

Enabling Autonomous Science for a Mars Rover

Tara Estlin¹, Rebecca Castano², Daniel Gaines³, Benjamin Bornstein⁴, Michele Judd⁵ and Robert C. Anderson⁶
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

The Onboard Autonomous Science Investigation System (OASIS) evaluates geologic data gathered by a planetary rover. This analysis is used to prioritize the data for transmission, so that the data with the highest science value is transmitted to Earth. In addition, the onboard analysis results are used to identify and react to new science opportunities. A planning and scheduling component of the system enables the rover to take advantage of identified science opportunities. In this paper, we provide an overview of the OASIS system and report on our experience testing this software with a Mars rover prototype. In particular we discuss how such capabilities can be enabled during ground operations planning and how this increased autonomy will affect downlinked data. We also introduce a new area of OASIS work, which is to provide autonomous targeting capabilities for the MER rovers.

I. Introduction

The Mars Exploration Rovers (MER 2003) have traveled many kilometers over the Martian soil and have outperformed all expectations by lasting an order of magnitude longer than their original mission goal. Both rovers have now lasted over 1400 sols (or Martian days), whereas their primary mission was for 90 sols. The longevity of these vehicles will have significant effects on future mission goals, such as objectives for the Mars Science Laboratory rover mission which is scheduled to fly in 2009 or for a Mars Sample Return Mission, which could fly as early as 2018. Surface rovers offer scientists the ability to move around a planetary surface and explore different areas of interest. The farther a rover can travel, the greater the opportunity exists for increased scientific discovery. Due to advances in rover navigation and other onboard software, traverse distances are increasing at a rate much faster than communications bandwidth. While the Sojourner rover traveled around 100m in the entire mission, the MER rovers have now traveled combined over 19km. As this trend in increased mobility continues, the quantity of data that can be returned to Earth per meter traversed is reduced. Thus, much of the terrain the rover observes on a long traverse may never be observed or examined by scientists. We present a system developed to maximize the quality of the science data transmitted to Earth by in situ missions that enables new science opportunities to be identified and handled onboard a rover platform.

The Onboard Autonomous Science Investigation System (OASIS) system¹ has been developed to evaluate, and autonomously act upon, science data gathered by in-situ spacecraft such as planetary landers and rovers. OASIS analyzes the geologic data gathered by the rover onboard. This analysis is used to identify terrain features of interest and additional science gathering opportunities. A planning and scheduling component of the system enables the rover to take advantage of the identified science opportunity by updating the command sequence to include the opportunistic measurements. OASIS currently works in a closed loop fashion with onboard control software (e.g., navigation and vision) and has the ability to autonomously perform the following sequence of steps: analyze gray scale images to find rocks, extract the properties of the rocks, identify rocks of interest, re-task the rover to take additional imagery of the identified target and then allow the rover to continue on its original mission.

This paper provides an overview of the OASIS system and report on our experience testing elements of this system with the MER rovers and with Mars rover prototypes at JPL. In particular we discuss how such capabilities can be enabled during ground operations planning and how this increased autonomy will affect downlinked data. We

¹ Senior Technical Staff, Artificial Intelligence Group, M/S 301-260, AIAA Member, Tara.Estlin@jpl.nasa.gov

² Assistant Section Manager, Instrument Software Section, M/S 168-527, Rebecca.Castano@jpl.nasa.gov.

³ Senior Technical Staff, Artificial Intelligence Group, M/S 301-260, Daniel.Gaines@jpl.nasa.gov.

⁴ Senior Technical Staff, Machine Learning Group, M/S 306-463, Benjamin.Bornstein@jpl.nasa.gov.

⁵ Research Manager, Science Division, M/S 183-335, Michele.Judd@jpl.nasa.gov.

⁶ Senior Technical Staff, Planetary Geosciences Group, M/S T1722, Robert.C.Anderson@jpl.nasa.gov.

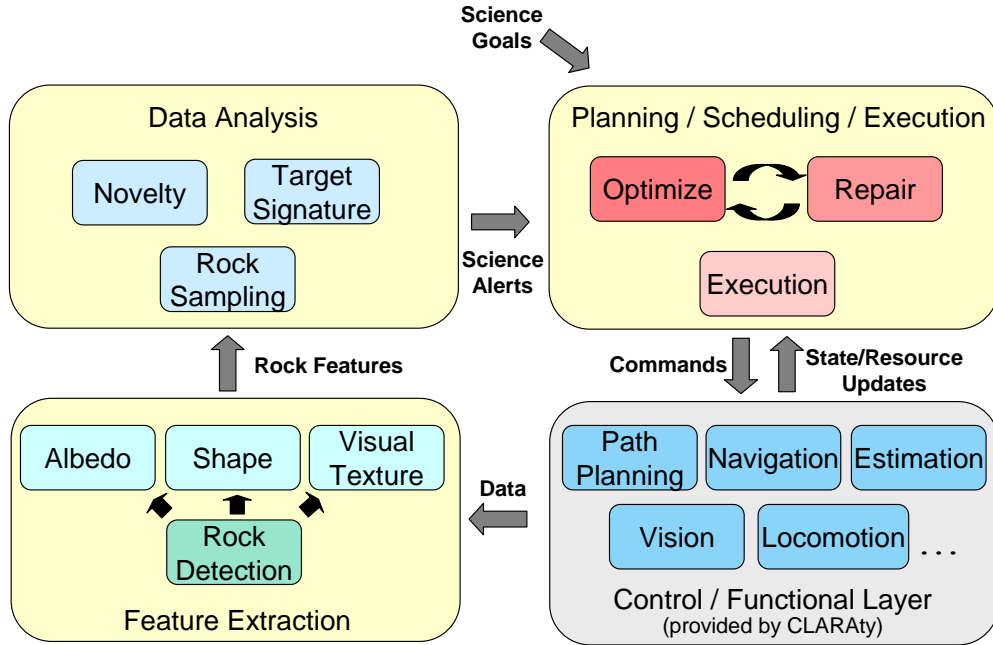


Figure 1. OASIS (Onboard Autonomous Science Investigation System) framework. This diagram shows how different decision-making capabilities interact within OASIS (shown in the yellow boxes) and how OASIS interacts with low-level robotic control software (shown in the gray box).

also highlight how OASIS can be used with the Maestro planning and data visualization tool², which enables user to command autonomous science operations using a well-known ground operations tool. Finally, we introduce a new area of OASIS work, which is to provide autonomous targeting capabilities for the MER rovers.

II. Overview of the OASIS System

OASIS is designed to operate onboard a rover identifying and reacting to serendipitous science opportunities. OASIS analyzes data the rover gathers, and then using machine learning techniques, prioritizes the data based on criteria set by the science team. This prioritization can be used to organize data for transmission back to Earth and it can be used to search for specific targets it has been told to find by the science team. If one of these targets is found, it is identified as a new science opportunity and a “science alert” is sent to a planning and scheduling system. After reviewing the rover’s current operational status to ensure that it has enough resources to complete its traverse and act on the new science opportunity, OASIS can change the command sequence of the rover in order to obtain additional science measurements. A large amount of work in OASIS has been done to identify and characterize rocks within images, which will be used as the primary science focus for the rest of this paper. However, OASIS has been applied to analyze data from other instruments (e.g., spectrometers) as well as to identify other terrain or atmospheric features. For example, OASIS can currently identify dust-devils and clouds in images as part of the MER rover onboard software.

A breakdown of OASIS capabilities is shown in Fig. 1. There are three major components that comprise OASIS:

- **Extract Features from Images:** Enables extraction of features of interest from collected images of the surrounding terrain. This component both locates rocks in the images and extracts rock properties (or features) including shape, texture, size, and albedo.
- **Analyze and Prioritize Data:** Uses the extracted features to assess the scientific value of the planetary scene and to generate new science objectives that will further contribute to this assessment. This component consists of three separate prioritization algorithms that analyze the collected data and prioritize the rocks. A new set of observation goals is generated to gather further data on rocks that either conform to the pre-set specifications by the science team, or are so novel in comparison to the other rocks, that another data measurement may be required.

- **Plan and Schedule New Command Sequence:** Enables dynamic modification of the current rover command sequence (or plan) to accommodate new science requests from the data analysis and prioritization module. A continuous planning approach is used to iteratively adjust the plan as new goals occur, while ensuring that resource and other operation constraints are always met.

III. Feature and Event Detection

OASIS uses a number of different methods to analyze data gathered by a rover. This data can include both terrain and atmospheric information. Though our techniques are applicable to a wide range of data modalities, our initial focus is on image analysis as images are commonly taken during rover activities and provide significant information about a scene.

A. Rock Detection

One important focus of OASIS over the last several years has been detecting interesting rocks in grayscale images. Several methods for identifying rocks have been developed. In one method, called Rockfinder, the detection of rocks is carried out by finding closed shapes in the image¹. The image is initially normalized, filtered with an edge preserving smoother and its edges are enhanced using unmask sharpening. The edges of the resulting image are detected using both a Sobel and a Canny edge detectors. For each result, the algorithm searches for closed shapes (which presumably correspond to relatively small homogeneous regions) using an edgewalker. The results from both detectors are combined and output as a list of contours of the identified shapes.

A second rock detection method, called Rockster, also focuses on intensity edges in grayscale imagery³. Rockster initially locates partial boundary contours of rocks using a procedure similar to the Canny edge detector. In particular, an intensity gradient is calculated over the image; ridges in the intensity gradient are linked together using non-maximum suppression, hysteresis thresholding and edge-following yielding a set of raw contours.

This initial set of contours does not directly provide a usable segmentation of the rocks from the background due to various problems, including: spurious contours from the sky-ground boundary (horizon line) and texture within individual rocks and the background. Rockster attempts to resolve these problems by splitting the initial contours into low-curvature fragments. Potential T-junctions that were missed by the edge detector are identified and used to further split fragments into even smaller pieces. A gap-filling mechanism is then applied to add new contour fragments between existing fragment endpoints. The final step is to regroup the edge fragments into coherent contours, which is accomplished through background flooding.

B. Rock Feature Extraction

Once rocks are identified in an image, the OASIS system classifies their physical properties. The properties that OASIS currently estimates are albedo, texture, size, and shape. The albedo of a rock is an indicator of the reflectance properties of a surface. OASIS measures albedo by computing the average gray-scale value of the pixels that comprise the image of the rock. The reflectance properties of a rock provide information about its mineralogical composition. Shadows and sun angle can both affect the gray-scale value of a pixel. Although this can be corrected by using the range data along with knowledge of both the sun angle and the camera orientation, the current system does not address these specific issues.

OASIS uses Gabor filters to estimate the visual texture of observed rocks⁴. Visual texture can provide valuable clues to both the mineral composition and geological history of a rock.

One of the most important properties of rocks on the surface is their size. Size can be used to identify sorting and geologic contacts. We model rocks as ellipses (if no range data is provided) or ellipsoids (if range data is available).

Although the shape of a rock is complex and often difficult to describe, significant geologic information can be extracted to better understand provenance (source of material) and environmental conditions. Various shape parameters are used to classify rocks in terrestrial studies, including elongation (or aspect ratio), ruggedness (or angularity), and surface area. In OASIS, an ellipse is fit to the outline of the rock. The eccentricity of the fit ellipse as well as the error is computed. The angularity of each rock is assessed using a measure of ruggedness. Further, another shape parameter identified is the orientation of the rock with respect to the ground.

C. Target Signature Prioritization for Rocks

Once rocks and their properties are identified in images, OASIS uses several methods to prioritize these identified features. One method for prioritization is called *target signature*. This algorithm enables scientists to efficiently and easily stipulate the value and importance to give each particular feature. Rocks are then prioritized as a function of the distance of their extracted feature vector from the specified weighted feature vector. Scientists can either manually specify a feature vector, or they may select a rock from among the set of rocks already identified and rank the rocks as a function of the distance of their feature vectors from the feature vector of the selected rock. Fig. 2 shows an example of using target signature to look for rocks of high albedo.



Figure 2. Image taken in response to a science alert on the JPL FIDO rover. In this test, science alerts were generated for “high albedo” (or light-colored) rocks.

D. Novelty Prioritization for Rocks

Another prioritization method is to highly rank identified rocks that are dissimilar from any rocks seen previously. The OASIS novelty detector utilizes the same historical traverse database as our target signature capability. However, instead of searching for specific rocks, as in the target signature case, the novelty detector i) groups rocks based on feature similarity, and then ii) searches for statistical outliers from those groups. Intuitively, this has the effect of identifying and prioritizing for downlink the most novel rocks. For example, a novelty detector that utilized albedo or VIS/NIR spectral signature would identify a calcite as novel in a sea of martian basalts. It’s important to note that our novelty detector operates without any prior knowledge from scientists or operators by design.

The novelty detector groups rocks based on feature similarity. To perform this grouping, rock features are treated as points in N-dimensional space and then clustered using the unsupervised *k*-means clustering algorithm. The number of clusters, *k*, signifies the number of geologic rock groups a geologist might expect to encounter on a traverse. For simple novelty detection, when operating on only a few rock features and/or on a relatively short traverse, it is sufficient to set *k* to 1. Once the N-dimensional points (representing rock features) have been clustered into one or more groups, statistical outliers are selected by choosing points (rocks) at least three standard deviations away from the cluster center(s). Rocks that pass this threshold test are considered novel and ranked by their distance from cluster centers (i.e. the farther away from a cluster center, the greater the degree of novelty)

E. Dust-Devil and Cloud Detection

A more recent effort on the OASIS task has been to develop two algorithms, which are now in use onboard the MER rovers, to observe and opportunistically identify atmospheric phenomena⁵. The atmosphere of Mars is highly dynamic with phenomena including clouds and dust devils observed by the MER rovers. Previously, the MER mission monitored for these events by performing observation campaigns where sequences of images were acquired with the hope of imaging the event in one or more of the frames. There is no guarantee that the phenomena of interest will be captured. For example, only around 10-25% of these cloud campaign images collected have clouds in them. Downloading these images without the phenomena of interest represents an inefficient use of limited bandwidth.

To identify dust devils, we are using motion detection using a temporal sequence. An example of this technique is shown in Fig. 3. Dust devils are high dust opacity features on a dusty background and often have a faint signature in an image. The main challenge for robust automated detection occurs when the difference in the intensity of the two images, at the location of the change, is comparable in magnitude to the noise of the image. The detection of faint dust devils in the image takes into account the noise of the image and uses the fact that a dust devil is bounded within a portion of the image. To reduce the noise in one image I_i we look at the difference between the average change in a series of images containing I_i , and the average change in the same series excluding I_i . Assuming that the major component of the image noise is zero-mean Gaussian noise, then the areas with no change tend to zero



Figure 3. Result of motion detection in an image. Two of the dust devils are evident (3rd and 5th red box), while the other three require the sequence to play out to become apparent.

while the areas with change do not. Thus, although the intensity of the motion information has been damped, the motion can be detected because the areas with no change tend to zero faster than those with change.

A second type of atmospheric phenomena of interest on Mars is clouds. In detecting clouds, it is assumed that large variations in the intensity of the sky in the image correspond to clouds. Our approach to automating the detection of clouds is to first locate the sky (equivalently, the horizon) in an image and then determine if there are high variance regions in the sky. In contrast to the dust devil detection, this algorithm operates on individual images. The time frame over which the clouds may change significantly is too long to require the rover to remain motionless on a regular basis, as would be necessary for effective application of image differencing.

Sequencing these detectors on MER has been integrated into the standard command uplink process. Algorithm execution on a rover is triggered by sequencing the WATCH command. Parameters for this command set various bounds for resource usage including total runtime, maximum number of images to save and maximum amount of megabits to save. If any of these resource limits are exceeded during execution, the WATCH command automatically terminates its run. Supporting commands set memory usage, various algorithm parameters (including whether to search for clouds or dust devils), and whether or not to enable image masking to allow for better bandwidth. For current nominal operations, there are two WATCH sequences, one used to search for clouds and another used to search for dust devils.

For dust devil detection, WATCH has two different image acquisition strategies. In the first, a single image buffer is used to acquire Navcam images. Since dust devil detection relies on image differencing to detect motion, at least two images must be acquired before image analysis can start. When only a single image buffer is used, the time between successive frames can be up to two minutes, as it takes one minute to acquire the image and one more to perform frame-to-frame analysis and search for differences between images. The second image acquisition strategy makes use of additional MER image buffers (up to 10) to more rapidly (every minute) acquire a series of images and then analyze the series in batch. While both acquisition strategies have been tested on the surface, the single image buffer mode is currently more common.

IV. Data Prioritization for Downlink

The results of image data analysis can then be used to prioritize image data for downlink, ensuring that the data with the highest science value is downlinked first. This strategy can be used with any of the previously mentioned prioritization algorithms, where high priority rocks or atmospheric events would be used to guide the prioritization assignment. Different strategies for prioritizing data for downlink based on autonomous science analysis can be used. For instance, all data of a certain category could be selected and/or prioritized for downlink based on analysis results or a certain percentage of downlink could be allocated for data selected by analysis results. For instance, standard downlink procedure could be used for 80% of downlink but the remaining 20% could be data marked as high-priority by autonomous science.

Currently, the primary output for the MER dust-devil detectors and cloud detectors is to downlink all images that have been identified as containing the target atmospheric phenomena and deleting any campaign images that do not. Rock prioritization has also been shown to be useful for prioritizing images where high priority images are those that contain the best match to a specified target signature or those that contain the most novel rock detected. Beyond image prioritization, identified terrain targets can also be used to direct additional science activities, such as taking additional observations of high priority targets.

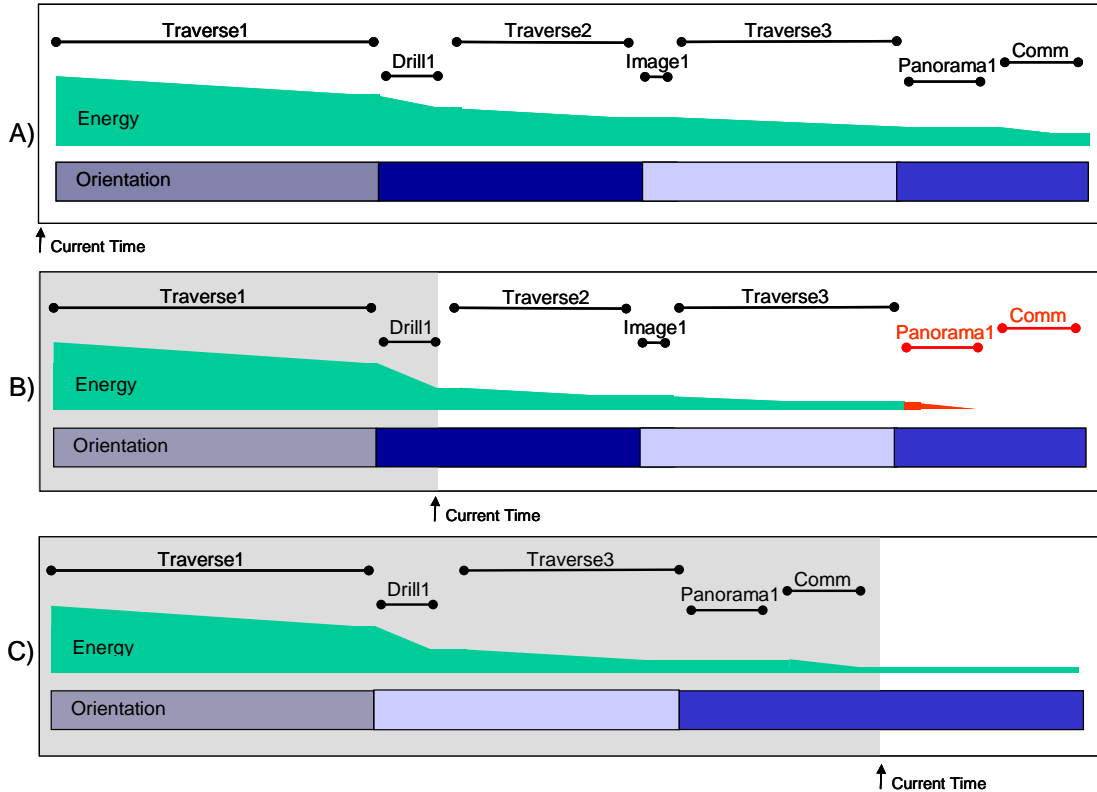


Figure 4. Example re-planning scenario. A) The initial plan generated by CASPER which drives the rover to three locations, performs a science operation at each location, and finally performs a critical end-of-day communications link with Earth. B) The plan after partial execution. The first science activity has taken more energy than expected, which has caused several resource conflicts with later activities. C) The plan after re-planning has occurred to resolve the conflict. The second science operation and drive were dynamically removed so the last (higher priority) science operation and communication activity can still be achieved.

V. Scheduling Opportunistic Science

When the OASIS data analysis software identifies science targets of interest (e.g., a novel rock), a science alert is generated. This results in a new science goal being passed to the planning and scheduling module which determines if the new measurement request can be accommodated. If it can be, the current rover command sequence is modified to collect new science data.

The OASIS planning and execution module⁶ is intended to run with little communication with ground. It accepts new science goals and then modifies the current rover command sequence (or plan) to try and achieve as many of the goals as possible while still respecting relevant state and resource constraints. This module also executes the current rover plan by dispatching commands to the rover's low-level control software and monitoring relevant state and resource information to identify potential problems or opportunities. If problems or new opportunities are detected, the system is designed to handle such situations by using re-planning techniques to add, move, or delete plan activities.

Planning capabilities for OASIS are provided by the CASPER continuous planning system⁷. Based on an input set of science goals and the rover's current state, CASPER generates a sequence of activities that satisfies the goals while obeying relevant resource, state and temporal constraints, as well as operation (or flight) rules. Plans are produced using an iterative repair algorithm that classifies plan conflicts and resolves them individually by performing one or more plan modifications. CASPER also monitors current rover state and the execution status of plan activities. As this information is acquired, CASPER updates future-plan projections. Based on this new information, new conflicts and/or opportunities may arise, requiring the planner to re-plan in order to accommodate the unexpected events. A simple re-planning situation is shown in Fig. 4. Here, a rover science activity (of drilling for sample collection) took more power than originally estimated, which causes a resource conflict. CASPER fixes

the conflict by deleting a low-priority science activity to ensure that enough time remains to execute a later, high-priority science activity and a required communications link.

To reason about science goal priorities and other plan quality measures we use the CASPER optimization framework to search for a higher quality plan. User-defined preferences are used to compute plan quality based on how well the plan satisfies these preferences. An overall plan score is computed based on the preference specification. Plan optimization works in an iterative fashion (similar to plan repair) and searches for plan modifications that could potentially improve the overall plan score.

To handle opportunistic science, we enabled the OASIS planning and execution module to recognize and respond to science alerts, which are new science opportunities detected by onboard data-analysis software. For example, if a rock is detected in navigation imagery that has a previously unseen shape or texture, a science alert may be generated to take additional measurements of that rock.

Science alerts can have different levels of reaction from the planning and execution system. 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 for further analysis at the next communication opportunity. The next level of reaction is to collect additional data at the current site before transmitting back to Earth. A further step is to have the rover alter its path to get closer to objects of interest before taking additional measurements. These operations provide new data that could not be obtained through analysis of the original image.

To handle a science alert that requests additional measurements, the planner must generate a plan that achieves the new goals without deleting existing activities or causing conflicts that cannot be resolved (e.g., scheduling more activities than can be executed over a certain time window). Depending on rover resources, not all new opportunities may be realized. Science alerts can also contain priorities, which are representative of their scientific value.

VI. User Interface for Data Visualization and Target Selection

The Maestro Operations Tool² provides user interface and data visualization capabilities for science operations and is used to support science on a number of Mars surface missions (e.g., MER, Phoenix, MSL). In particular, Maestro is used for the selection and viewing of science targets during a mission operations day. Maestro can display imagery taken by a mission spacecraft or rover, such as images taken with the MER panoramic or navigation cameras, and mission operators can define targets in those images by selecting an image pixel and assigning it a target name. This name is then translated into XYZ coordinates based and the target name and location is stored into a database that all mission operators and scientists can share. Maestro not only provides a flexible user interface, but enables mission personnel from different physical locations to easily view and share target data.

The Maestro Operations tool has been integrated with OASIS to provide a similar user interface. It can be used with to select science targets, emulate ground commanding of science goals, select parameters for using OASIS algorithms, and to view a number of collected data products, such as single pair of stereo images or panoramas of multiple images. This integration enables Maestro to be used as a user front-end for testing and demonstration of OASIS with rover hardware. For instance, it can be used to select and send ground science targets to an automated planner, to view opportunistic science images that were taken based on data analysis results, and to view any other collected image data products that were gathered during autonomous science tests. Fig 5 shows an example of target selection in Maestro during a demonstration with the FIDO rover. Further, Maestro has also been extended to support enabling autonomous science capabilities and selecting relevant parameters³. When planning a set of rover activities with Maestro, users can decide when to activate OASIS's autonomous science capabilities and specify parameters for the OASIS underlying algorithms. When a traverse is planned, for example, users specify whether or not to perform autonomous science during the traverse. If they select to activate autonomous science, they can also design an imaging sequence (e.g., that points the mast and takes images mast cameras) that can be run at a particular frequency to gather images of the passing terrain. Users can also select target features that can be used to prioritize images. For instance, users can select combinations of rock features (e.g., rock albedo levels, shape characteristics, size characteristics) that represent high priority rock targets.

The usage of Maestro to setup and run autonomous science activities and to view gathered data products illustrates how the OASIS system could be integrated into the missions operations process for future rover missions.

VII. CLARAty Robotic Architecture

The systems described in this paper are integrated with the Coupled Layered Architecture for Robotic Autonomy (CLARAty)⁸. CLARAty was developed to provide reusable robotic software and was designed to simplify the

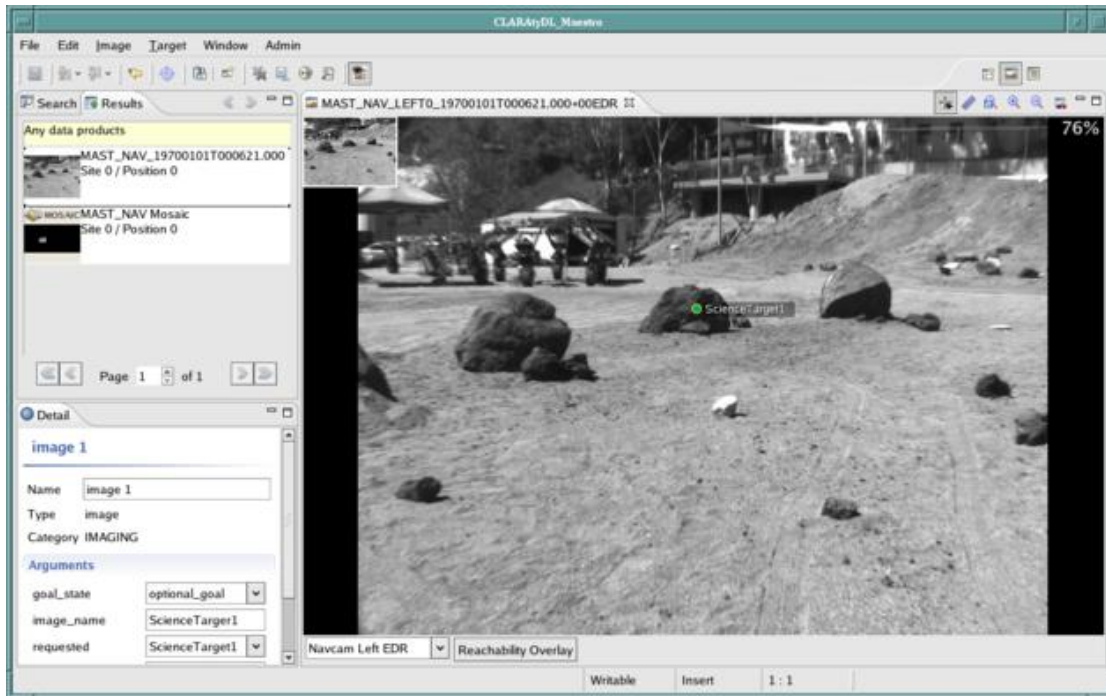


Figure 5: Example of target selection using the Maestro tool. Here, Maestro is used to display an image taken with the FIDO rover navigation cameras. The user has selected a science target, named *ScienceTarget1*, for the rover to drive towards and examine.

integration and testing of new robotic technologies on different robotic platforms. Currently, CLARAty provides a large range of robotic functionality from more basic capabilities, such as motor and camera control, to high-level autonomy capabilities, such as automated planning and scheduling and model-based fault diagnosis. CLARAty is divided into two primary layers. The lower tier is the Functional Layer (FL), which provides low- to mid-level robotic control functionality such as locomotion, vision, navigation, etc. The FL is organized primarily in a hierarchical manner, making it straightforward for new technology to plug in and interface needed capabilities. The upper tier is the Decision Layer, which provides higher-level autonomy systems that typically reason about system resources, high-level science goals, mission constraints, activity temporal and state requirements, etc. The DL is organized to enable plug-and-play of different autonomy systems and enable them to perform global reasoning over a large span of lower-level functionality and data. A primary distinction between technology in the DL and in the FL is that DL modules are characterized by their need for a supervisory interface to robotic functionality while FL modules tend to require tighter control and feedback. For example, an automated planner needs the ability to command the rover to navigate to a location, but only needs occasional updates (> 1 Hz) on the rover's progress. In contrast a motor PID controller, which would reside in the FL, needs a high-frequency control-feedback loop to correctly operate.

The overall OASIS system is primarily contained within the Decision Layer. Through CLARAty, the OASIS planning and execution system has been tested with a number of JPL rover platforms and a high-fidelity rover simulation tool. Many Functional Layer capabilities were also used during tests with rover hardware. This software includes the Morphin navigation system⁹, which enables a rover to avoid obstacles and navigate to specified waypoints, and a position estimation algorithm, which integrates IMU (Inertial Measuring Unit) measurements with wheel odometry to estimate rover position and attitude (roll, pitch and heading). Recent tests have involved the use of the MER Visual Targeting Tracking (VTT) algorithm¹⁰, which is also provided through CLARAty. VTT enables a rover to track a selected target as the rover drives and to closely approach a target location. Other algorithms used in most tests include rover locomotion, mobility and stereo processing as well as control functions for mast pan/tilt and camera operation.




Figure 6. Example of an opportunistic image taken of a rock identified as novel (based on the orientation property). After the initial target identification, the rover was driven to close range of the rock using the MER Visual Target Tracking algorithm. The two boxes show the particular rock feature that was tracked.

and consistently selected new science targets and/or new science target combinations that had not been previously tested.

A primary scenario element was dynamically identifying and handling opportunistic science events. For many of these tests, we used onboard data analysis software to generate science alerts based on a target rock signature. Various types of signatures were used, but they typically corresponded to a combination of target albedo (brightness) levels, shape characteristics, and size estimations. Other tests used a different analysis algorithm that searched for novel rocks (or outliers). If rocks were identified in rover camera imagery that had a high score for these features, then a science alert was created and sent to the planner. Science alerts often happened during rover traverses to new locations, but they were also used for testing autonomous targeting at the end of a traverse. If a science alert was detected, the planner attempted to modify the plan so an additional image of the rock of interest would be acquired. A sample image that was taken in response to a science alert (for a rock of high albedo) was shown in Fig. 2. Science alerts also had varying rover position constraints. Sometimes the rover was requested to take an image from its current position using the rover mast cameras. Other times the rover was commanded to turn towards the rock of interest and take an image with the front hazard cameras. A more complex request was to use Visual Target Tracking software to drive within 1-2 meters of the opportunistic science target and take a close-up



Figure 7. JPL FIDO rover

VIII. System Testing with Robotic Hardware

To evaluate our system we have performed a large series of tests using rover hardware. These tests covered a wide range of scenarios that included the handling of multiple, prioritized science targets, limited time and resources, opportunistic science events, resource usage uncertainty causing under or over-subscriptions of power and memory, large variations in traverse time, and unexpected obstacles blocking the rover's path^{1,3,6}.

A. Testing Setup

For many tests, our testing scenarios consisted of a random number of science targets specified at certain locations. A map was used that would represent a sample mission-site location where data would be gathered using multiple instruments at a number of locations. Targets were typically prioritized and constraints on time, power or memory would limit the number of science targets that could be handled. A large focus of these tests was to improve system robustness and flexibility in a realistic environment.

Towards that goal we used a variety of target locations

image. Fig. 6 shows a rock that was identified as having a novel orientation and then was tracked by VTT to bring the rover in for a close observation.

B. Rover Hardware

A large number of tests have been run in the JPL Mars Yard using different rover hardware platforms. For the past few years, the FIDO rover (shown in Fig. 7) was used for the majority of tests. FIDO¹¹ (the Field, Integration and Design Operations) is an advanced technology prototype rover similar to the MER rovers. FIDO's mobility sub-system consists of a six-wheel rocker-bogie suspension capable of traversing over obstacles up to 30 cm in height. FIDO currently contains a fixed mast that contains both panoramic and navigation camera pairs. FIDO also has front and rear hazard cameras that can be

used during driving to detect obstacles.

All OASIS software has been designed to run onboard the rover, however during testing, only functional-level CLARAty modules, such as navigation, vision and tracking, and the OASIS rockfinding software were run onboard FIDO. Other modules, including the planning and execution module and the analysis module, were run on offboard workstations that communicated with the rovers using Wireless Ethernet, since a port of these components to the onboard operating system (VxWorks) was not complete.

Tests in the Mars Yard typically consisted of 20-50 meter runs over a 100 square meter area with a range of obstacles that caused deviations in the rover's path. Science measurements using rover hardware were images from one of three sets of cameras on the rover (hazard cameras, navigation cameras and panoramic cameras). Other instruments, such as spectrometers, were not readily available and thus not directly incorporated into hardware tests. However different types of measurements were included when testing in simulation.

C. Demonstration Summary

A number of live demonstrations of our system have been performed, including a several hour long demonstration which showed the system successfully handling a random combination of science targets and science alerts (that had not been used in previous testing) and resulted in over 40 meters of autonomous driving. This demonstration consisted of several runs that showed scenario elements such as handling new science alerts, dynamically adding new ground-specified science when time became available, and deleting low priority science targets in a later run where more power was used than originally estimated. The OASIS software operated correctly in all tested cases.

Another live demonstration showed a combination of traverse science and autonomous targeting. Autonomous targeting was supported by adding a targeting subplan at the end of a drive that executes a FIDO navigation camera panorama that is analyzed online for new targets. Autonomous, new measurements were then scheduled and taken by the high-resolution FIDO panoramic cameras, which were being used as an example limited FOV instrument. More recent demonstrations highlighted the OASIS novelty detector and showed how the rover can be autonomously driven towards a novel terrain feature in order to collect close-range observations. To support these last tests and demonstrations, the OASIS planning components was required to execute and reason about drives that used both navigation capabilities and visual target tracking functionality to drive the rover towards the object of interest.

IX. Automated Targeting for the MER Rovers

One new area of application for OASIS is to provide automated targeting capabilities for the MER rovers as part of a new MER mission technology experiment. Specifically, OASIS will provide technology for *targeted* remote sensing science in an automated fashion during or after rover traverses. A number of remote sensing instruments for rovers have a very narrow field-of-view (FOV) and thus require very specific targets for correct sampling. Examples of such instruments include the MER mission Mini-Thermal Emission Spectrometer (or mini-TES) and the ChemCam spectrometer, which performs Laser-Induced Breakdown Spectroscopy and is planned to be flown as part of the Mars Science Laboratory (MSL) 2009 rover mission. Targeting of these instruments by mission personnel requires a lengthy planning process. The typical scenario for selecting targets is to manually identify the targets using data that has already downloaded on a previous sol. Thus after reaching an end-of-day location, the rover must sit and wait until images can be analyzed and new measurement commands can be uplinked, which at best will happen on the next sol. As a result, only untargeted remote sensing is possible immediately after or during a traverse, where typically wider field of view instruments are used such as the MER navigation cameras.

OASIS will enable the rover flight software to analyze imagery onboard in order to autonomously select and sequence targeted remote-sensing observations in an opportunistic fashion. Fig. 8 shows an example of selecting five rocks that could be sampled by a limited FOV remote-sensing instrument. This capability is especially useful for multi-sol plans where a drive is performed on the first sol. Currently only untargeted remote sensing can be done on the second and third sols since ground does not have imagery of the rover's new location. Using OASIS, targeted remote sensing can be performed on these sols, significantly increasing science gain for multi-sol plans. Further, targeted measurements could also be made during a traverse.

Onboard MER trials will be performed to demonstrate and validate the OASIS capability. OASIS will perform two main tasks: 1) identify science targets in imagery from the MER navigation cameras and 2) schedule an opportunistic response where a high-resolution, subframed panoramic camera image is taken of the selected target. Science targets will include rocks and other terrain features that match pre-set scientist criteria that will be determined during sequencing. Fig. 8 illustrates the additional targeted measurements that could be acquired during




Figure 8. OASIS selects five potential targets for a limited FOV instrument, such as the 2009 MSL ChemCam, to sample. This image was taken on the MER mission. Autonomously selecting targets vs. blind sampling greatly increase the chances of accurately targeting a rock.

selection, and activity planning and scheduling. Current approaches require human analysis to determine goals and to manually convert the set of high-level science goals into low-level rover command sequences. By integrating these components onboard, we enable a rover to function more autonomously in collecting valuable science data. This type of capability should dramatically increase the science return of future rover missions.

Acknowledgments

This work was performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

- ¹R. Castano, T. Estlin, R. C. Anderson, D. Gaines, A. Castano, B. Bornstein, C. Chouinard, M. Judd, "OASIS: Onboard Autonomous Science Investigation System for Opportunistic Rover Science," *Journal of Field Robotics*, Vol 24, No. 5, May 2007.
- ²M. W. Powell, T. Crockett, J. M. Fox, J. C. Joswig, J. S. Norris, K. J. Rabe, M. McCurdy, G. Pyrzak, "Targeting and Localization for Mars Rover Operations," *Proceedings of the 2006 IEEE Conference on Information Reuse and Integration*, September 2006.
- ³R. Castano, T. Estlin, D. Gaines, C. Chouinard, B. Bornstein, R. C. Anderson, M. Burl, D. Thomspson, A. Castano, and M. Judd, "Onboard Autonomous Rover Science," *Proceedings of the 2007 IEEE Aerospace Conference*, Big Sky, Montana, March 2007.
- ⁴R. Castano, T. Mann, and E. Mjolsness. (1999) "Texture Analysis for Mars Rover Images," in *Applications of Digital Image Processing XXII, Proc. of SPIE* Vol. 3808, Denver, CO, July 1999, pp. 162-173.
- ⁵A. Castano, A. Fukunaga, J. Biesiadecki, L. Neakrase, P. Whelley, R. Greeley, M. Lemmon, R. Castano, S. Chien, "Automatic detection of dust devils and clouds at Mars," *Machine Vision and Applications*, 2007.
- ⁶T. Estlin, D. Gaines, C. Chouinard, R. Castano, B. Bornstein, M. Judd, and R.C. Anderson, "Increased Mars Rover Autonomy using AI Planning, Scheduling and Execution," *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2007)*, Rome Italy, April 2007.
- ⁷S. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau "Using Iterative Repair to Improve Responsiveness of Planning and Scheduling," *Proceedings of the 5th International Conference on Artificial Intelligence Planning and Scheduling*, Breckenridge, CO, 2000.

or after rover drive activities. This automated targeting capability will be scheduled and run using a process similar to the one in place for the onboard dust-devil and cloud detectors. In this process, the ground team will schedule automated targeting blocks where time and resources are reserved to safely acquire new image data. Parameters can also be set during sequencing to specify science-target features, such as rock albedo or size.

The resulting Panoramic camera images will be downlinked with other standard MER data products. New images will also be relatively inexpensive to downlink since only subframed images will be taken. This result will show how a mission can receive high-quality opportunistic data without requiring a large amount of resources, such as downlink bandwidth, or onboard storage.

X. Conclusion

We have demonstrated an autonomous science system in the field conducting opportunistic science. By integrating data analysis and planning capabilities, the resulting system can operate in a closed-loop fashion. This framework enables new science targets to be addressed onboard with little or no communication with Earth. An important contribution of this work is closing the loop between the sensor data collection, science goal

⁸I.A. Nenas, R. Simmons, D. Gaines, C. Kunz, A. Diaz-Calderon, T. Estlin, R. Madison, J. Guineau, M. McHenry, I. Shu, and D. Apfelbaum, "CLARAty: Challenges and Steps Toward Reusable Robotic Software," *International Journal of Advanced Robotic Systems*, Vol. 3, No. 1, 2006.

⁹C. Urmson, R. Simmons, and I. Nenas, (2003) "A Generic Framework for Robotic Navigation," *Proceedings of the IEEE Aerospace Conference*, Montana, March 2003.

¹⁰W. Kim, R. Steele, A. Ansar, K. Ali, and I Nenas, Rover-based Visual Target Tracking Validation and Mission Infusion, *Proceedings of the American Institute of Aeronautics and Astronautics (AIAA) Space 2005 Conference*, Long Beach, CA, August, 2005.

¹¹P. Schenker, E. Baumgartner, P. Backes, H. Aghazarian, L. Dorsky, J. Norris, T. Huntsberger, Y. Cheng, A. Trebi-Ollennu, M. Garrett, B. Kennedy, A. Ganino, R. Arvidson, S. Squyres, "FIDO: a Field Integrated Design & Operations rover for Mars surface exploration," *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Autonomous for Space*, Montreal Canada, June, 2001.