Planning for Rover Opportunistic Science

Daniel M. Gaines, Tara Estlin, Forest Fisher, Caroline Chouinard, Rebecca Castaño, and Robert C. Anderson

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr., Pasadena CA 91109
{*firstname.lastname*}@jpl.nasa.gov

Abstract.
The Mars Exploration Rover Spirit recently set a record for the furthest distance traveled in a single sol on Mars. Future planetary exploration missions are expected to use even longer drives to position rovers in areas of high scientific interest. This increase provides the potential for a large rise in the number of new science collection opportunities as the rover traverses the Martian surface. In this paper, we describe the OASIS system, which provides autonomous capabilities for dynamically identifying and pursuing these science opportunities during long-range traverses. OASIS uses machine learning and planning and scheduling techniques to address this goal. Machine learning techniques are applied to analyze data as it is collected and quickly determine new science goals and priorities on these goals. Planning and scheduling techniques are used to alter the rover's behavior so that new science measurements can be performed while still obeying resource and other mission constraints. We will introduce OASIS and describe how planning and scheduling algorithms support opportunistic science.

1 Introduction

On Sol 36 of the Mars Exploration Rover mission, the Spirit rover successfully performed its first autonomous traverse using obstacle avoidance software. During these traverses, the rover acquires hazard camera (a.k.a. hazcam) images to look for obstacles and make decisions about the direction to travel. While there is valuable engineering information in these images, there is also a potential wealth of science data. However, by the time these images are downlinked, the rover may be far from the site of interest. In addition, future missions may not downlink these images, in which case the science data could be lost. As the length of autonomous drives increases so does the possibility of passing over some important scientific data. Furthermore, there are short-lived events, such as dust devils, that cannot be adequately studied if the rover must wait for instructions from Earth.

To avoid the potential loss of valuable science, we are developing onboard techniques to support opportunistic science. We are creating a system called OASIS (Onboard Autonomous Science Investigation System) that integrates science analysis and planning to enable the rover to detect potentially interesting science events and re-task the rover to respond appropriately. OASIS includes a science analysis unit that performs onboard processing of collected science data. When a science opportunity is detected, one or more requests are sent to the planning and execution system which attempts to accomplish these additional objectives while still achieving current mission goals.

Opportunistic science poses significant challenges for an autonomous planning and execution system. In many ways, the challenges of handling opportunistic science are similar to dealing with unexpected events and anomalies during plan execution. When an autonomous system detects an anomaly, such as a traverse taking longer than expected or a science activity consuming more power than predicted, the system must assess the impact this event will have on its ability to complete other mission objectives. If necessary, the system will revise the plan in an attempt to achieve as many of the remaining mission objectives as possible, or enter safe mode and wait for assistance from the ground team on Earth. Similarly, when an opportunistic science event arises, the autonomous system must assess the impact that accomplishing these new objectives would have on its current mission goals and, if feasible, alter its plan so that it can achieve these new objectives while still accomplishing the original mission goals.

As with anomalies, it is difficult to predict when a science opportunity will arise and, hence, what the state of the rover will be when the opportunity presents itself. As a result, the specific state of the rover, including its location, resource availability and resource constraints, will not be known ahead of time. It can also be difficult to predict what type of science operation will be called for ahead of time. For example, depending on the type of event, the science analysis software may request an additional image or a spectrometer measurement. The different observations types place different demands such as time, power and memory, on the rover. Because of these uncertainties, it is difficult to determine ahead of time which opportunities can be achieved and what steps need to be taken to accomplish the new science objectives.

Given the similarities between handling anomalous events and opportunistic science, we were able to leverage our previous work with the Continuous Activity Scheduling, Planning and Re-Planning (CASPER) system as a basis for our opportunistic science planning system within OASIS (Estlin et al.(2002)).

In contrast to anomalous events, CASPER has more control over opportunistic science. In particular, the system can decide whether or not to pursue the opportunities that are identified. Much of our work on extending CASPER to support opportunistic science has been in enabling the system to make decisions about whether or not to pursue opportunistic science and which opportunities to accomplish. Our extensions include a Science Alert protocol to enable the data analysis algorithms to communicate new science goals to the planner and a set of plan modification functions to assist the planner in reasoning about these new objectives.

In the next section we provide an overview of our integrated science analysis and planning system that supports opportunistic science. We will then describe a series of scenarios that we have used to develop and test our system in simulation and on rover prototype hardware in the Jet Propulsion Laboratory Mars Yard. These scenarios demonstrate our current capabilities in responding to opportunistic science events.

2 OASIS

Our initial emphasis in OASIS has focused on image analysis and the characterization of surface rocks. Rocks are among the primary features populating the Martian landscape and the understanding of rocks on the surface is a first step leading to more complex regional geological assessments. Figure 1 shows the main components of the OASIS system and how they interact to analyze images of rocks and re-task the rover to respond to opportunistic science events. OASIS consists of:

Feature Extraction: detects rocks in images and extracts rock properties (e.g. shape and texture).

Data Analysis: uses extracted features to assess the scientific value of the planetary scene and to generate new science objectives that will further contribute to this assessment.

Planning and Scheduling: dynamically modifies plan in response to new science requests.

The feature extraction and data analysis components of OASIS have been described previously in (Castano et al.(2004)). Here we will give a brief overview of these components and concentrate on the planning and scheduling unit and how it supports opportunistic science.

2.1 Feature Extraction

The first step in the OASIS system is analyzing rock features from images taken by rover cameras as the rover traverses. The image is segmented using a rock detection algorithm based on edge detection and tracing. Next, a set of properties is extracted from each rock. Our feature extraction priorities are based upon our knowledge of how a geologist in the field would extract information. Important features to look for and categorize include albedo (an indicator of rock surface reflectance properties), visual texture (which provides valuable clues to mineral composition and geological history), shape, size, color and arrangement of rocks. Currently our system identifies the first three of this set; future work will expand this to cover additional features.

2.2 Data Analysis

After features have been extracted from each rock, OASIS runs a set of data analysis algorithms to look for interesting rocks. Two of these algorithms can result in the generation of science alerts: key target signature and novelty detection.

Key Target Signature: enables scientists to efficiently and easily stipulate the value and importance of certain features. Scientists often have an idea of what they expect to find during a rover mission and/or are looking for specific clues that reflect signs of life or water (past or present). Using this technique, target feature vectors can be pre-specified and an importance value assigned to each of the features. Rocks are then prioritized as a function of the weighted Euclidean distance of their extracted features from the target feature vector.

Novelty Detection: detects and prioritizes unusual rocks that are dissimilar to previous rocks encountered. We have looked at three different learning techniques for novelty detection: distance-based using k-means clustering, probability-based using Gaussian mixture models and discrimination-based using kernel one-class classifier.

2.3 Science Alert Protocol

Using the above algorithms, the data analysis software can flag rocks that should be further analyzed and produce a new set of measurement goals. We call this capability the *science alert*, since it alerts other onboard software that new and high priority science opportunities have been detected.

A science alert may involve several different levels of reaction. The most basic reaction is to adjust the rover plan so that the rover holds at the current position and the flagged data is sent back to Earth at the next communication opportunity for further analysis. The next level of reaction would likely be to collect additional data at the current site before transmitting data to Earth. Further steps include having the rover alter its path to get closer to objects of interest before taking additional measurements. These operations would provide new data that could not be obtained through image analysis alone.

The data analysis unit uses the *science alert protocol* shown in Figure 2 to communicate opportunistic science requests to the planner. The protocol consists of two message types. When the rover receives a **Stop and Call Home** message, it responds by altering its plan so that it remains near the target location (tx, ty, tz) until the next communication opportunity. The **Data Sample Request** message represents a request for an additional science measurement. In addition to the target location, this message includes a priority.

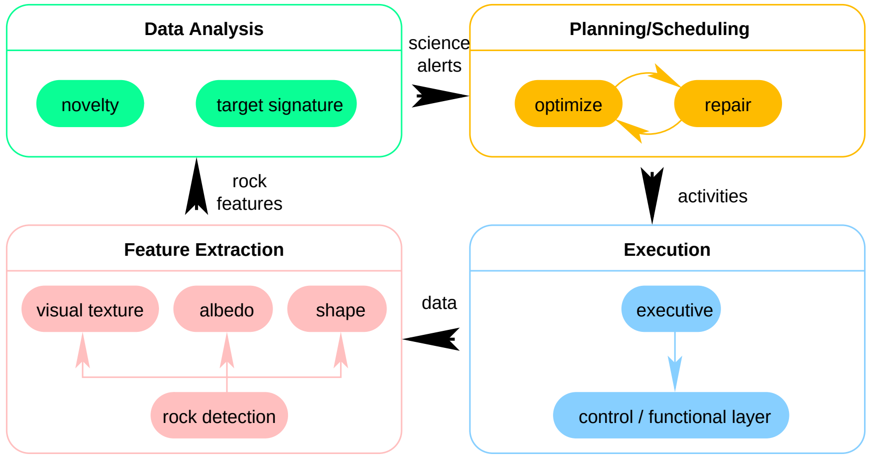

Figure 1: OASIS architecture.

If multiple data sample requests are received, the priority is used to decide which alerts to give preference to in the case that they cannot all be achieved within the current time and resource constraints. In future work, we will also use priorities to compare the value of new science opportunities with objectives already in the plan. Data Sample Requests also specify the type of measurement (e.g. image, spectrometer, . . .) that should be collected.

stop_and_call_home
tx = *real*; ty = *real*; tz = *real*;

(a) Stop and Call Home

data_sample_request
priority = *int*; data_type = {*image, spectra, panorama, libs, . . .*}; tx = *real*; ty = *real*; tz = *real*; heading = *real*;

(b) Data Sample Request

Figure 2: The science alert protocol.

2.4 Planning and Execution

The new science targets are passed to onboard planning and scheduling software that can dynamically modify the current rover plan in order to collect the new science data. This component takes as input the new science requests, the current rover command sequence (or plan), and a model of rover operations and constraints. It then evaluates what new science tasks could be added to the current plan while ensuring other critical activities are preserved and no operation or resource constraints are violated.

Planning and scheduling capabilities in OASIS are provided by CASPER (Estlin et al.(2002); Chien et al.(2000)), which employs a continuous planning technique where the planner continually evaluates the current plan and modifies it when necessary based on new state and resource information. Rather than consider planning a batch process, where planning is performed once for a certain time period and set of goals, the planner has a current goal set, a current rover state, and state projections into the future for that plan. At any time an incremental update to the goals or current state may update the current plan. This update may be an unexpected event (such as a new science opportunity) or a current reading for a particular resource level (such as power). The planner is then responsible for maintaining a plan consistent with the most current information.

A plan consists of a set of grounded (i.e., time-tagged) activities that represent different rover actions and behaviors. Rover state in CASPER is modeled by a set of plan timelines, which contain information on states, such as rover position, and resources, such as power. Timelines are calculated by reasoning about activity effects and represent the past, current and expected state of the rover over time. As time progresses, the actual state of the rover drifts from the state expected by the timelines, reflecting changes in the world. If an update results in a problem, such as an activity consuming more memory than expected and thereby over-subscribing RAM, CASPER re-plans, using iterative repair (Zweben et al.(1994)), to address conflict.

Plan optimization: CASPER includes an optimization framework for reasoning about soft constraints. User-defined preferences are used to compute plan quality based on how well the plan satisfies these constraints. Optimization proceeds similar to iterative repair. For each preference, an optimization heuristic generates modifications that could potentially improve the plan score.

OASIS uses this optimization framework to decide how to respond to science alerts. A science alert comes into the system as an *optional goals*, which is a soft constraint indicating that the plan's quality will be improved if the goal is achieved. Because it may not be possible to accomplish optional goals, CASPER protects the plan from corruption by saving a copy of the plan before optimizing. If the quality has not increased after a pre-defined number of iterations, the previous plan is restored, and CASPER tries optimizing again. To prevent CASPER from churning away endlessly on the same optional goal, we keep track of the number of times optimize attempts to satisfy it. After a certain number of times, CASPER gives up on the goal and throws it out.

Responding to science alerts: We created a set of plan modification functions that are invoked when the optimizer attempts to satisfy a science alert. How the plan is modified depends on the type of alert that is considered. When a Stop and Call Home alert is received, the planner alters the plan to remove any non-engineering critical activities. If the activities are already executing, the planner requests that the executive abort them. If the activities are scheduled in the future, the planner deletes them and resolves any inconsistencies created by these deletions.

To achieve a Data Sample Request, the system must generate a plan that achieves the new goal without deleting existing activities. It adds the new goal to the plan and attempts to resolve any conflicts that might arise. For example, the planner might need to add a traverse activity to get to the location of interest. As with a Stop and Call Home request, the planner must abort currently executing activities, but it does not delete activities that are scheduled to be executed in the future.

3 Testing and Evaluation

We created several scenarios to develop and test OASIS opportunistic science capabilities. We will use four scenarios to highlight the system's current abilities. All scenarios begin with the map show in Figure 3 (a). Here, the rover must take an image of the rock at the far right of the map and downlink the data to Earth. Figure 3 (b) shows an excerpt from the plan to accomplish this goal. The rover will traverse to the rock, take an image and perform the downlink. To facilitate testing, we created a science alert generator that sends science alerts at pre-determined times. We used this capability with the following scenarios to send different alerts as the planner executes the plan.

Scenario 1: Stop and Call Home In the first scenario, the planner receives a Stop and Call Home alert in the middle of the traverse. Figure 4 shows the results. The traverse is aborted and the image is deleted. The downlink is tagged as being engineering-critical, so it is preserved in the plan. It is not moved up in time as it is temporally constrained at a specific time.

Scenario 2: Data Sample Request This time the planner receives a Data Sample Request in the middle of the traverse. After the traverse is aborted, it is able to insert a new traverse to the new science target along with an image activity and a traverse to get back to the original goal.

In order to complete opportunistic science, the planner must be able to find sufficient time in the schedule to perform these extra activities. Extra time might be put into the schedule specifically to allow for opportunistic science. Alternatively, the rover may be able to take advantage of extra time put in the schedule due to conservative estimates of activity durations. This extra time is commonly added to mission schedules to help increase robustness. Often, temporal padding is added between activities in case it takes the rover longer than expected to complete tasks. The planner may also use over-estimates of how long activities take to perform. If, for example, the rover traverses faster than expected, the planner can make use of this spare time for opportunistic science. We have tested with both possibilities. The planner can move activities around, if permitted by temporal constraints, to take advantage of extra time in the schedule. We also experimented with traverses going better than expected by having the rover (either simulated or real) drive faster than the planner predicted. In this case, the planner takes advantage of the extra time gained during the traverse to insert opportunistic science activities. Figure 5 illustrates the latter approach in which extra time is gained because the rover traversed faster than expected.

It should be noted that a third possibility of extra time is to delete other activities in the schedule to make room for opportunistic science. We will explore this option in future work.

Scenario 3: Data Sample Request that cannot be achieved This scenario is identical to the previous except that the system begins in a slightly different state. In this case, the system has additional data in RAM. When the science alert is received, the planner finds that it cannot collect the extra data without over-subscribing RAM. Figure 6 illustrates the problem found by the planner. As a result, the planner gives up on the alert and continues on to the original goal.

Scenario 4: Multiple Data Sample Requests This scenario demonstrates the use of priorities when multiple alerts are received. The rover begins with an empty RAM buffer and receives three Data Sample Requests. Only two of these requests can be accomplished without exhausting memory. As seen in Figure 7 the two alerts with highest priority are included in the plan.

Testing environments: We have successfully tested the planning system's response to opportunistic science with real, prototype rover hardware. These hardware tests were performed in the JPL Mars Yard using the Fido rover.

4 Related Work

The idea of having a scientific discovery system direct future experiments is present in a number of other systems. Work on learning by experimentation, such as IDS (Nordhausen & Langley(1993)) and ADEPT (Rajamoney(1990)), varied certain quantitative and qualitative values in the domain and

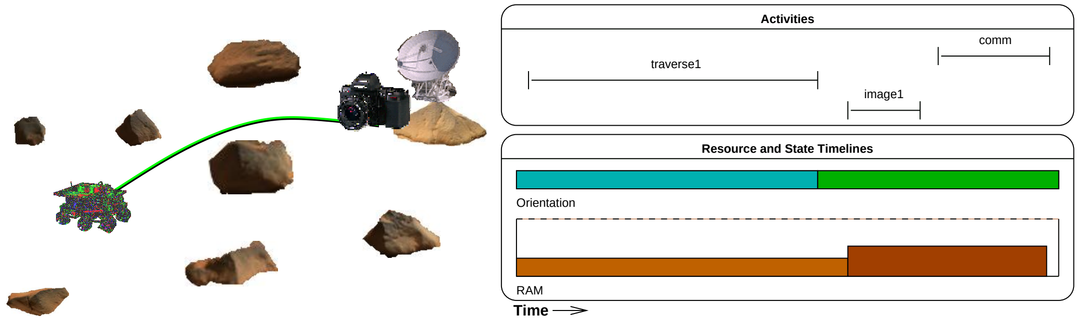

(a) map

(b) plan excerpt

Figure 3: Initial plan for scenarios.

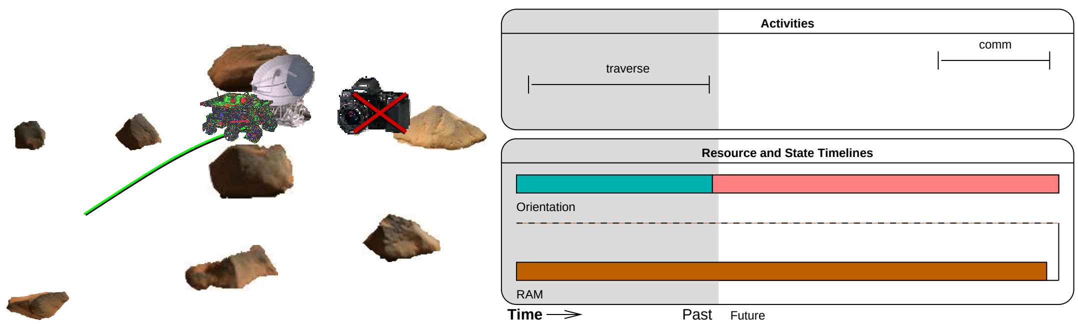

(a) map

(b) plan excerpt

Figure 4: Results for scenario 1: stop and call home.

then measured the effects of these changes. OASIS differs from these systems in that it interacts with the environment to perform experimentation, and it is specialized to address particular problems and scenarios in planetary science. OASIS is also integrated with a planning system, which constructs the detailed activity sequence needed to perform new science experiments.

Several researchers have addressed methods for extracting features from data with the intention of performing the operations onboard a spacecraft. (Gulick et al.(2001)) presented methods for locating rocks in an image using information about the sun angle, identifying the horizon and recognizing layers. There has also been work on developing a framework for feature extraction and event detection for use onboard Earth orbiting satellites (Tanner et al.(2001)). Our work has specifically focused on identifying and analyzing rocks in grayscale images thus far and, in contrast to the work mentioned here, takes the next step of using the feature extraction to determine desirable additional actions a rover could autonomously take.

The objectives of OASIS are similar to those of the Autonomous Sciencecraft Experiment (ASE) (Sherwood et al.(2003)) which also uses science analysis to generate additional goals for a planner. OASIS differs from ASE in the types of feature extraction and data analysis that are performed. In addition, while ASE has focused on planning for orbiter missions, the focus for OASIS has been on ground operations. To support this type of planning OASIS must deal with the high degree of uncertainty inherent in ground operations and integrate path planning into the planning and scheduling process. Finally, in OASIS it is often necessary to temporarily halt currently executing activities, such as a traverse, in order to accomplish new science goals.

A number of other systems have used planning methods to coordinate robot behavior (e.g. (Bonasso et al.(1997); Alami et al.(1998))). However, these systems generate plans with a batch approach where plans are generated for a certain time period and if re-planning is required, an entire new plan must be produced. In OASIS, plans are continuously modified in response to changing conditions and goals. The CPS planner generates contingent plans which are then executed onboard a rover and can be modified at certain points if failures occur (Bresina et al.(1999)). Since only a limited number of contingencies can be anticipated, our approach

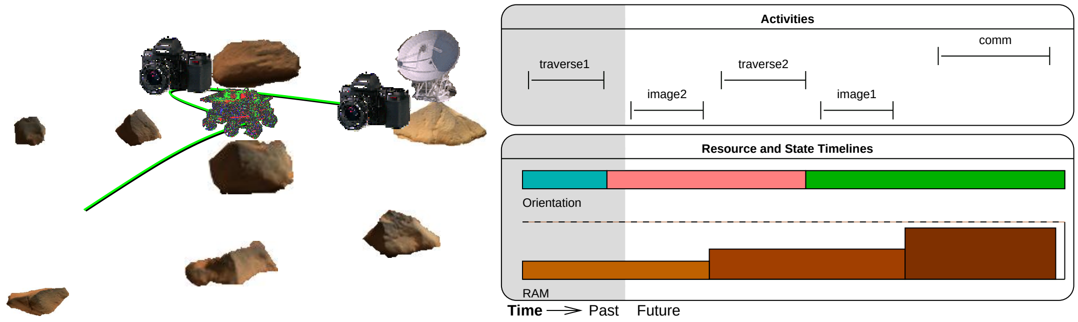

(a) map

(b) plan excerpt

Figure 5: Results for scenario 2: data sample request.

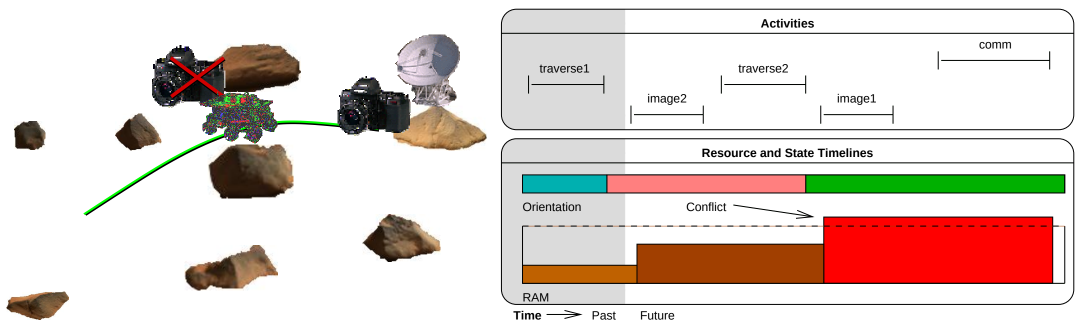

(a) map

(b) plan excerpt

Figure 6: Reason for not performing opportunistic science in scenario 3: cannot achieve request due to RAM over-subscription.

provides more onboard flexibility to new situations. If a situation occurs onboard for which there is not a pre-planned contingency, the rover must be halted to wait for communication with ground.

5 Conclusions

OASIS supports opportunistic science by integrating data analysis algorithms, which identifies potentially interesting science measurements, with planning and scheduling algorithms, which enables the rover to respond to these new requests. Our current system has been tested with several scenarios in simulation and on prototype rover hardware. In these scenarios we demonstrate a spectrum of responses to opportunistic science from halting activity and waiting for communication with Earth, to acquiring additional measurements and proceeding with the original mission objectives.

We are still developing OASIS and there are several capabilities we will be adding. One of the challenges in planning for opportunistic science is finding "spare" time in the schedule in which new activities can be inserted. Currently, we either have the planner add slack time to the schedule during initial plan creation or we allow the rover to make up time by traversing faster than expected. In future work we will consider allowing the system to plan ahead for opportunistic science so that it can decide to schedule extra time in some situations but not in others. This will have the benefit of allowing ground personnel to control when and how much opportunistic science is permitted.

Currently, the planner preserves the original mission goals when attempting to perform opportunistic science. We will relax this constraint and allow the system to use priorities to determine when it is appropriate to achieve opportunistic science at the cost of existing goals. There are significant challenges with introducing autonomous techniques into the mission operations culture. We are taking steps to address this by introducing MER scientists to off-line versions of our software.

Acknowledgments

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work was funded by the NASA Intelligent Systems program and the JPL Interplanetary Network

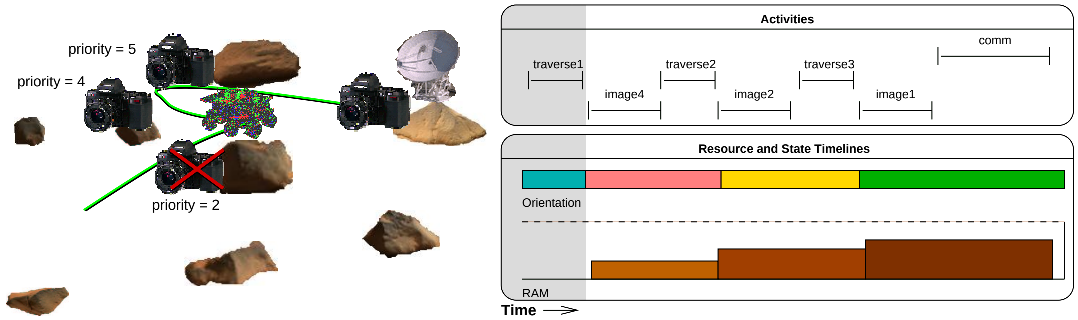

(a) map

(b) plan excerpt

Figure 7: Results for scenario 4: multiple data sample requests.

Directorate. We wish to thank the other members of the OASIS team for their work on feature extraction and data analysis and the CLARAty team, especially Issa Nesnas and Max Bajrachary, for their assistance with the rover.

References

Alami, R., Chautila, R., Fleury, S., Ghallab, M., & Ingrand, F. (1998). An architecture for autonomy. *International Journal of Robotics Research*, **17**(4).

Bonasso, R., Firby, R., Gat, E., Kortenkamp, D., Miller, D., & Slack, M. (1997). Experiences with an architecture for intelligent, reactive agents. *Journal of Experimental and Theoretical Artificial Intelligence Research*, **9**(1).

Bresina, J., Golden, K., Smith, D. E., & Washington, R. (1999). Increased flexibility and robustness for Mars rovers. In *Proceedings of the Fifth International Symposium on Artificial Intelligence, Robotics, and Automation in Space*, Noordwijk, Netherlands.

Castano, R., Judd, M., Estlin, T., Anderson, R., Scharenbroich, L., Song, L., Gaines, D., Fisher, F., Mazzoni, D., & Castano, A. (2004). Autonomous onboard traverse science system. In *IEEE Aerospace Conference*, Big Sky, Montana.

Chien, S., Knight, R., Stechert, A., Sherwood, R., & Rabideau, G. (2000). Using iterative repair to improve the responsiveness of planning and scheduling. In *Fifth International Conference on Artificial Intelligence Planning and Scheduling*, Breckenridge, CO.

Estlin, T., Fisher, F., Gaines, D., Chouinard, C., Schaffer, S., & Nesnas, I. (2002). Continuous planning and execution for a mars rover. In *Third International NASA Workshop on Planning and Scheduling for Space*, Houston, TX.

Gulick, V., Morris, R., Ruzon, M., & Roush, T. (2001). Autonomous image analyses during the 1999 marsokhod rover field test. *Journal of Geophysical Research-Planets*, **106**(E4), 7745–7763.

Nordhausen, B. & Langley, P. (1993). An integrated framework for empirical discovery. *Machine Learning*, **12**, 17–47.

Rajamoney, S. (1990). A computational approach to theory revision. In J. J. Shrager and P. Langley, editors, *Computational Models of Scientific Discovery and Theory Formation*, pages 225–254. Morgan Kaufman, San Mateo, CA.

Sherwood, R., Chien, S., Castano, R., & Rabideau, G. (2003). The autonomous sciencecraft experiment. In *Proceedings of the IEEE Aerospace Conference*, Big Sky, MT.

Tanner, S., Keiser, K., Conover, H., Hardin, D., Graves, S., Regner, K., Wohlman, R., Ramachandran, R., & Smith, M. (2001). EVE: An on-orbit data mining testbed. In *IJCAI-01 Workshop on Knowledge Discovery from Distributed, Heterogeneous, Autonomous, Dynamic Data and Knowledge Sources*, Seattle, Washington.

Zweben, M., Daun, B., Davis, E., & Deale, M. (1994). Scheduling and rescheduling with iterative repair. In M. Fox and M. Zweben, editors, *Intelligent Scheduling*, pages 241–256. Morgan Kaufmann Publishers Inc.