Planning in the Dark: Robotic Sorties into Lunar Cold Traps (A Preliminary Report)

Robert A. Morris
NASA Ames Research Center
Moffett Field, CA 94035
morris@email.arc.nasa.gov

Abstract

Future robotic missions are being planned in order to explore permanently dark regions of the moon, located in craters near the poles. There the rover will monitor for the presence of hydrogen concentration, and collect and analyze samples in order to *verify* the presence of water ice, and, in addition, and determine the *spatial distribution* of the ice, including location, depth, and concentration. Unlike the case with Mars, short communication delay to the moon makes safeguarded teleoperation of the rover's surface operations a viable strategy for navigation and control in general. Automated planning systems on the ground could also assist ground scientists in generating waypoint-based exploration routes. Nonetheless, the intermittent loss of direct line-of-sight communication near the poles justifies considering an approach that combines autonomous on-board decision making with teleoperation. This paper proposes an approach for autonomously constructing, executing and revising plans for multiple sorties into and out of cold traps on craters. The approach combines ground planning with on-board execution and plan revision.

Motivation

The depositing of water and other volatiles on planetary bodies is a fundamental aspect of the history of the solar system and is intimately connected to the origins of life on Earth. Evidence exists for the presence of ice on the Moon's permanently shadowed areas, called "cold traps", near the lunar poles. Lunar models propose different sources of lunar volatiles, including cometary impacts. These models suggest a subsurface in which ice in a cold trap appears in layers from the earliest (deepest) primal ice to the most recent that occurs near or on the surface. In addition to the scientific value to be gained in verifying the presence of ice, such deposits could be tapped for life support and rocket propellent, thus enabling In-situ Resource Utilization (ISRU) for maintaining permanent human outposts on the moon.

Although important information for confirming the presence of frozen water can be obtained by orbiting spacecraft, true confirmation can only be obtained by the proximity and scale offered by instruments on the surface. Although human explorers could be tasked with this goal, an efficient alternative would be to send robotic explorers, or mixed human-robotic teams.

This paper focuses primarily on the application of autonomous planning and scheduling technology to the problem of sortie exploration in cold traps. Challenging computational issues emerge in the integration of navigation and activity planning, as well as the need for ensuring rover safety while accomplishing mission goals. The next section describes the overall capabilities for instrumentation and mobility required by a rover. Next, the overall sortie exploration planning problem is discussed, including capabilities required for autonomous execution and replanning. The paper concludes with a description of the software components to be used in the implementation of these capabilities.

Rover Requirements and Instrumentation

Rovers on polar lunar sorties will enter cold traps, monitor for the presence of hydrogen concentration, and collect and analyze samples in order to **verify** the presence of water ice, and, in addition, determine the **spatial distribution** of the ice, including location, depth, and concentration. Instruments and tools on the rover that could be used for the unambiguous detection of water ice include:

- A *neutron detector* for monitoring the hydrogen concentration close to the surface around the rover;
- A *cryogenic drill and sample acquisition tool* for obtaining sub-surface cores up to 1 meter below the surface.
- *compositional instruments* such as Tuned Diode Lasers (TDLs) for the molecular identification of $H_2 0$.

A typical sortie will transition from a region of full sunlight, into one with partial darkness, and finally into complete darkness. Desirable timing for sorties will be based on minimizing the distance traversed in partial darkness in order to maximize useful battery life, as well as maintaining Earth line-of-sight. A strobe light synchronized with a camera will allow the rover to illuminate its path. A Neutron Detector will allow for a determination of drill locations based on suitable concentrations of hydrogen deposits.

The safety of the rover is the primary constraint on the length of a sortie; if the rover cannot escape from the cold trap, it is lost. The duration of the sortie is determined by the power requirements for rover mobility and in-situ sampling. For example, one conservative estimate states that a

---

fully charged 178 amp-hr lithium-ion battery allows for a 3.5 hour sortie into permanent darkness for drilling and analyzing samples.

The short communication delay between Earth and moon (approximately 1.5 second one-way time delay) makes safeguarded teleoperation of the rover's surface operations, perhaps combined with navigational and fault detection autonomy, a viable strategy for navigation and control. Automated planning systems on the ground could also assist ground scientists in generating waypoint-based routes. Nonetheless, at the poles, the Earth will not be in view for approximately 75% of the time. During these blackout periods, contact can be maintained through a lunar Orbiter for part of the time (roughly 52 minutes out of every 168-minute of a 1000km lunar orbit period). This still implies intermittent communication with Earth, making continuous control of the rover impossible. When mission costs are considered, autonomy could potentially offer an alternative to using an orbiter for maintaining a high degree of interactivity with a rover. These considerations motivate the discussion that follows.

Previous research in autonomous rover operations has led to demonstrations of autonomous navigation and odometry on Mer (Biesiadecki, Leger, & Maimone 2005), traverse of several kilometers in Mars-analogue settings (Wettergreen *et al.* 2002), instrument placement on science targets (Pedersen *et al.* 2003), sub-surface drilling (Paulsen *et al.* 2006), and opportunistic science (Washington *et al.* 1999). The cold trap lunar sortie scenario has new challenges for autonomy technologies not previously encountered. Specifically, the need to traverse in total darkness over partially or completely unknown terrain to find a promising site to drill below the surface makes fine-grained mission planning on the ground infeasible. On-board planning and execution to ensure safety while accomplishing mission exploration goals requires continuous incremental path planning, and monitoring and predicting of power consumption. The closest episode that resembles this kind of scenario is sol 109 of the MER mission, when Spirit was commanded to drive into terrain that was not previously imaged using its autonomous navigation capability (Biesiadecki, Leger, & Maimone 2005). Nonetheless, the sortie scenario extends the autonomy capabilities over the MER episode by requiring decision-making for drill site selection.

Mission Architecture and Sortie Command Cycle

For the sake of ease in presentation, we consider in this paper a simplified formulation of the sortie scenario, which assumes no communication is possible with the Earth during exploration of the dark region, but only while the rover is in daylight. Second, we assume a mission requirement for multiple sorties, where each sortie consists of a traverse into a dark region, followed by an (optional) drilling and sampling session, followed by a return to a lighted region in Earth view for recharging battery and data downlink. Mission goals of verifying the presence of water, and of mapping the concentration, stratigraphy and distribution of water ice, must eventually be accomplished by the end of some finite time.

Accomplishing a mission consisting of multiple sorties into cold traps for the purpose of verifying and mapping the spatial distribution of water ice requires automated capabilities distributed between ground operations tools and on-board rover capabilities. The capabilities are integrated into a system that allows for a *sortie command cycle* to be performed.

A single command cycle is illustrated in Figure 1. A *nominal plan* is built on the ground. A feasible nominal plan is a set of waypoints and drill activities. Data from previous sorties will offer guidance in the selection of waypoints within the cold traps to explore on the next sortie. However, in general, it is assumed that detailed geological or topological maps of the regions to be explored is scarce or non-existent. Over time the rover will enable the construction of such a map as it returns data acquired from the Neutron Detector and drilling activities.

A sortie plan is uplinked to the rover during the period in which it is recharging its battery in sunlight. The rover executes the plan by navigating towards the waypoints, potentially stopping to take pictures and sending additional data while direct communication with Earth is still possible. Each sortie consists of a period in which the rover is in sunlight, followed by a period in which it is in partial sunlight, followed by a period in total darkness and out of communication contact with the Earth. Once in darkness, the rover autonomously explores a designated region while looking for drill sites. selects a drill site, drills, collects a sample and analyzes it, returns to sunlight and Earth view, and downlinks the data acquired.

As this cycle repeats, a more detailed map of the distribution and contribution of water ice is developed, as well as a more detailed map of the terrain inside the cold trap. The increased knowledge will allow both for better estimates of where the promising exploration areas reside, and also for finding the most effective escape routes to ensure continued safety.

As noted in the Figure, associated with ground cycle activities is an integrated *ground mission operations* tool comprised of capabilities for modeling the terrain under investigation for planning purposes, automated planning, and execution monitoring. Similarly, to accomplish rover tasks a set of *On-board capabilities* for robust execution, incremental path planning, and autonomous drill site selection, are required. The next sections discuss these capabilities in detail.

Ground Planning Operations

From Figure 1 it is clear that the length of a single sortie mission command cycle is hard constrained by the duration of time for the battery on the rover to recharge to full capacity, plus the duration of time it can explore in the cold traps before it must return to the sunlight to recharge, plus the time it takes to downlink the data acquired during the sortie. Let us call the sum of these durations D. Thus, any difference between D and the actual length of the mission cycle must accrue as the result of time spent on the ground building the plan or analyzing the data (assuming for simplicity, as

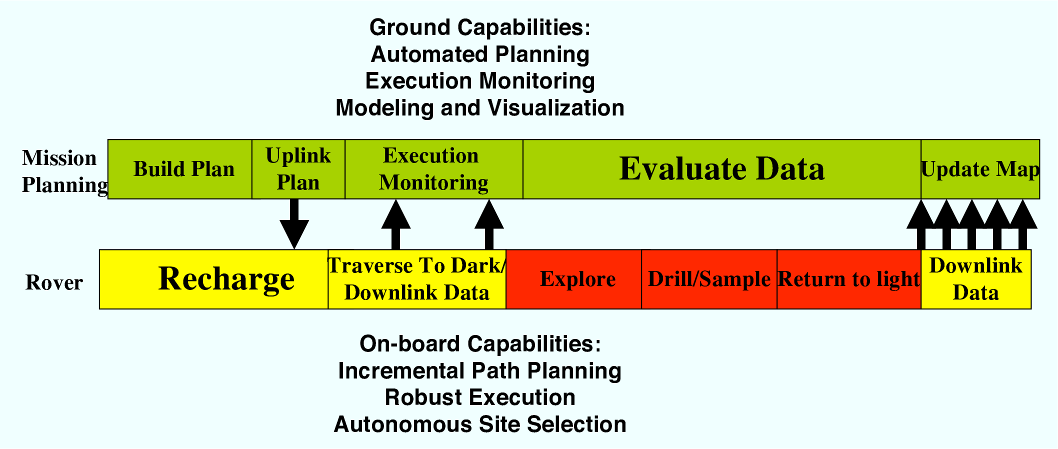

Figure 1: A sortie mission planning cycle and distributed architecture for cold trap exploration, showing timelines for ground and on-board activities

we are here, that no teleoperation of the rover occurs during the sortie mission planning cycle). Clearly, it is desirable to maintain a mission cycle length as close to D as possible; this increases the number of sorties in the mission and hence increases the chances for mission success as well as the overall amount of knowledge gained.

Automation on the ground in the form of mission planning and visualization tools has been demonstrated to be effective in reducing the amount of time devoted to mission planning on MER (Ai-Chang *et al.* 2003). Ground mission operations for multiple cold trap exploration sorties will require similar tools for integrating data acquired from previous sorties into a *visualization tool* to enable selection of waypoints for the next sortie, *automated baseline planning* that integrates temporal flexibility and uncertainty, and *execution monitoring*. This section describes a set of three capabilities for assisting ground teams to meeting these constraints.

Visualization Tool

For the purpose of exploiting the results of cold trap exploration on previous sorties to plan the next sortie, we propose a visual organization of geologic and topological information in graphical form, called a *sortie plan map*. The nodes of the graph are waypoints. Labeled edges between the nodes specify either terrain information or the results of mineralogical analysis.

An example of a sortie plan map for a four-day repeat sortie is found in Figure 2. The graph is divided into three regions: full sunlight, partial sunlight, and total darkness. Nodes indicate waypoints; smiley faces indicate drill sites; the square node indicates planned waypoints that were not reached during the traverse due to decisions made on-board (discussed in more detail below). Edges between nodes are labeled with an indicator of the cost of traversing the edge (e.g. battery resource comsumption). Values in the node are an indicator of the evidence for the existence of water ice in that area (e.g. a reading from the neutron detector or the results of analysis by the tuned diode laser). Dotted edges indicate paths that may have been generated by the plan but were not traversed due to decisions made on-board. The pair of values at the drill sites indicate depth of the drill and the result of the analysis. (It should be noted that these numbers are arbitrary and are meant simply to illustrate the concept of the sortie plan map.)

Ground mission planners will begin a sortie planning cycle by examining the current sortie plan map, i.e., the one recently updated by data acquired from the previous sortie. Each subsequent sortie will extend an existing planning map by adding a new path. There are different strategies for growing a planning map over time, based on a tradeoff between exploring new areas and ensuring safety by re-using data acquired from previous sorties. For example, a *breath-first strategy* favors exploration over exploitation by covering a broader area at shallower levels. A *depth-first strategy* favors exploitation by incrementally extending a single path. More intelligent strategies generate generate plans based on the geological or topological data downlinked from previous sorties. These data will indicate areas to exploit or avoid based either on terrain features or expected science gain.

Baseline Mission Planning

Robotic planning problems, unlike standard planning problems involve:

- *Integration of planning decisions with acting, sensing, and resource monitoring*. In particular, data acquired during the sortie will lead to decisions involving
    - Time allocation. e.g., when to end exploration or drilling;
    - Whether to perform a task or not, e.g. whether to perform a sampling operation at all, given the concentrations of hydrogen detected during an exploration;
    - Path selection, e.g., whether to return to a given area earlier traversed during exploration, or what path to use to return to the light.
- *uncertainty and partial models of the environment*. Initially, there is at best only coarse data available about either the terrain features or the most promising areas to

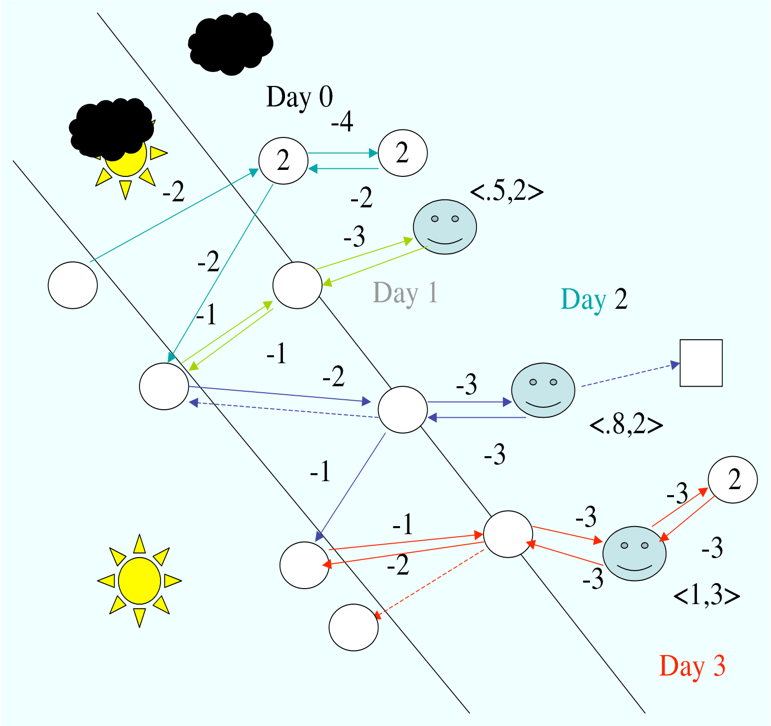

Figure 2: A sortie plan map

explore for water. Over repeated sorties this knowledge is refined, but on the whole the knowledge remains partial.

A baseline sortie plan should satisfy three critical constraints:

1. it should contribute to *rover safety* by being resource constrained and temporally flexible;
2. it should contribute to *accomplishing mission goals* by utilizing all available relevant data about promising water exploration sites acquired by previous sorties; and
3. it should exploit on-board decision-making capabilities by allowing opportunistic replanning of the region explored by the rover, of the site selected for drilling, and of the path chosen by the rover for return.

A *baseline sortie plan* consists of five high-level activities: a traverse to darkness, an exploration in darkness, a (optional) site selection and drill sequence, a return to sunlight and a downlink of data. (This paper ignores issues related to planning and execution of communication downlinks and focuses primarily on issues related to cold trap exploration.)

Baseline sortie plans will address constraint (1) by utilizing environment and rover models to reason about the cost of accomplishing the component actions. Environment models (Tompkins, Stentz, & Whittaker 2004) include terrain models derived from previous traverses or other sources (e.g. orbital data), and lighting models (to determine the durations of light and dark traverse times). Rover models include those for battery (minimum and maximum storage of charge), locomotor (power load for mobility), and drill (power load required for sampling).

To address constraint (2), sortie plans will use the results of previously acquired data from the rover Neutron Detector or composition sensors to plan traverses through promising exploration sites, while ignoring unpromising regions,

Finally, constraint (3) is addressed by deferring the choice of path for part of the sortie to on-board decision-making. Thus, a sortie plan can be viewed as a combination of a set of waypoints and a ”way region”. Formally, a way region is an area defined by a waypoint and a radius around this point. The rover is constrained to explore anywhere within the defined region, but can plan the path through the region on its own. An example of a sortie plan is found in Figure 3. The rover plan consists of a set of waypoints. The way region for the plan is the area within the circle. The path inside the circle was generated incrementally on-board by a path planner.

Ground Monitoring and Control

As noted, it is assumed here that prior to entering cold traps, the rover will undertake part of its traverse in Earth sight. This allows for teleoperated commanded updates to the baseline plan to occur dynamically during this phase of the traverse. This also allows for the rover to acquire and downlink images of the surface features while traversing, either for science purposes or to build higher resolution terrain maps. Dynamic changes to the baseline plan during this phase may have consequences for the amount of time that can be spent exploring in the dark region. Because the baseline plan is refined on-board during execution in the dark region, more command flexibility can be tolerated during execution.

Sortie Plan Execution

A nominal sortie plan will be partially executed in a highly unknown environment in complete darkness without the option of human supervised control. Figure 3 illustrates a complete sortie scenario. The top of the figure shows a set of waypoints. At the time of the plan uplink, the rover is in position A. Waypoints B and C were generated on the ground, incorporating, if possible, data acquired from previous sorties. B is the point at which the rover enters the cold trap (note: this scenario is a simplification insofar as the transition into the cold trap will include an intermediate region in which the rover is partly in sun and partly in dark).

The circle in the figure defines a region within which the rover is allowed to explore autonomously in order to select a drill site; C is the entry point into this region. Hence, in this sortie, waypoints D, D' and E were generated on board.

After entering the cold trap, the rover will turn on its strobe light to enable obstacle avoidance in the dark. It will also begin to monitor of the values returned by the neutron detector (the H-count) as well as the remaining battery charge. At any moment, the energy required to perform the drilling and return to sunlight must not be more than the remaining charge. We define a state variable, the *Drill-Return (DR) value*, as the estimated amount of charge required to drill and return to the light.

The middle graph in Figure 3 shows the interaction between the current battery charge and the DR-value. While in the light (from A to B) the charge does not diminish; also the DR-value is clearly 0 because the rover is already in sunlight. At point B (entry into the cold trap) the DR-value undergoes a jump in value to indicate the amount of charge required simply to perform the drilling to some desired depth and sampling. This value is derived from the drill model. Clearly, if the DR-value ever crosses above the remaining charge value, the rover will not be able to perform the drill and return to light. The on-board executive will monitor this interaction so that this condition never arises. If for some reason these lines become close or touch without crossing, the executive may still decide for safety reasons not to perform the drill and simply return.

Different strategies for autonomously selecting a drill site can be distinguished based on two criteria: first, how to select a path through the region, and second, the condition that must be satisfied for exploration to be terminated and the drill site selected. For example, traversing a ”breadth-first” zig-zag pattern through the exploration region allows for more unknown region to be explored while maintaining a relatively close proximity to the light. In figure 3, by contrast, the rover choses a strategy of traversing a straight line from its entry point C to the center of the region.

The bottom of Figure 3 shows a hypothetical reading of the Neutron Detector as the rover traverses the dark region. Between C and D the H-count increases steadily, then decreases to point E. One simple strategy for terminating exploration would be to traverse until the H-count reaches a predetermined value. Another approach would be to explore until the distance between the remaining battery charge and the DR-value reaches some predetermined threshold.

In Figure 3 a slightly more sophisticated strategy is followed. The rover keeps track of H-values obtained at different locations, and thus has a record of the location at which the H-count reached its highest value (D). After some distance is explored while the count decreases (to E), the rover returns to the location of the highest count ($D = D'$). Notice that on the return to the DR-value decreases proportionally to the decrease in charge, indicating that the rover is on a return path to the light.

After the drill site is selected, the executive executes commands for the drilling and sampling. Based on the remaining charge at the onset of drilling, the drilling duration may be shortened or lengthened. High level control of drilling decisions such as the desired weight on the bit, penetration rate, and rotation rate, can be made based on telemetry received from the drill actuators or down-hole sensors (Paulsen *et al.* 2006). Detecting and recovering from drilling failures will also be performed. If subsurface cores are collected, the data analysis can be conducted while in the darkness, or saved until the return to light (this decision can itself be made during execution).

After drilling, the rover will decide on the return path. The simplest approach would be to backtrack over the path it took from its original position (viz. $D' \rightarrow C \rightarrow B \rightarrow A$ in Figure 3). However, this simple approach is not always optimal. For one thing, the terrain may make the return on the same path prohibitively unsafe. Alternatively, if the rover chose a zig-zag or other non-linear pattern of traverse, the return path could be much longer than a path that followed a straight line back to the region entry point.

Implementation

A complete system for mission planning operations and on-board execution and replanning is being developed using a suite of development and planning tools developed at Ames Research Center and JPL. This section briefly describes each component of the overall system.

Mission Operations Software

For engineering nominal plans for traversal and drilling, we are employing a combination of the EUROPA activity planner (Smith, Jonsson, & Frank 2000) and the TEMPEST mission path planner (Tompkins, Stentz, & Whittaker 2004). A EUROPA model of activities and constraints related to time and resources will be used to develop flexible activity plan sequences. TEMPEST adds the capability to build and utilize rich terrain and resource models to generate optimal traversal paths.

For visualization of data from, and regions explored by previous sorties, we are employing Viz. Viz is a user interface that provides an integrated 3-D display for exploration. Viz provides a mission team with a 3-D model of the remote environment and the ability to simulate robot operations in its virtual 3-D world using an accurate kinematic simulator before issuing commands to the ”real” robot. Viz will be used to select paths towards sortie targets. Viz allows for

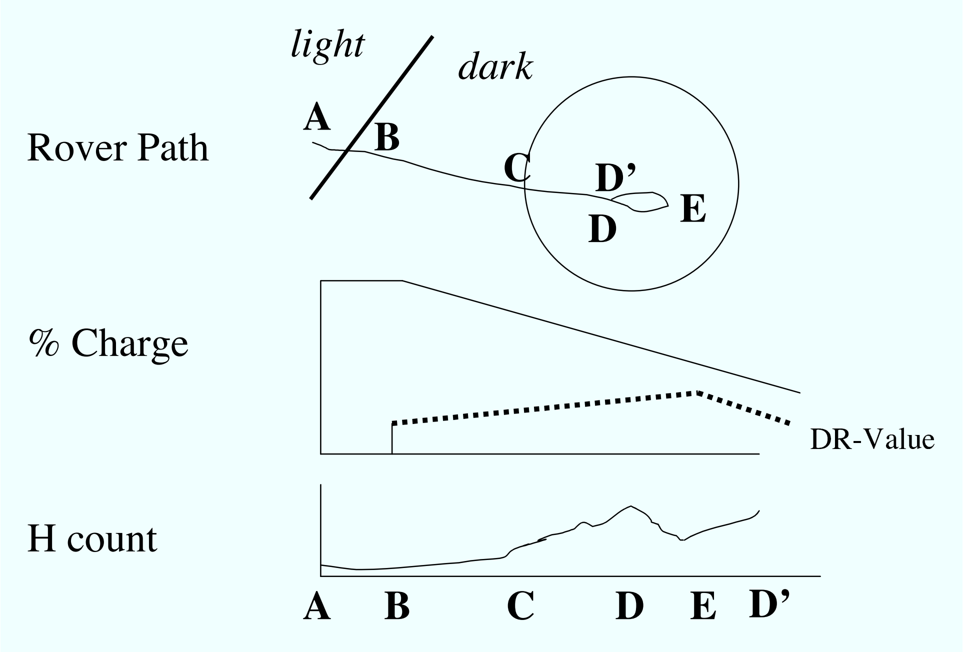

Figure 3: An Autonomous Site Selection Scenario.

the accurate simulation of shadows, which will be important during sortie planning to find routes to the cold traps that maximize time in sunlight.

Finally, we are employing Ensemble, a platform for developing, integrating and deploying mission operations software. Ensemble allows for the data visualization, execution monitoring and planning software components to share the same representation and interfaces. Ensemble is based on open source Eclipse Rich Client Platform(RCP). Using Eclipse's notion of perspectives, users of the Ensemble GUI product are able to downlink and manipulate data, and create new targets using either the downlink or data browsing perspective, and build plans using the planning perspective.

Rover hardware and Software

The autonomy technology will be demonstrated using the cold trap sortie scenario on an experiment using Gromit, an RWI ATRV Jr at NASA Ames. Gromit's functional layer is built out of the LAAS architecture (Alami *et al.* 1998). Gromit uses stereo vision to continuously build a model of its environment. Robust execution will be realized through the integration of the IDEA model-based autonomy architecture and the Plexil-based universal executive.

IDEA (Muscettola *et al.* 2002) supports the development of robust control systems through a uniform agent-based model that uses a reactive planner (RP) as its core reasoning engine. The RP uses an activity model to generate control procedure invocations, based upon the same semantics as the EUROPA activity model. Response-time guarantees are achieved by limiting the time horizon of the RP. The RP will generate path trajectories during the exploration phase of the sortie.

The RP will be invoked by the Plexil universal executive (Verma *et al.* 2005). Plexil is a language for designing execution systems. Control is specified as a set of execution nodes, arranged in a hierarchy, where leaf nodes are command invocations. Attached to each node are conditions that drive node execution. The Plexil Universal Executive interprets a Plexil representation of an execution control instance. Plexil also allows for monitoring resources and the status of executing commands.

Conclusion

The sortie mission scenario described in this paper will be demonstrated in an analogue setting using facilities at NASA Ames Research Center. Future work will document the results of this demonstration.

This paper has described a plan and software architecture for demonstrating autonomy for lunar exploration of cold traps at the poles for the purpose of verifying and mapping concentrations of water ice. The architecture partitions the autonomy capabilities to allow mission planners to control and monitor the development of a detailed map of the features of the explored area.

Acknowledgements

The lunar sortie mission scenario described in this paper was an extension of previous work by Michael Sims, Paul Tompkins, and Brian Glass. Ari Jonsson, Jeremy Frank, Paul Morris, Vijayakumar Baskaran. Kevin Greene, Matthew Boyce, and Lina Khatib contributed to the development of the planning architecture and implementation plan.

References

Ai-Chang, M.; Bresina, J.; Charest, L.; Jonsson, A.; Hsu, J.; Kanefsky, B.; Maldague, P.; Morris, P.; Rajan, K.; and Yglesias., J. 2003. Mapgen: Mixed initiative planning and scheduling for the mars 03 mer mission. In *Proceedings of iSAIRAS-2003*.

Alami, R.; Chatila, R.; Fleury, S.; Ghallab, M.; and Ingrand, F. 1998. An architecture for autonomy. *International Journal of Robotics Research, Special Issue on Integrated Architectures for Robot Control and Programming* 17(4).

Biesiadecki, J.; Leger, C.; and Maimone, M. 2005. Tradeoffs between directed and autonomous driving on the mars exploration rovers. In *Proceedings of ISRR-2005*.

Muscettola, N.; Dorais, G. A.; Fry, C.; and Plaunt, R. L. C. 2002. Idea: Planning at the core of autonomous reactive agents. In *Proceedings of the 3rd International NASA Workshop on Planning and Scheduling for Space*.

Paulsen, G. L.; Mumm, E.; Kennedy, T.; Chu, P.; Davis, K.; Frader-Thompson, S.; Petrich, K.; and Glass, B. 2006. Development of autonomous drills for planetary exploration. In *Lunar and Planetary Science XXXVII*.

Pedersen, L.; Sargent, R.; Bualat, M.; Kunz, C.; Lee, S.; and Wright, A. 2003. Single-cycle instrument deployment for mars rovers. In *Proceedings of i-SAIRAS*.

Smith, D.; Jonsson, A.; and Frank, J. 2000. Bridging the gap between planning and scheduling. *Knowledge Engineering Review* 15.

Tompkins, P.; Stentz, A. T.; and Whittaker, W. R. L. 2004. Mission-level path planning for rover exploration. In *Proceedings of the 8th Conference on Intelligent Autonomous Systems (IAS-8)*.

Verma, V.; Estlin, T.; Jnsson, A.; Pasareanu, C.; Simmons, R.; and Tso, K. 2005. Plan execution interchange language (plexil) for executable plans and command sequences. *International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*.

Washington, R.; Golden, K.; Bresina, J.; Smith, D.; Anderson, C.; and Smith, T. 1999. Autonomous rovers for mars exploration.

Wettergreen, D.; Dias, M.; Shamah, B.; Tompkins, P.; Urmson, C.; Wagner, M. D.; and Wittaker, W. R. L. 2002. First experiment in sun-synchronous exploration. In *International Conference on Robotics and Automation*. 3501–3507.