkamp, B. Pell and D. Schreckenghost, "Adjustable autonomy for human-centered autonomous systems," Working notes of the 16th Intl. Joint Conf. on Artificial Intelligence, Workshop on Adjustable Autonomy Systems, Stockholm, Sweden, 1999, pp. 16-35.

[33] P.C. Leger, R.G. Deen and R.G. Bonitz "Remote image analysis for Mars Exploration Rover mobility and manipulation operations," Proc. IEEE International Conference on Systems, Man, and Cybernetics, Waikoloa, HI, October 2005, pp. 917-922.

[34] J.S. Norris, M.W. Powell, M.A. Vona, P.G. Backes and J.V. Wick, "Mars Exploration Rover operations with the Science Activity Planner," Proc. IEEE Conf. on Robotics and Automation, Apr. 2005.

[35] A. Howard and H. Seraji, "Vision-based terrain characterization and traversability assessment," *Journal of Robotic Systems*, Vol., 18, No. 10, pp. 577-587, 2001.

[36] A. Talukder, R. Manduchi, K. Owens, L. Matthies, A. Castaño, R. Hogg, "Autonomous terrain characterization and modelling for dynamic control of unmanned vehicle," Proc. IEEE/RSJ Conf. on Intelligent Robots and Systems (IROS), Sept. 2002.

[37] R.L. Castaño, T. Mann and E. Mjolsness, Texture analysis for Mars rover images," Proc. Applications of Digital Image Processing XXII, SPIE Vol. 3808, July, 1999.

[38] A. Shirkhodaie, R. Amrani and E. Tunstel, "Visual terrain perception modeling of space planetary robotic systems based on soft computing classifiers," Proc. 14th IEEE Intl. Workshop on Robot and Human Interactive Communication, Nashville, TN, Aug. 2005.

[39] R.A. Lindemann and C.J. Voorhees, "Mars Exploration Rover mobility assembly design, test and performance," Proc. IEEE International Conference on Systems, Man, and Cybernetics, Waikoloa, HI, October 2005, pp. 450-455.

[40] A. Stentz, "Optimal and efficient path planning for partially-known environments," Proc. IEEE Intl. Conf. Robotics and Automation, 1994, pp. 3310-3317.

[41] D. Ferguson and A. Stentz, "Field D*: An interpolation-based path planner and replanner," Proc. 12th Intl. Symp. on Robotics Research, San Francisco, CA, Oct. 2005.

[42] R. Volpe, "Rover functional autonomy development for the Mars Mobile Science Laboratory," Proc. IEEE Aerospace Conf., Big Sky, Montana, Mar. 2003, Paper #1289.

[43] E. Tunstel, A. Howard, M. Maimone and A. Trebi-Ollennu, "Mars Exploration Rover baseline for flight rover autonomy technology assessment," Proc. 8th Intl. Symposium on Artificial Intelligence, Robotics, and Automation in Space, Munich, Germany, Sept. 2005.