Rover Functional Autonomy Development for the Mars Mobile Science Laboratory

Richard Volpe
Space Exploration Technology Program
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91103
volpe@jpl.nasa.gov

Abstract—This paper provides an overview of the rover technology development, integration, and validation process now being used by the Mars Technology Program. Described are the relevant mission scenarios of long traverse and instrument placement, and the enabling algorithmic components that are being captured into a common software environment for demonstration and validation. As discussed, these components come from the ongoing 2003 rover mission, funded MTP research, and other complementary sources. All are providing software elements integrated into the CLARAty software system, enabling test and validation on a group of Mars rover experimental platforms.

TABLE OF CONTENTS

1. INTRODUCTION

Beyond the 2003 Mars Exploration Rovers (MER) Mission, NASA plans to send a larger, longer life Mobile Science Laboratory (MSL) in the 2009 timeframe. As envisioned in Figure 1, this rover is planned to survive 500 days, travel approximately ten kilometers, and demonstrate autonomous capabilities that reduce the number of communication cycles now needed to achieve successful completion of activities on the surface. Specifically, there are two primary categories of activity now being addressed by technology development efforts in The Mars Technology Program (MTP): long range traverse, and instrument placement.

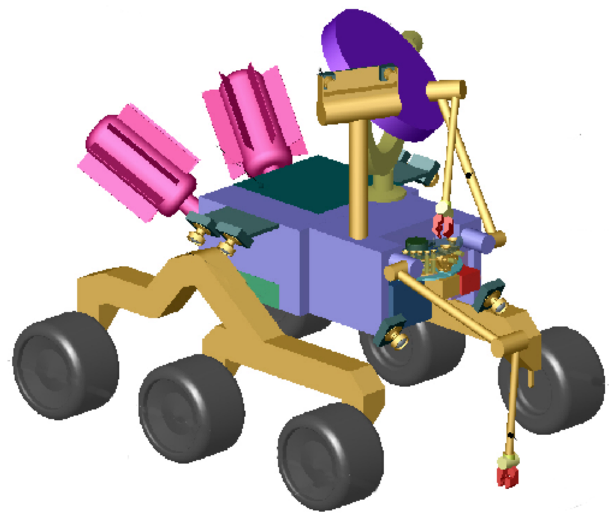

Figure 1 – Pre-decisional concept drawing of the 2009 MSL rover. For scale, the height is 2m.

Long range traverse for the 2009 mission is defined as driving hundreds of meters per day, and venturing safely and effectively through terrain not previously seen by operators via rover panoramic imagery. To achieve this, improvements are being developed in onboard algorithms for estimation of the rover position, estimation of the surrounding terrain qualities, and navigation decision making for driving to the goal in a safe and more optimal manner. Specifically, MTP is funding research in next generation position estimation using: visual odometry, soil sinkage and slippage estimation from wheel current and visual evidence, novel methods of inertial sensor placement and data processing, and integrated estimation software for

[1] 0-7803-7651-X/03/$17.00 © 2003 IEEE
[2] IEEEAC paper #1289, Updated December 1, 2002

combining all information. Next generation environment estimation is being addressed by new techniques for wide baseline stereo, correlation of surface and overhead imagery, and soil property estimation and correlation with imagery. Using this improved knowledge of the rover position and the qualities of the surrounding terrain, improved navigation is being provided by new power-sensitive global path planning software, and experimental evaluation of existing local navigation solutions.

Instrument Placement for the 2009 mission is defined as single-day positioning of an instrument on a rock selected by operators and scientists in panoramic imagery. Typically, this target would be at most ten vehicle lengths away, and no smaller than a single pixel in the panoramic imagery. For example, in MER the target on a rock will be selected from at most ten meters away, and be as small as one centimeter in size. However, MER will perform instrument placement in three days minimum, with strict oversight by operators. To move to single day capability, long traverse technologies must be augmented with others specific to the instrument placement problem: visual servoing on selected environmental features, autonomous recognition of scientific properties of the terrain, elevation map seaming from panoramic imagery, and onboard manipulator motion planning.

At this time, the described technology components are being integrated into the CLARAty (Couple Layer Architecture for Robotic Autonomy) software environment by participating researchers of MTP. From this suite of capabilities, long range traverse and instrument placement validation efforts will mix and match relevant capabilities to quantify their performance – both for near-term iterative improvements, as well as mid-term documented software delivery to the MSL flight project. This work will continue through FY05, thereafter expanding its scope to other potential rover-based missions.

2. ROVER TECHNOLOGY INFUSION

As more surface missions are anticipated for Mars, with elevated expectations of mobility and autonomy, it becomes important to develop a process for capture of advanced research capabilities in flight systems. Up to the present, this has often been accomplished by having technology developers assume positions on the flight team, and bring their technology components with them. However, such a process is not always feasible or desirable, and is biased against technology developers not located at the institution of the flight project.

The Mars Technology Program has attempted to remedy this situation by developing a process by which technology providers infuse their component technologies into a coherent whole, where they may be leveraged by other participants, compared with competing techniques, and validated for capture by upcoming missions. This process is designed to be distinctly different than its predecessors in the way it organizes participants, captures their technology products, and experimentally validates the resulting system capabilities prior to infusion into the mission. A diagram of the process flow is shown in Figure 2, and portions of it will be described throughout the remainder of this paper.

First, all technology providers are competitively selected through proposal calls and technical evaluation [1]. The content of the call is based on specified mission needs, currently provided by MSL. Resultant proposals must demonstrate that the technology to be provided is reasonably mature, addresses mission needs, lives within mission constraints, and can be transferred to MTP within the period of funding. Maturity should be at Technology Readiness Level (TRL) four at the start of funding, and demonstrated in the integrated MTP system at level six by the conclusion of funding [2].

Second, the product of these efforts is not just journal papers and documented results. Rather, the primary product of all providers is software delivered to MTP by integration into a common software environment. This software system, the *Coupled Layer Architecture for Robotic Autonomy* or CLARAty [3], is being actively developed and provided by MTP, and its support team actively assists technology providers with integration of their software products. Further, the CLARAty infrastructure itself if being developed through multi-institutional collaboration between JPL, NASA Ames Research Center (ARC), and Carnegie Mellon University (CMU). CLARAty will be reviewed more in Section 6.

Third, with the research products integrated into a common software system, they may be combined, compared, and quantified in their performance. Further, since CLARAty provides abstraction of, and support for, numerous test platforms, the performance may be elucidated independent of single platform particularities. Included amongst the platforms is a rover simulation system, ROAMS [4], which will allow for test trial repetition not possible by slower experimentation with physical rovers. Therefore, quantification of software and algorithm performance will not only be based on experimentation, but statistical results from simulation. Further, the experimentation will validate the simulation fidelity. The documented results will quantify the performance of the individual technology products, as well as their integrated configurations, in mission relevant scenarios.

Based on these results, the flight projects may make well informed decisions about which subset of technology software products will be used in the missions. Since mission infusion will come from a single, validated source of software, the complexity of the process is drastically reduced. Also, since the technology products have been

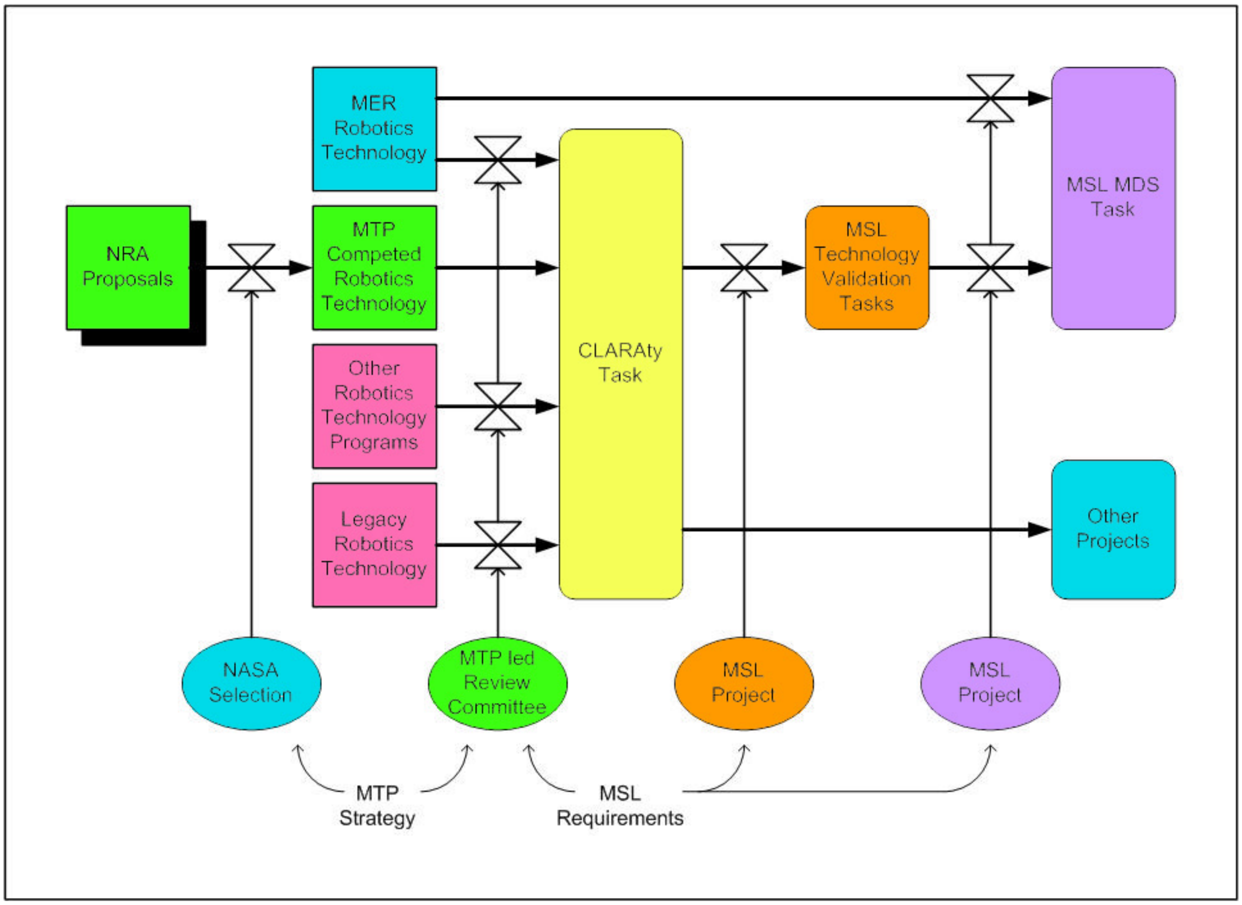

Figure 2 – Rover Functional Autonomy Technology Flow

decoupled from the individual providers, there is no implicit need to bring the developers into the mission to ensure success. This last point is especially important since the extended community of technology providers assumed by the initial competitive selection is distributed throughout the nation, and not readily available for mission support roles. But their technology components can and will greatly enhance the mission performance.

The following sections will describe in more detail examples of technology components going through this infusion process, the software architecture that binds them together, and the validation scenarios used in measuring their performance.

3. 2003 Mars Exploration Rovers

In May and July of 2003, the MER twin rovers will be launched, arriving at Mars in early 2004. (Figure 3 shows one of these rovers being tested in late 2002.) Once safely reaching the surface, a number of new capabilities will be utilized to drive to science targets and place instruments against them. The robotic capabilities are the product of previous NASA funding in the research program, transferred to the mission through an inconsistent process of software infusion. But once validated and used by the mission, they necessarily become the de facto standard for future mission performance comparison. Therefore, it is important to migrate these into CLARAty so that new technology products can be directly compared against them.
There are several robotic capabilities that define the baseline onboard MER:

a. Stereo vision

Each MER rover has three pairs of cameras available for onboard left/right image correlation resulting in depth perception and terrain elevation models. These will be use primarily for autonomous navigation of the vehicle, but also for manipulation and instrument placement [5,6].

b. Obstacle Detection and Navigation

To avoid obstacles and navigate to the goal, MER uses a software package called GESTALT [7], which estimates the local terrain traversability, and steers the rover to avoid nearby sensed obstacles while trying to get to the specified goal. GESTALT is derivative of navigation software from CMU [8].

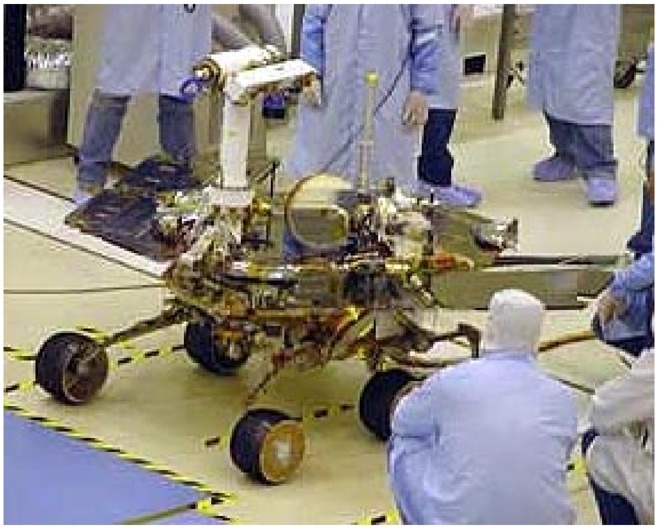

Figure 3 – One of the twin Mars Exploration Rovers.

c. Vehicle Kinematics

Kinematic computations are used to determine both steering angles and wheel rotations needed to effect the desired movement of the vehicle. The inverse is utilized to estimate actual motion of the center of the vehicle based on individual wheel motions. Similar mapping between joint angles and end effector motion is computed for the arm on each rover.

d. Position Estimation

In addition to position estimation of the vehicle based on solely on the wheel motion, inertial sensors are used. Integration of angular rate sensors provides an estimate of heading, which is much better than that obtained from wheel measurements. In addition, sun sensing is used at the end of the day to obtain an independent measurement of the vehicle orientation, both for navigation planning and communications antenna pointing.

In addition, there are several complementary off-board capabilities that are intended for use either within, or in conjunction with CLARAty:

e. Science Activity Planner

SAP is a product of the Mars Technology Program which has been adopted for MER mission use in collaboratively selecting science targets and establishing science activity sequences within mission resource constraints [9]. This software package is actually capable of acting as an entire ground data system for rover technology development, and will be interfaced with CLARAty to provide this functionality.

f. ROAMS

A subset of the ROAMS rover simulation environment is being used in MER for previewing rover commanded actions, and post-viewing telemetry [4].

g. Calibration

Calibration techniques for camera models and arm kinematics will be captured for off-board use. Validation will determine the accuracy these techniques, as well as possible improvements forthcoming [10].

h. Motion Planning

Motion planning for MER is of the nature of operator assistance tools. These assist manual selection of vehicle and arm motions, highlighting rough terrain or possible arm collisions [11]. For future missions, autonomous path planning will typically replace this functionality, both off-board and onboard.

4. MARS TECHNOLOGY PROGRAM

The Mars Technology Program, in conjunction with the Mars Science Laboratory Mission, is funding three complementary infrastructure elements: ROAMS, WITS, and CLARAty. In addition, MTP is funding eight competitively selected technology providers, and will be adding to this number through upcoming NASA Research Announcements (NRAs) [1].

Software Infrastructure

Rover Analysis Modeling and Simulation Software, or ROAMS [4], is a high fidelity rover simulation environment built upon a Dynamics and Real-time Simulation engine (DARTS) which was the 1997 recipient of NASA software of the year. The same underlying DARTS software is used for Entry, Descent, and Landing Simulation, thereby providing a complete simulation system for MSL needs.

ROAMS provides simulation services for off-line analysis, as well as acting as a virtual rover platform for CLARAty control software. In the latter mode, actuators, sensors, and environment are simulated at different levels of resolution appropriate for the controls problem. For instance, if control of vehicle slippage is being tested, then simulation of wheel-soil interaction is required. In this case, the interface to ROAMS is done at the level of individual wheels. Alternatively, when planning and execution algorithms are tested using the simulator, then connectivity is performed at the vehicle level, and wheel-soil interactions need not be explicitly calculated, instead modeled statistically if at all. This flexibility matches the simulation to the level of fidelity needed for the problem being addressed. It also, allows for increases in the speed of simulation, permitting more testing of lower frequency system control loops.

The Web Interface for Tele-Science, or WITS [9], is a operations software environment for perusal of rover telemetry and construction of sequences for rover control. A subset of its capabilities is used for the MER Science Activity Planner (SAP). It has also been demonstrated to provide a goal specification interface to planning and scheduling systems such as CASPER [12]. CASPER, in turn, has served as the prototype Decision Layer for the CLARAty architecture. Further interfacing between CLARAty and WITS will occur in FY03, tying WITS to the CLARAty Functional Layer, both for execution and telemetry.

Both the Decision Level and Functional Level of CLARAty will be described in Section 6.

Competitively Selected Rover Technology Components

As previously described, technology software algorithm developers have been competitively selected to address the needs of the MSL mission. The software products from these teams are being integrated into CLARAty for use with rovers and simulation surrogates, and access by the WITS operations interface. Currently, there are eight funded research teams, with more expected in the near future:

a. Driving on Slopes, JPL/Caltech

This research is improving the vehicle controls performance while driving on sloped and soft soils. Three specific issues are being addressed: visual estimation of position changes to overcome inaccurate odometry due to slippage [13], estimation techniques for visual and other estimates of rover position, and wheel steering and drive control techniques to keep the vehicle moving in the desired orientation and direction when disturbed by slippage.

b. Visual Servoing, JPL/Caltech

This work is combining previously demonstrated techniques in monocular and stereo visual tracking of terrain features [14,15]. The combined capability is expected to be more robust than either technique alone. The primary value of these algorithms are to track science targets selected by operators, enabling the rover to move robustly to them, and place instruments on them.

c. Autonomous Science, NASA Ames

Two of the limiting factors on the accomplishments of any remote spacecraft is the restricted communications bandwidth through which science data is returned, and the time consumed by the cycle of ground analysis and subsequent commanding. One solution to these problems is to analyze science data onboard the rover, enabling prioritization of science data telemetry, or immediately guiding rover actions in response to measurements. This research team is providing analysis algorithms for visual and spectrographic data, enabling onboard detection of rocks, layered terrains, and carbonate signatures [16,17].

d. Fault Diagnosis, NASA Ames

Similar to the science telemetry bottleneck, engineering analysis of system health is limited by communications bandwidth. This work is developing algorithms for onboard fault detection and diagnosis using particle filters. The first application of these techniques is toward manipulator health determination in the presence of motor failures and unexpected environment contact [18].

e. Vehicle Planning, Carnegie Mellon University

Leveraging on previous accomplishments in mobile robot path planning [19], this work is developing improved algorithms that add other system constraints into the state space during solution search. Primary amongst these is power, both its production and expenditure and their relation to the terrain. Algorithms developed to calculate solar power production have dual use for determining view angles for communications windows, and science imaging opportunities, which will also be factored into resultant plans [20].

f. Mapping, University of Washington

This work provides correlation of imagery from multiple sources to develop improved elevation maps of the environment around the rover [21]. One form of the research product is a capability for wide-baseline stereo, using images taken by the rover before and after a motion of several meters. Another form of the research provides elevation map seaming for panoramic stereo image data.

g. Terrain Estimation, MIT

This research is concentrating on non-visual techniques to estimate soil properties experienced by the rover platform. By measuring wheel torque, rotation, translation, and sinkage, terrain cohesion and internal friction angle can be accurately estimated [22].

h. Position Estimation, University of Michigan

This work is investigating improved position estimation based solely on inertial and encoder measurements. Through novel configuration of the sensors, and fuzzy logic based processing of the data from them, improvements over current position estimation techniques are anticipated [23].

5. LEGACY AND OTHER TECHNOLOGY

The active flight and research software development described above represents only a subset of technology available for capture and use on future rover missions such as MSL. There has been over 15 years of autonomous rover research funded by NASA, and many of the products of that funding do not have software implementations available, or the implementations are in heterogeneous systems [24,25,26,27,28,29]. To enable quantified performance according to metrics, and qualified performance by comparison to competitive techniques, it is necessary to bring these legacy technology products into a common software environment. There are several issues to be considered when reviewing and prioritizing legacy technology products:

- applicability to currently planned missions
- overlap with currently funded or integrated products
- level of maturity previously achieved
- completeness and quality of documentation
- ease of software capture or re-creation

All of these factors translate into a cost/benefit ratio that must be developed and prioritized. Such an effort is currently underway through a recently formed inter-institutional team formed by MTP, and it is anticipated that initial efforts of capturing legacy products will begin in FY03.

In addition to legacy products, which by definition have no current funding base, there are also complementary technology products being developed in other programs. A case in point is the NASA Code R Intelligent Systems Program (IS). In the recent past, IS has largely concentrated on Decision Layer technology such as planning and scheduling technology. There has already been some progress in incorporating these results by interfacing to the resulting software, but the bulk of MTP efforts have concentrated on Functional Layer controls technology. More recently IS funding has begun to cover areas of control, making the projects very complementary to those of MTP. Also, maturing controls infrastructure in MTP has led to the desire to interface to more of IS Decision Layer products. Details of this interaction are still under development at the time of this writing.

6. CLARATY

The 'Coupled Layer Architecture for Robotic Autonomy', or CLARAty, has been developed to serve as the technology integration software architecture for MTP [3]. From the beginning, it has been designed to satisfy multiple objectives:

1. Provide a common software environment for heterogeneous rover research platforms, and transparently include simulated versions.
2. Provide a generalized, modular, and reusable software framework, that spans existing and past robotics research.
3. Provide tight coupling of the traditional artificial intelligence (AI) fields of planning, scheduling, and execution, with the traditional robotics fields of sensing, estimation, and control.
4. Satisfy the design and usage objectives of participating institutions, including JPL, ARC, and CMU.
5. Utilize contemporary development tools such as object-oriented programming, UML documentation, distributed and collaborative design and development, comprehensive version control, etc.

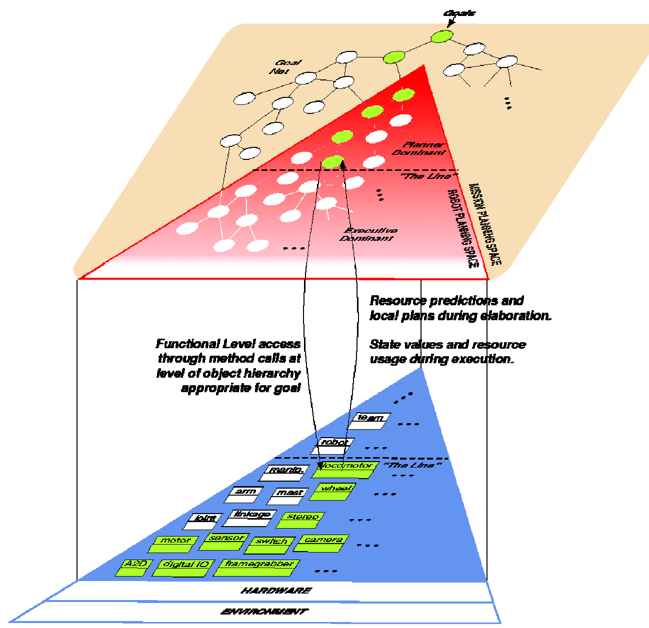

Figure 4 -- The CLARAty architecture with top Decision Layer and bottom Functional Layer.

Figure 4 shows the resultant design as a dual layer architecture with a Decision Layer (DL) for AI software, and Functional Layer (FL) for controls implementations. Implicit in the design is the concept of granularity, which increases for each layer, moving into the figure. FL granularity allows for the nesting of capabilities and the hiding of system details, often through the use of polymorphism. DL granularity allows for variability in the planning system time quanta, and conditional goal expansion.

As described by Figure 2, CLARAty serves as the integration environment primarily for MTP funded research, but also for capture of MER flight capabilities, IS program software products, and other legacy software relevant to MSL. Through abstraction of the hardware layers, it currently enables these software products to be transparently used on 4 custom research rovers (Rocky 7, Rocky 8, K9, and FIDO), one commercial platform (ATRV Jr.), and benchtop duplicates of these systems' avionics.

Integration of technology products to the software architecture, instead of the individual platforms, has a number of advantages:

- Timesharing of platforms for development and testing.
- Experimental comparison of similar techniques on a single platform.
- Experimental demonstration of the robustness of a single algorithm on differing platforms.
- Distribution of parts of the whole rover control problem across multiple research teams, with integration of new products later into the whole.
- Leveraging of the integrated products of others by all teams, thereby reducing overhead and duplication of effort by all.
- Centralization of the final resultant software system, providing a single source of technology products for infusion to flight systems.

This final point provides a pathway to flight, but doesn't necessarily provide the needed information by which the flight project can properly select amongst all technology components available in the research software environment. Therefore, validation of each is needed to provide the information for this decision process, as described next.

7. Technology Validation

After research technology products have been integrated into CLARAty and been verified by the providers to perform as expected, there is still a need for additional extensive testing. This is to validate the technology, by using it with multiple rovers and numerous conditions, and quantifying its performance. There are a number of reasons for this need:

- To provide independent verification that the technology providers have delivered what was claimed, and quantify the performance.
- To provide possible feedback to technology providers enabling fixes or improvements of their products.
- To test single technology components interacting with each other, and confirm there are no algorithmic or architectural problems.
- To test combinations of technology components grouped to achieve a single mission designated capability.

Specific to this last item there are two primary mission capabilities designated by MSL: long traverse, instrument placement. In addition, research is also addressing the enhancing capability of autonomous science data processing.

Long traverse requires autonomously driving distances on the order of 100 times the vehicle length. Many terrain features of significance, such as obstacles, will typically not be apparent in panoramic imagery provided to operators by the rover from its starting location. High resolution imagery from orbit may help map large scale terrain qualities, and may be used by operators or onboard the rover for global path planning. However, determining the original position of the rover and maintaining an accurate estimate during the traverse become important issues. This is especially true in soft terrains which cause slippage, or featureless terrains where visual correlation is difficult.

Instrument placement requires approaching a terrain feature designated by scientists from up to 10 vehicle lengths distant, and reliably placing a instrument on the feature. An important facet of this capability is keeping track of the target even while traversing toward it through rough terrain. A continuous line of sight may not possible, and differences in lighting or view angle may complicate the process. Also, the rough terrain expected for rock fields of interest can make navigation and position estimation difficult. As the desire target becomes close to the vehicle, another complication may be introduced by the necessity to use cameras with different focal length, stereo separation, field of view, and vehicle mount position. Finally, once the target is within the workspace of the manipulator with science instrument, the arm must be deployed safely and reliably to the target location. This last operation may require repetition for surface preparation steps, require force control for grinding operations or surface compliance, and must handle contingencies through lighting changes and thermal cycles during long deployments.

Finally, autonomous science data processing is seen as a mission enhancing capability that will be extremely important during 500 day missions. Three types of data processing are possible:

Data Compression – This provides passive categorization, or compression of data collected for other purposes. Examples might be as simple as cropping sky from images taken for geology, or using navigation imagery to quantify rock distributions during traversal.

Activity Suspension – This requires detection of known features using periodic measurements, and aborting current plans if specified conditions are met. An example of this type of capability would be to monitor periodic spectral readings and abort the remainder of a traverse if a carbonate signature is detected.

Conditional Activity Initiation – This is similar to above, except rover activities are initiated without further review by ground operators. An example would be suspension of a long traverse and initiation of an instrument placement operation, based on data collected during the traverse. While this level of capability is a goal for the technology program, it is currently considered by many as too aggressive for MSL.

Currently two funded activities are in progress to perform validation for long traverse and instrument placement. It is planned that a third activity will address autonomous science data processing validation beginning in FY04.

8. SUMMARY

This paper has provided an overview of the MTP technology development, integration, and infusion process for the upcoming 2009 MSL mission. A review of pertinent MER robotics capabilities has been provided, as has an overview of ongoing competitively selected technology development in MTP. These sets of technology are being captured into the CLARAty software environment, to leverage each other in performance of mission scenarios, and enable quantified validation of their performance. Results will be provided to the mission so that informed selections may be made for technology inclusion in the mission flight software.

9. ACKNOWLEDGEMENTS

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract to the National Aeronautics and Space Administration.

REFERENCES

[1] http://research.hq.nasa.gov/

[2] http://www.hq.nasa.gov/office/codeq/trl/

[3] R. Volpe, I.A.D. Nesnas, T. Estlin, D. Mutz, R. Petras, H. Das, "The CLARAty Architecture for Robotic Autonomy," proceedings of the *2001 IEEE Aerospace Conference*, Big Sky Montana, March 10-17 2001. http://robotics.jpl.nasa.gov/~volpe/papers/aerospace01.pdf

[4] J. Yen and A. Jain, "ROAMS: Rover Analysis Modeling and Simulation Software," in International Symposium on Artificial Intelligence, *Robotics and Automation in Space (i-SAIRAS'99)*, Noordwijk, Netherlands, June 1999.

[5] L. Matthies, ``Stereo Vision for Planetary Rovers: Stochastic Modeling to Near Real-time Implementation'', *International Journal of Computer Vision*, 8(1), July 1992.

[6] Y. Xiong and L. Matthies, "Error Analysis of a Real-Time Stereo System", *IEEE Conference on Computer Vision and Patter Recognition (CVPR),* Puerto Rico, June 17-19, 1997. http://www.cs.cmu.edu/~yx/papers/StereoError97.pdf

[7] S. Goldberg, M. Maimone, and L. Matthies, "Stereo Vision and Rover Navigation Software for Planetary Exploration," *2002 IEEE Aerospace Conference*, Big Sky Montana, March 2002, pp 2025-2036. http://robotics.jpl.nasa.gov/people/mwm/visnavsw/aero.pdf

[8] S. Moorehead, R. Simmons, D. Apostolopoulos, and W.L. Whittaker, "Autonomous Navigation Field Results of a Planetary Analog Robot in Antarctica," *International Symposium on Artificial Intelligence*, Robotics and Automation in Space, June, 1999. http://www.ri.cmu.edu/pub_files/pub1/moorehead_stewart_1999_1/moorehead_stewart_1999_1.pdf

[9] P. Backes, K. Tso, J. Norris, G. Tharp J. Slostad, R. Bonitz, and K. Ali, "Group Collaboration for Mars Rover Mission Operations," *IEEE International Conference on Robotics and Automation*, Washington DC, May 2002. http://wits.jpl.nasa.gov:8080/WITS/publications/2002-group-icra.pdf

[10] D. Gennery, "Least-squares Camera Calibration Including Lens Distortion and Automatic Editing of Calibration Points," in *Calibration and Orientation of Cameras in Computer Vision*, A. Gruen and T.S. Huang, Editors, Springer Series in Information Sciences, Vol. 34, Springer-Verlag, pp. 123–136, July 2001.

[11] C. Leger. "Efficient Sensor/Model Based On-Line Collision Detection for Planetary Manipulators," *IEEE International Conference on Robotics and Automation*, Washington DC, May 2002.

[12] T. Estlin, R. Volpe, I.A.D. Nesnas, D. Mutz, F. Fisher, B. Engelhardt, S. Chien, "Decision-Making in a Robotic Architecture for Autonomy," *The 6th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS),* Montreal Canada, June 18-21 2001. http://robotics.jpl.nasa.gov/tasks/claraty/reports/publications/isairas01.pdf

[13] C. Olson, L. Matthies, M. Schoppers, M. Maimone, "Stereo Ego-motion Improvements for Robust Rover Navigation," *IEEE International Conference on Robotics and Automation*, Seoul Korea, May 2001, pp 1099-1104. http://robotics.jpl.nasa.gov/~mwm/papers/icra01.ps.gz

[14] M. Maimone, I.A.D. Nesnas, H. Das, "Autonomous Rock Tracking and Acquisition from a Mars Rover," *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS),* Noordwijk, The Netherlands, June 1999, pp. 329-334. http://robotics.jpl.nasa.gov/tasks/pdm/papers/isairas99/

[15] D. Wettergreen, H. Thomas, M. Bualat, "Initial Results From Vision-based Control of the Ames Marsokhod Rover," *IEEE Int'l Conf on Intelligent Robots and Systems,* France, 1997, pp. 1377-1382.

[16] P. Gazis and T. Roush, "Autonomous identification of carbonates using near-IR reflectance spectroscopy during the

February 1999 Marsokhod field tests," *J. Geophys Res.-Planets,* February 2001.

[17] V. Gulick, R. Morris, M. Ruzon, and T. Roush, "Autonomous Image Analyses During the 1999 Marsokhod Rover Field Test," *Journal of Geophysical Research--Planets,* 106(E4), 25 April 2001, pp. 7745-7764.

[18] V. Verma, R. Simmons, D. Clancy, and R. Dearden, "An Algorithm for Non-Parametric Fault Identification," *AAAI Spring Symposium on Robust Autonomy*, Stanford CA, March 2001.

[19] Stentz, A., "Optimal and Efficient Path Planning for Partially-Known Environments," *IEEE International Conference on Robotics and Automation*, May 1994. http://www.frc.ri.cmu.edu/~axs/doc/icra94.pdf

[20] P. Tompkins, A. Stentz, and W. Whittaker, "Automated Surface Mission Planning Considering Terrain, Shadows, Resources and Time," *The 6th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS),* Montreal Canada, June 18-21 2001. http://www.ri.cmu.edu/pub_files/pub2/tompkins_paul_2001_1/tompkins_paul_2001_1.pdf

[21] C. Olson, F. Xu, K. Di, R. Li, and L. Matthies, "Automatic Feature Registration and DEM Generation for Martian Surface Mapping," in proceedings of the ISPRS Commission II Symposium, August 2002.

[22] K. Iagnemma, H. Shibly, and S. Dubowsky, "On-Line Terrain Parameter Estimation for Planetary Rovers," *IEEE International Conference on Robotics and Automation*, Washington DC, May 2002. http://robots.mit.edu/people/Karl/ICRA_02.pdf

[23] L. Ojeda, and J. Borenstein, "FLEXnav: Fuzzy Logic Expert Rule-based Position Estimation for Mobile Robots on Rugged Terrain," *IEEE International Conference on Robotics and Automation*, Washington DC, May 2002. http://www-personal.engin.umich.edu/~johannb/Papers/paper89.pdf

[24] "Robotic Vehicles for Planetary Exploration," *IEEE International Conference on Robotics and Automation,* Nice, France-May 1992.

[25] R. Volpe, J. Balaram, T. Ohm, R. Ivlev. "The Rocky 7 Mars Rover Prototype," *IEEE International Conference on Intelligent Robots and Systems (IROS),* November 4-8 1996, Osaka Japan. http://robotics.jpl.nasa.gov/~volpe/papers/rocky7.ps.gz

[26] P. Schenker, et al., "FIDO: a Field Integrated Design & Operations Rover for Mars Surface Exploration," *The 6th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS),* Montreal Canada, June 18-21 2001.

[27] D. Christian, D. Wettergreen, M Bualat, K. Schwehr, D. Tucker, E. Zbinden, "Field Experiments with the Ames Marsokhod Rover," *International Conference on Field and Service Robotics, Canberra Australia, 7-10 Dec 1997.* http://www.frc.ri.cmu.edu/~schwehr/Papers/mars-fsr97.pdf

[28] E. Krotkov and R. Simmons, "Performance of a six-legged planetary rover: power, positioning, and autonomous walking," *IEEE International Conference on Robotics and Automation*, May 1992, pp. 169-174. http://www.ri.cmu.edu/pub_files/pub2/krotkov_eric_1992_1/krotkov_eric_1992_1.pdf

[29] D. Wettergreen, D. Bapna, M. Maimone, and H. Thomas, "Developing Nomad for Robotic Exploration of the Atacama Desert," *Robotics and Autonomous Systems,* Vol. 26, No. 2-3, February, 1999, pp. 127-148. http://www.ri.cmu.edu/pub_files/pub2/wettergreen_david_1999_1/wettergreen_david_1999_1.pdf

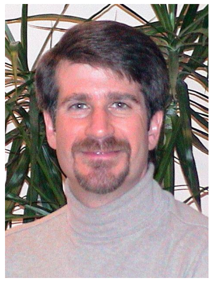

Richard Volpe, Ph.D., *is Manager of the Mars Regional Mobility and Subsurface Access Office of the JPL Space Exploration Technology Program. In addition to guiding technology development for future robotic exploration of Mars, he is actively involved in 2003 & 2009 rover mission development. His research interests include natural terrain mobile robots, real-time sensor-based control, manipulation, robot design, and path planning.*

Richard received his M.S. (1986) and Ph.D. (1990) in Applied Physics from Carnegie Mellon University, where he was a US Air Force Laboratory Graduate Fellow. His thesis research concentrated on real-time force and impact control of robotic manipulators. Since December 1990, he became a Member of the Technical Staff at the Jet Propulsion Laboratory, California Institute of Technology. Until 1994, he was a member of the Remote Surface Inspection Project, investigating sensor-based control technology for telerobotic inspection of the International Space Station. Starting in 1994, he led the development of Rocky 7, a next generation mobile robot prototype for extended-traverse sampling missions on Mars. In 1997, he received a NASA Exceptional Achievement Award for this work, which has led to the design concepts for the 2003 Mars rover mission. In 1999 and 2000 he served as the System Technologist for the Athena-Rover, then part of a 2003 Mars Sample Return Project. Until mid 2001 he was the Principal Investigator for the Robotic Autonomy Architecture and Long Range Science Rover Research Projects.

.