Safe and Efficient Robotic Space Exploration with Tele-Supervised Autonomous Robots

Alberto Elfes*, John M. Dolan, Gregg Podnar, Sandra Mau, Marcel Bergerman

*Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA, 91109-8099
elfes@jpl.nasa.gov

Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, 15213-3890
jmd@cs.cmu.edu

Abstract

A successful plan for space exploration requires the commissioning of fleets of robots to prospect, mine, build, inspect and maintain structures, and generally assist astronauts, rendering the overall mission as safe as reasonably achievable for human beings, the most precious resource. The authors are currently developing, under the support of NASA, a Robot Supervision Architecture (RSA) which will allow a small number of human operators to safely and efficiently telesupervise a fleet of autonomous robots. This represents a significant advance over the state of the art, where currently one robot is overseen by a group of skilled professionals. In this paper we describe some aspects of this work, including the architecture itself for coordination of human and robot work, failure and contingency management, high-fidelity telepresence, and operation under limited bandwidth. We also present highlights of our first application: wide area prospecting of minerals and water in support of sustained outposts on the Moon and on Mars.

Introduction

NASA has initiated the implementation of its Vision for Space Exploration by planning to return human beings to the Moon by 2018 and then on to Mars by 2030. This bold, risky, and costly enterprise will require that all possible actions be taken to maximize the astronauts' safety and efficiency. The authors believe that this can be facilitated by fleets of robots autonomously performing a wide variety of tasks such as in-space inspection, maintenance and assembly; regional surveys, mineral prospecting and mining; habitat construction and *in-situ* resource utilization (ISRU); etc. These robots will be telesupervised by a small number of human ground controllers and/or astronauts, who will be able to share control with and teleoperate each individual robot whenever necessary, all from a safe, "shirtsleeve" environment.

In this paper we present the Robot Supervision Architecture (RSA), a multilayered architecture for human telesupervision of a fleet of mobile robots. This research is supported by the Advanced Space Operations Technology Program of NASA's Exploration Systems Mission Directorate. Our objective is to demonstrate that the RSA enhances the telesupervisor's efficiency in a real-world mineral prospecting task (see Figure 1) while allowing supervision of the robot fleet from the relative safety of a lander, orbiter or ground station. The architecture is general enough to accommodate a wide range of applications and to span the entire spectrum of so-called sliding autonomy levels, from "pure" teleoperation to supervised control to "fully" autonomous [6].

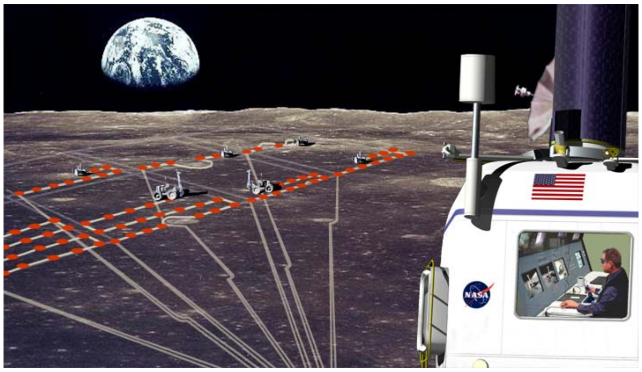

Figure 1. Artist's conception of telesupervised wide-area robotic prospecting on the Moon.

Our Robot Supervision Architecture addresses the following important challenges:

1. Coordination of human and robot work: The RSA is designed to both minimize the need for humans to perform costly and risky extra-vehicular activities (EVA), in which astronauts in space suits execute tasks in orbit or on lunar/planetary surfaces; and multiply their efforts beyond direct constant teleoperation of one robot. It does so by allowing the robots to operate as autonomously as technically possible, receiving assistance whenever necessary; and by allowing the human to assume partial or total control of a robot or its task-specific instrumentation whenever appropriate. While we do not explicitly address issues of tightly-coupled coordination between collocated humans and robots (i.e., joint task performance), the architecture supports augmenting the robots in these scenarios when autonomy is insufficient and continued assistance is critical.

2. Enhanced human safety and efficiency: Allowing a human to telesupervise a fleet of autonomous robots from

the safety of a "shirtsleeve" environment is a prudent and efficient approach to future space exploration. As our research progresses, we will test the efficiency gains over the baseline of a single human in a space suit performing EVA by tasking a fleet of robots to execute wide-area prospecting; and we will compute system performance metrics that take into account the area covered, prospecting accuracy, terrain difficulty, human safety, number of robotic rovers, human effort defined as the degree of human task intervention in full-time equivalents, and task completion time.

3. Failure and contingency management: The architecture explicitly defines a Hazard and Assistance Detection subsystem which operates on multiple levels. At the robot level, it assesses deviations from standard or expected operating conditions, both with respect to the robot's health and its assigned task. At the workstation level, the Hazard and Assistance Detection subsystem is responsible for queuing and prioritizing all robot alerts and requests, based on a criterion that takes into account the urgency of the request and a predicted amount of time the teleoperator needs to assist that particular robot.

4. Remote operations with bandwidth constraints: Our system relies on both high-bandwidth radio links for geometrically-correct, high-fidelity telepresence and teleoperation of any robot in the fleet, and lower-bandwidth links for command and telemetry exchange between the robots and the telesupervisor workstation. Should the high-bandwidth video link malfunction, the system design provides for graceful fallback to lower-bandwidth communications.

It is important to note that our approach is optimal when the human telesupervisor and the robot fleets are "near" each other, meaning that they are separated by a roundtrip communication delay of no more than about 300 milliseconds (28,000 mi / 45,000 km distant). It is also applicable with short telecommunication delays, such as between the Earth and Moon.

The paper is structured as follows. In the next section we discuss the novelty of our work with respect to the state of the art. In the following section we describe the architecture itself, as an overarching paradigm under which all other subsystems reside. It addresses specifically challenges #1 and #2 above. In the sequence we address challenge #3, presenting the Hazard and Assistance Detection subsystem and its underlying protocols. In the next-to-last section we present details of our first application area, the wide-area mineral prospecting task, including a task-specific performance metric. The last section presents conclusions and future work. For completeness, we note that our approach to challenge #4 is presented in more detail in another paper [10].

Related Work

The most important aspect of our work, namely, creating a Robot Supervision Architecture that allows a human safely and efficiently to telesupervise a fleet of autonomous robots, encompasses a variety of robotic technologies. We review here only research focussed on human-robot interaction for space exploration, or which is strongly applicable to the area.

The state of the art in robotic space exploration are the Mars Exploration Rover (MER) missions [8]. Spirit and Opportunity combined have logged over 10 km and operated for over 1200 sols (Martian days). This is achieved by assigning a large team of highly skilled professionals to download telemetry and imagery, interpret the data, plan the rover activities, and program and upload commands every sol, in addition to a large science team to select science targets and tasks. In contrast, we are multiplying one human's capability to telesupervise a large number of robots, while still allowing the human to perform other tasks. Another difference between MER and this work is that, because of the long communication delays between Earth and Mars, the only possible way of operating the rovers is via batch command sequences which are executed in autonomous mode; whereas the RSA accommodates a large variety of operation modes.

The sliding autonomy aspect of space exploration is one of great importance. Heger et al. [6], in particular, have developed an architecture geared towards humans and robots "jointly performing tightly coordinated tasks." They focus on "how small teams of robots in space and a few humans on Earth could work together to assemble large orbital structures," while we focus on maximizing an astronaut's efficiency by coordinating a large fleet of robots.

From the point of view of direct assistance to and collaboration with astronauts, a relevant project is Robonaut [1]. Robonaut's focus is a space robot "designed to approach the dexterity of a space suited astronaut." From the point of view of RSA, a Robonaut would be another robotic device whose operation could be coordinated using our architecture. When teleoperated, Robonaut's main similarity with our work is the telepresence capability implemented with stereo cameras. However, Robonaut's use of a head-mounted display and converged cameras differs from our geometrically-correct remote viewing system.

Other human-multirobot architectures are those of Nourbakhsh et al. [9] and Sierhuis et al. [12]. The former focuses on urban search and rescue operations; their architecture allows for simultaneous operation of real-world and simulated entities. The latter have created a Mobile Agents Architecture (MAA) integrating diverse mobile entities in a wide-area wireless system for lunar and planetary surface operations. Our work is conceptually similar to these, but it differs in the fact that we focus on human safety and efficiency in a planetary exploration environment by providing high-fidelity telepresence and a hazard and assistance detection methodology that seeks to optimize the use of human attention resources given a set of robot assistance requests.

Finally, we note the work of Fong et al. [4], where the authors also develop an architecture for supervision of multiple mobile robots. Their work and ours differ in the assistance request protocols and our use of stereoscopic telepresence.

The Robot Supervision Architecture

The RSA is implemented as a multilayered multi-robot control and coordination architecture that can accommodate different configurations of robotic assets based on previous work by Elfes [2], [3]. Here, “multilayered” means that robot system control is performed at multiple levels of resolution and abstraction and at varying control rates. Likewise, “replicated” means that the fundamental activities of perception, decision-making and actuation occur at each layer of the architecture. A diagram of the overall RSA architecture is shown in Figure 2 and explained below.

The Autonomous Navigation System (ANS) is replicated on each robot for local rover navigation. In the same way, each robot's Hazard and Assistance Detection (HAD) system is tightly coupled with the local ANS, and is supported up through the layers of the RSA architecture for high-level decision-making and handover to the telesupervisor.

The **Human Telesupervisor** oversees the entire mission planning and execution, being able to assume a wide range of roles – from “pure” supervision while monitoring the progress of the assigned tasks; to monitoring the performance of the fleet of autonomous robots; to “pure” teleoperation of any robot vehicle or its subsystems. This means that the RSA covers the entire sliding autonomy spectrum as defined in [6]. The reader should note that this is not to be confused with the so-called levels of interaction engagement [11], as we are not dealing with the issue of robots interacting with humans in a “social” way.

Task Planning and Monitoring and **Robot Fleet Coordination** lie at the core of the Robot Supervision Architecture. The high-level mission plans are created and edited with the Planning Tools, and are then assigned to the Robot Fleet Coordination module, which decomposes them into tasks and assigns these to the individual robot controllers (see Figure 3).

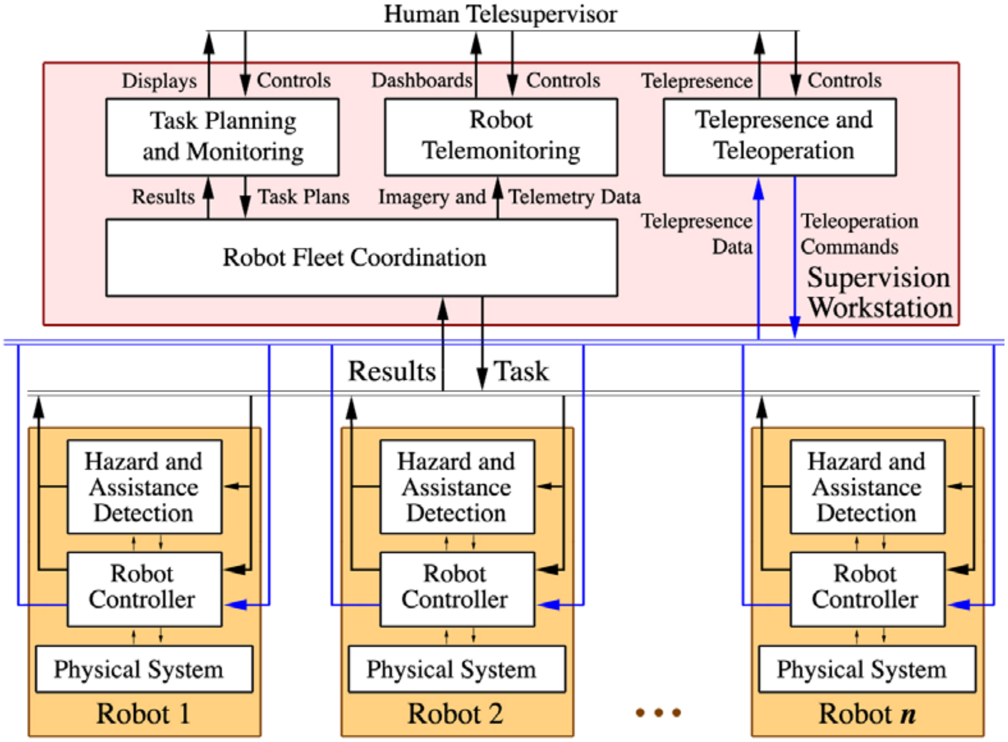

Figure 2. RSA system-level block diagram and main data paths.

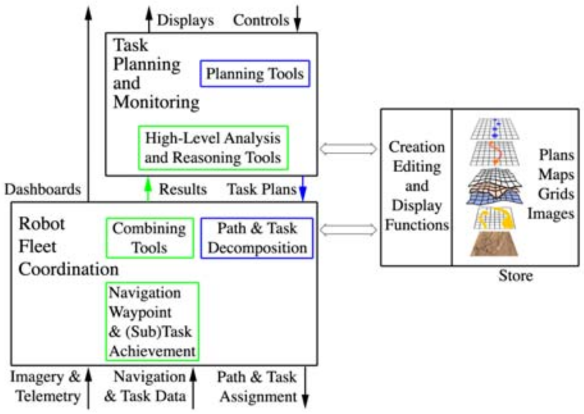

Figure 3. RSA Planning, Monitoring, and Coordination.

Robot Fleet Coordination also collects operational results from all robots and integrates them for convenient human monitoring and analysis. The Monitoring Level also includes Perception — Decision-Making — Actuation sequences to monitor multi-robot operations at the system-level, and to analyze for high-level hazard and assistance detection. Robot Fleet Coordination imagery and telemetry are combined for building regional imagery and maps, and are also presented graphically to the telesupervisor, as shown in Figure 4.

Figure 4. Telesupervisor workstation (concept).

A suite of editing tools for plans and maps, as well as manipulation tools are available at both the Planning and Coordination levels. The telesupervisor uses these tools and other display-oriented tools such as overlay maps, grids, and images for viewing. All these data structures are maintained in a data store as depicted at the right of Figure 3.

As depicted in Figure 5, each individual **Robot Controller** subsystem is responsible for receiving a collection of tasks from the Robot Fleet Coordination and monitoring its execution. The robot controller has direct access to all of the robot's subsystems to drive actuators and read sensor data. When a robot is a relatively complex combination of mobility, manipulation, and other engineering or science subsystems, the corresponding robot controller may be implemented as a collection of modules responsible for each one of them. Each robot controller is currently subdivided into the Autonomous Navigation System and Prospecting Task Support modules, respectively responsible for controlling the mobility and prospecting subsystems. The ANS is responsible for decomposing the navigation path assigned to it and reporting its progress. The initial Prospecting Task Support software at the robot level is very simple: it merely supports the control and monitoring of the prospecting tools.

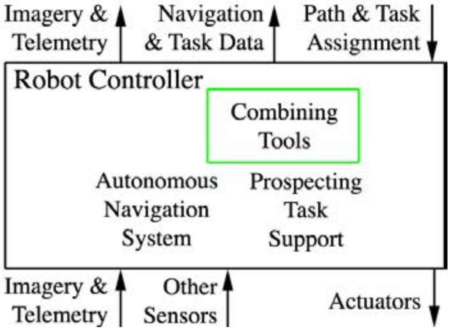

Figure 5. Robot Controller.

Just like the Robot Fleet Coordination module, the Robot Controller includes a set of robot-specific Combining Tools. These include functions such as generating composite elevation maps and images along the path of the individual robot by combining the separate depth maps and images captured by the robot's cameras.

Each Hazard and Assistance Detection subsystem is responsible for the robot's and the mission's integrity, including vehicle monitoring and health assessment and failure detection, identification, and recovery. It assesses the robot's sensor data to infer potential hazards or the machine's inability to complete a task, and communicates with the Robot Fleet Coordination, which must take the appropriate action(s).

Robot Telemonitoring allows each robot to be constantly monitored at low bandwidth by the human telesupervisor with imagery and data updated regularly from each prospecting robot vehicle. As shown on the right in Figure 4, each robot has a "dashboard" which includes streaming images from one of the robot's cameras, and graphical depictions of status data such as battery charge, attitude, motor temperatures, and any other monitored telemetry.

The Hazard and Assistance Detection (HAD) system (see next section) automatically monitors each robot. When the telesupervisor should be made aware of a hazardous condition, it is on that robot's dashboard that it is indicated. This is exemplified for Robot #2 in Figure 4 by an orange surround of the robot view.

Telepresence and Teleoperation subsystems: We make a distinction between monitoring the operation of each robot, and telepresently taking control of a robot. Where monitoring is supported by simultaneous low-bandwidth data streams from each of the robots, telepresence is supported by high-bandwidth stereoscopic imagery and other telesensory modalities one robot at a time. It provides not only the stereoscopic visual, but also aural, and attitude proprioceptive feedback that allows for more immersive telepresence.

Teleoperation involves direct human control of a single robot when a vehicle must be remotely driven rather than operating under its Autonomous Navigation System; and when the task-specific tools must be operated manually. Joystick, keyboard, and task-specific input devices support this. This subsystem is implemented over a dual-path data communication infrastructure, where the low-bandwidth path is used for communication of commands and telemetry data, and the high-bandwidth path is used for stereoscopic video. In addition to the described functionality, each subsystem is a source for data which are both archived for later analysis, and also provided in part as a stream for access by a Distant Expert who can consult as required.

Hazard and Assistance Detection

The Hazard and Assistance Detection (HAD) subsystem is responsible for the following high-level capabilities:

- Single-rover hazard and object-of-interest detection.
- Single-rover hazard and assistance assessment.
- Multi-rover assistance request prioritization and management.

The overall goal of the HAD assessment protocol is to take into consideration both the relative urgency of the hazard and the availability of a human telesupervisor to deal with it. The detection of hazards involves a fusion of various sensor inputs to generate a comprehensive picture of the situation. Both the lower-level HAD located on each rover and the higher-level HAD located in the workstation are implemented as a perception – decision-making – actuation sequence.

At the lower level, the perception aspect includes receiving inputs from the various sensors and subsystems, including motor current, path waypoints, and a number of HAD flags in the assistance queue on the telesupervisor side. The decision-making aspect involves assessing the whole spectrum of available sensor data and determining whether there is a potential for hazard. If such a hazard is identified, the urgency of the situation is evaluated and compared to the urgency tolerance threshold and range. The actuation aspect involves responding to the detection of a hazard by halting the actuators on the robot and alerting the supervisor with a hazard flag. If the urgency is merely within the urgency tolerance range but not above the threshold, then the relevant data are passed to the supervisor as a caution, but the operation of the rover remains uninterrupted.

At the higher level on the supervisor side, HAD perception receives cautionary telemetry, and hazard flags and their related telemetry. The decision-making aspect prioritizes the hazard flags in the assistance queue in order of urgency (passed from the rovers). The resulting actions (actuation) are to inform the teleoperator of the hazard or potential hazard through the control panel.

Since a telesupervisor oversees multiple rovers, it is not feasible for the operator to be fully aware of each rover's situation at all times. When a hazard is flagged, the operator should be made aware of exactly what hazard was detected and why, rather than having to figure it out. To get a better picture of the hazard situation, data from other sensors in addition to the video feed will be considered by the HAD algorithm. To address these requirements, the HAD subsystem was designed as in Figure 6. It is explained in detail in the sequel.

Various factors come into play when determining whether to ask the operator for assistance. These include the type of hazard, its associated risk and time to fix, and the availability of the telesupervisor. To determine when assistance is needed, time and risk are weighted and combined to determine the "urgency" of the situation. This urgency value is compared to the urgency tolerance threshold (based on operator availability) and urgency tolerance range (a parameter set by the operator). Above the threshold, the operator is notified of the hazard and the rover is halted. Within the tolerance range, the operator is cautioned, but the rover may still act autonomously. If the urgency value is within a given range below the threshold value (not urgent enough to stop the rover, but urgent enough to bear watching) the information detailing the possible hazard will still be passed to the operator indicating a "warning", but no hazard flag is queued.

Risk Associated with Hazard [5]
This includes both risk to the physical well-being of the rover (physical hazard) and risk to the rover's ability to execute its assigned task (mission hazard). Physical hazards are more critical, since physical damage also affects ability to complete the mission. For example, a rover that is about to tip over would likely require immediate attention, as opposed to a rover with navigation difficulty. Some risk values vary depending on the degree of possible physical "discomfort" for the rover. For example, a path over rough terrain will have a higher risk value than a smooth path. The Autonomous Navigation System will provide the waypoints for the suggested path as well as the risk in taking that path (which considers factors such as roughness and tilt.) The risk value is normalized to range between 0 and 1, with 1 being most critical and 0 being minimum risk.

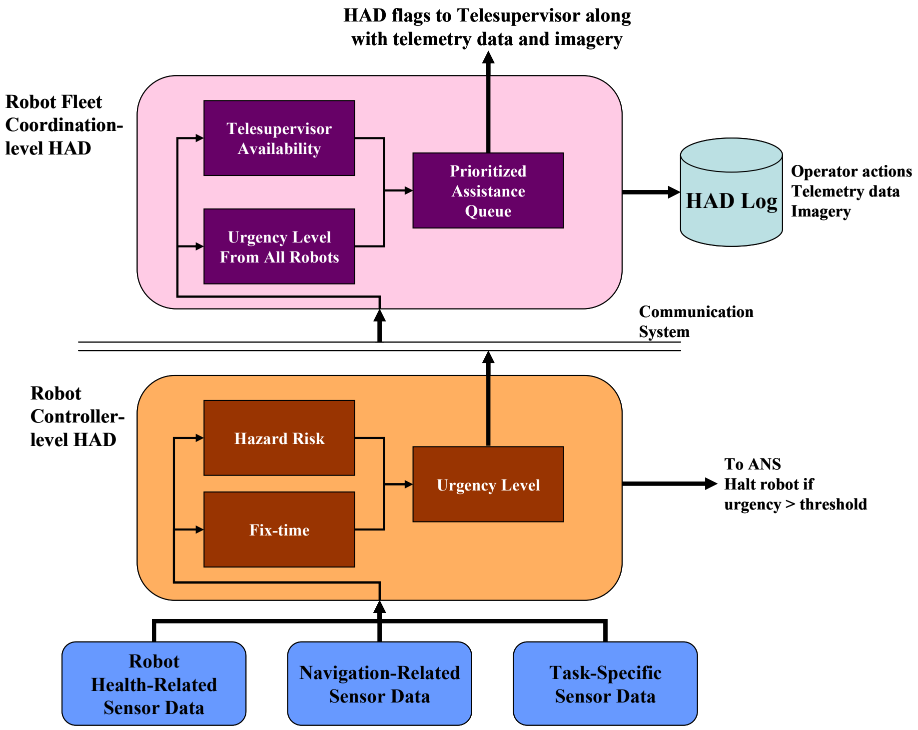

Figure 6. Hazard and Assistance Detection block diagram.

Time Associated with Hazard

The estimated amount of time necessary to remedy a hazard ("fix-time") is relevant to hazard urgency. If one non-critical situation requires a lot of time to overcome, it may be more efficient to tackle another hazard with shorter fix-time first before turning to the longer one. This decreases the backlog of inoperative rovers in the case where many rovers need assistance. Hazards which take less time to fix are therefore be weighted with a preference for being addressed first. Initial fix-times will be based on average data from user studies, and then updated by monitoring of specific user actions. For an unknown hazard, the fix-time associated will likely be the least amount of time possible since caution dictates addressing an unknown case quickly, rather than slowly. The fix-time value is also normalized to range between 0 and 1, with 1 corresponding to the least amount of time to resolve a hazard or an unknown hazard, and 0 corresponding to the (thresholded) longest time.

Operator Availability (Urgency Tolerance Threshold and Urgency Tolerance Range)

Since the RSA deals with multiple rovers operating simultaneously, it is likely that more than one will ask for assistance at the same time. If the operator is very busy, there may be a queue of inoperative rovers waiting for attention. Instead of waiting, some of the rovers making decisions with lower risk (for example, when the only hazard is slightly rough terrain) may choose to act autonomously. In other words, depending on the number of rovers in the assistance queue and their urgency, the urgency tolerance threshold will dynamically increase to attempt to ease backlog. One simple way to do this is to have several set threshold values based on number of hazard flags and average urgency. The threshold is a number between 0 and 1, with 1 being highest urgency and 0 being no urgency.

When a rover is nearing the urgency threshold, entering a tolerance range below the threshold value causes the operator to receive a warning as well as the relevant telemetry data associated with it. The tolerance range is a value less than 1 and is set as a parameter by the operator.

Robot Action: Assistance Request/Urgency vs. Urgency Tolerance Range and Threshold

The time and risk associated with the hazard are multiplied by weights that the operator can set. The two weighting values add up to 1. If an operator works very quickly, he may place less emphasis (lower weighting value) on the time it takes to address hazards. The linear combination of the weighted risk and time factors is the urgency, which is then compared to the urgency tolerance threshold described above. If the urgency is greater than the threshold, the rover sends a hazard flag with the relevant information (urgency, hazard case, rover name, telemetry) to be placed in the assistance queue on the supervisor side, and the rover is immediately halted. This hazard flag to the operator is also continuously updated to monitor changes in the situation (i.e., the rover will continue to monitor itself for increasing/decreasing urgency.) If the urgency falls in the tolerance range under the threshold, the rover continues its operations, but the relevant telemetry is passed to the supervisor to signal a caution warning on the control panel. This allows the operator to address a potential hazard before it occurs.

The use of user-adjustable weights for time and risk to determine urgency provides flexibility. For example, setting the time weight to zero creates the simple case in which hazards are dynamically prioritized in the queue according to risk alone, so that highest-risk hazards receive attention first without regard to fix-time. Conversely, setting the risk weight to zero creates the simple case in which hazards are prioritized in the queue according to fix-time alone, so that the most rapidly addressed hazards receive attention first without regard to risk. Experiments will allow us to explore a suitable balance between these extremes.

Assistance Queue & Prioritization

The most obvious method of prioritization of hazard flags in the assistance queue is to rank in order of urgency. The higher the urgency, the higher the priority in the assistance queue.

Hazard Logs

Scientists and engineers may want later to study the detected hazards, so all relevant data in the hazard situation must be recorded and saved. This log includes a sequence of events starting when the HAD flag was raised and lasting until the operator hands back control to the rover. Each event receives a time stamp. The relevant telemetry from vision and sensors are saved, as well as a record of the actions taken. The operator is asked to briefly state what and why s/he is taking such an action either through text or voice recording.

Wide-Area Prospecting

Our first-year test for the RSA is to autonomously search an area for *in situ* resources with assistance from a human telesupervisor when needed. A set of onboard instruments for each rover represents an analysis system that will function as a stand-in for a suite of instruments for the Moon. It is not the purpose of this project to design a complete, integrated, chemical analysis system, but rather to demonstrate the interactions between human telesupervisors and prospecting robots, and validate their performance in identifying resources in the field. In the future, the prospecting instruments will be expanded to include sampling tools.

A chosen area will be prospected using a predetermined search algorithm defined prior to the test (see Figure 7). This simple grid-search algorithm is akin to any initial terrestrial prospecting task where no a priori mineral information is known about the area. It is also analogous to sub-surface sample prospecting, such as core boring, which is one of our target prospecting tasks in the future.

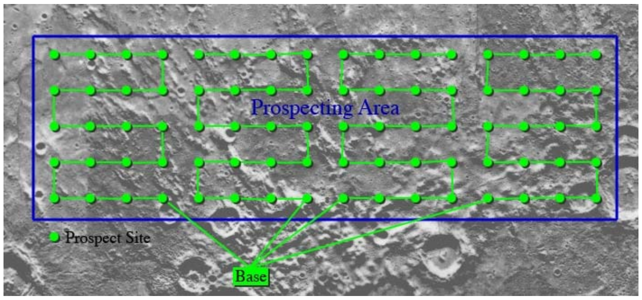

Figure 7. Robot prospecting paths within prospecting area.

To provide a tractable prospecting task for development and validation, we have elected to do a surface study with limited manipulation and non-contact sensor-based sampling. Because in the future the prospecting task will involve discrete prospect sites with physical soil and/or core samples collected and returned to base for analysis, the initial study is designed to be an analogue, with surface sensor data being taken in a similar fashion at discrete prospect sites. This surface sensor will provide data that can be analyzed without the need to deploy expensive sampling tools (e.g., a mass spectrometer) on each of the robots. This will be accomplished by adapting a video camera to collect visual sample data for an area of approximately 0.25 m^2 at each sampling site.

The target material density measurements collected at each sampling site together with the coordinates of each site will be used to build a map of material density over the prospecting area. An algorithm that infers a distribution over the whole prospecting area will be employed. This map will then be compared to the resource map maintained by our geologist who seeded the prospecting area to characterize the accuracy of the prospecting aspect of the system.

To quantitatively assess how well the system performs in the prospecting task, we propose a basic performance metric based on the following notions: 1) greater area, accuracy, terrain difficulty, and safety of prospecting coverage mean increased performance; 2) greater effort and time required mean decreased performance. Given these factors, we propose the following metric:

$$P = \frac{ACTS}{(R / w + H_E)t}$$

where:

P: performance in units of (area accurately prospected)/(effort-time).
A: area covered.
C: prospecting accuracy; $C = 1$ corresponds to the highest possible accuracy and $C = 0$ corresponds to the lowest possible accuracy.
T: terrain difficulty factor ($T \geq 1$) with $T = 1$ corresponding to the easiest terrain (a flat surface without obstacles).
S: safety factor ($S \geq 1$) with $S = 1$ corresponding to the least safe task performance, i.e., via 100% EVA.
R: number of robotic rovers (integer).
H: number of humans (integer, but does not occur in the performance formula); note that although our project's focus is on a system in which a single human astronaut controls multiple rovers, the metric is general enough to allow for multiple humans.
H_E = human effort defined as the degree of human task intervention in full-time equivalents ($0 \leq H_E \leq H$); e.g., if one human intervenes 30 min. during a 1-hr. task, $H_E = (30/60) = 0.5$; if three humans intervene 15, 30, and 45 min. respectively during a 1-hr. task, $H_E = (15/60) + (30/60) + (45/60) = 1.5$.
w: factor allowing commensurability of human and rover time by giving the relative value of the former to the latter; e.g., $w = 4$ sets human time to be four times as valuable as rover time
$R/w + H_E$: combined human-rover effort.
t = time required to cover A.

We will report on the results obtained with this metric after we conclude our indoor and outdoor tests in the Fall of 2005.

Conclusion

The work presented in this paper summarizes parts of a larger technology development effort being undertaken by the authors under NASA support and in cooperation with NASA centers. Other aspects of this effort include the Autonomous Navigation System, based currently on standard binocular vision [5]; the Telepresence and Teleoperation System [10]; and other task-specific elements. Our ultimate goal is to deliver the entire Robot Supervision Architecture to NASA at technology readiness level 6 [7], after extensive field tests where one human will telesupervise a fleet of eight to ten autonomous robots performing mineral prospecting and core sample drilling.. Specifically with respect to HAD, in the near future more sophisticated assessment protocols will be implemented, possibly using a Bayesian framework for dynamic state estimation [13]. This will improve the ability to identify obstacles and will also aid in the performance of opportunistic science in that features of interest can be detected with greater accuracy and frequency.

References

[1] R.O. Ambrose, R.T. Savely, S.M. Goza, et al. "Mobile Manipulation Using NASA's Robonaut." *Intl. Conference on Robotics and Automation*, New Orleans, USA, April 2004, pp. 2104-2109.

[2] A. Elfes. "Incorporating spatial representations at multiple levels of abstraction in a replicated multilayered architecture for robot control." in *Intelligent Robots: Sensing, Modelling, and Planning*, R. C. Bolles, H. Bunke, H. Noltemeier (eds.), World Scientific, 1997.

[3] A. Elfes. "Dynamic control of robot perception using multi-property inference grids." *Intl. Conference on Robotics and Automation*, Nice, France, May 1992, pp. 2561-2567.

[4] T. Fong, C. Thorpe, and C.; Baur. "Multi-robot remote driving with collaborative control." IEEE Transactions on Industrial Electronics, Vol. 50, No. 4, August 2003, pp. 699-704.

[5] S. Goldberg, M. Maimone, and L. Matthies. "Stereo vision and rover navigation software for planetary exploration". *IEEE Aerospace Conference*, Montana, USA, March 2002, pp. 2025-2036.

[6] F. Heger, L. Hiatt, B.P. Sellner, R. Simmons, and S. Singh, "Results in Sliding Autonomy for Multi-robot Spatial Assembly." 8th *International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Munich, Germany, September, 2005.

[7] J.C. Mankins. *Technology Readiness Levels: A White Paper*. NASA Office of Space Access and Technology, April 1995. http://advtech.jsc.nasa.gov/ downloads/ TRLs.pdf.

[8] *NASA Mars Exploration Rover Mission*. http://marsrovers.jpl.nasa.gov/home/index.html.

[9] I. Nourbakhsh, K. Sycara, M. Koes, et al. "Human-Robot Teaming for Search and Rescue." *Pervasive Computing*, Vol. 4, No. 1, Jan-Mar 2005, pp. 72-78.

[10] G. Podnar, J. Dolan, A. Elfes, M. Bergerman, H.B. Brown and A.D. Guisewite. "Human Telesupervision of a Fleet of Autonomous Robots for Safe and Efficient Space Exploration." Submitted to *Conference on Human-Robot Interaction*, Salt Lake City, USA March 2006.

[11] T.E. Rogers, J. Peng, and S. Zein-Sabatto. "Modeling Human-Robot Interaction for Intelligent Mobile Robotics." *IEEE International Workshop on Robots and Human Interactive Communication*, Nashville, USA, August 2005, pp. 36-41.

[12] M. Sierhuis, W.J. Clancey, R.L. Alena, et al. "NASA's Mobile Agents Architecture: A Multi-Agent Workflow and Communication System for Planetary Exploration." 8th *International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Munich, Germany, September 2005.

[13] A. Soto and P. Khosla. Probabilistic adaptive agent based system for dynamic state estimation using multiple visual cues. Carnegie Mellon University, Pittsburgh, PA, 2001.