

Human Telesupervision of Very Heterogeneous Planetary Robot Teams

Gregg Podnar* and John Dolan†
Carnegie-Mellon University, Pittsburgh, Pennsylvania, 15213

Alberto Elfes‡
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, 91109

and

Marcel Bergerman§
Carnegie-Mellon University, Pittsburgh, Pennsylvania, 15213

Lunar and planetary surfaces are the most hostile working environments into which humans can be sent. The protective spacesuit is massive and cumbersome, with EVA mission time limited by both the suit's resources and the astronaut's stamina. To maintain human presence on the Moon and to expand it to Mars requires enormous investments in transportation and life support for each human. Therefore, successful and sustainable space exploration and operations must maximize the efficiency of every astronaut and keep them "as safe as reasonably achievable".

Towards this goal, tasks for which current robotic autonomy technologies are effective should be offloaded from the astronauts. However, whenever the limits of autonomy are reached, a human will need to intervene, preferably by telesupervising the robotic assets (thus reducing EVAs). Employing an effective telesupervision architecture to augment the ingenuity of a human supervisor with state-of-the-art autonomous systems results in a manifold increase in the human's performance and a significant improvement in safety. This completely changes the risk profile of a mission, and allows astronauts to perform substantial amounts of hazardous work from a well-supplied operations base, such as an orbital station, a CEV, or a Lunar or Martian habitat.

Telesupervised robotic systems have been identified as a key technology by the NASA Exploration Systems Mission Directorate, and are crucial to the success of the Vision for Space Exploration. However, very little applicable work has been done in the design of telesupervised system architectures, the appropriate mix of autonomy and remote control, and context switching between them, or in the testing and deployment of such systems.

This paper focuses on the development of an advanced telesupervision system architecture that will provide a highly efficient approach to human-robot interaction while allowing very heterogeneous robotic assets to be deployed. These assets include exploration rovers and climbers; large autonomous miners and transporters; stationary ISRU processing plants, materials fabricators, and power stations; and construction and maintenance robots.

We argue that for the telesupervisor to acquire the state of each varied robot and its environment involves not only telemetry and high-fidelity telepresence (including proprioceptive cues), but also a sensorial "playback" of the recent history of autonomous operation that will reveal the issues that led to the crisis that now requires assistance. Providing the framework within which this history and context are acquired and reproduced is crucial to a viable telesupervision architecture.

* Program Manager, Robotics Institute, AIAA Member.

† Senior Systems Scientist, Robotics Institute.

‡ Principal Member of Technical Staff, Mobility and Manipulation Group, M/S 198-235, AIAA Member.

§ Project Scientist, Robotics Institute.

Our philosophy of maximizing the efficiency and safety of humans through telesupervision of autonomous robotic systems is applicable across all anticipated operational phases of the Vision for Space Exploration, including: telesupervising Lunar robots from Earth, Lunar orbit, or the Lunar surface; and telesupervising Martian robots from Mars orbit, or the Martian surface. The described architecture also applies to on-orbit assembly, inspection, and maintenance operations. Finally, the telesupervision architecture is relevant for earth science applications such as ecological forecasting, water management, carbon management, disaster management, coastal management, and homeland security.

We describe two applications of our Multilevel-Autonomy Robot Telesupervision Architecture: to planetary mineral prospecting using multiple semi-autonomous rovers based on work conducted under a past project funded by the NASA Exploration Systems Mission Directorate; and to Harmful Algal Bloom detection and characterization by multiple semi-autonomous ocean vessels based on work conducted under an ongoing project funded by the NASA Earth Science Technology Office.

By addressing the real problem of human and robot cooperative effectiveness, the system described in this paper is responsive to the goals of NASA's Global Exploration Strategy and the Lunar Exploration Program Architecture. The Lunar activities to which this architecture is applicable are an essential testbed for refining the technology for subsequent deployment to Mars. Applying this telesupervision architecture will save thousands of hours of astronaut time, as well as thousands of tons of mass due to fewer astronauts needing support to achieve the Lunar and Martian objectives.

I. Introduction

IN January 2004, NASA established a long-term program to extend human presence across the solar system. It encompasses a broad range of human and robotic missions. The key goal of this policy and strategic direction is that future space exploration activities will be enabled by combining human and robotic capabilities in order to achieve a long-term and well-orchestrated campaign of space exploration.¹

Efficient methods for human-machine cooperation are necessary for effective space exploration and operations both on orbit and on planetary surfaces. Our approach is to combine the strengths of the state of the art in autonomous systems with the ingenuity of a nearby human supervisor. We contend that augmenting a human's expertise with multiple autonomous robots results in a manifold increase in the human's performance and a significant improvement in life safety.

To support our approach we have defined, designed, and are continuing to develop an advanced telesupervision system architecture that supports highly efficient human-robot interaction while allowing very heterogeneous robotic assets to be deployed. Demonstrating the architecture's flexibility, we have applied these concepts to two very different task domains: adaptive investigation of harmful algal blooms in Earth's oceans; and planetary mineral prospecting for *in situ* resource utilization.

This flexible architecture, developed, hardened, and proven by real-world tests, becomes the operational framework for a wide variety of tasks where teams of autonomous robots are telesupervised by fewer humans. Such tasks include space assembly, inspection, and maintenance; and planetary prospecting, mining, and construction, while keeping humans "as safe as reasonably achievable" (ASARA) and increasing the effectiveness of humans on orbit, the Moon, and planetary surfaces.

II. Telesupervision Architecture

Our flexible Multilevel-Autonomy Robot Telesupervision Architecture (MARTA) is designed as a general architecture for telesupervision of teams of robots. It supports human telesupervision of robotic assets with a range of autonomy capabilities. Implemented as a multilevel, multi-robot control and coordination architecture, it can accommodate different configurations of stationary and mobile robotic assets. Here, "multilevel" means that robot system control is performed at multiple levels of resolution and abstraction.

The lower-level robot-specific functions of MARTA operate under typical Perception – Decision-Making – Actuation sequences. While the higher-level functions are for multi-robot control and monitoring. These follow a similar paradigm, also having Perception – Decision-Making – Actuation sequences but at a higher level of abstraction. For example, stationary robotic assets have low-level modules for operational autonomy, and mobile assets (e.g., rovers and boats) each have a low-level Autonomous Navigation System. These robot-specific functions are abstracted at a higher level in MARTA for team coordination, monitoring, decision-making, and when required,

handover to the telesupervisor. A diagram showing the modules and levels of MARTA is presented in Figure 1 with each of the subsystems explained below.

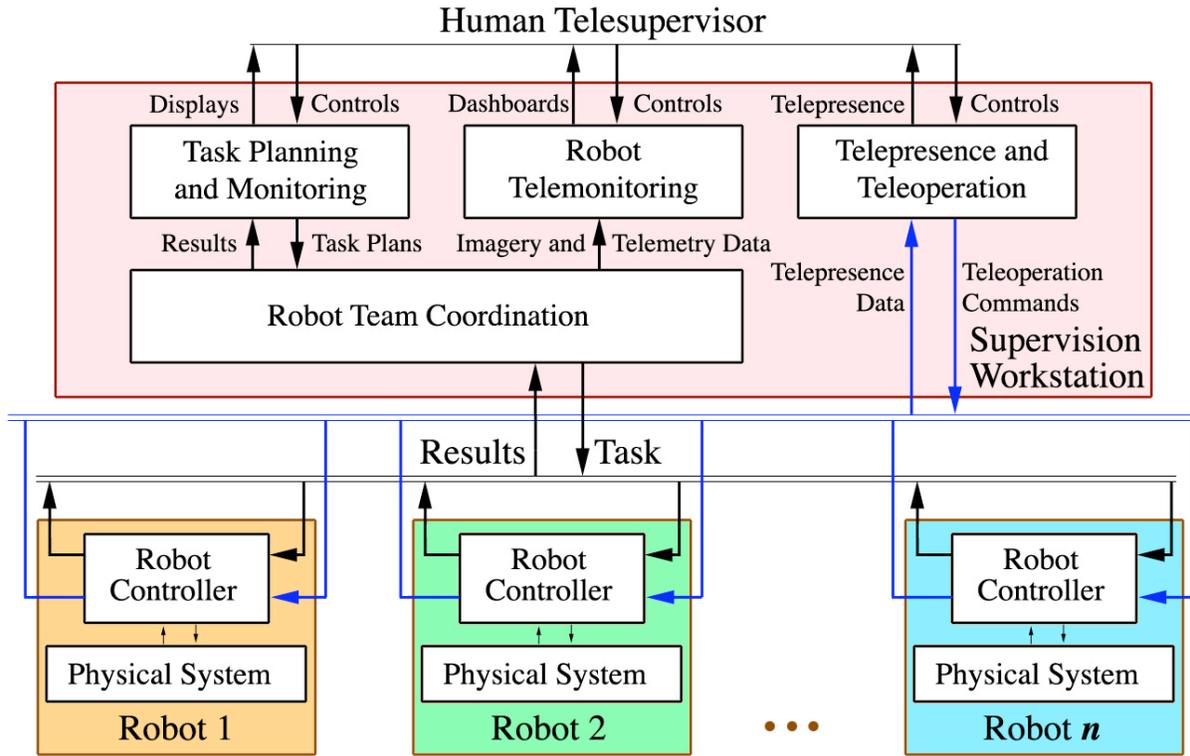


Figure 1. Multilevel-Autonomy Robot Telesupervision Architecture (MARTA) block diagram and main data paths.

A. Human Telesupervisor

A human telesupervisor oversees the entire mission planning and execution, being able to assume a wide range of roles – from “pure” supervision at the highest level while monitoring the progression of accomplishing the assigned task; to monitoring the performance of the team of autonomous robots; to “pure” teleoperation at the lowest level of any robot vehicle or its subsystems. The telesupervisor devises the high-level mission plans and assigns them to the Robot Team Coordination module for further task decomposition and assignment to the individual robots. The telesupervisor may also drive the robots directly via the Telepresence and Teleoperation subsystem. All results, in the form of images and sensor data, science data, and other task-specific information are provided at the Telesupervisor Workstation, one conception of which is shown in Figure 2.

B. Task Planning and Monitoring

To the left in Figure 2 are displays that allow interaction with the high-level task planning and monitoring. Functionally, this is represented by the Task Planning and Monitoring module in Figure 1. Maps representing different aspects of the area of operation are displayed and overlaid with navigation information such as planned paths, waypoints achieved, current locations, and other robot and team specific data. These displays provide the graphical user interface for mission task design, which for HAB ocean sensing may include a marine navigational chart and areas of interest, while for the mission of wide-area mineral prospecting it may include terrain map and a list of prospecting sites. This high-level mission monitoring also provides the feedback needed for successful replanning as necessary. As one example, the lower left monitor in Figure 2 shows each of four robots, their navigated paths, and the points at which prospecting data have been collected.

Where *a priori* and sensor data are integrated, Inference Grid² representations are used. Inference Grids (IG) are a multi-property lattice-based Markov-Random-Field (MRF) model where sensor information is stored in spatially and temporally registered form, and is used for both scientific inferences and for vehicle mission planning. The

information in Inference Grid cells is represented as a stochastic vector, and metrics such as entropy are used to encode the uncertainty in the IG. Measured data fields such as mineral concentrations can be manipulated together with generated maps such as *terrain roughness* that are calculated from composite image analysis of rover image sequences built up into high-resolution regional maps.

These high-level functions are adapted to the task domain: the methods for coordinating prospecting rovers are instantiated for the prospecting task; and the methods for sensor-based adaptive fleet control are instantiated for the marine platforms investigating Harmful Algal Blooms. The Monitoring level also includes Perception — Decision-Making — Actuation sequences to monitor multi-robot operations as a system, and to analyze them for high-level hazard and assistance detection.

C. Robot Team Coordination

The Robot Team Coordination module (Figure 1) decomposes high-level commands and generates robot-specific commands. This provides a separation from the high-level planning that adds flexibility to the system by supporting coordination of robots with very different physical characteristics and even different command protocols. Imagery and telemetry are received from the various robotic assets and conformed to a common data storage schema which can be archived and is also presented graphically to the telesupervisor in the form of dashboards, as explained in more detail in the next subsection.

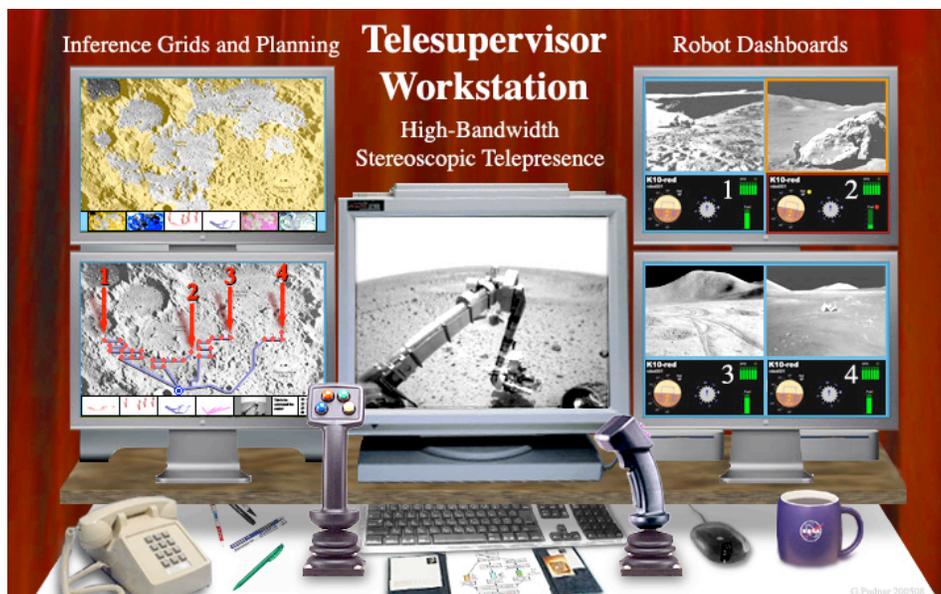


Figure 2. Planetary Telesupervisor Workstation concept.

D. Robot Telemonitoring

The team of robotic assets is constantly monitored at low bandwidth by the human telesupervisor with imagery and data updated regularly from each robot. The Robot Telemonitoring module manages the image and telemetry archives for each robot and presents these data for human and automatic monitoring and analysis. As shown at the right in Figure 2, each robot has a dashboard that includes a regularly-updated image from one of the robot's cameras, and engineering data such as battery charge, attitude, motor temperatures, and other telemetry provided to the operator in both tabular and graphical form. The archive function allows for later mission analysis, but also allows for a *replay* to allow the telesupervisor to quickly review the robot's recent activities, thus providing the context leading up to its current situation. This attention to human factors allows the telemonitor to rapidly assess conditions when a hazardous situation is detected or assistance is required.

E. Telepresence and Teleoperation

We make a distinction between monitoring the operation of every robot, and telepresently taking control of a robot. Whereas 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 for one robot at a time, such as 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 controlled rather than operating under its Autonomous Navigation System; and when the task-specific tools must be operated manually. Appropriate joystick, keyboard, and task-specific human interface devices support this. Robot actions and sensor data continue to be reported through the Robot Controller up to the Robot Team Coordination Module for maintaining a model of the robot's state which is necessary when returned to autonomous control.

This subsystem is ideally 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. This is represented by the parallel data path of Figure 1 (colored blue).

F. Robot Controller

As depicted in Figure 1, the Robot Controller subsystem on each robot is responsible for receiving a collection of lower-level task commands from the Robot Team Coordination module and monitoring its execution on that robot. The Robot controller has direct access to all of the robot's physical systems 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. For example, in addition to a navigation module, a vehicle *health* module can be monitoring at a low level — battery charge and temperature, vehicle attitude, motor currents — and when an out-of-range condition occurs, appropriate action is taken, including sending a message to the Telesupervision Workstation where higher-level Hazard and Assistance Detection functions are monitoring.

III. Application Examples

The Multilevel-Autonomy Robot Telesupervision Architecture has been applied to two very different real-world tasks: investigation of marine-borne Harmful Algal Blooms; and Planetary Prospecting for exploration and *in situ* resource utilization (ISRU). Each of these task application areas is described in the following subsections.

A. Harmful Algal Bloom Investigation

Interest in Harmful Algal Bloom (HAB) detection has grown in recent years for scientific, commercial and public health reasons. The Woods Hole Oceanographic Institute (WHOI) has mapped the distribution of *Alexandrium fundyense* cysts on the sea floor of the Gulf of Maine for both 2005 and 2006.³ The Florida Fish and Wildlife Research Institute has commissioned statisticians to analyze a historical database of concentrations of the HAB dinoflagellate *Karenia Brevis* in Florida waters.⁴ There is also an ongoing effort by the Northwest Fisheries Science Center and collaborators to develop probes for detecting toxins produced by each species of *Pseudo-nitzschia*.⁵

Earth science research requires information obtained from space, the atmosphere, and the ocean to foster understanding of the Earth and its natural processes. Developing a better understanding of ocean processes in particular, is crucial for global warming, meteorological, and ecological studies. Ocean sensing is typically done with satellites, buoys, airborne assets and research vessels. Satellites and airplanes are limited by cloud cover and temporal/geographical coverage and resolution; while research vessels are expensive to deploy, and buoys cannot be self-deployed to specific areas of interest.

The National Oceanic and Atmospheric Administration (NOAA) is addressing some of these constraints through the development of robotic ocean vessels for weather-related ocean monitoring. The OASIS (Ocean-Atmosphere Sensor Integration System) vessels are long-duration solar-powered autonomous surface vehicles (ASVs) designed for global open-ocean operations (Figure 3). Using our Multilevel-Autonomy Robot Telesupervision Architecture to investigate HABs enhances the science value of the OASIS sensing assets by coordinating their operation, and adapting their activities in response to sensor observations. MARTA supports a human telesupervisor monitoring these and other sensing assets such as Earth observing satellites.⁶



Figure 3. Concept of OASIS Platforms investigating an HAB

The Telesupervised Adaptive Ocean Sensor Fleet (TAOSF, *tao'sef*) is an instantiation of MARTA that takes advantage of the integrated approach to multi-vehicle coordination and multilevel human-vehicle autonomy. TAOSF allows the networking of a fleet of autonomous ocean vehicles to study the surface and sub-surface characteristics and the dynamics of such ocean phenomena as coastal pollutants, oil spills, hurricanes or harmful algal blooms. It supports mobile and fixed sensing assets (including buoys and satellites) to function in a cooperative fashion, and the operating mode of different vessels to vary from autonomous control to teleoperated control.

In TAOSF, MARTA supports the following features:

- Multilevel autonomy, allowing an operator to control the vehicles by setting high-level goals, such as specifying an area to monitor, or by taking direct control of the vehicles via teleoperation, or at other autonomy levels in between.
- Adaptive replanning of the activities of the OASIS vessels based on sensor inputs (“smart” sensing) and sensorial coordination between multiple assets, thereby increasing data-gathering effectiveness while reducing the effort required for tasking, control, and monitoring of the vehicles.
- Web-based communications permitting control and communications, and the sharing of data with widely-dispersed remote experts for higher-level scientific analysis and mission planning.
- Autonomous hazard and assistance detection, allowing automatic identification of hazards that require human intervention to ensure the safety and integrity of the robotic vehicles, or of science data that require human interpretation and response.
- Science analysis of the acquired data in order to perform an initial onboard assessment of the presence of specific science signatures of immediate interest.

The TAOSF instantiation of MARTA integrates and extends subsystems developed by the participating institutions. The major TAOSF subsystems are:

1. The OASIS Autonomous Surface Vehicle (ASV)** system includes the vessels themselves as well as the land-based control and communications infrastructure developed for them. Each OASIS Robot Controller module directly controls the hardware of its platform (sensors, actuators, etc.), and also provides a low-level waypoint navigation capability.
2. A task-specific Robot Team Coordination module, named the Adaptive Sensor Fleet (ASF) and developed by the Goddard Space Flight Center, provides platform assignment and automated path planning for area coverage, as well as monitoring of mission progress. Command protocols and collected data are conformed and passed here.
3. The Task Planning and Monitoring module provides high-level planning, monitoring, and telesupervision, as well as analysis of science data from both the OASIS platforms and external sources such as satellite imagery and fixed sensors. These data are used in planning of vessel navigational trajectories for data gathering.
4. The Robot Telemonitoring module provides an operator interface for those occasions when a scientist desires to exert direct monitoring and control of individual platforms and their instruments. Both the Task Planning and the Robot Telemonitoring module are presented to the telesupervisor through an instantiation of the Multi-Robot Operator Control Unit (Figure 4) developed by SPAWAR Systems Center San Diego.⁷

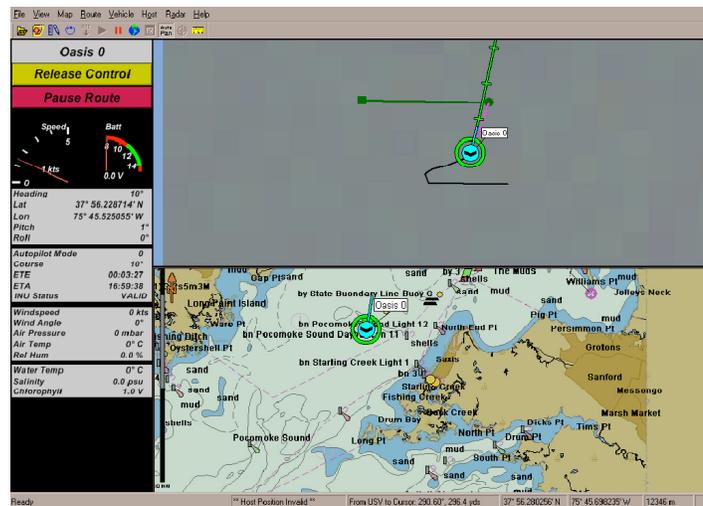


Figure 4. Multi-robot Operator Control Unit.

B. Planetary Mineral Prospecting

Space exploration is one of the costliest and riskiest activities human beings pursue. It is estimated that, for each astronaut who walks on the surface of Mars, it will be necessary to lift some 500,000 tons of cargo, at a cost of

** The OASIS platforms have been developed by Emergent Space Technologies (EST) working with the team members at the NASA GSFC Wallops Flight Facility (WFF), where the platforms are physically built and maintained.

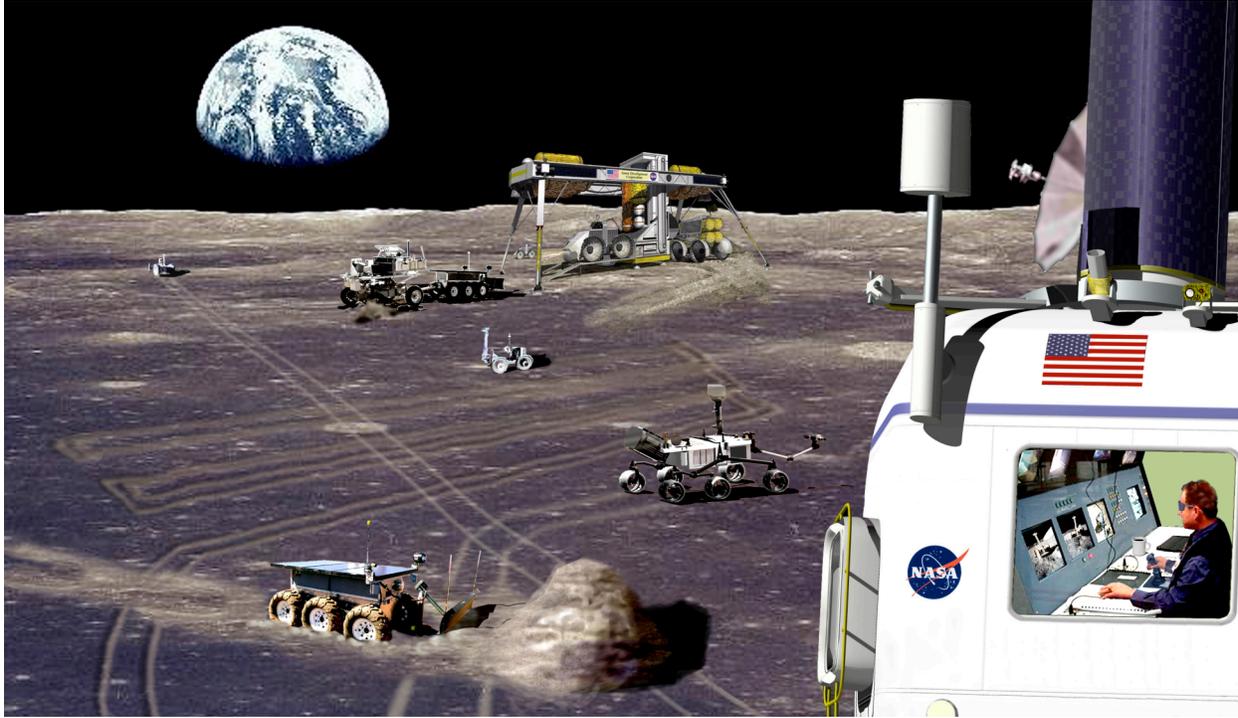


Figure 5. Telesupervision of a very heterogeneous team of ISRU robots.

several hundred billion dollars. Because these missions require extended stays, the use of *in situ* resources is compulsory to provide building materials, breathable atmosphere, and even fuel.

Once there, an astronaut will explore and work, often performing extra-vehicular activities (EVA) wearing a space suit that is heavy and cumbersome. The duration of an EVA is limited by the amount of life-support consumables that can be carried, such as energy and oxygen. More importantly, for experiences both on orbit and on the Moon it is commonly reported that after about five hours of work in a space suit, astronauts must return to their base habitat, exhausted by the intensive effort required.

Robots, on the other hand, have been proven to work very well in extraterrestrial environments, being able to travel autonomously, explore, photograph, collect samples, and even perform limited scientific analysis of materials. The Spirit and Opportunity Mars Exploration Rovers, for example, combined have logged more than 2,400 Martian days, and over 18 km, and are still operating. However, human interaction with the Mars Exploration Rovers is restricted by round-trip communication delays ranging from six to forty-five minutes.

When humans and robots are near each other the communication delays become less significant, and a human telesupervisor can take more direct control of a robot when necessary as if the human were immediately present. By exploiting this scenario of human-supervised robotic exploration and work, we can optimize the human's time by deploying a fleet of robots that operate *as autonomously as possible*, and have the human supervisor provide direct control only when necessary. This paradigm has significant benefits: the robot fleet multiplies the effectiveness of one human; the human supervisor remains in a relatively safe, shirtsleeve environment; and the weight lifted from Earth can be reduced by orders of magnitude.

The Prospecting⁸ instantiation of MARTA follows the same framework as outlined in *Section II. Telesupervision Architecture*, and are analogous to the TAOSF subsystems while being applied to a very different task domain. The major Prospecting subsystems are:

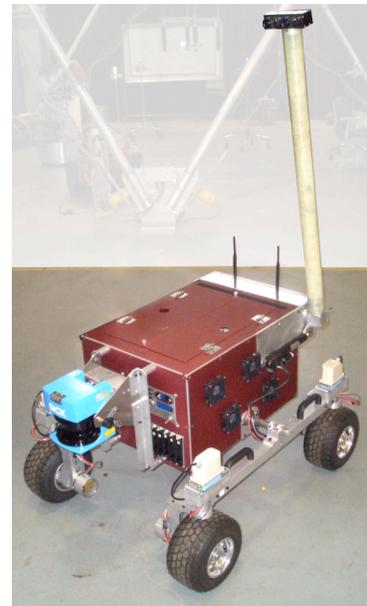


Figure 6. CMU-modified K-10.

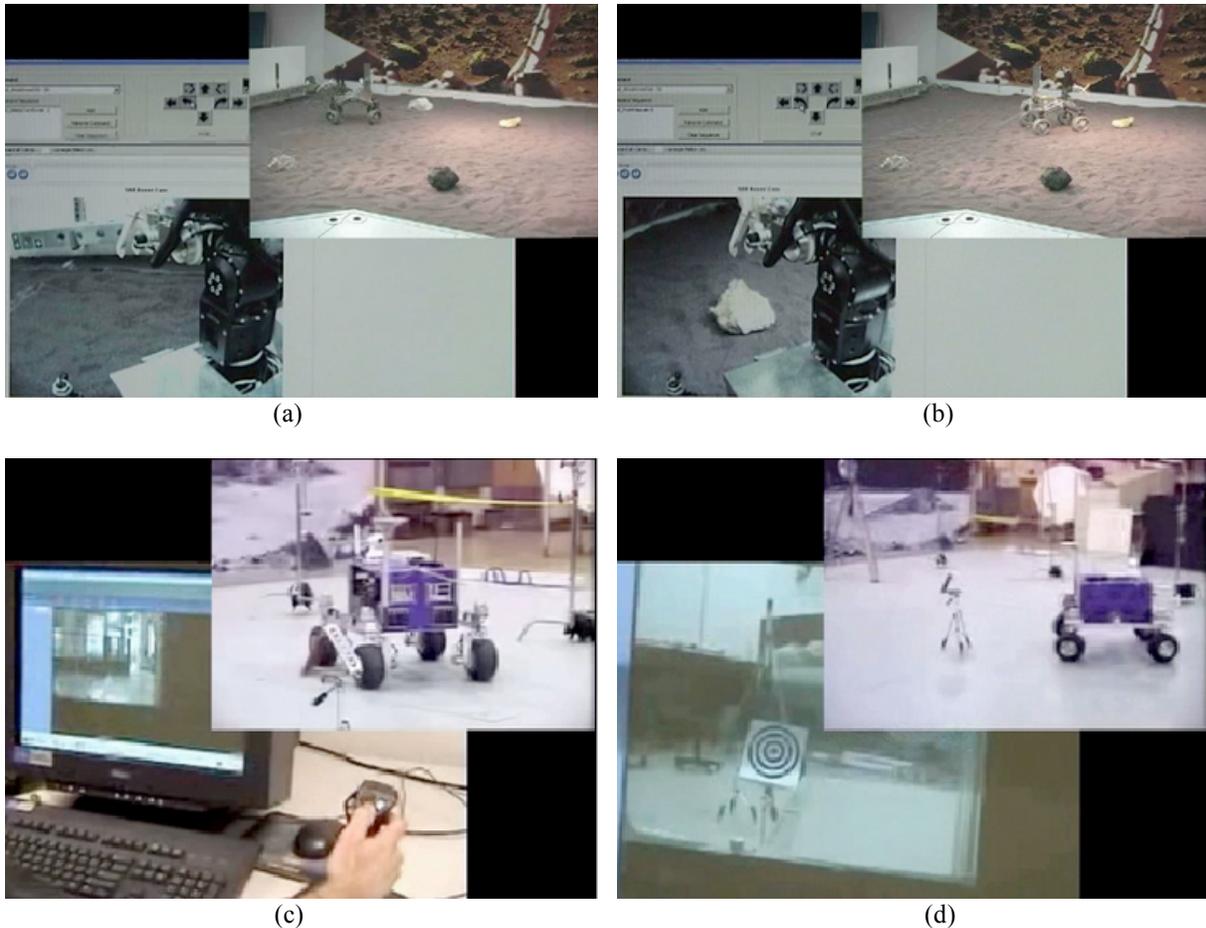


Figure 7. Teleoperation experiments.
(a) Between CMU and NASA JPL, initial position. (b) At “prospecting” site.
(c) Between CMU and NASA ARC, initial position. (d) At “prospecting” site.

1. The Autonomous Prospecting Vehicles included testing with two different physical platforms: the NASA Ames K-10 and the JPL Sample Return Rover (SRR). We augmented the baseline K-10 at Carnegie-Mellon with a scanning laser ranger, and a geometrically-correct stereoscopic telepresence camera system mounted on a deployable arm (Figure 6) that allowed both a high view for teleoperation and a low view for investigating surface rocks. While we were unable to conduct tests on the Moon, we did remotely telesupervise a K-10 at Moffett Field, California, and the SRR in Pasadena, California, both from CMU in Pittsburgh, Pennsylvania (Figure 7).
2. A task-specific Robot Team Coordination module supported efficient multi-robot area coverage algorithms. The prospecting task required that each robot autonomously traverse to the next prospecting site and then request assistance, where the human telesupervisor would perform the actual prospecting task. This is part of the high-level Hazard and Assistance Detection monitoring.
3. The Task Planning and Monitoring module was designed to include functionality for generating high-resolution maps by building up robot-acquired imagery on top of satellite imagery as the prospecting robots traversed the terrain. This was provided using Inference Grid technology. As part of this module’s development, significant scheduling for efficiency work was also done.^{9,10}
4. The Robot Telemonitoring module included a graphical Dashboard which relayed vehicle telemetry (Figure 8).

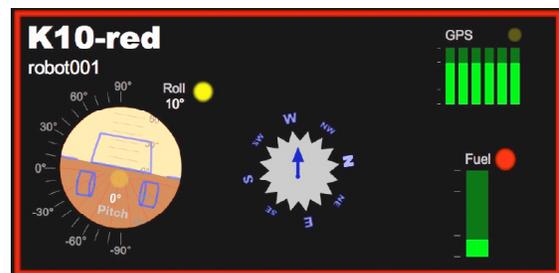


Figure 8. Prospecting vehicle “Dashboard”.

The Telesupervisor Workstation provided a live video display for teleoperation. This was demonstrated by setting up bullseye targets at NASA Ames' "Moonscape" lab, and teleoperating their K-10 with a telepresence camera and a long pointing tool (Figure 7 (c) and (d)). These same experiments using CMU's K-10 with a geometrically-correct stereoscopic camera allowed the teleoperator to approach the targets with much greater accuracy and then to just barely touch them (Figure 9).

For a task such as prospecting, being able to autonomously navigate to a prospecting site, then turn control over to the telesupervisor allows the human to make decisions for precise placement of the sampling tool and to closely monitor delicate operations such as core boring and retrieval.



Figure 9. Teleoperating Prospecting robot using geometrically-correct stereoscopic system.

IV. Telesupervision of Planetary Robot Teams

The summaries of these two multi-robot telesupervision tasks, HAB investigation and Planetary Prospecting, demonstrate the flexibility of MARTA working within these very different domains. While the ongoing program for the Telesupervised Adaptive Ocean Sensor Fleet is focused on this planet, NASA Goddard Space Flight Center's Innovative Partnerships Program¹¹ recognizes that:

Within the context of the Vision for Space Exploration, this system will have applications for lunar, orbital, or planetary construction and inspection; lunar and planetary *in situ* resource utilization; and prospecting, mining, transport, and construction.

Telesupervision of robot teams using MARTA, with its support for multiple levels of autonomy, is significantly different than exploration programs that have been conducted on the Moon and Mars to date: our human exploration of the Moon was conducted by humans primarily out on the surface in space suits; and our remote exploration of Mars is being conducted from the Earth with a roomful of people for each rover communicating for short periods twice per Martian sol to download science data and to upload the next sol's program.

A. Human and Robot Teams

The Human and Robotic Technology Formulation Plan¹ envisioned intelligent multi-agent systems of human and robot teams "...working in close proximity (i.e., shoulder to shoulder), and humans and robots working together but distributed across space and/or time".

Our telesupervision approach to the relationship between humans and robots working together addresses each of these interaction modes with one reservation. In reverse order:

- Working together across *time* involves human analysis of data archived during an autonomous robotic mission, decisions being made for the next desired mission, and providing the plans or programming for the next mission. (This is how the Mars Exploration Rovers are currently being operated.) It is also a useful mode where scientists on Earth analyze archived data from previous missions to suggest future missions.
- Working together across *space* involves humans being some distance from the robotic team members. When this distance gets to be relatively great (again, as is the case with MER) adaptations for the communications time delay are necessary.
- Working together in *close proximity* deserves very great attention now precisely because the plans are to "extend human presence across the solar system". Humans *will* be in close proximity to the robotic assets, unlike most of our previous experiences. However, the concern is over how one defines "close".

Our reservation arises from the conflict between two goals of sustainable space exploration systems as stated in the Formulation Plan: "when astronauts are involved, missions must be *as safe as reasonably achievable* (ASARA)"; and the concept of having humans and robots working "shoulder to shoulder". The half-century that artificially intelligent systems have been researched and developed has given researchers a healthy skepticism about entrusting human safety to machines that can interact physically with the world. Experienced programmers of real-world artificially intelligent systems would most probably not prefer to be in a situation where a robot anticipates that the human needs to be handed a sharp tool (e.g., a screwdriver) while the human is standing in vacuum in a spacesuit.

The goal is certainly to develop such reliable artificially intelligent systems eventually. But their modest abilities now do not preclude taking full advantage of the current state of the art. This has led to our pragmatic approach of combining the strengths of autonomous systems with the ingenuity of a *nearby* human supervisor. Augmenting a human’s expertise with teams of autonomous robots will result in efficient use of the human’s reasoning ability and attention while simultaneously keeping the human in the safest possible environment.

B. Robot Teams with a Human *Nearby*

We suggest that “nearby” means the communication time delay is relatively short and, depending on distance, interactions can be relatively quick or practically immediate. A human in a space station telesupervising robot teams just “outside” in the same orbit is “nearby”, as is a human in a habitat telesupervising robot teams out on the surface.

Pursuing a further definition of “nearby” begins with the distances depicted in Figure 5. Anywhere within geostationary Earth orbit is technically nearby^{††} the Earth’s surface with less than a 250 millisecond round-trip communications delay. These nearby scenarios allow the best of telesupervision, where the robots operate as autonomously as possible, but when one is in need of direct human assistance, the human can efficiently teleoperate the robot through a high-fidelity, low-temporal-delay telepresence system.

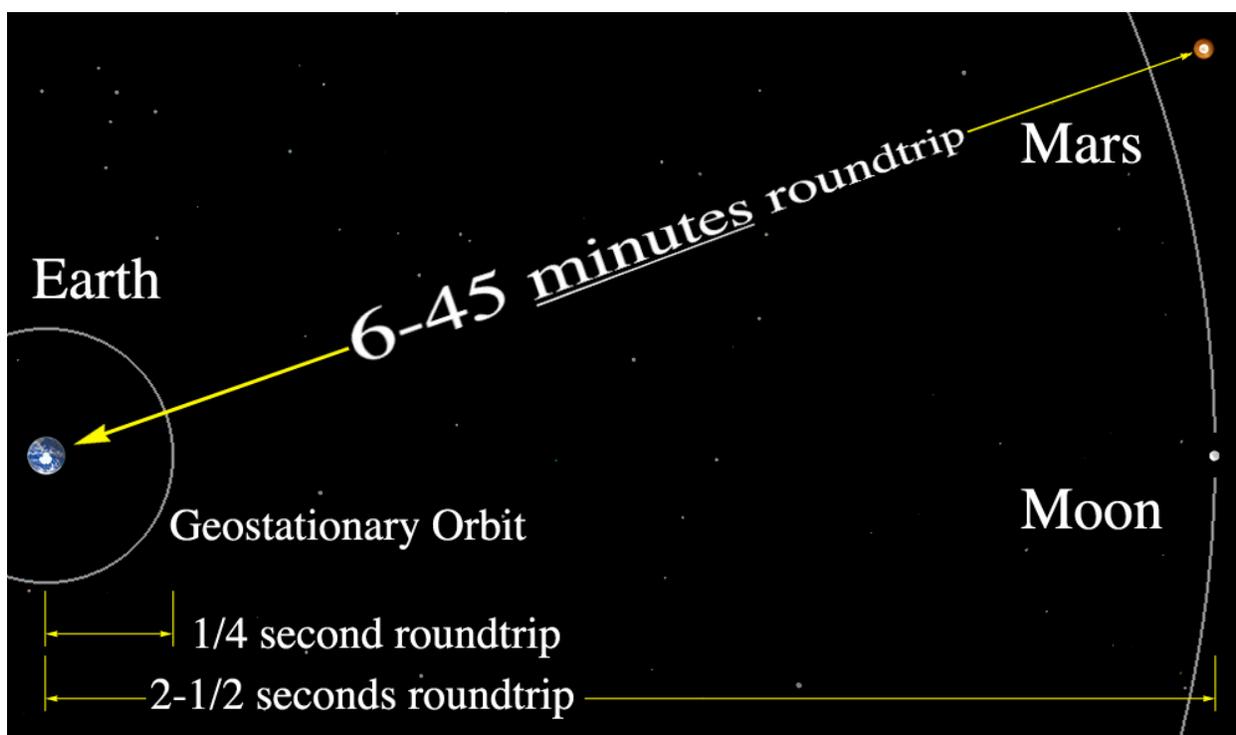


Figure 10. Roundtrip communications delays.

Telesupervision of robotic assets on the Moon from the Earth becomes far less efficient with its order-of-magnitude increase in communications delay. And the long communication delays between the Earth and Mars require that human intervention take the form of teleprogramming.

C. Practically Mars

Taking the approach of developing robotic technology for humans working on the Moon in preparation for sending humans to Mars is an excellent way to prove these technologies (where backup systems can be available) before taking them so far away that we must be able to depend on them unequivocally. By employing

^{††} Practically, communications from the Earth to low Earth orbit (LEO) and geostationary Earth orbit (GEO) for teleoperation currently entail multiple hops through geostationary communications satellites multiplying the roundtrip delay by two or three times or more. As the delay increases, the temporal immediacy decreases and the efficiency of working telepresently through the delay is greatly reduced. Should remote repair of communications satellites become economically interesting for example, it is entirely feasible to upgrade the communications infrastructure to circumvent these delays.

telesupervised-autonomy and high-fidelity telepresence technologies we will be able to mount a telerobotic precursor mission to Mars which puts humans *nearby* Mars without landing them on the surface of the planet. Interestingly, the roundtrip communications delay between a Mars geostationary orbit and the surface is only about one-eighth of a second (Figure 11). This is *nearer* than geostationary orbit above Earth.

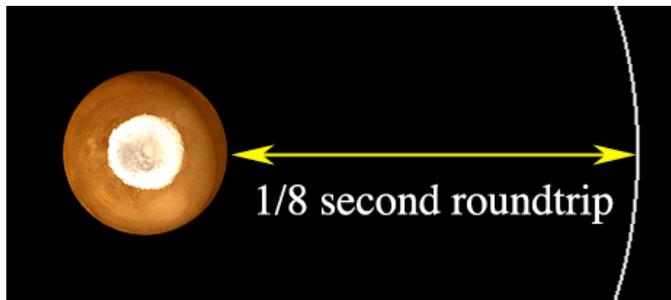


Figure 11. Mars geostationary orbit — very *nearby*.

Geologists on an orbital station might actually have a better experience through telepresence than if working on the surface in bulky spacesuits. Samples can be lifted on small rockets for direct analysis by scientists, and only the best and most varied samples need be returned to Earth for further study, dramatically increasing the potential science return.

By postponing the need to build a launching center and heavy-lift rockets down in the gravity well of Mars, the amount of hardware needed to be transported is reduced by orders of magnitude, and the probability of a safe return is dramatically increased.

D. Heterogeneous Robot Teams

An example of terrestrial testing of a very heterogeneous group of robots is being conducted in the Arizona high desert.¹² These tests by the Desert Research and Technology Studies (RATS) team are investigating human and robotic interaction with a collection of potential extraterrestrial equipment.

A homogeneous team of two K-10 robots in the Haughton Crater Site Survey Field Test¹³ conducted detailed surveys in the high Canadian Arctic of a kind that will be needed at a variety of sites on the Moon and Mars to lay out infrastructure locations, prospect for resources, and plan astronaut excursions. This NASA Ames Intelligent Robotics Group investigation was carried out this year.

These are excellent efforts by dedicated teams braving inhospitable terrestrial conditions to test a wide variety of discrete robotic systems for future space explorations. The ratio of humans-to-robots for these tests is large. Integrating these and nascent robotic technologies by applying our Multilevel-Autonomy Robot Telesupervision Architecture will allow a few astronauts to efficiently telesupervise these very heterogeneous assets.

V. Conclusion

For a sustained human presence on the Moon, and human exploration of Mars, *in situ* resources must be used. This entails substantial and sustained surface operations that must be largely conducted with autonomous systems. Each flight to the Moon and to Mars will bring robotic assets with very different capabilities: prospecting, mining, processing, transporting, construction. The most efficient use of the time resources of the nearby humans is necessary to make tractable the goals of this human exploration.

The flexibility of our multilevel architecture supports integration of higher levels of autonomy and efficient coordination of robot teams. As more intelligent agents are developed they can be directly integrated, thus further reducing the demand on the astronauts' attention, and making these systems-of-systems more efficient.

The most precious resources in space are the human beings. To multiply their effectiveness and keep them safe is fundamental to an affordable long-term future for space exploration. Telesupervised robotic systems have been identified as a key technology by the NASA Exploration Systems Mission Directorate, and are crucial to the success of the Vision for Space Exploration. The systems challenges, for both in-space and on-surface space systems and operations, can only be addressed with higher-level integration and intelligent supervision such as that supported by our Multilevel-Autonomy Robot Telesupervision Architecture.

Acknowledgments

The work herein described was supported by NASA award NNX06AF27G, "Telesupervised Adaptive Ocean Sensor Fleet", granted under the Advanced Information Systems Technology program of NASA's Earth Systems Technology Office. The TAOSF project is a collaboration among Carnegie Mellon University (CMU), NASA Goddard Space Flight Center (GSFC), NASA Goddard's Wallops Flight Facility (WFF), Emergent Space Technologies, Inc. (EST), and the Jet Propulsion Laboratory (JPL). Work on the OASIS platforms is conducted by Emergent Space Technologies, Inc., EG&G, and Zinger Enterprises under award NA03NOS4730220 from the

National Oceanic and Atmospheric Administration (NOAA), U.S. Department of Commerce. Earlier work herein described was supported by NASA under Cooperative Agreement No. NNA05CP96A, "Wide Area Prospecting Using Supervised Autonomous Robots", granted under the Exploration Systems Research and Technology program of NASA's Exploration Systems Mission Directorate. The Prospecting project was a collaboration among Carnegie Mellon University (CMU), NASA Ames Research Center (ARC), and the Jet Propulsion Laboratory (JPL). The statements, findings, conclusions, and recommendations are those of the authors and do not necessarily reflect the views of NASA, NOAA or the U.S. Department of Commerce.

References

- ¹ *Human and Robotic Technology (H&RT) Formulation Plan*, Development Programs Division, Office of Exploration Systems, National Aeronautics and Space Administration, Version 5.1, 29 July 2004 (Amended 13 September 2004).
- ² A. Elfes. "Robot Navigation: Integrating Perception, Environment Constraints and Task Execution Within a Probabilistic Framework". In *Reasoning with Uncertainty in Robotics*, Springer-Verlag, Berlin, Germany, 1996.
- ³ Woods Hole Oceanographic Institution. <http://www.whoi.edu/>
- ⁴ Fish and Wildlife Research Institute. <http://research.myfwc.com/>
- ⁵ Northwest Fisheries Science Center: Harmful Algal Blooms Program. <http://www.nwfsc.noaa.gov/hab/index.html>
- ⁶ Earth Observing-1 (EO-1) <http://eo1.gsfc.nasa.gov/>
- ⁷ Bruch, M. H., "The multi-robot operator control unit (MOCU)", SPIE Proc. 6230: Unmanned Systems Technology VIII, Defense Security Symposium, Orlando, FL, April 18-20, 2006.
- ⁸ Podnar, G.; Dolan J.M.; Elfes, A.; Bergerman, M.; Brown, H.B.; Guisewite, A.D. "Human Telesupervision of a Fleet of Autonomous Robots for Safe and Efficient Space Exploration". *Proceedings of the First Annual Conference on Human-Robot Interaction*, pp. 325-326, Salt Lake City, Utah, March 2006.
- ⁹ Mau, S. and Dolan, J., "Scheduling to Minimize Downtime in Human-Multirobot Supervisory Control", *Proceedings of the 5th International Workshop on Planning and Scheduling for Space*, pp. 268-277, Baltimore, MD, October, 2006.
- ¹⁰ Mau, S. and Dolan, J., "Scheduling for Humans in Multirobot Supervisory Control", to appear in the *Proceedings of the 2007 IEEE International Conference on Intelligent Robots and Systems*, San Diego, CA, November 2007.
- ¹¹ Pekar, N. "2006 Innovative Partnerships Program Updates", *GoddardView – Year in Review*, Vol. 3 No. 2, February 2007.
- ¹² Desert Research and Technology Studies (RATS). <http://www.nasa.gov/vision/space/preparingtravel/desertrats.html>
- ¹³ Haughton Crater Site Survey Field Test. http://ti.arc.nasa.gov/projects/haughton_field/