Overview of the Mars Exploration Rovers' Autonomous Mobility and Vision Capabilities

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

Abstract— NASA's Mars Exploration Rovers have set the standard for autonomous robotic exploration of planetary surfaces. Their abilities to detect and avoid geometric hazards, and measure and compensate for slip or heading changes, have made it possible to drive farther and in highly sloped areas, increasing the science return of the mission. Software updates that took place during the more than three year mission have increased their abilities even further, raising the bar for the remainder of their mission and all that will follow.

In this paper we summarize the autonomous capabilities available on the Mars Exploration Rovers following the September 2006 software update.

1. Background

All spacecraft include a high degree of autonomy by necessity. Capabilities included in spacecraft launched from around the world over the past five decades include responses to faults, high gain antenna pointing, orientation assessment using star trackers, and onboard data storage and retransmission strategies.

But the detailed exploration of the surface of another world brings with it even more challenges than those. Spacecraft that travel through space or in orbit about a planetary body might be commanded to perform navigation corrections tens or hundreds of times during multi-year missions, but the Mars Exploration Rover (MER) vehicles have performed over 60,000 coordinated motions (powering either steering or drive motors continuously) as of January 2007. That much motion requires nearly constant attention, and the uncertainty of interaction with new terrain means that little drive progress could be made in challenging terrain without help from onboard autonomy.

Many constraints limit what each rover can accomplish on any given Martian solar day (or *sol*): the need for human confirmation of vehicle state before beginning certain activities, and the availability of power, data volume, and execution time for each activity. Onboard autonomy enables improved performance along one or more of these dimensions, typically the elimination of the need for human confirmation of the vehicle state. Human operators generally interact with the rovers at most once per sol, so each human confirmation step that can be eliminated from a procedure can potentially reduce the overall activity time by at least one sol [Mishkin et al., 2006], [Biesiadecki et al., 2007].

Several driving modes were used during the first three years of operation. *Directed* drives executed a planned course without any onboard compensation for position or attitude drift. *Visual Odometry* drives updated the rover's position but did not check for obstacles. *Terrain Assessment* drives looked for geometric hazards (e.g., rocks, ditches), but did not measure any slip. *Local Path Selection* drives corrected for heading changes and anything measured by Visual Odometry and Terrain Assessment, instead of just blindly following a directed drive.

There have been very few successful rover missions to other worlds. Not counting NASA's human-driven lunar rovers and Soviet teleoperated Lunakhod 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 successfully 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. Future rover missions currently being planned include NASA's Mars Science Laboratory (launch in 2009) and the European Space Agency's ExoMars mission (launch in 2013) [Baglioni et al., 2006].

In this paper we outline the primary autonomous mobility and image processing capabilities available to the MER science and engineering teams.

Onboard Computing

The MER vehicles have a single general purpose processor available for autonomy processing, a 20 MHz RAD6000 CPU with 128 Mbytes of RAM and 256 Mbytes of flash memory. The flight computer uses the VxWorks operating system, and runs dozens of parallel tasks that implement the flight software [Reeves and Snyder, 2005]. During normal operations, at most 75 percent of the CPU time is available for autonomy software, but telemetry processing can sometimes reduce that amount even more. Dynamic memory allocation is strongly discouraged by the coding standards document, but some non-system RAM is available in a set of dedicated memory pools of 4 Mbytes, 9 Mbytes, and up to 10 additional 2 Mbyte blocks (use of these 2 Mbyte blocks reduces the memory available for image processing and is also discouraged).

2. Primary Mission Capabilities

The MER vehicles landed with several Mobility and Vision technologies already incorporated into their flight software [Maimone et al., 2006b]. Although incremental improvements were made over the mission lifetime, their basic functionality was already demonstrated during the first 90 sols comprising the Primary Mission.

Descent Image Motion Estimation System

Safely landing on Mars required complete autonomy during entry, descent, and landing since the round-trip time for telemetry and commands was much longer than the time between entering the atmosphere and landing. One part of this autonomy was DIMES, the Descent Image Motion Estimation System. MER's radar could not measure horizontal velocity relative to the surface, but atmospheric models predicted that sustained winds could impart a horizontal velocity to the MER entry system that was greater than the design rating of the airbag-based landing system. DIMES was tasked with estimating horizontal velocity from three images taken 2.5 seconds apart starting at a target altitude of 2000m. This estimate was then fed to three solid rocket motors that could change the horizontal velocity of the lander just a few seconds before impact.

DIMES uses radar and attitude measurements to project the three images onto the nominal ground plane, then choose and track 2 surface features between each image pair. This, along with knowledge of the time between images, allowed DIMES to compute two horizontal velocity measurements, and check the difference between them – the average acceleration – against that determined from accelerometers. The main challenges for DIMES were threefold:

- Develope a flight-qualified system with less than two years to go before launch.
- Incorporate enough robustness into DIMES so that it could compute accurate measurements and – equally important – know when its measurements were unreliable.
- Acquire 3 images (at 3.75 seconds each) and analyze them within 17 seconds while using only 40% of the 20MHz CPU.

The DIMES team was able to leverage existing computer vision expertise and algorithms for much of the low-level image processing, allowing the focus to be on robustness, optimization, and sound engineering analysis of every part of DIMES, from camera lenses and filters, to image smear from slew rates up to 60 deg/sec, to correction of radiometric and frame-transfer artifacts, and end-to-end performance assessment and error analysis.

DIMES successfully produced horizontal velocity measurements for both Spirit and Opportunity. While Opportunity's horizontal velocity was low enough that it would have been within the design envelope of the airbags, Spirit experienced higher winds and the estimate produced by DIMES was essential in reducing the horizontal velocity to within the airbags' rating. [Cheng et al., 2004], [Johnson et al., 2007].

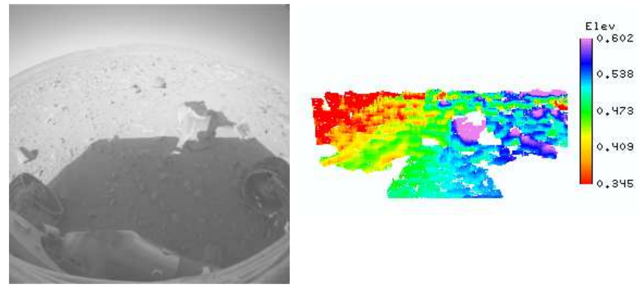

Figure 1. On Sol 118, Spirit autonomously detected a wheel-sized rock in the middle of a 54 meter autonomous drive. The left image is the original (unrectified) image, the right is the disparity image color-coded by elevation (recomputed on Earth from the original images). The rock, somewhat obscured by the rover shadow on the left, is clearly visible as the large violet area on the right.

Absolute Orientation Sensing

MER vehicles rely on the PANCAM cameras to provide an absolute heading measurement during surface operations [Ali et al., 2005]. The onboard position estimate can accumulate several degrees of drift after integrating the Inertial Measurement Unit (IMU) gyros for thousands of seconds. This drift is eliminated by combining a vector pointed at sun, knowledge of the current local solar time, and the direction of gravity measured by the IMU accelerometers. The Sun vector is found by pointing the camera where the Sun is expected to be, then processing the image to re-locate its center [Eisenman et al., 2002].

As of January 2007, Sun-locating commands have run more than 100 times on each rover.

Stereo Image Processing

One of the most capable sensors on the MER vehicles is the combination of any available camera pair and software for measuring the shape of the surrounding terrain from those images [Goldberg et al., 2002]. Four stereo pairs of cameras with varying fields of view are available [Maki et al., 2003]. Our stereo vision software relies on the cameras' geometric lens calibration [Gennery, 2006] and performs a windowed 1D search using the Sum-of-Absolute-Differences metric to generate full 3D measurements of points in the stereo images. It has been a critical component of the onboard mobility system, providing the Terrain Assessment software the geometric information it needs to predict vehicle safety, expressed as a cloud of 3D points.

The primary use of this stereo vision capability has been in support of Terrain Assessment operations as of January

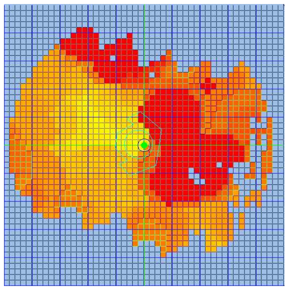

Figure 2. Spirit's onboard Terrain Assessment map associated with the image in Figure 1. Cells are colored based on whether the center of rover would be safe there; red areas are impassable obstacles, yellow areas are not entirely flat yet are safe enough to traverge. The rover is in the center of the map (north is up), facing the obstacle to the east. The onboard map clearly shows the rock as a large red obstacle and includes information from earlier locations as well.

2007. Images are typically acquired at "binned" resolution (1024x256 pixels^2) and then software-downsampled and rectified to 256x256 pixels^2. Stereo processing results in a 256x256 Disparity image, where each pixel is either unknown or encodes the 3D location of the terrain shown at that pixel in the rectified left image (see Figure 1). Spirit typically uses 125^o Field of View (FOV) front and rear HAZCAMs to compute 15,000 3D measurements per stereo pair on average, while Opportunity uses 45^o FOV NAVCAMs and computes 48,000 3D measurements on average [Maimone et al., 2006a].

Onboard stereo is also used as part of the Instrument Placement autonomous technology described in Section 3. In that mode, one or two pairs of front HAZCAM images are taken at 1024x1024 resolution (but downsampled to 512x512 and then cropped to 420x412) to generate a coarse view of the terrain, followed by a full-resolution view of a relatively small patch of terrain in a 241x241 subframed view. This most recent application of stereo was still being checked out in January 2007, and Figure 6 shows the results of one of the first Instrument Placement tests using it.

Terrain Assessment

The GESTALT (Grid-based Estimation of Surface Traversability Applied to Local Terrain) system detects geometric hazards (e.g., rocks, ditches, cliffs) in the area around the rover. GESTALT processes the cloud of 3D points computed by Stereo Vision by fitting rover-sized patches of data to a plane. It looks for Step Obstacles (large deltas in elevation from the best fit plane), Tilt Hazards (large angle between surface normal and the Up vector), and Roughness hazards (residual of the planar fit) and maintains a traversability map of the terrain immediately around the rover (10 x 10 meters^2 on Spirit, 12 x 12 meters^2 on Opportunity, each with 0.2 meter cells) [Goldberg et al., 2002], [Biesiadecki and Maimone, 2006]. See Figure 2 for an example. As many as 10 separate point clouds can be collected before performing the traversability analysis (e.g., on sol 378 Opportunity used two pairs of NAVCAMs at each step), but normally only a single stereo pair is taken.

Terrain Assessment can be used in combination with other drive modes. It is primarily used to extend the distance driven by the rovers beyond the edge of what can be reliably seen in the NAVCAM and PANCAM stereo image data. It functions both as a simple predictive safety check prior to executing one or more directed commands (so-called *Guarded Motion*, in which the presence of an obstacle would preclude all further driving that sol), and also in combination with Local Path Selection to provide a fully autonomous obstacle avoidance capability (see the next section below).

Terrain Assessment also provides the foundation for a simple "Turn toward a Rock" command, in which the goodness map is searched for the nearest obstacle which then becomes the drive goal. A similar capability was found useful on the Sojourner rover [Wilcox and Nguyen, 1998], but the combination of MER's accurate position estimation, the desire to end each drive at a particular heading (to aid communication efficiency), and the limitation that it can only track something that looks like an obstacle means that this capability is not used operationally (though it was checked out on Opportunity during sols 103 and 352). The Visual Target Tracking technology in Section 3 describes a related capability.

Spirit's six month (and more than 2 kilometer) trek toward the Columbia Hills made much use of its Terrain Assessment capability. As a result, nearly one third of its traverse to the hills was driven using Terrain Assessment, enabling Spirit to reach the hills 50% sooner than would have been safely possible using only directed drives. Spirit's longest commanded drive was 124 meters on Sol 125, in which it drove 62 meters using directed driving and then another 62 meters using Terrain Assessment and Local Path Selection. Opportunity's longest planned drive took place over a holiday weekend in February 2005. After an initial 106 meter Local Path Selection drive without any vision processing, it covered an additional 284 meters using Local Path Selection with Terrain Assessment (70 meters, 104 meters, 110 meters on sols 383–385; although humans reviewed the vechile telemetry each sol, no futher commands were sent after the first day).

As of 15 August 2005, onboard terrain assessment was used on Spirit for 1354 meters of 4,798 total meters traversed, and on Opportunity for 1379 of 5,947 total meters traversed, or 25% of the overall distance for both rovers. [Biesiadecki and Maimone, 2006]

Local Path Selection

One of the most straightforward autonomous driving modes is Local Path Selection, in which the rover corrects its path as it drives toward some pre-specified goal location. This is in contrast to so-called *directed* drives, wherein during ARC drives the rover follows precise pre-planned motor rotations

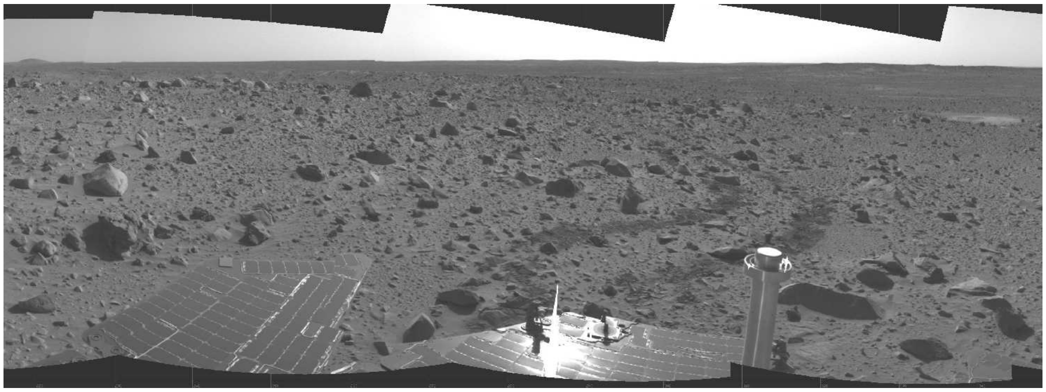

Figure 3. On Sol 107, Spirit avoided obstacles in previously-unseen terrain.

while disregarding all position measurement sensors other than the wheel encoders.

In the simplest instance of this mode, no vision processing is performed, and the vehicle chooses its next step based on the best position estimate derived solely from wheel encoders and the gyros in its Inertial Measurement Unit (IMU) [Ali et al., 2005]. Although this mode cannot measure slip in distance, it can detect and compensate for unplanned changes in yaw (which can occur when driving across sandy slopes, for instance).

Combining this mode together with Terrain Assessment allows the rover to not only detect obstacles, but also steer around them (see Figure 3). This combined mode was used extensively to allow both rovers to go "beyond the horizon", driving safely into terrain that had never been seen at sufficiently good resolution to assess the presence of obstacles. For example, Opportunity drove 390 meters during Sols 383–385, 284 meters of which was driven autonomously in this mode.

Combining Local Path Selection with Visual Odometry and clever sequences of commands (i.e., conservative driving strategies including Keep-out zones [Biesiadecki et al., 2007]) allows the rover to navigate safely on high slopes or in sandy areas. Drives in this mode were limited to shorter distances, since the vehicle could only be driven as far as the human rover drivers could see potential obstacles. Both Spirit and Opportunity have used this mode extensively in highly sloped and sandy areas (or if not exactly this mode then its equivalent, a series of directed ARC drives whose execution was conditional on the current vehicle position estimate).

Visual Odometry

On relatively level ground, the MER vehicles demonstrated remarkably good position estimation using only wheel encoders and the IMU (Li et al report that Spirit's position estimate was only off by 3% even after driving more than 2 kilometers [Li et al., 2006]). But on steep hillsides, in mixed sand/rock terrains inside craters, and even when crossing sandy ripples in the otherwise flat plains of Meridiani, very

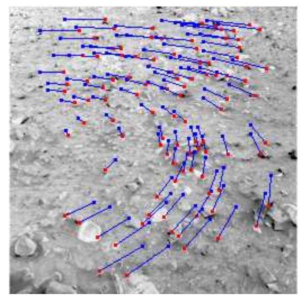

Figure 4. Illustration of autonomous feature selection and tracking by Visual Odometry. Red dots indicate the locations of 80 features in the given image, blue vectors point to where the same feature was found in the previous image. On Sol 520, Spirit drove along the side of the Columbia Hills at tilts ranging from 7 to 22 degrees. This image shows one tracking step in which the rover was commanded to drive 0.6 meters forward, but its 10 – 11 degree roll and/or the terrain caused it to slide 0.12 meters downslope (indicating 20% slip).

large slips were found (e.g., 100% slip during the first exit attempt from Eagle Crater, 99.9% slip getting stuck in Purgatory ripple, and 125% slip during one climb in the Columbia Hills). The only sensor on the MER vehicles capable of detecting position slip is the combination of NAVCAM images and Visual Odometry software.

The MER Visual Odometry software compares pairs of NAVCAM images of nearby terrain, and autonomously detects and tracks features between those images (see Figure 4). The 2D and 3D motions of those features is used to update the vehicle's onboard position estimate according to the algorithm described in [Maimone et al., 2007]. Originally included as an "extra credit" capability, Visual Odometry has become a critical component of the rovers' safety systems.

Several styles of driving that developed during the mission are only possible thanks to Visual Odometry. In sandy terrain that might bog down the vehicle, short directed drives (around 5 – 10 meters) are followed by Slip Checks to ensure the it is not yet stuck. On high slopes with large rocks or deep sand, manually-specified Keep-out zones keep the rover

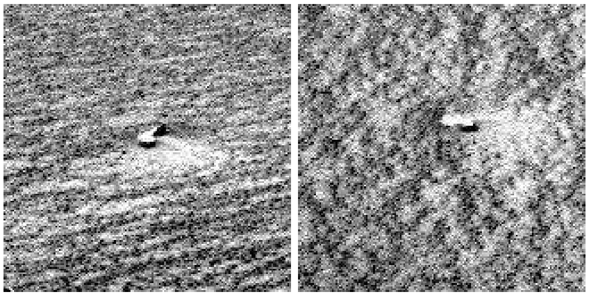

Figure 5. On Sol 992, Opportunity tested Visual Target Tracking for the first time. The feature on the left was tracked through 8 updates as the rover drove around it. The right image shows the final view, indicating successful tracking.

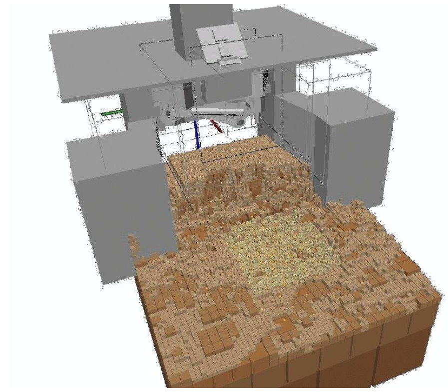

Figure 6. On Sol 1068, Spirit tested its Instrument Placement ability to process the terrain onboard. This rendering shows a model of the rover collision volume in grey, the terrain model that was generated as cubes, and the potential instrument placement targets as orange spheres.

safe while opportunistically trying to drive longer distances (getting too close to an obstacle precludes further driving for that sol). And after leaving an area studied in detail by the instrument arm, Visual Odometry enabled accurate pointing of the remote sensing mast instruments (PANCAM, miniTES spectrometer). These styles are discussed in more detail in [Biesiadecki et al., 2007], [Maimone et al., 2007].

As of 15 August 2005, Visual Odometry was used on Spirit for 792 meters of 4,798 total meters traversed, and on Opportunity for 682 of 5,947 total meters traversed, or 14% of the overall distance for both rovers. [Biesiadecki and Maimone, 2006]

3. Capabilities Added during the Extended Missions

MER flight software updates took place once en route to Mars, and three times during the surface mission in April 2004, January 2005, and September 2006. Earlier updates addressed incremental enhancements, but the most recent one incorporated several new robotic autonomy capabilities.

Visual Target Tracking

When planning a drive toward a precise, pre-specified target, two kinds of precision must be taken into account: target specification precision, and position estimation uncertainty. Throughout their missions thus far, MER vehicles have been driven using precise, metrically specified goal locations (e.g., "go to location x,y $\pm$ 2 meters", "drive 1.3 meters forward"). But target locations are often determined from stereo correlation, whose precision falls off with the square of the target's distance from the cameras, so targets more than a few meters away have nontrivial uncertainty (and hence less precision) in their specification. And unless Visual Odometry is enabled, the rovers might be unaware of slip that occurs while driving toward the target. Both of these effects can result in the rover not quite reaching its target, but both are mitigated by Visual Target Tracking (also known as Visual Servoing), which allows scientists and engineers to specify a feature by its *appearance*, rather than its predicted location.

The MER flight software now incorporates the ability to track a target using the pointable NAVCAMs. While driving near a feature of interest, a series of images is processed to update the rover's knowledge of the feature's location using a correlation function [Kim et al., 2005]. The current location of the target is maintained at each step, and can be used in many ways: as the imaging target of other mast-mounted sensors, as part of a sequence of conditional commands, or as the goal to be reached at the end of a drive. And because it is tracking only a single target, rather than the hundreds of targets found by Visual Odometry, each step can execute much more quickly than Visual Odometry.

During the first part of its checkout, this technology successfully tracked a feature as Opportunity drove by it on sol 992 (see Figure 5).

Instrument Placement

Each MER vehicle has several science instruments mounted on a robotic arm called the Instrument Deployment Device (IDD) [Trebi-Ollennu et al., 2005], [Baumgartner et al., 2005]. During most of the mission, MER project procedures dictated that the IDD should never be deployed onto a target without human confirmation so that its safety would not be compromised. Such confirmation was based on detailed processing of high-resolution stereo data to determine appropriate placement points on or near the terrain. The image data had to be acquired from the rover's current position to guarantee knowledge of terrain geometry relative to the rover, which required a downlink-and-command cycle (an additional sol) between the end of a drive and deployment of the IDD.

An autonomous instrument placement capability called AutoPlace was developed to enable safe instrument placements after driving a short distance to a target within a single command cycle. This requires stereo imaging, target selection, workspace safety analysis, and trajectory generation. Many of these capabilities were previously incorporated into the planning tools used by rover operators [Leger, 2002], [Leger et al., 2005].

The most important aspect of AutoPlace is safety: ensuring that any commanded trajectory is free of self-collisions and collisions with the terrain. Safety is assessed by testing each "via" point along a potential trajectory for collisions with a terrain model built from stereo data (see Figure 6). This terrain model also explicitly models volumes that are unknown due to occlusions and stereo dropouts, so that any volume not confirmed to be free of obstacles is considered unsafe. A tolerance several times greater (3cm) than the expected stereo-to-arm error ($<$ 1 cm) is used to ensure that the only collisions with terrain are those that are intentional–placing the commanded instrument in contact with a target.

Target selection requires finding the closest point on the terrain to the commanded target and ensuring that it meets workspace, surface orientation, and surface roughness requirements. The preferred target can be specified as relative to the rover (wherever it ends up), relative to a coordinate frame fixed to the surface, or relative to the target position determined by Visual Target Tracking. Multiple targets can be stored and attempted in series, and if the closest point to the commanded target is not safely reachable, alternate nearby targets can be autonomously selected.

Trajectory generation is based on the method used when manually building IDD command sequences: first, a joint-space move brings the tool over the target in the desired kinematic configuration; then a cartesian move to a point 10cm over the target and aligned with the surface normal; and finally, doing a guarded cartesian move to a position either over the target (for Microscopic Imager and APXS placements) or with a small positive overdrive (for the MB) to ensure contact with the surface. (A guarded move indicates that the contact switch on the active instrument should cause motion to stop without triggering a fault; in contrast, a contact switch trip during a free-space move will cause a fault.) The trajectory generation software can consider trajectories in multiple kinematic configurations (as human operators do) to avoid joint limits, terrain collisions, kinematic configuration flips, and other faults. The trajectory generation software can also make slight, bounded modifications to a surface normal to avoid faults, as is common in manually-generated command sequences.

Autonomous Science

Several types of science observations are done in an opportunistic fashion, taking multiple sensor readings in an attempt to accomplish one type of observation. Two examples of this are attempts to image dust devils and clouds. Throughout most of the mission, these types of observations would collect dozens of images at particular times of day, in the hope that those atmospheric activities could be detected. Many of the images did not show any such activity, yet still had to be transmitted to Earth so humans could assess their utility. This requires the transmission of a large volume of images, even the unchanging ones, reducing the overall amount of useful science returned on any given sol.

The latest software enables the rover to detect the presence of these features on its own, and selectively transmit only those images deemed to be interesting. This will result in overall increased science return, by greatly reducing the data volume to just those images (or subframed images) that demonstrate the desired criteria [Castano et al., 2006], [Castano et al., 2007].

Checkout activities so far have shown that the software successfully detected the absence of dust devils and the presence of clouds. Because dust devils and clouds tend to be seasonal events, positive dust devil feature detection may only become feasible some months later.

Global Path Selection

The combination of Terrain Assessment and Local Path Selection is sufficient to navigate around occasional small obstacles using a "greedy," goal-seeking behavior. But the MER vehicles have also encountered larger "extended" obstacles (e.g., large rocks 1 meter or more in length, multiple parallel ripples, fractures, and craters). The Local Path Selector is unable to plan more than one or two steps ahead, and as a result Spirit has occasionally failed to make progress beyond an extended obstacle (e.g., on sol 108). In such situations, the vehicle is not mechanically stuck, it just fails to backtrack far enough to allow it to continue toward its goal. Instead, it moves back and forth until the command times out or software checking for the lack of forward progress halts the drive.

Now a Global Path Planner has been incorporated into the onboard flight software. The Field D* planner developed at Carnegie Mellon University maintains a much larger world map (typically 50 x 50 meters^2 with 0.4 meter cells) and provides the ability to plan arbitrary paths through its map [Ferguson and Stentz, 2006], [Carsten et al., 2007].

This capability was first demonstrated on Opportunity during sol 1014 in a passive mode, computing its map and the desired path but not yet actively selecting the next drive command (see Figure 7).

Status of New Autonomy Technologies

Uplink of the new autonomy technologies took place just before the start of winter on Mars, and therefore the rovers had limited solar energy available for extensive testing. So although each of these technologies has completed the formal Earth-based Verification and Validation process required of all flight software, as of January 2007 all were still going through their "check out" procedures on Mars.

4. Combining Autonomous Capabilities

Even greater benefits of autonomy can be realized when different capabilities combine to provide even higher level ca-

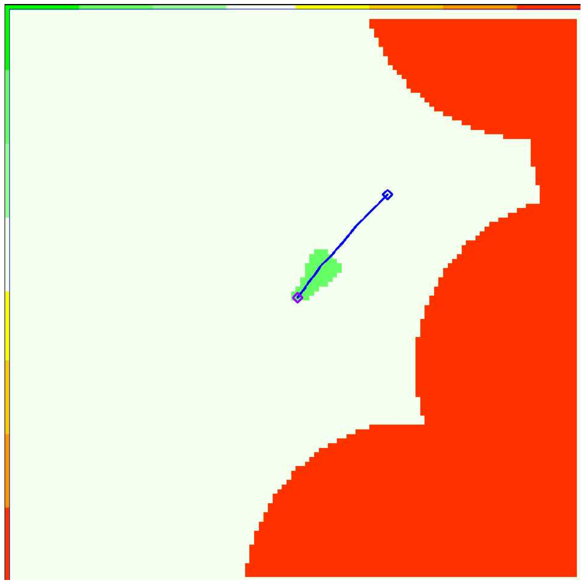

Figure 7. On Sol 1014, Opportunity tested Field D* for the first time. This frame shows the first of seven successful updates; red areas are human-specified Keep-out zones to avoid getting near the edge of Victoria crater, white areas are unknown, green is the known-to-be-safe area processed by NAVCAM stereo images. The blue line shows the current optimal path from the rover (lower left diamond) to its goal (upper right diamond).

pabilities. For example, Terrain Assessment alone provides a useful "safe or not safe" indication, but combined with Local (and soon, Global) Path Selection enables fully autonomous driving around obstacles to reach a goal point on flat ground; add Visual Odometry processing and it would become possible to navigate around obstacles in slippery areas as well. Criteria for choosing between the original autonomy modes are presented in [Biesiadecki et al., 2007], and statistics summarizing the actual use of combinations of those modes can be found in [Biesiadecki and Maimone, 2006].

Once the new technologies have been checked out, even greater levels of autonomy will become possible. As of January 2007, a minimum of two sols has been required to drive to a nearby target and deploy the *in situ* sensors of the IDD onto it. This minimum has been achieved several times, most notably during Opportunity's sol 304 8.7 meter drive at 20 – 24 degree slopes along the edge of Endurance crater [Maimone et al., 2007]). But the combination of (at least) Visual Target Tracking and Instrument Placement could enable safe instrument deployment in a single sol.

The main restriction on the actual use of combinations of MER autonomous capabilities is processing time. The relatively slow speed of its space-qualified CPU, an architecture that prevents full benefits of the processor cache from being realized, and limited development time left us with very capable autonomy that can take minutes to process a single set of images.

5. Conclusion

The MER vehicles have successfully demonstrated several autonomous capabilities during their first three years of operation. New software is making them even more capable, which will not only benefit MER operations but also raise the baseline expected of future missions.

6. Acknowledgments

Many thanks to the developers of these technologies, NASA's Mars Technology Program office, Software Intelligent Systems and Modeling, and the New Millennium Program for funding much of the new technologies, MER project management for supporting their inclusion in flight software, and the Validation and Verification testbed teams. Thanks also to FSW lead Khaled Ali for valuable suggestions on this paper.

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract to the National Aeronautics and Space Administration.

References

[Ali et al., 2005] Ali, K. S., Vanelli, C. A., Biesiadecki, J. J., Maimone, M. W., Cheng, Y., Martin, M. S., and Alexander, J. W. (2005). Attitude and position esimation on the Mars Exploration Rovers. In *IEEE Conference on Systems, Man and Cybernetics*, pages 20–27, The Big Island, Hawaii, USA.

[Baglioni et al., 2006] Baglioni, P., Fisackerly, R., Gardini, B., Gianfiglio, G., Pradier, A., Santovincenzo, A., Vago, J., and van Winnendael, M. (2006). The Mars exploration plans of ESA. *IEEE Robotics and Automation Special Issue (MER)*, 13:83 – 89.

[Baumgartner et al., 2005] Baumgartner, E., Bonitz, R., Melko, J., Shiraishi, L., and Leger, P. (2005). The mars exploration rover instrument positioning system. In *IEEE Aerospace Conference*.

[Biesiadecki et al., 2007] Biesiadecki, J. J., Leger, P. C., and Maimone, M. W. (2007). Tradeoffs between directed and autonomous driving on the Mars Exploration Rovers. *International Journal of Robotics Research*, 26(1):91–104. http://ijr.sagepub.com/cgi/content/abstract/26/1/91.

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

[Carsten et al., 2007] Carsten, J., Rankin, A., and Ferguson, D. (2007). Global path planning on board the mars exploration rovers. In *IEEE Aerospace Conference*, Big Sky, MT, USA.

[Castano et al., 2006] Castano, A., Fukunaga, A., Biesiadecki, J., Neakrase, L., Whelley, P., Greeley, R., Lemmon, M., Castano, R., and Chien, S. (2006). Autonomous detection of dust devils and clouds on Mars. In *IEEE Intl. Conf. Image Processing*, Atlanta, USA. http://trs-new.jpl.nasa.gov/dspace/bitstream/2014/39920/1/06-1553.pdf.

[Castano et al., 2007] Castano, A., Fukunaga, A., Biesiadecki, J., Neakrase, L., Whelley, P., Greeley, R., Lemmon, M., Castano, R., and Chien, S. (2007). Automatic detection of dust devils and clouds on Mars. *Machine Vision and Applications Journal*.

[Cheng et al., 2004] Cheng, Y., Goguen, J., Johnson, A., Leger, C., Matthies, L., Martin, M. S., and Willson, R. (2004). The mars exploration rovers descent image motion estimation system. *IEEE Intelligent Systems*, 19(3):13–21. http://www-robotics.jpl.nasa.gov/publications/Andrew_Johnson/DIMES-ai-space.pdf.

[Eisenman et al., 2002] Eisenman, A. R., Liebe, C. C., and Perez, R. (2002). Sun sensing on the Mars Exploration Rovers. In *IEEE Aerospace Conference*, volume 5, pages 2249–2262.

[Ferguson and Stentz, 2006] Ferguson, D. and Stentz, A. (2006). Using interpolation to improve path planning: The field d* algorithm. *Journal of Field Robotics*, 23(2):79–101. http://gs2045.sp.cs.cmu.edu/downloads/fdstar-journal.pdf.

[Gennery, 2006] Gennery, D. B. (2006). Generalized camera calibration including fish-eye lenses. *International Journal of Computer Vision*, 68(3):239 – 266.

[Goldberg et al., 2002] Goldberg, S. B., Maimone, M. W., and Matthies, L. (2002). Stereo vision and rover navigation software for planetary exploration. In *IEEE Aerospace Conference*, volume 5, pages 2025–2036, Big Sky, Montana, USA. http://robotics.jpl.nasa.gov/people/mwm/visnavsw/.

[Johnson et al., 2007] Johnson, A., Willson, R., Cheng, Y., Goguen, J., Leger, C., SanMartin, M., and Matthies, L. (2007). Design through operation of an image-based velocity estimation system for mars landing. *International Journal of Computer Vision*.

[Kim et al., 2005] Kim, W. S., Steele, R. D., Ansar, A. I., Ali, K., and Nesnas, I. (2005). Rover-based visual target tracking validation and mission infusion. In *AIAA Space Conference*.

[Leger, 2002] Leger, C. (2002). Efficient sensor/model based on-line collision detection for planetary manipulators. In *International Conference on Robotics and Automation*.

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

[Li et al., 2006] Li, R., Archinal, B. A., Arvidson, R. E., Bell, J., Christensen, P., Crumpler, L., Marais, D. J. D., Di, K., Duxbury, T., Golombek, M., Grant, J., Greeley, R., Guinn, J., Johnson, A., Kirk, R. L., Maimone, M., Matthies, L. H., Malin, M., Parker, T., Sims, M., Thompson, S., Squyres, S. W., and Soderblom, L. A. (2006). Spirit rover localization and topographic mapping at the landing site of Gusev Crater, Mars. *JGR-Planets, Special Issue on Spirit Rover*, 111(E02S06).

[Maimone et al., 2006a] Maimone, M., Biesiadecki, J., Tunstel, E., Cheng, Y., and Leger, C. (2006a). *Surface navigation and mobility intelligence on the Mars Exploration Rovers*, chapter 3, pages 45–69. TSI Press, San Antonio, TX, USA. http://www.intelligentspacerobotics.com/.

[Maimone et al., 2007] Maimone, M., Cheng, Y., and Matthies, L. (2007). Two years of visual odometry on the Mars Exploration Rovers. *Journal of Field Robotics*, accepted for publication.

[Maimone et al., 2006b] Maimone, M., Johnson, A., Cheng, Y., Willson, R., and Matthies, L. (2006b). *Autonomous Navigation Results from the Mars Exploration Rover (MER) Mission*, volume 21, pages 3–13. Springer Berlin / Heidelberg.

[Maki et al., 2003] Maki, J. N., III, J. F. B., Herkenhoff, K. E., Squyres, S. W., Kiely, A., Klimesh, M., Schwochert, M., Litwin, T., Willson, R., Johnson, A., Maimone, M., Baumgartner, E., Collins, A., Wadsworth, M., Elliot, S. T., Dingizian, A., Brown, D., Hagerott, E. C., Scherr, L., Deen, R., Alexander, D., and Lorre, J. (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 et al., 1998] Mishkin, A., Morrison, J., Nguyen, T., Stone, H., Cooper, B., and Wilcox, B. (1998). Experiences with operations and autonomy of the Mars pathfinder microrover. In *Proceedings of the 1998 IEEE Aerospace Conference*, Snowmass at Aspen, Colorado. http://robotics.jpl.nasa.gov/people/mishkin/papers/ IEEE_aerospace98.pdf.

[Mishkin et al., 2006] 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)*, pages 46 – 53.

[Reeves and Snyder, 2005] Reeves, G. E. and Snyder, J. F. (2005). An overview of the Mars Exploration Rovers flight software. In *IEEE Conference on Systems, Man and Cybernetics*, The Big Island, Hawaii, USA.

[Trebi-Ollennu et al., 2005] Trebi-Ollennu, A., Baumgartner, E., Leger, C., , and Bonitz., R. (2005). Robotic arm in-situ operations for the mars exploration rovers surface mission. In *IEEE Conference on Systems, Man and Cybernetics*, The Big Island, Hawaii, USA.

[Wilcox and Nguyen, 1998] Wilcox, B. and Nguyen, T. (1998). Sojourner on Mars and lessons learned for future planetary missions. In *Proceedings of the 28th ICES International Conference on Environmental Systems*, Danvers, Massachusetts. SAE Society of Automotive Engineers.