# Supporting Increased Autonomy for a Mars Rover

Tara Estlin, Rebecca Castano, Dan Gaines, Ben Bornstein,
Michele Judd, Robert C. Anderson and Issa Nesnas
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-393-5375
Tara.Estlin@jpl.nasa.gov

*Abstract*— **This paper presents an architecture and a set of technology for performing autonomous science and commanding for a planetary rover. The MER rovers have outperformed all expectations by lasting over 1100 sols (or Martian days), which is an order of magnitude longer than their original mission goal. The longevity of these vehicles will have significant effects on future mission goals, such as objectives for the Mars Science Laboratory rover mission (scheduled to fly in 2009) and the Astrobiology Field Lab rover mission (scheduled to potentially fly in 2016). Common objectives for future rover missions to Mars include the handling of opportunistic science, long-range or multi-sol driving, and onboard fault diagnosis and recovery. To handle these goals, a number of new technologies have been developed and integrated as part of the CLARAty architecture. CLARAty is a unified and reusable robotic architecture that was designed to simplify the integration, testing and maturation of robotic technologies for future missions. This paper focuses on technology comprising the CLARAty Decision Layer, which was designed to support and validate high-level autonomy technologies, such as automated planning and scheduling and onboard data analysis.[*]**

## I. INTRODUCTION

The Mars Exploration Rovers (MER) continue to make history as they drive over new Martian terrain and send back valuable scientific data. Both rovers have now survived over 1100 Martian days (sols) and combined have traveled over 17 kilometers. The longevity of these vehicles and the plans for several new rover missions to Mars has sparked significant interest in developing and testing new rover autonomy technology. New autonomous capabilities have already been uploaded to the MER rovers, including technology for automatically recognizing atmospheric events, such as dust-devils, and technology for visual tracking an object as a rover drives. These capabilities are also scheduled to fly on the 2009 Mars Science Laboratory rover mission and further autonomy will likely be desired for the MSL extended mission phase. Requirements for future rover missions, such as the 2016 Astrobiology Field Lab Mission, include a need for further autonomous capabilities, including increased onboard science data analysis, automated planning for sequence modification,

and techniques to support long-range traverse or "over the horizon" driving.

The CLARAty (Coupled Layer Architecture for Robotic Autonomy) architecture was developed to provide reusable robotic software. It was designed to simplify the 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 or model-based fault diagnosis. CLARAty is divided into two primary layers. The lower tier is the Functional Layer, which provides low- to mid-level robotic control functionality organized in a hierachical manner, making it straightforward for new technology to plug in and interface needed capabilities such as locomotion, vision, navigation, etc. 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.

One important goal of Decision Layer technologies is to support autonomous science capabilities. Systems for onboard data analysis and automated planning can enable this goal by analyzing data gathered onboard the rover (such as images gathered for navigation or science pursposes), prioritizing data for downlink, and providing opportunistic data collection when new science opportunities are identified. This paper will in particular discuss how the CLARAty Decision Layer supports such capabilities and how they have been tested and demonstrated on various JPL research rovers.

## II. OVERVIEW OF CLARATY DECISION LAYER

The CLARAty architecture [1] provides a framework for generic and reusable robotic components that can be adapted to a number of heterogeneous platforms. It also simplifies the integration, testing and maturation of new robotic autonomy technologies. As mentioned in the previous section, CLARAty is a two layer architecture with a Decision Layer (DL) and a Functional Layer (FL). This paper focuses on the Decision Layer and in particular, DL technologies that provide autonomous science capabilities.

The Decision Layer provides a number of capabilities for reasoning about higher-level goals and robotics resources. Autonomy technologies that are currently integrated within the DL include automated planning/scheduling, executive control, model-based fault diagnosis, and onboard science data analysis. These capabilities can be used individually or in an integrated fashion to provide high-level decision making for a robotic application. Automated planning/scheduling and executive technology provides functionality for breaking down high-level goals into a plan of activities that successfully coordinate FL capabilities in achieving the goals. These technologies also provide execution monitoring, dynamic sequence modification, as well as resource and constraint management. Onboard science data analysis provides capabilities for analyzing gathered data (e.g., image data, spectral data), prioritizing and/or summarizing data for downlink, and identifying opportunistic science events. Model-based fault diagnosis provides functionality for diagnosing system faults by monitoring sensor data and command execution status.

Work is currently progressing for providing generic infrastructure for DL components. Most past work for the CLARAty architecture has concentrated on developing generic infrastructure for the CLARAty Functional Layer and supporting low- to mid-level robotics capabilities such as vision, locomotion, navigation, etc. In the past few years, some key infrastructure elements have been designed and developed for the DL; though much work remains to fully support all desired DL capabilities.

*A. Interface between Decision Layer and Functional Layer*

One infrastructure element that is currently in place is a flexible interface between the DL and the FL, which enables components in the DL and FL work closely together. Communication using this interface is typically in one of two categories:

- Commands are sent from DL modules to FL modules that cause the rover hardware to execute a certain activity (e.g., drive, image).

- Updates are sent from FL modules to DL modules that report the current state of the rover hardware, command execution state and various telemetry data that the DL needs to access.

The developed interface is intended to provide a generic mechanism for communication between the DL and FL and was designed to be robust to errors and easily extendible. The interface was designed and built using the CORBA TAO C++ orb, which provides a distributed object approach to communication between objects where communication can be across multiple machines.

A number of features have been added in this interface, including strong error checking for all communication mechanism, supporting a wide range of command parameters and update values, enabling the DL and FL to communicate when on the same machine and when on different machines, and enabling the DL to communicate with FL modules at different hierarchical levels. Further, a key aspect of this interface is a data update service that provides a flexible and straightforward approach for DL modules to subscribe and receive various updates and other relevant data from FL modules.

*B. Generic Framework for Data Analysis*

Another infrastructure element that has been added is a generic framework for easily integrating multiple data analysis technologies. This framework was designed to enable different types of data analysis software to be easily integrated with other relevant rover technologies and be exercised on rover hardware. In particular, this framework was designed to support the "plug-n-play" of various data analysis algorithms and components.

This framework uses a pipe and filter architecture where data can be piped through different data analysis algorithms (or filters) and algorithms can be organized in a flexible manner to support varying analysis needs. The framework enables easy and flexible access to relevant data at different stages of analysis. It also allows different types of detectors and analysis approaches to be easily configured. Often, varying data analysis detectors are available that can be used individually or linked together with other detectors. For instance, a rock detector or layer detector could be used on their own to detect areas of interest. However, other detectors could also be used in conjunction and in different organizations. For instance, an albedo (or brightness) detector could be used with a rock detector to detect light or dark rocks and could also be used with a layer detector to detect light or dark layers.

Currently, a number of analysis algorithms for image processing are integrated within the framework, such as components for rock detection, rock shape analysis, cloud detection and dust-devil detection. The algorithms can be used on various images that could be collected onboard a rover, such as images taken for navigation by hazard cameras or images taken for science purposes by higher-resolution panoramic cameras. Further the framework is designed to support a large spectrum of instrument data analysis and a plan and design for integrating several detectors that use spectral data is already in place.

Further plans for the CLARAty Decision Layer include the design and development of generic infrastructure to support automated planning and scheduling systems and the integration of the Maestro command sequencing and data visualization tool [2] to provide a user interface for selecting

plan goals (or targets) during testing and visualizing collected science data.

## III. CLARATY AUTONOMOUS SCIENCE CAPABILITIES

A large focus of the CLARAty Decision has been to integrate and demonstrate technologies for autonomous science, and in particular for supporting opportunistic science. Many of these capabilities are provided by the OASIS system. OASIS is designed to operate onboard a rover identifying and reacting to serendipitous science opportunities. Science opportunities can include detection of novel rocks that the rover has not seen before [3] and/or the detection of atmospheric events, such as dust devils and clouds [4]. OASIS analyzes data the rover gathers, and then 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 the planning and scheduling component of OASIS. 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 central 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 [3].

When a science alert is found, the rover typically is instructed to stop its current traverse, locate the rock that triggered the science alert, and take additional data (e.g. color image, closer grayscale image, spectrometer reading) on that rock. The system enables the rover to take that data through several methods, which include pointing the rover mast and using either navigation or panoramic cameras to image the target, turning and pointing the rover to take images with the hazcams and driving to the target so that closer measurements can be collected. Once the rover has completed this additional measurement, it reverts back to its original plan and continues on its traverse.

OASIS also provides an end-of-day capability to increase rover science measurements. After the rover has completed its required image captures of its final location at the end of the sol, OASIS can run these images through its onboard rock detection program(s) and identify targets for a remote sensing instrument (such as a close-up panoramic image or a Laser Induced Breakdown Spectrometer, or LIBS measurement). As resources permit, OASIS can instruct the rover to point the mast instrument to these targets, take data and return the data to Earth at its next communication cycle.

A breakdown of OASIS capabilities is shown in Figure 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 module both locates rocks in the images and extracts rock properties (features) including shape, texture 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 module 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 met.

The OASIS Feature Extraction and Data Analysis modules are responsible for analyzing onboard image data from rover cameras, identifying interesting or novel parts of that data, and generating new science goals to gather additional measurements when new science targets have been identified. Analyzed images could be taken by either rover science or engineering cameras and could be shot during a variety of rover activities including obstacle avoidance imaging, high-resolution science imaging, and imaging done for long-range route planning.

As shown in Figure 1, new science data is first processed by the Feature Extraction component. Images are broken down by first locating individual rocks, and second, by extracting a set of rock properties (or features) from each identified rock. Extracted rock properties (e.g., shape, albedo, visual texture, size) are then passed to the Data Analysis component of the system. This component consists of several prioritization algorithms, which analyze the data by searching for items such as rocks with features that match pre-known signatures of interest (as identified by scientists on Earth), or novel rocks (i.e., outliers) that have not been seen in past activities. If the analysis component detects new science opportunities of significant interest, it will generate a *science alert* that is sent
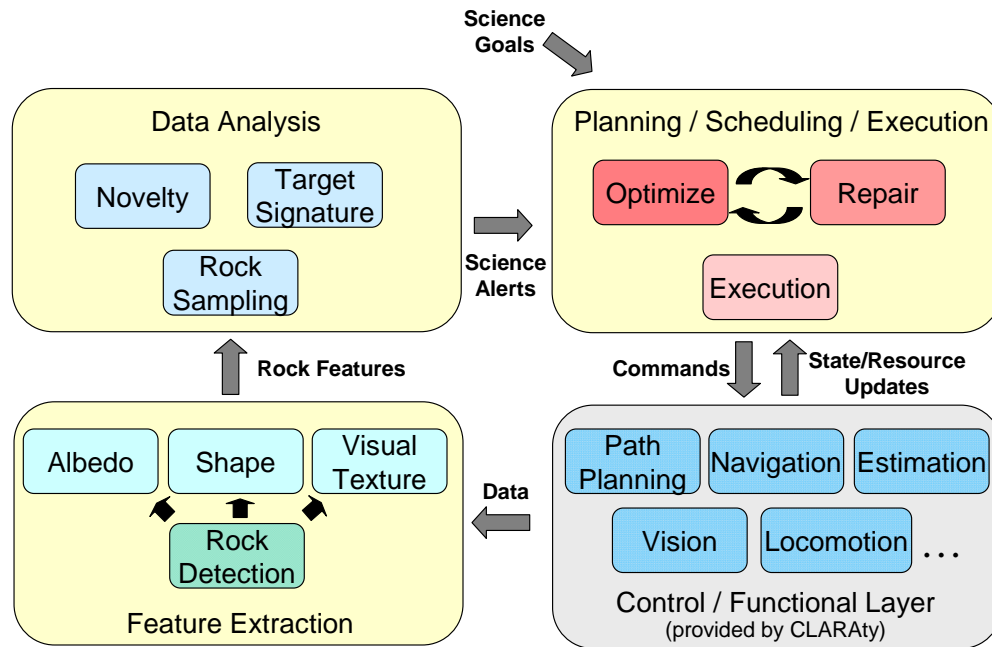
Fig 2. OASIS (Onboard Autonomous Science Investigation System) system and data flow. This diagram shows how different decision-making capabilities interact within OASIS (shown in the yellow boxes) and how OASIS interacts with CLARAty low- and mid-level control software (shown in the gray box).

to the planner, which signifies that new measurements should be taken of the identified target.

In OASIS, planning, scheduling and execution techniques are used to provide rover-plan generation, execution, monitoring, and re-planning in response to problems and new science opportunities. To provide these capabilities, the framework closely integrates the CASPER (Continuous Activity Scheduling, Planning, Execution and Re-planning) planning system [5] and the TDL (Task Description Language) executive system [6]. CASPER is given a model of rover operations and an input set of science and engineering goals (e.g., take a panoramic image at a certain azimuth and elevation). CASPER generates a sequence of activities that satisfies as many goals as possible while obeying relevant resource, state and temporal constraints. CASPER also monitors current rover state and the execution status of plan activities. As this information is acquired, CASPER updates future-plan projections. These updates may cause new conflicts and/or opportunities to arise, requiring CASPER to replan to accommodate the new data.

TDL was designed to handle intermediate control activities for a robotic system and to mediate between a planning system and low-level robot control software. It expands activities into lower-level commands, executes the commands, and monitors their execution. For this application, TDL handles expanding high-level activities into lower-level commands (based on current state information), executing commands by dispatching them to functional-level control software, relaying updates on activity status and state/resource data to CASPER, tracking plan execution progress, and relaying identified problems to CASPER.

## IV. ROBOTIC HARDWARE TESTING

Through CLARAty, the OASIS planning and execution system has been tested with several JPL rover platforms, including Rocky 7, Rocky 8 and the FISO rovers, which are shown in Figure 2. FIDO is a MER-class rover that supported early MER field trials and has been our primary testing platforms for the last several years.

To run realistic scenarios with rover hardware, a number of supporting pieces of CLARAty software were used, which could run on the relevant hardware platforms. These software components include the Morphin navigation system, which enables a rover to avoid obstacles and navigate to specified waypoints, a position estimation algorithm, which integrates IMU (Inertial Measuring Unit) measurements with wheel odometry to estimate rover position and attitude (roll, pitch and heading), stereo vision processing, basic mobility locomotion, as well as control functions for mast pan/tilt and camera operation.

A number of live demonstrations of OASIS autonomous science techniques with other CLARAty capabilities 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

**Figure 2: Rocky 8 rover (left), FIDO rover (middle), Rocky 7 rover (right)**

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 that starts with 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.

## V. CONCLUSION

The CLARAty Decision Layer was design to support high-level autonomy algorithms and to enable their easy integration and testing with low- to mid-level robotic capabilities and various robotic hardware platforms. Currently, a number of technologies have been integrated within the Decision Layer. This paper focused on capabilities for autonomous science, which have been tested on a number of scenarios with rover hardware and with relevant Functional Layer components, such as navigation, stereo vision, and mast manipulation. Future work is focused on developing new DL infrastructure for supporting automated planning systems and for providing a flexible user interface to support system testing on different mission scenarios.

## ACKNOWLEDGMENT

## REFERENCES

[1] I.A. Nesnas et al., "CLARAty: Challenges and Steps Toward Reusable Robotic Software," International Journal of Advanced Robotic Systems, Vol. 3, No. 1, pp. 023-030, 2006.

[2] J. Norris, M. Powell, M. Vona, P. Backes and J. Wick, "Mars Exploration Rover Operations with the Science Activity Planner." In proceedings of IEEE International Conference on Robotics and Automation, 2005.

[3] Rebecca Castano, Robert Anderson, Tara Estlin, Dennis Decoste, Forest Fisher, Daniel Gaines, Dominic Mazzoni, and Michele Judd, "Rover Traverse Science for Increased Mission Science Return," *Proceedings of the 2003 IEEE Aerospace Conference.* Big Sky, Montana, March 2003.

[4] Rebecca Castano, Tara Estlin, Daniel Gaines, Andres Castano, Caroline Chouinard, Ben Bornstein, Robert C. Anderson, Steve Chien, Alex Fukunaga, and Michele Judd, "Opportunistic Rover Science: Finding and Reacting to Rocks, Clouds and Dust Devils," *Proceedings of the 2006 IEEE Aerospace Conference*, Big Sky, Montana, March 2006.

[5] T. Estlin, D. Gaines, C. Chouinard, R. Castano, B. Bornstein, M. Judd, I. Nesnas, 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.

[6] R. Simmons and D. Apfelbaum, "A Task Description Language for Robot Control," Proceedings of the Intelligent Robots and Systems Conference, Vancouver, CA, October 1998.