SPACE 2006

Human-Robot Site Survey and Sampling for Space Exploration

Terrence Fong*, Maria Bualat, Laurence Edwards,
Lorenzo Flückiger, Clayton Kunz, Susan Y. Lee, Eric Park, Vinh To,
Hans Utz, Nir Ackner, Nicholas Armstrong-Crews, and Joseph Gannon
Intelligent Systems Division
NASA Ames Research Center, Moffett Field, CA 94035, USA

NASA is planning to send humans and robots back to the Moon before 2020. In order for extended missions to be productive, high quality maps of lunar terrain and resources are required. Although orbital images can provide much information, many features (local topography, resources, etc) will have to be characterized directly on the surface. To address this need, we are developing a system to perform site survey and sampling. The system includes multiple robots and humans operating in a variety of team configurations, coordinated via peer-to-peer human-robot interaction. In this paper, we present our system design and describe planned field tests.

I. Introduction

A CENTRAL concept of NASA's Vision for Space Exploration is that mission activities must be sustainable over the long-term.[1] To achieve this, space exploration systems must be affordable, reliable, and effective. In particular, human-robot teams must be able to operate effectively and safely on planetary surfaces. A key aspect of this will be for humans and robots to efficiently map and characterize sites of operational and scientific interest.

Our objective is to develop and demonstrate tools and techniques to support human-robot site survey and sampling. Specifically, we are developing methods that combine information from orbital and descent imagery with surface sampling by humans and robots. Two key topics are being addressed: (1) techniques for robot teams to perform resource mapping using a variety of instruments and (2) techniques to enable effective human-robot interaction for a range of team configurations and interfaces in order to improve survey performance and effectiveness.

With our approach, robotic survey tasks can be coordinated from ground-control (for pre-cursor exploration missions), as well as from inside surface habitats or nearby worksites (for short-term stay missions). A typical work scenario involves multiple survey robots mapping a region for resources while human operators assess reported finds and provide support (physical and cognitive intervention). Coordination and dialogue between ground control, crew (both EVA and IVA), and mobile robots is performed through peer-to-peer human-robot interaction.

*terrence.w.fong@nasa.gov

II. Survey Architecture

Our survey architecture is shown in Figure 1. The architecture supports three processing phases: preparation, execution, and analysis. In the preparation phase, we perform terrain modeling by first collecting stereo imagery (from aerial flyover or satellite) of the survey terrain. The stereo images are then used to create a digital elevation map (DEM), comprising a grid of cells each with an elevation value. Next, a traversability analysis is performed to determine hazardous terrain for each survey robot.

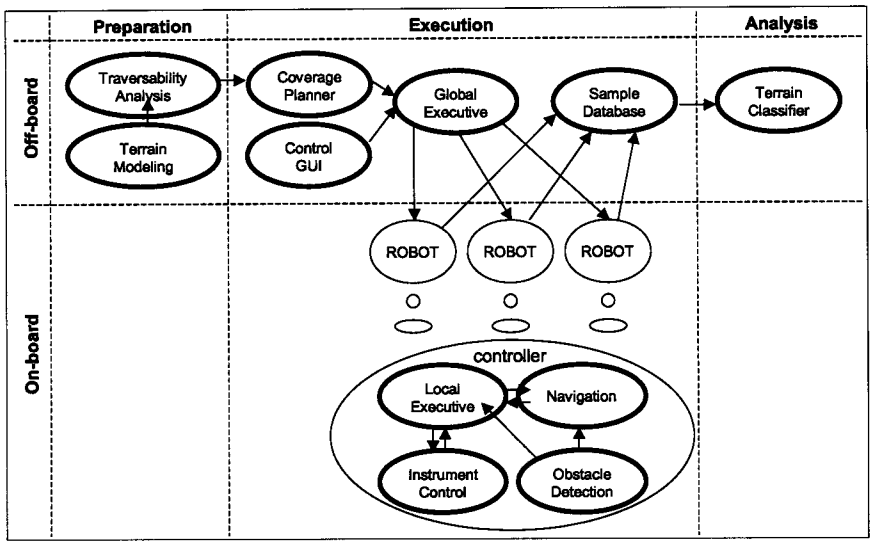

Figure 1. System architecture

During the execution phase, processing occurs off-board and on-board multiple survey robots. The traversability map is processed by a coverage planner, which dynamically computes survey points as the robots acquire them. A central, global executive coordinates task assignment and monitors task execution. The data acquired by each robot is collected into a central sample database for post-processing.

In the final phase, the acquired data are analyzed to characterize the types and distribution of resources in the survey site. Because the suite of survey instruments and number of survey robots will depend on site-specific requirements (e.g., resources being mapped), the architecture does not constrain how analysis is to be performed, nor does it attempt to provide a data fusion framework.

Figure 2 shows the flow of data through the system: from the initial acquisition of stereo terrain images, through terrain modeling and traversability analysis, to coverage planning and task assignment to multiple (heterogeneous) robots.

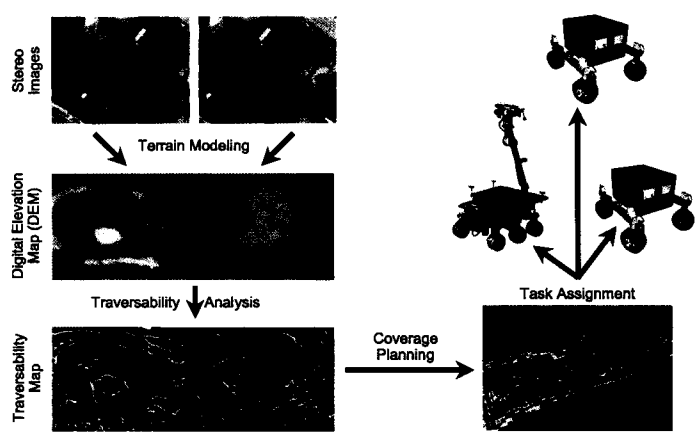

Figure 2. Dataflow

One challenge is that the time required

for a robot to acquire a measurement varies from survey point to survey point. This is due primarily to incomplete map information, i.e., the resolution of the traversability map is limited. Thus the robot's navigation and instrument control algorithms must adaptively avoid obstacles, select safe locations for instrument placement, etc. Moreover, it is not obvious how to plan against robots crossing paths, or otherwise interfering with each other, because it is not possible to always predict where the robots will be at a given point in time.

For this reason, we have chosen to plan for each robot independently, resolving interference at run-time via the global executive. In addition, we use a simple scheme for collision resolution: if two robots detect that they are getting to close to each other (via stereo vision or localization), the one with a lower, pre-assigned priority will yield the right of way.

III. Terrain Modeling

We compute digital elevation maps (DEM) of terrain using the Ames Stereo Pipeline (ASP), which we have been developing since 1998.[2] The ASP uses computer vision to reconstruct 3D models from a variety of stereo image data sources (rover mounted cameras, orbital imagers, etc) and camera types (both area and line scan imagers).

The ASP employs a fast image-based correlation method to compute positional disparities between corresponding points in left and right images of a stereo pair. From the image disparities, 3D points are calculated using camera geometry. The 3D points are then used to create a 3D triangle mesh, which can be output in a variety of file formats. In addition, various pre- and post-processing modules in the pipeline can be used to improve the mesh (e.g., NURBS-based surface interpolation for noisy image regions).

The resulting DEM's are of high quality: they are generally free of stereo correlation artifacts and noise. Moreover, in many cases, the DEM's can be automatically validated against other elevation information (e.g., laser altimeter readings), or ground truth models, when that data is available for comparison. Figure 3 shows a DEM constructed from a pair of overhead images.

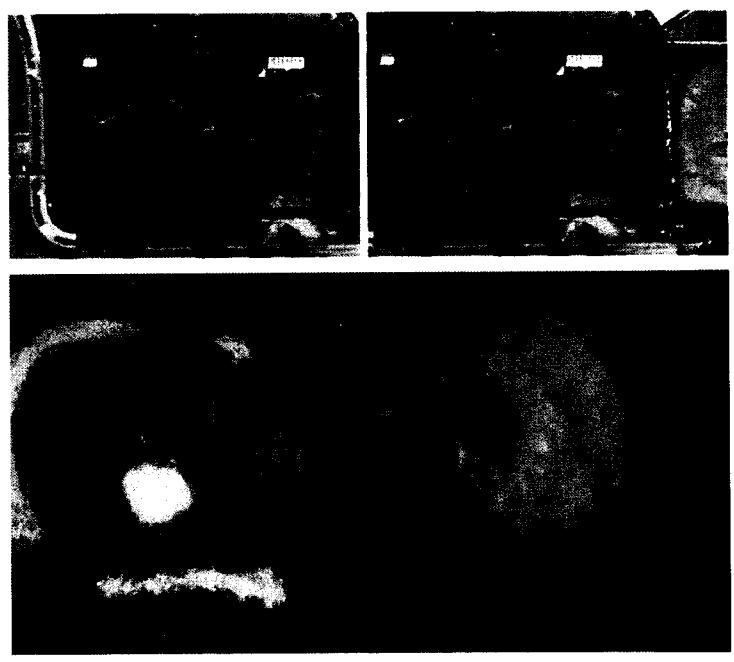

Figure 3. Stereo images (top) are used to construct a Digital Elevation Map (bottom). White points on the DEM are higher elevation.

IV. Traversability Analysis

A survey rover needs to be able to navigate a survey site safely and reliably. In particular, the rover has to avoid regions that are hazardous to traverse due to slope, obstacles, depressions, etc. For navigation purposes, therefore, we need to analyze the traversability of the site, assessing how easy it is to drive through different locations. This information is useful both for local obstacle avoidance and global path planning.

We perform local traversability analysis of DEM's using the Morphin[3] algorithm. Morphin operates by generating statistical metrics of terrain. Map cells are combined into overlapping robot-sized patches and traversability is determined by computing

three metrics: slope (roll and pitch), roughness, and "step height". Morphin estimates slope using a least-squares method to fit a plane to the elevation data points that cover each patch. Roughness is the residual of the plane fit. The "goodness" of a patch is then determined by taking the worst of the (normalized) values in comparison to rover-specific safety parameters. Goodness indicates how easy it would be for the rover to drive through a particular cell. Table 1 lists the parameters that we use for the NASA Ames K9 and K10 rovers.

Table 1. Traversability analysis parameters for NASA Ames rovers.

Parameter	Description	K9	K10
min_goodness	The minimum acceptable goodness level (normalized)	0.5	0.5
fav_goodness	The minimum desirable goodness level (normalized)	0.75	0.75
good_step height	Maximum desirable step height within a cell	0.05 m	0.07 m
bad_step height	Maximum acceptable step height within a cell	0.3 m	0.3 m
good_pitch	Maximum desirable pitch	3 deg	5 deg
bad_pitch	Maximum pitch to maintain rover stability	5 deg	10 deg
good_roll	Maximum desirable roll	5 deg	5 deg
bad_roll	Maximum roll to maintain rover stability	10 deg	10 deg

Figure 4 shows the traversability map produced by Morphin using the DEM from Figure 3. The map is color coded as follows: green indicates regions easy to traverse, yellow indicates regions of moderate difficulty and red indicates regions to be avoided.

Figure 4. Traversability map of survey site.

V. Coverage Planning

Survey coverage planning involves dividing a site into regions (e.g., a regular grid of uniformly sized cells) and determining the order to visit each region. Common methods include: line transects (the site is traversed in a series of parallel lines); zig-zag coverage (similar to line transects but the lines are not parallel); and Morse decomposition (the site is divided into small sections that can be covered trivially).

In our work, the basic task is to acquire as much information about a site as possible in a given period. When there is sufficient time to sample each "good" cell in the traversability map, we use a "full coverage" algorithm to plan paths that systematically sweep the entire map. When the allotted time is insufficient (e.g., due to resource or mission constraints), we employ a "partial coverage" planner to choose a prioritized set of sample points that maximize survey coverage in that period.

A. Full Coverage Planner

We define a "full coverage path" as one that allows a robot to survey all traversable regions in a systematic and efficient manner. To compute such paths, we use the "path transform" method, which is a grid-based path planning approach.[4] In the path transform approach, a wave front (a weighted sum of the distance from the goal and a measure of the discomfort of moving too close to obstacles) is propagated through the region to be covered. The resulting path forms contour patterns, which slope towards the goal and which

traces the shape of obstacles.

Figure 5 shows a full coverage path for the traversability map from Figure 4. One advantage of the path transform over potential fields is that it does not suffer from local minima. In addition, the method can accommodate multiple cost functions, which allow the production of "adventurous" and "conservative" paths in addition to the "optimized" (i.e., shortest distance) behavior.

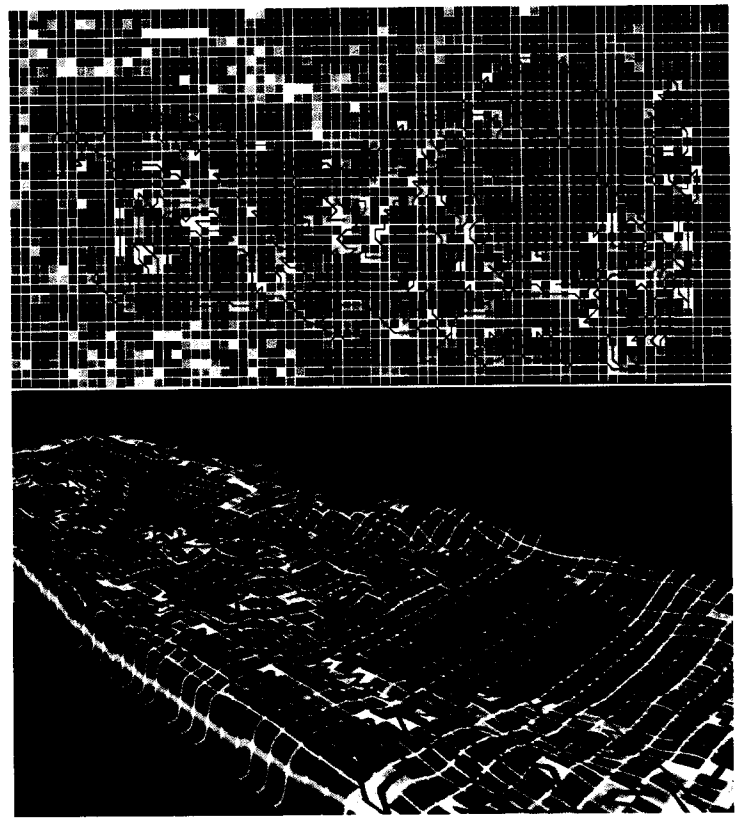

Figure 5. Full coverage path (top), overlay on DEM (bottom). The circles indicate start and end (goal) points.

B. Partial Coverage Planner

We define a "partial coverage plan" as set of sample points that maximizes coverage for a time-limited survey. To plan for partial coverage, we consider three factors. First, the number of possible sample points depends on their locations and order, since the robot must drive from point to point.

Second, in the absence of *a priori* information, we assume that survey samples are spatially correlated: a sample at one point is indicative of surrounding (nearby) points, but less so for distant points. Thus, for good coverage, we want sample points to be distributed and for all cells to be near at least one point.

Finally, to order the points, we consider the relative value and priority of each survey point. Specifically, we want to ensure that if a robot is not able to visit all the planned points (e.g., the robot fails in the middle of executing the survey) that we will have acquired the maximum information possible to that point in time.

Given a set of cells $c \in C$ in a m x n grid and sample points $s \in S \subseteq C$, we define "coverage" as:

$$coverage(C, S) = 1 - \frac{1}{k} \sum_{c \in C} \min_{s \in S}(distance(c, s))$$

where k is a normalizing constant:

$$k = n \left[(n-1)(\frac{2n-1}{3} + \frac{1}{2}) + (m-n)(\frac{m+n-1}{2}) + m \right]$$

and *distance* is the number of moves in an 9-connected grid. If each cell has a sample point, then the minimum distance from each cell to a sample point is 0 and $coverage = 1$. However, if there are no sample points, we assume that a sample exists just outside the grid (i.e., in the corner), leaving the sum of minimum distances to be exactly k and $coverage = 0$.

Our algorithm for partial coverage planning works then is:

1. Compute the "brushfire" distance from the robot to all traversable cells in the traversability map.

2. Choose the cell that maximizes sample utility.

3. Starting from the goal point, compute and return the shortest path by repeatedly choosing the neighboring cell of least brushfire distance (until we reach the robot's current position).

The sample utility for a set of cells $c \in C$ comprising a m x n grid ($m > n$) and sample points $s \in S \subseteq C$ at time t is defined as:

$$utility(c, t, S) = (1 - \lambda(t))reward(c, S) - \lambda(t)cost(c, t)$$

where *reward* is the mean distance to sample points of the cell:

$$reward(c, S) = \frac{1}{|S|} \sum_{s \in S} distance(s, c)$$

and *cost* is the brushfire distance from the robot's current position to the cell. The weighting function $\lambda(t)$ is any function whose output is in the interval $[0, 1]$. Its purpose is to define the trade-off between moving further to get a better sample, or staying close and getting a less useful sample.

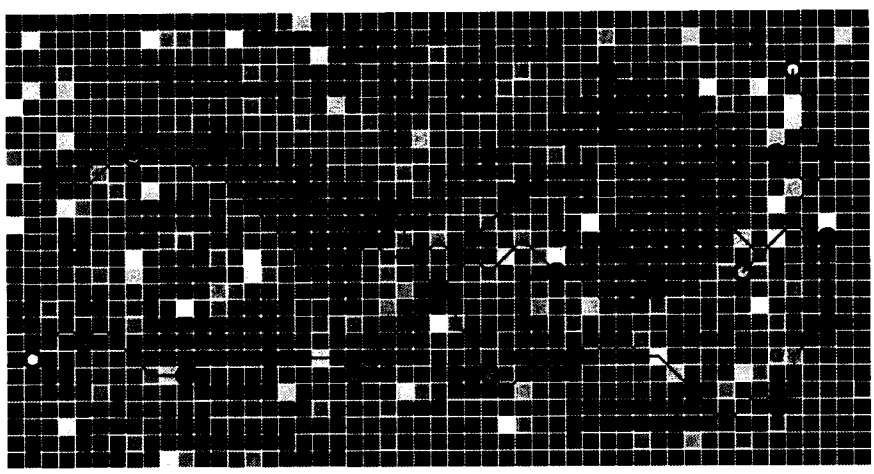

Figure 6. Partial coverage plan.

Figure 6 shows a partial coverage plan for the traversability map from Figure 4. This plan was generated for a short duration survey, assuming that: (1) the robots require a constant amount of time to move between cells and (2) a constant time is need to sample at each survey point.

VI. Task Executives

Multi-robot site survey requires a task execution system to assign survey tasks to the robots, monitor execution of those tasks, and resolve conflicts that may arise. Although centralized control is efficient for global coordination, it is also more vulnerable than distributed schemes, particularly in the case of partial or temporary system failures.

Thus, to ensure robust, co-ordinated fleet behavior, we employ a split execution system: (1) a central (global) executive performs overall co-ordination, and (2) a local executive, which runs on-board each robot,

handles reactive execution of assigned, individual tasks. Both the global and local task executives are currently implemented using the PLan EXecution Interchange Language (PLEXIL).[5]

The primary function of the global executive is to coordinate execution of survey plans produced by coverage planning. In particular, the global executive monitors robot pose and resolves conflicts using a fixed priority scheme. For example, in if two robots try to sample the same area with different instruments, or if two robots try to navigate through the same narrow passage from opposite directions, the executive halts the robot with lower priority until the robot with higher priority completes its task.

The local executive is responsible for performing individual survey tasks. At run time, it receives a target survey point and a list of intermediate waypoints as input. It then commands the robot to navigate to the point and perform the survey measurement. During execution of a single surveying task, no additional communication is required (i.e., the local executive can function autonomously). This enables survey operations to be robust in the presence of data network failures, which can be caused by poor network geometry in the survey site.

VII. Survey Tests and Robot Configuration

During Fall 2006, we will conduct multi-robot survey tests in the NASA Ames Marscape. The Marscape is a Mars surface analog that incorporates a variety of terrain including a dry lakebed, impact crater, and a volcanic zone. In these tests, we will use three mobile robots to map resources using three instruments: a microscopic imager, a terrain camera, and a subsurface sampler. All the robots are equipped with onboard computing, navigation sensors, and a software-based controller built on the NASA Coupled Layer Architecture for Robotic Autonomy (CLARAty).[6]

A. Rock Survey Robot

K9 (Figure 7, left) is a planetary rover based on a FIDO (NASA JPL) chassis.[7,8] K9 has a 6-wheel steer, 6-wheel drive rocker-bogey configuration, a 5-DOF instrument arm, a suite of mast-mounted steerable cameras, hazard cameras overlooking the arm workspace, and a variety of navigation sensors (odometry, compass/inclinometer, inertial measurement unit, and carrier-phase differential GPS unit).

Affixed at the end of K9's arm is the CHAMP (Camera Hand-lens MicroscoPe) microscopic camera[9] (Figure 7, right). CHAMP has a movable CCD image plane, allowing it to obtain focused images over a wide depth of field, from a few millimeters up to several meters. K9's arm allows arbitrary 3-DOF instrument placement as well as pitch and yaw control within the arm's workspace.

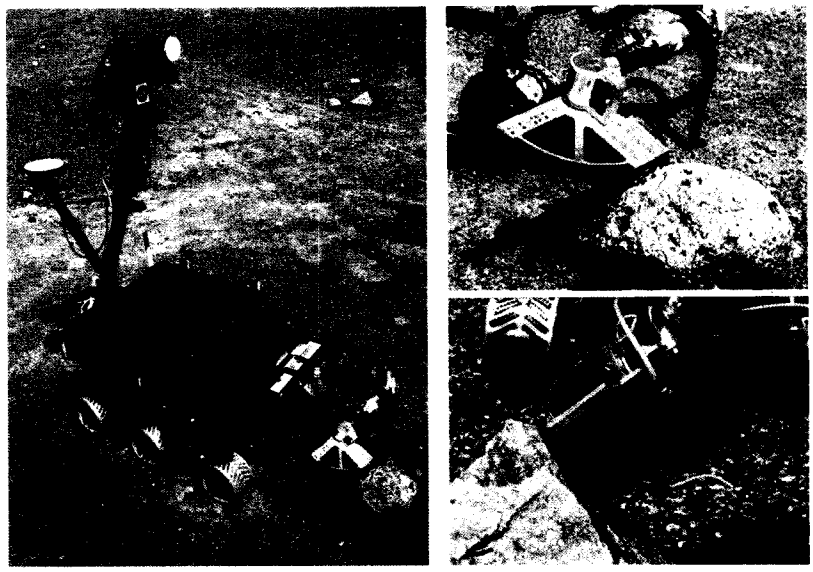

Figure 7. K9 rover (left) and CHAMP microscopic imager (right).

During the Marscape survey, K9 will autonomously place CHAMP against nearby rocks to acquire microscopic images of surface features to support physical characterization of rock geology. To do

this, K9 will use a visual tracking system, which combines a 2D interest point based tracker and a 3D-shape alignment technique to keep track of features and navigate to them.[10,11] Once K9 arrives at a rock, vision will also be used to find a place on the rock that is safe for CHAMP placement. K9 will then plan and perform an arm motion to safely position the instrument.

B. Soil Survey Robot

The K10 family of mobile robots is designed to operate in a wide range of environments, from high-friction indoor to moderate natural outdoor (30 deg slope, hard-pack dirt), at human walking speeds (up to 90 cm/s). K10 has four-wheel drive and all-wheel steering with a central rocker suspension, which allows it to traverse 20 cm step obstacles. K10's avionics duplicate those of K9 wherever possible, including power system and navigation sensors.

One K10 (Figure 8) is equipped with a downward-facing Point Grey Scorpion color camera. During the Marscape survey, this camera will acquire images of the terrain at different locations. Such images can then be used to classify soil. For example, image texture and can be used as to efficiently classify terrain images.[12] Moreover, color matching (e.g., based on color histograms) can be used to rapidly compare images.

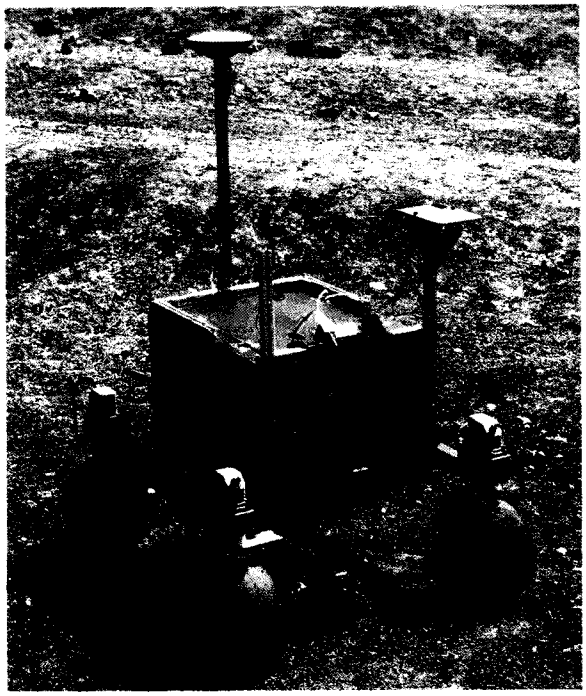

Figure 8. K10 rover and downward facing terrain camera.

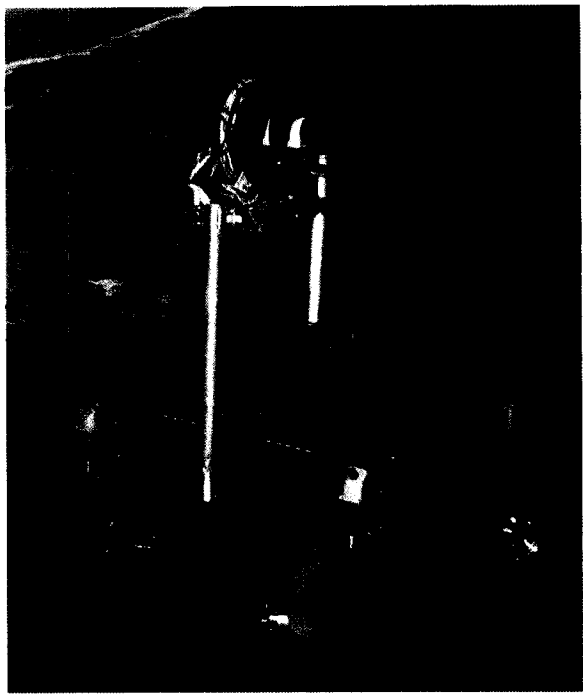

Figure 9. K10 rover and Mars Underground Mole (MUM).

C. Subsurface Sampler

A second K10 (Figure 9) is equipped with the NASA Mars Underground Mole (MUM).[7,13] MUM is a mobile subsurface penetrometer designed to deploy and retrieve itself, to burrow to depths of up to 5 meters (in Mars regolith), to detect subsurface mineralogy and retrieve soil samples. MUM uses an internal spring-and-mass impact system to produce a series of hammering movements that result in forward or reverse movement. MUM's design is based on a mole previously developed for the European Space Agency's Beagle-2 mission.

During the Marscape survey, the MUM will carry the Dual Spectroscopic Sensor (DSS), a subsurface sensing package that combines a Raman and near-infrared reflectance spectrometer. The DSS can be used to perform *in-situ* underground measurements to determine the composition of subsurface materials. As MUM burrows, the DSS views soil samples through a sapphire window and can be used to determine mineralogy (carbonates, iron oxides, etc.) as well as detect the presence of organic compounds and water at various depths.

VIII. Peer-to-Peer Human-Robot Interaction

In our work, we are investigating how *peer-to-peer* interaction can facilitate communication and collaboration between humans and robots.[14] We use the term "peer-to-peer" not because we expect humans and robots to have equal capabilities, but to emphasize that idea that humans and robots should work as partners.

To facilitate human-robot teaming, we have developed an interaction infrastructure called the "Human-Robot Interaction Operating System" (HRI/OS).[15] The HRI/OS allows humans and robots to work in a manner inspired by human work crews. In our system, for example, robots are able to ask task-oriented questions of the human in order to obtain assistance when they need help.

Site survey and sampling provides numerous opportunities for dynamic and flexible interaction. Humans, for example, may remotely interact with the survey robots: providing assistance for navigating through cluttered regions, helping assess instrument readings, etc. Robots, in turn, may need to communicate with EVA crew working nearby to request physical intervention or to coordinate sampling tasks (i.e., some operations may require both human and robot activity).

IX. Future Work

During the next year, our work will focus on survey and sampling activities needed for lunar *in-situ* resource utilization (ISRU). In particular, identifying and mapping key resources (minerals for oxygen production, water ice, and high glass-content regolith) will need to be done to optimize ISRU extraction and production. These lunar resources are likely to be distributed in variable quantity over a range of a few km, and thus well suited to rover-based surveying.

To do this, we plan to add additional sensors to our survey robots. This may include sensors for contact measurement (impedance spectrometer, conductivity meter, etc.), for surface composition assessment (raman spectrometer, visible-IR point spectrometer), and for near-surface measurements of stratigraphy (ground penetrating radar) and hydrogen (neutron spetrometer).

In the long term, we plan to conduct field tests to validate our approach in planetary analog environments, such as Meteor Crater (Arizona) and Haughton Crater (Devon Island, Canada). Particular emphasis will be placed on assessing: (1) human workload (both EVA and IVA) using NASA Task Load Index (TLX),[16] (2) Mean Time Between Interventions (MTBI) and Mean Time to Intervene (MTTI),[17] and (3) task-oriented human-robot interaction metrics.[18]

Acknowledgments

We would like to thank Matt Deans and Liam Pedersen for testing and improving K10's localization system. We would also like to thank Mark Allan, Leslie Keely and David Lees for adapting the Viz Explorer 3D visualization system for multi-robot monitoring. This work was supported by the NASA Robotics Academy Program and was sponsored by the NASA Exploration Systems Technology Development (ETDP) Program as part of the "Human System Interaction, Surface Handling and Surface Mobility Systems" project.

References

[1]NASA, "The vision for space exploration," Tech. Rep. NP-2004-01-334-HQ, NASA, Washington, DC, 2004.

[2]Lawrence, L., Sims, M., et al., "Photo-realistic terrain modeling and visualization for Mars exploration rover science operations," *Proc. 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, 2005.

[3]Singh, S., Simmons, R., et al., "Recent progress in local and global traversability for planetary rovers," *Proc. IEEE International Conference on Robotics and Automation*, 2000.

[4]Zelinsky, A., Jarvis, R., et al., "Planning Paths of Complete Coverage of an Unstructured Environment by a Mobile Robot," *Proc. International Conference on Advanced Robotics*, 1993.

[5]Verma, V., Estlin, T., et al., "Plan Execution Interchange Language (PLEXIL) for Executable Plans and Command Sequences," *Proc. 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, 2005.

[6]Volpe, R., Nesnas, I., et al., "The CLARAty architecture for robotic autonomy," *Proc. IEEE Aerospace Conference*, 2001.

[7]Bualat, M., Kobayashi, L., et al., "Flexible Rover Architecture for Science Instrument Integration and Testing," *Space 2006*, No. AIAA 2006-7420, AIAA, 2006.

[8]Park, E., Kobayashi, L., and Lee, S. Y., "Extensible Hardware Architecture for Mobile Robots," *Proc. IEEE International Conference on Robotics and Automation*, 2005.

[9]Lawrence, G., Boynton, I., et al., "CHAMP: Camera HAndlens MicroscoPe," *The 2nd MIDP Conference, Mars Instrument Development Program*, No. JPL Technical Publication D-19508, 2000.

[10]Pedersen, L., Bualat, M., et al., "Integrated Demonstration of Instrument Placement, Robust Execution and Contingent Planning," *International Symposium on Artificial Intelligence and Robotics in Space*, 2003.

[11]Pedersen, L., Deans, M., et al., "Multiple-Target Single Cycle Instrument Placement," *International Symposium on Artificial Intelligence and Robotics in Space*, 2005.

[12]Meyer, C., "Classification and content based retrieval of images for planetary exploration," Tech. Rep. M.S. Thesis, Microengineering, Swiss Federal Institute of Technology, 2006.

[13]Stoker, C., Richter, L., et al., "The Mars Underground Mole (MUM): A subsurface penetration device with in situ infrared reflectance and raman spectroscopic sensing capability," *Sixth International Conference on Mars*, 2003.

[14]Fong, T., Nourbakhsh, I., et al., "The Peer-to-Peer Human-Robot Interaction Project," *Space 2005*, No. AIAA 2005-6750, AIAA, 2005.

[15]Fong, T., Kunz, C., Hiatt, L. M., and Bugajska, M., "The Human-Robot Interaction Operating System," *Proc. First Human-Robot Interaction Conference*, 2006.

[16]Hart, S. and Staveland, L., *Human Mental Workload*, chap. Development of NASA TLX (Task Load Index): results of empirical and theoretical research, North-Holland Elsevier Science, 1988.

[17]Shah, J., Saleh, J., and Hoffman, J., "Analytical basis for evaluating the effect of unplanned interventions on the effectiveness of a human-robot system," *Reliability Engineering and System Safety*, in submission.

[18]Steinfeld, A., Fong, T., et al., "Common Metrics for Human-Robot Interaction," *Proc. Conference on Human-Robot Interaction*, 2006.