Mars Exploration Rover Mobility and Robotic Arm Operational Performance

Edward Tunstel, Mark Maimone, Ashitey Trebi-Ollennu, Jeng Yen, Rich Petras, Reg Willson
NASA Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA, 91109 USA
tunstel@robotics.jpl.nasa.gov

Abstract - *Increased attention has been focused in recent years on human-machine systems, how they are architected, and how they should operate. The purpose of this paper is to describe an actual instance of a practical human-robot system used on a NASA Mars rover mission that has been underway since January 2004 involving daily interaction between humans on Earth and mobile robots on Mars. The emphasis is on the human-robot collaborative arrangement and the performance enabled by mobility and robotic arm software functionality during the first 90 days of the mission. Mobile traverse distance, accuracy, and rate as well as robotic arm operational accuracy achieved by the system is presented.*

Keywords: Mars Exploration Rovers, space robotics, human-robot systems, mobility, performance assessment.

1 Introduction

As part of its Mars Exploration Rovers (MER) mission, NASA landed twin rovers, named *Spirit* and *Opportunity*, on Mars in January of 2004. These rovers are explicitly required to use robotic mobility and manipulator arm positioning functionality to achieve exploration objectives by serving as surrogate robotic field geologists for a science team on Earth. Software functionality that enables these robotic tasks includes wheel motion control and vision-guided autonomous navigation functions of varying complexity for traversing the Martian surface, as well as robotic arm motion control functions for accurate placement of scientific instruments onto rocks and soil. Mobility and robotic arm software runs onboard the rovers' computers to perform various exploration tasks. Tasks are specified in command loads unlinked to the rovers by engineers who plan their daily robotic activities on Earth.

The MER surface operations began in January 2004 when *Spirit* and *Opportunity* egressed their spacecraft landing systems and set all wheels on Martian soil at their respective landing sites on opposite sides of the planet. The planned duration of their *prime missions*, the baseline mission for which they were designed, was 90 Martian days (*sols*). As of this writing, both rovers have far out-lived their projected lifetimes and continue to explore the surface of Mars for over 1.5 years beyond their landing dates. The mission represents one of the longest deployments of field mobile robots in remote natural environments. In addition to establishing a landmark in planetary *in situ* scientific exploration [1, 2], MER represents a new benchmark in field robotics and human-robot systems. As such, it is important to capture and document the rovers' performance to facilitate later comparison of future robotic systems [3, 4] with the state-of-the-art established by this benchmark (as well as to gauge advancement relative to past accomplishments [5]). This is the motivation for this paper and related works that will follow.

The emphasis herein is on the human-robot collaborative arrangement within the MER surface system and the design-/mission-related performance enabled by the robotic software functionality. The scope is limited to the rovers' performance on Mars during their prime missions, or first 90 sols, for which we present mobile traverse distance, accuracy, and rate in addition to robotic arm operational accuracy achieved by the human-robot system. Section 2 sets the context by providing a high-level view of the mission operations system that runs the daily surface exploration. In Section 3, we briefly describe the relevant robotic capabilities of the rovers and related software-based functionality. Section 4 describes the approach used by the mobility and robotic arm engineers on the mission operations team to assess and analyze rover performance. Section 5 presents selected performance results against requirements for *Spirit* and *Opportunity*. Finally, Section 6 provides a discussion and conclusions with a glimpse at future work on rover performance assessment beyond the MER prime missions.

2 MER Surface Operations System

To explore the surface of Mars the rovers work in collaboration with a team of scientists and flight control engineers on Earth who plan and assess performance of daily mission operations for each rover. As a whole, this MER surface operations system represents a distributed human-robot system for *semi-autonomous* planetary surface exploration. The prefix "semi" connotes remote planning, command sequencing and visualization of rover activity sequences and related data products by the Earth-based science-engineering team, all under extreme time delay and intermittent communication afforded by daily uplink and downlink cycles of deep space networks.

The MER tactical mission operations system is a complex system broken down into two teams of human flight controllers on Earth that interact daily with two rover systems on Mars (all with supporting communications infrastructures on Earth and in Mars orbit). Each team of flight controllers includes a team that performs uplink command sequencing and a team that performs downlink telemetry analysis. The uplink and downlink teams are guided through the exploration process by a team of scientists to achieve the mission's science goals and objectives. The uplink team is further broken down into smaller teams that plan rover activities and create the supporting command sequences that govern the rovers' daily activities on the martian surface. The downlink team is further broken down into smaller teams that monitor daily command sequence execution and rover performance as well as ensure safe and proper remote operation of the rovers. These smaller teams are organized by engineering subsystem to cover focused discipline areas represented in the spacecraft and rover design (e.g., power, thermal, telecommunications, attitude control, flight software, etc).

One engineering subsystem of the downlink team is responsible for the health, safety, and performance of the mobility and robotic arm subsystems and, in particular, the related software functionality of *Spirit* and *Opportunity*. Hereafter, we shall refer to this sub-team as the *Mobility Engineers*. This paper focuses on performance of the mobility and robotic arm functional software subsystems as fostered and assessed by the Mobility Engineers in close collaboration with Rover Planners — a subsystem of the uplink team that plans and creates the rovers' mobility and robotic arm command sequences [6-8].

3 Rover Robotic Functionality

Spirit and *Opportunity* are 6-wheeled robots that employ an articulated rocker-bogie mechanical suspension system for rough terrain mobility (Fig. 1). Solar panels and batteries provide power for the rovers. Each rover is equipped with a robotic arm beneath the frontal area of its solar panel that carries a suite of instruments used for *in situ* science investigation of terrain surface materials. The robotic arm is also known as an Instrument Deployment Device since the end-effector is essentially a rotating turret of scientific instruments for field geology that must be accurately positioned near or against rocks and soil to acquire scientific measurements as part of the mission.

In this section we briefly describe the software-controlled functionality of the onboard mobility and navigation system as well as the robotic arm. This software runs on the rovers' computer — a 20 MHz RAD-6000 processor (radiation-hardened version of a PowerPC chip) running the VxWorks real-time operating system, with 128 MB of DRAM and 256 MB flash memory and EEPROM, embedded in a VME chassis. The robotic software functionality enables the sensing, perception, and actuation needed to achieve closed-loop and open-loop motion control and is responsible for the operational robotic system performance.

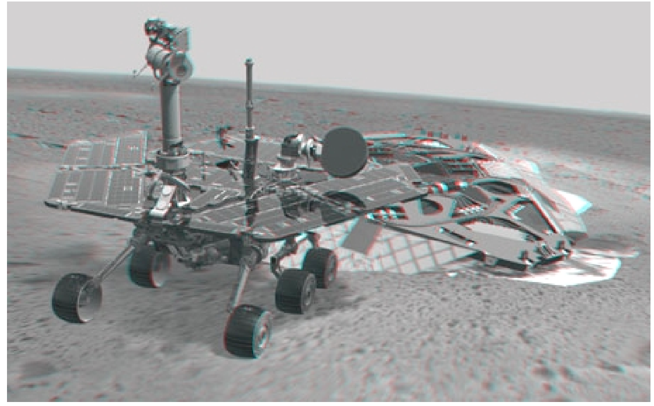

Figure 1. *Spirit* and lander (computer models combined with Mars 3-D surface data acquired by *Spirit's* cameras).

3.1 Mobility

With high torque, all-wheel drive, and double-Ackerman steering the rovers are designed to negotiate rough and rocky terrain. The rocker-bogie suspension endows the mobility system with the capability to traverse such terrain while surmounting rocks of heights as high as one wheel diameter (~ 25 cm) above the ground plane without high-centering or causing significant resulting roll of the vehicle chassis.

Sensing for control of mobility and navigation includes wheel encoders (for dead-reckoned odometry), potentiometers for articulated suspension kinematic state, inertial attitude sensing, celestial (sun) sensing for absolute heading determination, and several stereo camera pairs for navigation. Each rover has body-mounted, front and rear stereo camera pairs used for local terrain hazard detection and avoidance during autonomous navigation as well as stereo cameras for global path planning that are mounted on a fixed mast at a height of about 1.3 meters above the ground plane. Images from these cameras acquired during a traverse are also used on occasion to perform visual odometry. Visual odometry is a means to estimate position changes by estimating displacement of many image features tracked between successive image captures of overlapping scenes [9-10]. This is typically done when wheel odometry is deemed to be highly unreliable due to non-deterministic wheel-terrain interactions in low traction regimes (loose sand/soil, steep slopes, and rocky terrain in which wheels slip on/off of rocks).

Controlled motion is achieved via commands to onboard software functions that exploit the rover sensing and kinematics. The primary functions include three basic driving capabilities for translation and rotation in the plane (while the suspension system conforms to the 3-D terrain topography) as illustrated in Fig. 2. 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 as well as specific Cartesian coordinates of a location to face towards.

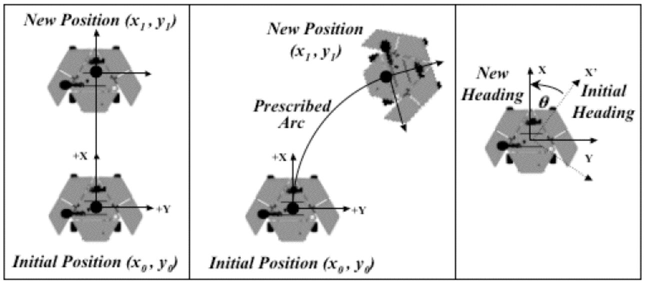

Figure 2. Basic surface mobility functionality.

Autonomous navigation is commanded in terms of desired destinations or explicit surface coordinates to be reached using onboard sensing, perception, and motion planning. Rover Planners provide a global path plan in the form of a series of waypoints. Execution is enabled by onboard hazard (obstacle) detection and avoidance consisting of autonomous selection and execution of incremental paths toward waypoints and, ultimately, navigation goal locations. This process is depicted in Fig. 3. Hazard detection is achieved using passive stereo vision to build and maintain a local terrain map onboard that is used onboard to infer traversability of the local terrain. This traversability map is then used to make automated selection of the best incremental path towards a waypoint or goal while avoiding obstacles. The rover drives until its estimated position is within a specified radial distance (tolerance) of its commanded goal location, or until a specified timeout period expires. Underlying details of the algorithm can be found in [11].

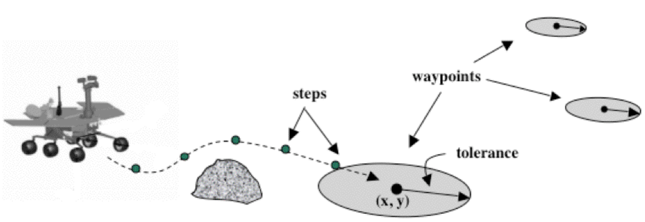

Figure 3. Autonomous surface navigation functionality.

Command sequences for traverses on the martian surface by *Spirit* and *Opportunity* have used the following types of mobility commands or a combination of them all: (a) basic mobility commands alone (*manual* driving without hazard avoidance enabled); (b) basic mobility commands 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). 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.

3.2 Robotic Arm

The robotic arm (Fig. 4) has five revolute degrees-of-freedom supporting its deployment (from a nominally stowed configuration) and 3-D fine positioning required to achieve accurate instrument placement onto rocks and soil. The instrument arm includes a microscopic imager to capture extreme close-up images, a Mössbauer spectrometer to detect composition and abundance of iron-bearing minerals, an Alpha-Particle-X-Ray Spectrometer to determine the elemental chemistry of surface materials, and a Rock Abrasion Tool for exposing fresh material beneath dusty or weathered rock surface layers via controlled-force loading and physical abrasive action. The arm is also used to position the spectrometers for physical placement on an instrument calibration target and science-related magnets mounted at different locations on the rover body.

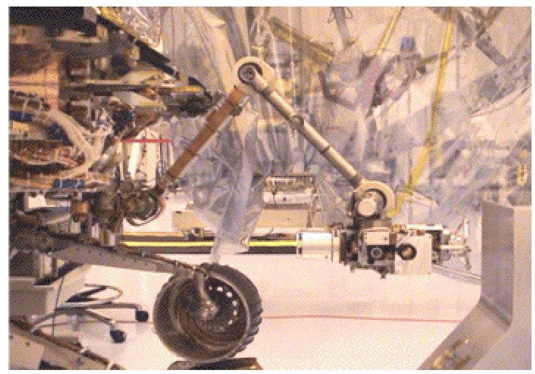

Figure 4. Robotic arm (Instrument Deployment Device)

Sensing for control of the robotic arm kinematic configuration includes joint angle encoders primarily, while controlled contact of the end effector is facilitated by redundant sets of contact sensors on each instrument. The contact sensors provide tactile feedback used by the software to halt arm motion upon expected or unexpected contact with parts of the terrain or the rover. Joint angle set points for achieving the arm configurations to reach locations of a body-mounted calibration target and magnets are taught manually before the rovers' launch from Earth. Several additional robotic arm configurations are taught prior to launch that are frequently used in its operations sequences. In order to reach arbitrary coordinates within the kinematic work volume of the robotic arm, the onboard software provides a variety of functions. A subset of these includes forward and inverse kinematics, straight-line end-effector trajectory generation, and model-based self- and rover-collision prediction. The frontal pair of body-mounted stereo cameras used for terrain hazard detection, is also used to specify 3-D target positions at which to place any of the arm-mounted instruments. As such, the absolute positioning accuracy of the arm when reaching stereo image-designated targets is a partial function of any stereo ranging errors associated with these cameras.

Controlled motion of the robotic arm and instruments is achieved via commands to onboard software functions that enable a set of gross and fine motions. Several modes of motion are possible including free-space motion for unintended contact movements, guarded motion for intended contact, pre-load motion for force-controlled contact against surfaces, and retracting motion away from reached targets after scientific measurements are acquired [8]. In addition, software control functions to unstow the arm, change tools (instruments), and move the arm back into its stowed configuration when not in use are provided.

4 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. Within the complex MER surface mission operations system, the Mobility Engineers, Rover Planners, and the rovers form a closed-loop human-robot control system (notwithstanding the NASA's Deep Space Network and supporting teams and systems beyond the scope of this paper). 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. Mobility Engineers effectively function in the feedback loop of the human-robot system (Fig. 5) 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.

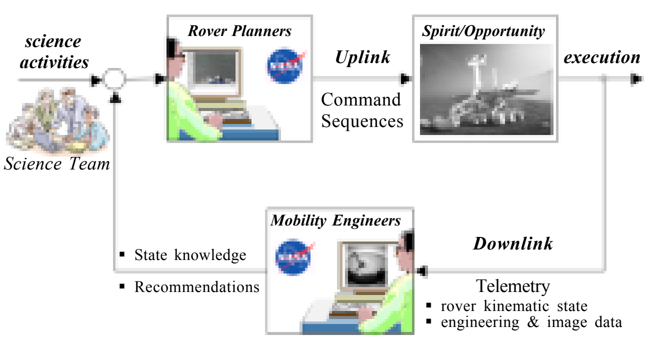

Figure 5. Simplified human-robot control system for remote mobility and robotic arm surface operations.

Within this closed-loop human-robot system the science instrument, image, and engineering data telemetered to Earth on a given sol, from either rover, drive its exploration plan for the next sol. Typically, the next sol's planning cannot begin without certain critical data and necessary rover state knowledge representing the last known state of the rover at the termination of the previous sol's activity. 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 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 [6, 7]. 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 daily and executes uplinked sequences throughout a period of 3-6 hours around local noon (with occasional nighttime communications/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.

4.1 Rover state determination

Each communications session in which *Spirit* or *Opportunity* transmits data received at Earth includes instrument, image, and/or engineering data acquired or periodically sampled and logged by the rover computer during execution of command sequences. The MER Ground Data System (GDS) delivers these data in various forms to computer file systems, monitors, and mission operations consoles. The GDS is a comprehensive collection of computer systems, scientific data and information and, in particular, software tools that aid decision-making for mission operations. Particular data of interest to Mobility Engineers includes many items of health and state data displayed on or accessible via a common console. These data typically represent the final state values of raw telemetry from the rovers' sensors and derived data produced by the onboard software. Typical key rover state data include (among much more): estimated position with respect to an established local coordinate frame on the martian surface, attitude with respect to the local Mars gravity vector, heading with respect to true north (as measured by celestial sensing), positions of wheel and robotic arm actuators, suspension articulations, and instrument contact switch states.

In addition to final state values, telemetry also delivers histories of the same states mentioned above, and others, which enable playbacks of state evolution during executed robotic motions via graphical and animated visualization tools [12]. Stereo image browsing, processing, and analysis tools are used to assess downlinked images acquired by the rover as well as generate 3-D synthetic virtual representations the rover surroundings as captured in actual digital imagery [12]. The GDS software tools feature several means to immerse articulated rover models within the photo-realistic synthetic terrains to yield high-fidelity 3-D representations of kinematic motions executed by the rover on the terrain. These same tools are used by Rover Planners to plan rover activities, and they feature a means to simulate rover predicted behavior and final state upon animated execution of planned motion sequences. This facility enables Rover Planners to provide predicted states that the rover should go through and terminate at for each mobility and robotic arm sequence executed on Mars. State predictions are utilized by Mobility Engineers as references to compare against the actual state values after execution and upon receipt in the communications downlink on Earth. As such, we are able to assess actual performance of daily activities relative to expected performance as predicted by command sequences and their simulations. This provides a basis for computing position estimation errors for rover traverses and robotic arm positioning errors associated with instrument placements.

4.2 Rover telemetry downlink assessment

Health and kinematic state determination is an integral part of the surface tactical downlink assessment necessary to understand the operational readiness of the rovers for the next sol. Additional tasks are allocated to the Mobility Engineers who operate in the feedback loop of the human-robot system. Thus far we have mentioned the careful examination of the latest received raw and derived data as well as motion history data for consistency with expectations and predictions. This activity establishes the basic rover health, performance, and readiness for subsequent mobility and robotic arm activities in the absence of any fault or failure conditions. In the event of any mobility, navigation, or robotic arm sequence execution errors, Mobility Engineers diagnose the cause(s) and prescribe the corrective action(s) to be commanded in the next uplink. Any serious anomaly in sequence execution prompts the most thorough analyses and may lead to prescription of operational constraints or restrictions on how the mobility or robotic arm functionality is utilized until full resolution of the anomaly.

To gain a richer understanding of rover performance of the most recently executed sequences, Mobility Engineers utilize the GDS image browsing, graphics, and visualization tools. These tools facilitate detailed analysis of the execution of each gross mobility and robotic arm motion in the commanded sequence(s) — essentially providing a replay of what the rover actually did on Mars. In this case, a 3-D model of the rover is immersed within synthetic terrain renderings or superimposed atop stereo images of the local surrounding terrain [12]. Mobility Engineers make further use of these tools to assess traversability of the surrounding terrain while considering rover mobility system capabilities and/or operational constraints. In a like manner, tool features are employed to assess the kinematic ability of the robotic arm to reach scientifically interesting rock or soil targets within its dexterous work volume.

On occasion, due to non-deterministic wheel-terrain interactions that are not modeled by the kinematics-based rover sequence simulators, actual mobility performance may produce markedly different state outcomes relative to the state predictions (primarily when onboard visual odometry is not enabled). If rover reported position after a traverse is inaccurate (as discerned from rover imagery or other means) and not within the prescribed position error tolerance, Mobility Engineers utilize all available data to derive a best estimate of the rover state. This often involves deriving a better localized position estimate than the onboard position estimate reported in telemetry from the rover based on available sensor data and feature matching between images acquired prior and after the traverse. Rover reported state knowledge might thus be augmented by derived knowledge resulting from further analysis by humans. Best knowledge of rover position and other key state information mentioned above is fed back to Rover Planners and the uplink team, along with any resulting engineering recommendations for subsequent use of the mobility and robotic arm subsystems. These final conditions for sol n become the critical initial conditions required to begin planning of activities for sol $n+1$.

The approach described above is a working human-robot system model in which humans apply perception, localization and navigation knowledge to supplement the rovers' capabilities and limited onboard intelligence to achieve mission objectives. This collaboration is enabled through the use of sophisticated ground-based software tools that help to bridge the gap between robot and human perceptions. Beyond this tactical operations approach, Mobility Engineers perform strategic analysis and trending of performance changes over multiple sols with respect to flight hardware condition and tunable flight software parameter sets. The intent is to identify and understand system performance in order to maintain operability and predictability as well as maximize performance throughout the mission given the evolving hardware condition and operational constraints. Over many sols, the same infrastructure, processes, and procedures that enable daily performance assessment also enable performance assessment and determination or prognosis of trends in rover health and system behavior. In the end, we can establish overall rover performance with respect to design and mission requirements.

5 Performance on Martian Surface

The *Spirit* and *Opportunity* rovers were designed to meet a collection of low-level flight system and software design requirements, many of which were derived from high-level mission requirements. In general, surface mobility and navigation software capabilities were required that would enable autonomous mobility, including 3-D stereo range mapping, hazard detection, local avoidance path selection, and position estimation while satisfying certain performance related requirements. In brief, the rovers had to be able to safely navigate at a low average rate of 35 m/hr in autonomous mode in rocky terrain (~7% rock abundance) 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). The required rate of autonomous traverse was derived from an early desire to be able to navigate 100 m per sol of operations. The primary constraint on achievable distance is the available time for driving (versus other activities), which is limited by combined thermal, power, communications, and science activity constraints. 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.

Robotic arm software capabilities were required that would enable fine positioning and control as well as accurate and repeatable placement of science instruments onto targets of interest. Each rover's robotic arm had to be capable of positioning each instrument to within 10 mm of a science target with position repeatability of +/- 4 mm. The flight system design requirements were verified and validated by testing prior to launch of each rover. Below we briefly summarize each prime mission and report on selected items related to performance requirements as achieved during the rovers' separate missions on Mars.

5.1 Spirit in Gusev Crater

Spirit's surface mission began when she was commanded to egress from her spacecraft lander on sol 12. On sol 1, the lander bounced and rolled to a stop within its target landing ellipse on the floor of Gusev Crater, a large crater of 160 km in diameter. The surrounding terrain was somewhat rocky and flat with most prominent features being 100 meter high hills over 2.5 km away dubbed the Columbia Hills complex. While traversing to explore the landing site, *Spirit* stopped frequently to perform *in situ* science using its robotic arm to place instruments on scientist-selected rocks and soil targets. She visited a large crater dubbed Bonneville Crater encountering significant sloped, rocky, and sandy terrain while climbing towards the crater rim. *Spirit's* next long range goal was the Columbia Hills towards which she encountered increasingly rough and undulated terrain (Fig. 6).

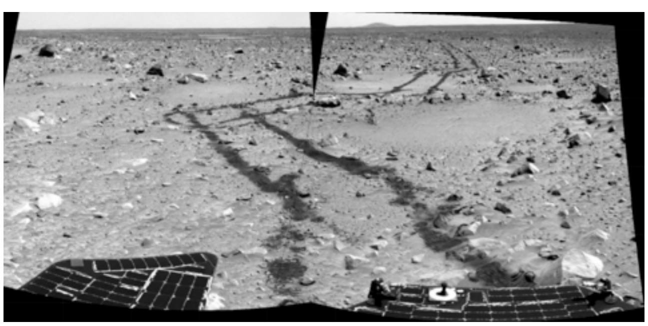

Figure 6. *Spirit* wheel tracks in Gusev Crater (solar panel in view at bottom of navigation camera image mosaic).

In order to cover significant distance safely during the long 2 km journey the Rover Planners and Mobility Engineers used the following strategy. The initial segments of a sol's drive would be sequenced using manual and guarded driving modes out to distances beyond which the available pre-drive range data accuracy was insufficient for designating global path waypoints. These distances were typically tens of meters. After reaching such "sensor horizons" the autonomous navigation mode would take over and drive through previously unseen terrain while keeping the rover safe. Spirit's prime mission ended en route to Columbia Hills (later reached during her extended mission). Fig. 7 depicts her total traverse from the lander in 2-D; dark-shaded segments of the graph correspond to manual driving and light-shaded to autonomous driving.

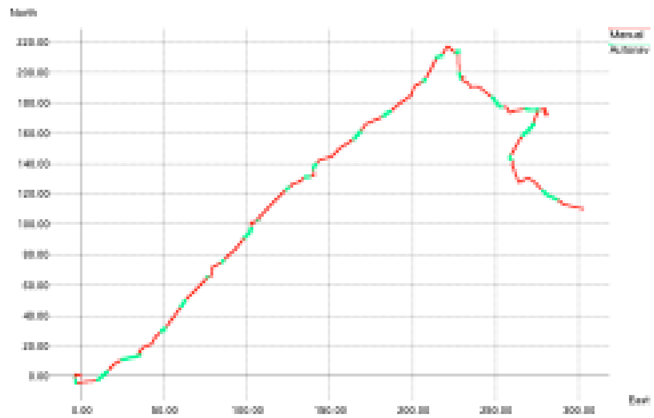

Figure 7. *Spirit's* prime mission total traverse (meters).

The total integrated distance traveled by *Spirit* during her prime mission was 635 meters based on onboard odometry. Of that distance, 73% was driven in manual driving mode and 27% with hazard detection enabled, i.e., guarded and autonomous driving modes. The best ground truth measurements for true distance traversed are derived from global landmark-based and radiometric localizations performed by science team members and spacecraft navigation engineers, respectively [13]. They performed triangulation on landmarks visible in images acquired by NASA spacecraft in Mars orbit and images taken by the MER spacecraft during landing, together with a formal least-squares bundle adjustment method using features in overlapping rover images, along with rover mobility and navigation sensor data. They also performed localization of the rover on the martian surface using two-way Doppler radio transmissions between Earth and Mars [13]. The ground truth of *Sprit's* integrated traverse distance, thus derived, reveals an accumulated difference between odometric and localized positions of 2% of total distance traveled [13].

The flexibility of the autonomous navigation software design allowed for constructive parameter adjustments to be sequenced by Rover Planners to improve the rover performance as needed. As an example, on certain sols early in the mission deep shadows cast by the rover caused problems for the navigation functionality due to shadows being interpreted as terrain hazards after onboard image processing. The design enabled parameter adjustments to be made that had the effect of correcting for shadows such that they would no longer be misinterpreted. The top speed achieved by *Spirit* during autonomous navigation was 34.9 m/hr. The Gusev Crater landing site was characterized to have 8% rock abundance.

5.2 Opportunity on Meridiani Planum

Opportunity's surface mission began when she was commanded to egress from her spacecraft lander on sol 7. On sol 1, the lander bounced and rolled to a stop on the floor of a small crater 20 meters in diameter, later dubbed Eagle crater. Her first 57 sols of mobility and robotic arm operation were spent exploring high-priority science targets within Eagle crater making extensive use of the robotic arm to position instruments for *in situ* science. During that period the rover traversed approximately 160 meters mostly up, down, and across crater walls (Fig. 8) of soil-covered slopes up to 20° and encountering wheel slippage up to 20% on slopes of up to 10°. Rover Planners and Mobility Engineers gained valuable experience with better compensating for wheel slippage using slip-predictive planning based on results of slope traversal tests performed on Earth and estimated slips from earlier sols within Eagle Crater on Mars. However, in one instance upon a first upslope attempt to exit the crater on sol 56 *Opportunity* encountered 100% slip on a soil slope of about 17° despite this. On the following sol, Rover Planners sequenced a cross-slope and up-slope traverse to successfully egress Eagle Crater having completed an exploration campaign within the crater that yielded conclusive evidence that liquid water flowed on the surface in that region.

Figure 8. *Opportunity* tracks in Eagle Crater (rear hazard camera view with lander in background on crater floor).

Outside of Eagle Crater is a flat and smooth expanse of barren landscape covered with dark-colored, fine-grained soil without many apparent mobility hazards except sparsely distributed terrain depressions and craters, both smaller and larger than Eagle Crater. Opportunity spent the rest of her prime mission driving and using the robotic arm to explore scientifically interesting rocks, soils, and features on the terrain of Meridiani Planum. This was done en route to a much larger crater later dubbed Endurance Crater. Her prime mission ended about 200 meters shy of that crater (later explored extensively in her extended mission). Fig. 9 depicts *Opportunity's* total traverse from the lander in 2-D; dark-shaded segments of the graph correspond to manual driving and light-shaded to autonomous driving. Darkest and smallest segments of the graph correspond to driving using visual odometry.

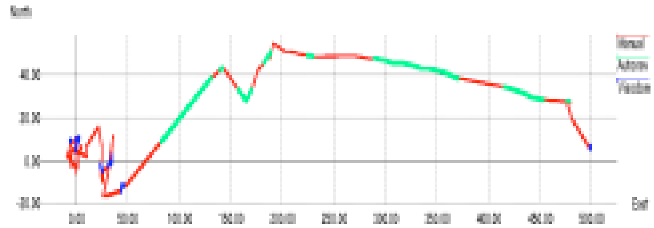

Figure 9. *Opportunity's* prime mission traverse (meters).

The total integrated distance traveled by *Opportunity* during her prime mission was 772 meters based on odometry. Of that distance, 68% was driven in manual driving mode and 28% in guarded and autonomous driving modes; the remaining 4% of traverse distance was driven with visual odometry enabled. The highly concentrated portion of the graph in Fig. 9 corresponds to mobility within Eagle Crater, a subset of tracks of which are apparent in Fig. 8. The ground truth of *Opportunity's* integrated traverse distance reveals an accumulated difference between odometric and localized positions of just over 20% of total distance traveled [13]. This is due almost entirely to the significant accumulated wheel slippages encountered within Eagle Crater during the first two-thirds of the prime mission. Exploitation of the autonomous navigation software design flexibility, in *Opportunity's* case, was driven by an early realization that images of the terrain captured by the body-mounted hazard cameras (with 128x128 pixel nominal image resolution) lacked sufficient texture needed to achieve good stereo correlation and produce useful 3-D range maps for hazard detection. Instead, Mobility Engineers prescribed necessary adjustments to permit the rover to use its mast-mounted navigation cameras (at a resolution of 256x256) for hazard detection. The navigation cameras were used to better process images of the smooth and nearly featureless terrain at that landing site. Even though their field of view is significantly narrower than that of the hazard cameras, the navigation algorithm was still able to perform its intended functions. The top speed achieved by *Opportunity* during autonomous navigation was 35.6 m/hr. The Meridiani Planum landing site was characterized to have a rock abundance of only a few percent.

In summary, the *Spirit* and *Opportunity* rovers performed well throughout their respective prime missions. Over the course of their first 90 sols, there were 5 and 2 sols, respectively, on which *Spirit* and *Opportunity* reported mobility/navigation software errors. These errors were each due, to some extent, to some form of human error; that is, command sequence or sequencing errors, unaccounted for system operational behaviors not experienced during Earth testing, or isolated shortfalls of the planning tools or processes. All of these were later rectified as lessons learned. Whenever enabled, the autonomous navigation software kept both rovers safe from terrain hazards while making progress toward commanded goals. For robotic arm operations, no functional anomalies or faults were encountered on either rover that could be attributed to the robotic arm software. The arm on both rovers operated well within requirements and performed extremely well throughout the prime mission. Instrument positioning and placement operations by the robotic arms were achieved within design and mission performance requirements. Positioning errors were consistently within 1 mm and instrument contact-placement errors were in the 1 cm range.

6 Conclusions

The long-term deployment of robotic systems that rely on human guidance for successful scientific exploration, as exemplified by the MER mission, is a significant achievement. The human-robot collaboration experiences gained are definitive of the state of the art in human-robot systems. The MER experience is a working example and a data point against which ideas for advancements in human-robot systems can be compared and contrasted. Descriptions of such actual human-robot systems that work effectively in real applications are needed contributions to the literature, and will be instrumental in further advancing the state of the art. As of this writing, the rovers continue to perform well as surrogate explorers and robotic geologists on our behalf. Working in concert with a functionally diverse team of flight controllers/mission operators on Earth, this complex human-robot system has established a benchmark in robotic planetary surface exploration missions and human-robot collaboration. It enabled exploration of different terrain regions and acquisition of key measurements from which new scientific knowledge was gained [1, 2].

Both rovers have far out-lived their projected 90-sol mission duration and continue to explore for over 1.5 years beyond their landing dates. *Spirit* is climbing and exploring a hill complex in Gusev Crater of older rock than she traversed during her prime mission, which may hold evidence of an ancient body of water thought to have once filled the crater. Since the end of her prime mission *Spirit* has traversed over 4 km. *Opportunity* continues to explore the sparse distribution of rocks and craters on the open plains of Meridiani Planum. Since the end of her prime mission she has traversed over 4.6 km. Future work will consist of rating their extended-mission performances to those of the prime missions using metrics to be formulated. This will allow effective comparison of baseline and improved software loads later uplinked to the rovers that provided enhanced mobility and sequencing. We will also apply metrics to compare current and future prototype and flight rovers using MER as a reference.

Acknowledgments

The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Engineering contribution and early leadership from Diana Darus of JPL is acknowledged.

References

[1] *Science*, Special Issue: Spirit at Gusev Crater, v. 305, 5685, AAAS, pp. 793-845, Aug. 2004.

[2] *Science*, Special Issue: Opportunity at Meridiani Planum, v. 306, 5702, AAAS, pp. 1697-1756, Dec. 2004.

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

[4] P. Schenker, T.L. Huntsberger, P. Pirjanian, E.T. Baumgartner, E. Tunstel, "Planetary rover developments supporting Mars exploration, sample return and future human-robotic colonization," *Autonomous Robots*, Vol. 14, pp. 103-126, 2003.

[5] A. Mishkin, et al., "Experiences with operations and autonomy of the Mars Pathfinder Microrover," Proc. IEEE Aerospace Conference, Aspen, CO, March 1998.

[6] C. Leger, et al, "Mars Exploration Rover Surface Operations: Driving Spirit at Gusev Crater," 2005 Proc. IEEE Intl. Conf. on Systems, Man, and Cybernetics, Waikoloa, HI, this proceedings.

[7] J. Biesiadecki, et al, "Mars Exploration Rover Surface Operations: Opportunity at Meridiani Planum," 2005 Proc. IEEE Intl. Conf. on Systems, Man, and Cybernetics, Waikoloa, HI, this proceedings.

[8] A. Trebi-Ollennu, C. Leger, E.T. Baumgartner, R.G. Bonitz, "Robotic arm in-situ operations for the Mars Exploration Rovers surface mission," 2005 Proc. IEEE Intl. Conf. on Systems, Man, and Cybernetics, Waikoloa, HI, this proceedings.

[9] C.F. Olson, L.H. Matthies, M. Schoppers, M.W. Maimone, "Robust stereo ego-motion for long distance navigation," Proc. IEEE Intl. Conf. Computer Vision and Pattern Recognition, Hilton Head, SC, pp. 453-458, 2000.

[10] Y. Cheng, M.W. Maimone and L. Matthies, "Visual Odometry on the Mars Exploration Rovers," 2005 Proc. IEEE Intl. Conf. on Systems, Man, and Cybernetics, Waikoloa, HI, this proceedings.

[11] S.B. Goldberg, M.W. Maimone, L. Matthies, "Stereo vision and rover navigation software for planetary exploration," Proc. IEEE Aero. Conf, Big Sky, MT, 2002.

[12] F. Hartman, B. Cooper, S. Maxwell and J. Wright, "Data visualization for effective rover sequencing," 2005 Proc. IEEE Intl. Conf. on Systems, Man, and Cybernetics, Waikoloa, HI, this proceedings.

[13] R. Li, et al., "Initial results of rover localization and topographic mapping for the 2003 Mars Exploration Rover mission," *Journal of Photogrammetric Engineering & Remote Sensing*, in press.