

Autonomy Needs and Trends in Deep Space Exploration

Dr. Richard J. Doyle

Manager, Information Technology and Software Systems Division
Leader, Center for Space Mission Information Systems and Software
Jet Propulsion Laboratory
California Institute of Technology
MS 126-221/4800 Oak Grove Drive
Pasadena, California 91109-8099 USA

rdoyle@jpl.nasa.gov

<http://it.jpl.nasa.gov>, <http://csmis.jpl.nasa.gov>, <http://cs.jpl.nasa.gov>

Abstract: The development of onboard autonomy capability is the key to a set of vastly important strategic technical challenges facing NASA: increased efficiency in the return of quality science products, reduction of mission costs, and the launching of a new era of solar system exploration characterized by sustained presence, in-situ science investigations and missions accomplished via multiple, coordinated space platforms.

Autonomy is a central capability for enabling missions that inherently must be accomplished without the benefit of ongoing ground support. This constraint may arise due to control challenges, e.g., small-body rendezvous and precision landing, or may arise due to mission planning challenges based in the difficulty of modeling the planetary environment coupled with the difficulty or impossibility of communications during critical or extended periods. A sophisticated Mars rover, a comet lander, a European under-ice explorer, and a Titan aerobot are examples of missions, some unprecedented, which typify these challenges.

This paper describes the set of NASA missions that aim to utilize autonomy and recent developments in the creation of space platform autonomy capabilities at NASA.

1. THE NASA MISSION CHALLENGES

NASA is embarking on a new phase of space exploration. In the solar system, an initial reconnaissance of all of the planets except Pluto has been accomplished. In the next phase of planetary exploration, the emphasis will be on direct, i.e., in-situ scientific investigation in these remote environments. In the next phase of exploration relating to astrophysics, the emphasis is on new observing instruments – often based on principles of interferometry – to achieve unprecedented resolution in remote observing. A theme that runs through all of these science missions is the search for life.

The development of onboard autonomy capability is on the critical path to addressing a set of vastly important strategic technical challenges arising from the NASA mission set: increased efficiency in the return of quality science products, reduction of mission costs, and the launching of a new era of solar system exploration characterized by sustained presence, in-situ science investigations and missions accomplished via multiple, coordinated space platforms. These new classes of space exploration missions, as a rule, require new capabilities and technologies.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 01 JUN 2003	2. REPORT TYPE N/A	3. DATES COVERED -			
4. TITLE AND SUBTITLE Autonomy Needs and Trends in Deep Space Exploration		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S)		5d. PROJECT NUMBER			
		5e. TASK NUMBER			
		5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Information Technology and Software Systems Division, Center for Space Mission Information Systems and Software Jet Propulsion Laboratory California Institute of Technology MS 126-221/4800 Oak Grove Drive Pasadena, California 91109-8099 USA		8. PERFORMING ORGANIZATION REPORT NUMBER			
		10. SPONSOR/MONITOR'S ACRONYM(S)			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
		12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited			
13. SUPPLEMENTARY NOTES See also ADM001519. RTO-EN-022, The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1.1 Mars and Autonomy

Mars is a primary target for future exploration, and certainly has captured the interest of the general public. The set of Mars missions under development differ from previous space exploration in one important aspect: they are being conceived as a collective whole, with the establishment and evolution of infrastructure at Mars as an important sub-goal.

The next rover mission to Mars will be the Mars Exploration Rover (MER) mission, which will place two rovers on the surface of Mars in 2004. This mission will extend the accomplishments of the famous 1997 Mars Pathfinder Sojourner rover mission in several ways. The larger size of the physical rover platforms allows for significantly upgraded payloads of half a dozen sophisticated instruments and tools on each rover. The traverse objectives for the rovers will be 100 meters per day, in contrast to Sojourner, which did not depart the vicinity of the lander. Finally, the nominal mission lasts 90 sols (Martian days), compared to 30 for Mars Pathfinder.

The Mars mission following MER will be the Mars Reconnaissance Orbiter (MRO) mission, in 2005. Although MRO will not contribute new autonomy capabilities, it is an example of how missions in the Mars program are related and support each other. MRO is expected to be a communications relay for future Mars surface missions, and more specifically, it will collect high resolution surface images, and these images will be utilized to plan landing sites for the more sophisticated Mars rover mission to follow, called Mars Smart Lander (MSL), in 2009.

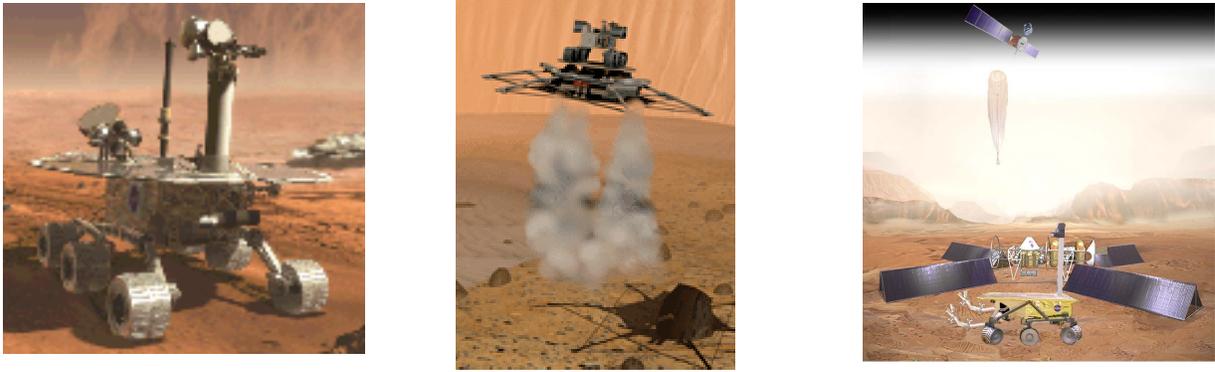
In the words of the MSL project manager, this mission will be “the most software-intensive, autonomy-dependent mission” to date. The emphasis on autonomy capabilities derive directly from the report generated by the study of the project’s Science Definition Team (SDT). Scientists are desirous of greater efficiency and reliability in commanding planetary surface rovers to generate science return, following the exciting, but often tedious experience with Sojourner on Mars Pathfinder, and anticipating that improvements on Mars Exploration Rover may be of an incremental nature.

The autonomy capability requirements generated by scientists on the SDT generally fall into three areas: the ability to command the rover, in a single command cycle, to move to a destination beyond the visible horizon; the ability, in a single command cycle, to reliably place an instrument on a designated target; and, for the first time, the ability to perform some limited science activities while traversing from one identified science investigation site to another.

The key phrase in the foregoing is “in a single command cycle.” Scientists and mission operators together have experienced frustration at the number of iterations sometimes required to reach a desired location or to successfully place an instrument. When the complicated constraints of light-time delay, in-view communication periods, and solar illumination cycles on Mars are factored in, such inefficiencies translate into serious limitations on overall science return. Conversely, if such actions can be accomplished autonomously and reliably in a single command cycle, overall science return is boosted greatly.

Another autonomy need on MSL is for safe and precise landing. The mission plans to utilize active hazard detection and avoidance capabilities during descent not only to achieve safe landing, but also to ensure that the target science sites are accessible from the landing site.

Looking further in to the future, infrastructure on Mars may include permanent science stations on the surface, propellant production plants, and a network of communications satellites in orbit to extend internet-like capability to Mars, and to enable the coordination of an array of heterogeneous, autonomous explorers: rovers, balloons, airplanes, even subsurface devices. No longer would each mission be conceived and executed in isolation, but through a combination of in situ and multiple platform mission concepts humanity’s presence at Mars would continually expand, culminating in the arrival and safe return of the first human explorers. See Figure 1.



**Figure 1. Future Mars Missions:
Mars Exploration Rover, Mars Smart Lander, Mars Outposts**

1.2 Small Planetary Bodies and Autonomy

Small planetary bodies – comets and asteroids – pose different kinds of challenges to space platforms that would investigate them in situ. Their gravitational environments are more difficult to model, particularly if the shape of the body is well off the spherical, and if it is tumbling rather than rotating. Achieving stable orbit is a complex task, much less landing. The environment of a comet is the most unpredictable: gas and dust and perhaps larger chunks of material can present hazards as well as make it difficult to select and track scientific targets of investigation.

The next mission to grapple with a cometary environment will be the Deep Impact mission, which will rendezvous with Comet Tempel 1 in 2005. The space platform for this mission consists of a flyby spacecraft and a detachable impactor. The impactor is designed to excavate below the surface of the cometary nucleus and reveal primordial material thought to be unchanged from the origin of the solar system. The targeting of the impactor has to be autonomous, because there is insufficient time to return images, analyze them, and send appropriate commands to the impactor. The image processing to guarantee successful impact is complicated by lighting, uncertain surface topography, and the inherent uncertainties of the environment.

The proposed Comet Nucleus Sample Return (CNSR) mission will involve autonomous landing and the return of samples. The unpredictable and volatile cometary environment will amplify the requirements for hazard detection and avoidance capabilities during descent and while on the surface. This mission concept calls for multiple site investigations, i.e., multiple takeoffs and landings in the course of the mission. The engineering and science considerations of autonomy merge in small body missions, as the same phenomena which represent potential engineering hazards (e.g., the formation of jets on comets), are also the phenomena of scientific interest. See Figure 2.



Figure 2. Future Small Body Missions: Deep Impact and Comet Nucleus Sample Return

1.3 Other Planetary Targets and Autonomy

The proposed Titan Organics Explorer would utilize a combination of platform concepts to conduct pre-biological chemistry and atmospheric investigations at Saturn's intriguing satellite, long known to possess an atmosphere, organic materials, and the possibility of a non-water ocean in an entirely different temperature regime. The platform concepts for this mission include an orbiter combined with an aerobot with detachable rover deployables. Aerobots utilize natural thermal cycles to periodically go aloft to sample multiple sites over a wide range of territory. When worthy science sites are found, the in situ investigation capabilities of surface explorers like rovers are then utilized. This type of exploration has a random element however, since landings can be only semi-directed, using direct control only of the vertical dimension, along with knowledge of wind patterns.

Europa is a notable focus for future exploration, second only to Mars as a target of interest within the solar system. The reason, of course, is the possibility that a liquid water ocean may exist beneath its surface, with obvious implications for the search for life. Three mission concepts for Europa exploration have been studied: an orbiter mission that can resolve the question of whether the subsurface ocean exists or not, followed by a lander mission, and ultimately, a cryobot/hydrobot mission. The orbiter and lander would have the additional challenge of survivability in the intense radiation environment at Europa, deeply embedded in the Jovian magnetosphere. If the European ocean does indeed exist, the cryobot/hydrobot mission concept involves melting through the ice surface of Europa and releasing an underwater submersible to reach and explore the ocean floor, looking for signs of life. The submersible would require high degrees of autonomy to perform its mission successfully, including onboard algorithms embodying knowledge of possible biosignatures. See Figure 3.



Figure 3. Other Future Planetary Missions: A Titan Explorer, A European Submersible

1.4 Astrophysics and Autonomy

Looking beyond the solar system, NASA has a series of next-generation deep sky observing missions planned in its Origins Program, whose end goal is the capability to image Earth-like planets around nearby stars, even to resolve features and perform spectroscopic investigations of such planets. The hallmark mission in this series is known as the Terrestrial Planet Finder, a deep-space-based interferometer consisting of multiple elements. These elements are guided at unprecedented precision to first null the light, via interference effects, coming from the primary star in these distant stellar systems, then collect the precious photons coming from any planetary companions it may possess. The autonomy challenge in this mission involves guidance and control of the interferometer itself during observations, and detection of and compensation for faults and performance degradations in the elements such that the collective capability of the multiple-platform interferometer is maintained. See Figure 4.

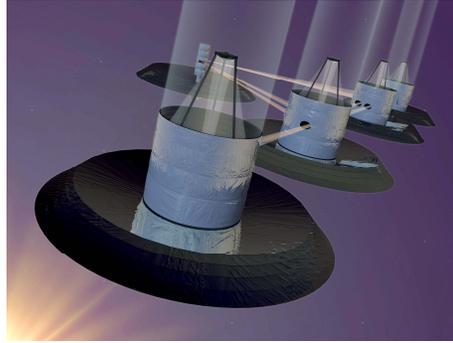


Figure 4. A Future Astrophysics Mission: Terrestrial Planet Finder

2. THE EMERGENCE OF AUTONOMY

Intelligent, highly autonomous space platforms will evolve and deploy in multiple phases. The first phase involves automation of the basic engineering and mission accomplishment functions of the space platform. The relevant capabilities include mission planning and resource management, health management and fault protection, and guidance, navigation and control. Stated differently, these autonomous capabilities will make the space platform *self-commanding* and *self-preserving*. Some of the relevant technologies include Artificial Intelligence (AI)-based planning & scheduling, model-based reasoning, and intelligent agents. In this initial *engineering-directed autonomy* phase, NASA space platforms will achieve onboard automated closed loop control among: planning activities to achieve mission goals, navigating, maneuvering, and deploying instruments and sensors to execute those activities, and detecting and resolving faults to continue the mission without requiring ground support. Also in this phase, the first elements of *science-directed autonomy* will appear. However, the decision-making capacity to determine how mission priorities should change and what new mission goals should be added in the light of intermediate results, discoveries and other events would still reside with scientists and other analysts on the ground.

Work on automating the spacecraft will continue into challenging areas like greater onboard adaptability in responding to events, closed-loop control for small body rendezvous and landing missions, and operation of the multiple free-flying elements of space-based telescopes and interferometers. In the next phase of autonomy development and deployment, a portion of the scientist's awareness will begin to move onboard, i.e., an observing and discovery presence. Knowledge for discriminating and determining what information is important would begin to migrate to the space platform. The relevant capabilities include feature detection and tracking, object recognition, and exploratory data sampling. At this point, the space platform begins to become *self-directing*, and can respond to greater uncertainty within the remote operating context. Ultimately, a significant portion of the information routinely returned from platforms would not simply and strictly match features of stated prior interest, but would be deemed by the onboard software to be "interesting" and worthy of further examination by appropriate science experts on the ground. At this point, limited communications bandwidth would then be utilized in an extremely efficient fashion, and "alerts" from various and far-flung platforms would be anticipated with great interest.

For surveys of NASA autonomy technology activities, see [1,2].

2.1 The Remote Agent

The most notable and successful effort in spacecraft autonomy development at NASA to date has been the Remote Agent, a joint technology development project by NASA Ames Research Center and the Jet

Propulsion Laboratory (JPL) [3]. The Remote Agent Experiment was conducted on the New Millennium Deep Space One (DS1) mission in May 1999, whose primary goal was to flight validate new technologies.

The Remote Agent consists of a Smart Executive [4], a Planning and Scheduling module [5], and a Mode Identification and Reconfiguration (MIR) module [6]. The onboard system receives mission goals as input, which is translated to a set of spacecraft activities free of resource and constraint violations by the Planner/Scheduler. The Smart Executive provides robust, event-driven execution, with runtime monitoring and decision-making. MIR continuously monitors representations of sensor data, identifying current spacecraft modes or states, and when these are fault modes, selects recovery actions. Other functions such as guidance, navigation and control, power management, and science data processing are domain-specific functions that can be layered on top of this basic autonomy architecture, and are developed or modified for each new mission. The Remote Agent was designed to serve as a general spacecraft autonomy architecture.

The demonstration objectives of the Remote Agent Experiment (RAX) on DS1 included nominal operations with goal-oriented commanding, closed-loop plan execution, onboard failure diagnosis and recovery, onboard planning following unrecoverable failures, and system-level fault protection. All of the technology validation objectives for RAX were accomplished. Additional details may be found in [7]. The Remote Agent was a co-winner of the NASA Software of the Year Award in 1999.

2.2 Some Definitions

There is often confusion on the differences between automation and autonomy. Acknowledging that the definitions given here may not be regarded as the final word on the subject, we make the following distinctions:

Automation applies to the creation of functionality (typically via algorithms), which can be fully defined independent of the context in which the functionality will be deployed, or when the context (e.g., the remote environment) can be modeled with sufficient confidence that the required functionality is well understood.

Autonomy, on the other hand, applies to the creation of functionality (typically via reasoning or inference capability), which is designed to be effective when context is important, and when the ability to model context (again, e.g., the remote environment) is limited. Autonomy specifically includes the capability to assess context and to support decisions based on knowledge that will only be available when the functionality is accessed, not when it is created. Knowledge and importance of context is the key consideration for distinguishing the need for automation vs. autonomy.

Handling of context is the central difference, but there is also a way in which automation and autonomy are similar: the functionality being created, whether based in algorithms or reasoning, is in both approaches fixed at deployment time (typically, launch time).

In the remainder of this section, we examine the research and technology investment areas that are contributing to the creation of autonomy capabilities for NASA missions. All of the work described is sponsored either by NASA's Intelligent Systems Program or NASA's Mars Technology Program.

2.3 Planning and Execution

Automated planning, scheduling, resource management and execution form, arguably, the core of system-level autonomy. These capabilities, when integrated, provide the basis for a space platform to perform engineering functions in closed-loop fashion onboard.

Planning involves reasoning about activities that will accomplish a desired set of goals. Automated planners work with models of tasks, resources and constraints, determining a set of activities which when executed, moves the space platform from its current state to a state where the desired goals have been achieved. In

addition to the basic task of planning (determining the relevant tasks and their ordering), the planning system must also determine the absolute or relative timing of the tasks (scheduling), ensure that no relevant resources are oversubscribed (resource management) and that no flight constraints are violated. The mission plan provided by the planning system is then passed to an execution system, or executive. The executive issues the specific commands contained in the plan, and monitors their execution. Execution violations can be handled in several ways: the executive typically has mechanisms to accomplish local recoveries, global recoveries are the domain of system-level fault protection, and finally, the planning system itself can generate a new plan when it is determined that assumptions in the plan have been violated, e.g., the conditions of the operating environment, or the availability of a resource or capability.

The planner from the Remote Agent, known now as EUROPA, has continued to mature. Other topics, relating to the interaction of planning and execution, are also under investigation [8]. There is work on contingent planning, which is examining the partial expansion of alternate paths of execution in a mission plan [9]. The idea is to anticipate the kinds of failures that may occur during plan execution, and to predetermine recovery actions for those failures, thereby providing both for greater robustness of execution in the plan, and obviating the need to engage full-up, computationally expensive replanning by default when execution failures occur. Another benefit is that ground personnel have the opportunity to validate alternate execution pathways before they may be engaged.

Another successful planning system is ASPEN/CASPER [10]. This planner is based on a continuous planning model where a plan is always being modified, either to optimize it further, or to incorporate the latest information available. This process can begin from a description of an initial state. ASPEN has been deployed on several applications, notably for the Antarctic Mapping Mission, where it achieved an order-of-magnitude reduction in the time required to create schedules for downlink activities, as well as supported “what-if” negotiations between the science and operations teams. CASPER is the high-performance version of ASPEN designed for onboard use, which can achieve plan turnarounds in the tens of seconds.

2.4 Science Data Understanding

Scientists are naturally skeptical of an autonomous capability that purports to be able to interpret raw science data in an informed fashion. The key to progress in this area is working closely with scientists to properly define the scope of any such onboard capability so that it is realistic and value-added. As described above, scientists associated with the Mars Smart Lander mission are interested in a capability to collect science data during traverses and to alert scientists on the ground when certain pre-defined phenomena have been detected.

Science-directed autonomy will typically rely on technologies for pattern and object recognition and machine learning techniques. The system known as QuakeFinder was successful in the use of a change detection capability [11]. QuakeFinder applies a registration technique to before and after images. Rather than insisting on global registration, the technique instead focuses on successful registration of local regions in the images, and as a side effect detects change along the boundaries of those regions. The technique is particularly suited to the detection of linear displacements. QuakeFinder has been successfully applied to the detection of ground movement due to earthquakes, detecting such motion at 1/10 pixel resolution. This system is now being applied to Voyager and Galileo images of Europa, looking for evidence of surface changes that may be consistent with the existence of a subsurface ocean.

Object recognition techniques previously applied successfully to the automated detection of volcanic features in the Magellan SAR image set of Venus are now being applied to the more general problem of crater detection on planetary bodies [12]. The crater detection problem is harder because craters may overlap, be in various stages of degradation, and be observed in diverse lighting conditions. Nonetheless, this crater detection capability has been successfully trained on Mars images and tested on lunar images. Putting this kind of capability onboard would be value-added from a planetary scientist’s viewpoint: crater studies are more about determining the number and distribution of craters on a planetary surface, rather than examining images of each individual crater.

Another activity is examining the interaction of onboard science data analysis with mission planning and execution for a Mars rover [13]. The data analysis capability here is a combination of texture analysis and spectral analysis looking to detect mineralogical signatures. The system summarizes and prioritizes science targets based on data collected and analyzed during traverse activities, and generates an update to the mission plan. This new plan can be communicated to scientists on the ground for modification and approval or executed autonomously.

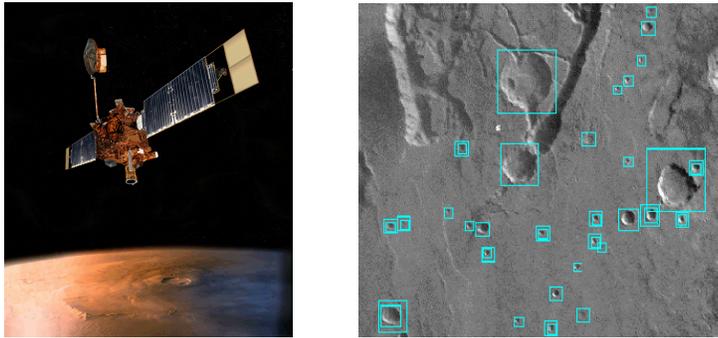


Figure 5. Craters Detected in Viking Data

2.5 Safe and Precise Landing

Many future mission concepts call for the ability to land space platforms on planetary bodies, ranging from small and hard-to-model asteroids and comets, to larger targets with more predictable environments, such as Europa and Mars. Several machine vision-based techniques have matured sufficiently to be the basis for safe and precise landing capabilities [14].

For example, spacecraft motion can be estimated accurately by precisely tracking features during descent. These image-based techniques can estimate motion to 1% accuracy of the total distance traveled. Feature tracking also can be used to accurately estimate position relative to defined landmarks, such as the centers of craters. These feature-based navigation techniques use different methods than, e.g., the crater recognition algorithms described earlier. When the objective is navigation, the task is to select features that can be reliably tracked under different viewing angles and lighting conditions. Features of scientific interest, however, are selected by different criteria, and may or may not line up with features that are useful to support navigation.

The stereo techniques that support surface navigation for rovers can provide critical information during the late stages of descent, when elevation knowledge of the topography directly below can be the key to final guidance to a safe landing spot. These techniques can compute elevation information to greater than one part in a hundred accuracy relative to altitude.

These vision-based techniques are being combined into an integrated hazard detection and avoidance capability for missions such as Mars Smart Lander, to locate and guide the landing spacecraft to a safe zone with local surface roughness of less than ten centimeters and local slopes of no greater than ten degrees. Path planning algorithms also can guarantee that there is accessibility from the landing site to the pre-defined science investigation sites. Benchmarking indicates that all of the required computation can be achieved in real-time and near real-time during the critical and intense entry, descent and landing sequence.

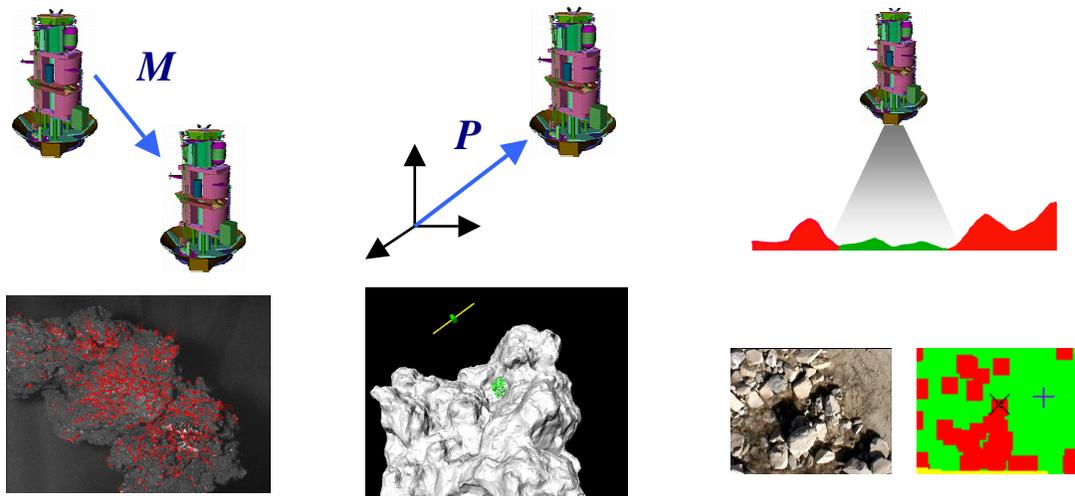


Figure 6. Vision-based Motion Estimation, Position Estimation, Hazard Detection and Avoidance

2.6 Mobility

Mobility is a defining feature of planetary rovers. Future rover missions are being planned now around the ability to go further and through rougher terrain than missions to date have accomplished. The four salient questions of mobility are: “Where am I?”, “What is the environment surrounding me?”, “Where should I be?”, and “How do I get there?”. These are the challenges of, respectively, position estimation, terrain estimation and obstacle detection, goal selection and tracking, and trajectory generation and path planning [15].

Addressing the challenge of position estimation has proven to be most difficult in the barren, rough, natural terrain known to exist on Mars. While current capabilities are able to provide error bounds down to 3% of the distance traveled, newer techniques promise to reduce that margin by an order of magnitude. These techniques include visual odometry for tracking the motion of features in navigation imagery to independently measure displacement of the vehicle, and visual servoing for tracking the motion of a selected visual objective to ensure motion toward that objective. Sensor fusion can play an important role to merge multiple measurement sources through statistical and logical filters, improving an overall position estimate. Related terrain estimation techniques address this challenge by estimating soil properties through visual and contact methods to determine sinkage and slippage of the vehicle. Other techniques match displacement of terrain map features on an intermittent basis to enhance estimates of vehicle position.

Obstacle detection via stereo processing is by now a standard technique, but is limited currently to providing surface elevation estimates only in the immediate vicinity of the rover. Continuing work in this area includes elevation map seaming to merge elevation maps provided by separate sets of stereo image data, and wide baseline stereo for the correlation of imagery taken from separate locations, providing elevation maps out to much greater ranges than is currently possible. Also, multi-resolution mapping techniques correlate surface and orbital imagery.

Mission operators currently perform the tasks of goal selection and tracking. Moving this function onboard the vehicle will require maturation of the position and terrain estimation techniques described above along with maturation of onboard techniques for planning and execution.

Finally, the challenge of trajectory generation and path planning concerns methods for navigating from the current location to the goal location, given the terrain. At this time, it appears that the techniques to be utilized on the MER mission will be sufficient for missions like MSL.

2.7 Distributed Autonomy

Mission concepts involving multiple space platforms are becoming increasingly common. The first examples of so-called spacecraft constellations will likely be in Earth orbit, where fleets of satellites carrying different sensors and instruments may coordinate observations and responses to provide global coverage for the detection of events such as volcanic eruptions or forest fires. Some future space-based observatories will consist of multiple elements, carefully arrayed and maintained to create apertures of unprecedented scale. Finally, heterogeneous assets deployed at Mars will collectively perform science missions, as well as count as the first true example of off-planet infrastructure.

The engineering capabilities of a single space platform do not scale simply to appear as a coordinated capability across multiple platforms. This is most apparent for mission planning, considered by many to be the core capability of system-level autonomy. An important consideration is architecture. Three different architectures for distributed planning are under study [16].

Perhaps the most straightforward is a centralized architecture, where a specially designated space platform performs the planning for all platforms. This is conceptually indistinguishable from mission planning on a single platform but has distinct disadvantages: intense communications to provide timely relevant data from all of the platforms to the central platform, and the need to update the models used by the central planner whenever another platform is added to the configuration.

In a distributed architecture, a central planner only allocates goals to the separate platforms, each of which has a planner of its own to autonomously achieve the goals assigned to them. Less communications are required in this architecture, but planners do appear everywhere, and some models used by the planners still need to be shared, so that the central planner can make informed choices about assigning goals. Whenever another platform is added to the configuration, the central planner requires a minimal model of that platform's capabilities, so that it can continue to assign goals appropriately.

In a market-based architecture, each platform has a model of its own capabilities, and bids to accomplish goals posted in auctions managed by a central platform. This architecture is the most scalable, because no platform requires models of the capabilities of other platforms, which can be added to the configuration at any time. There is global computational overhead however, for all planners examine all goals.

The first example of distributed autonomy in flight will occur in 2004, when the Autonomous Sciencecraft Constellation, a technology experiment on the New Millennium ST-6 mission, will test an integrated capability for onboard mission planning and onboard science data analysis aboard the Air Force's three-satellite constellation known as Techsat-21 [17].

2.8 The Role of Software Architecture

The Remote Agent technology experiment on DS-1 made clear the central role that software engineering must play in deploying autonomy capabilities in flight. In particular, it is highly desirable that the software architecture used in flight provides direct support for autonomy capabilities. The Mission Data System (MDS) is a flight/ground/test software architecture conceived and designed to be such an autonomy-friendly architecture [18]. In this goal- and state-based architecture, goals are defined to be constraints on the values of state variables over specified time intervals. Estimators interpret measurement and command evidence to estimate state. State variables hold state values, annotated with estimates of uncertainty, critical for informed, autonomous decision-making. Controllers issue commands, striving to achieve goals. Explicit models express specific relations among states, commands and measurements. Goals and states are the atomic concepts of autonomy, and MDS thereby provides a head start on implementing autonomy capabilities.

The MDS approach offers other important benefits, such as providing a common language of discourse for systems engineers and software engineers, who otherwise interact imperfectly at best across a pile of textual requirements. MDS also utilizes a components-based software architecture, and designs out certain classes of software defects at the architectural level, such as certain race conditions, along with units or dimensionality errors.

The Mars Smart Lander mission has adopted the MDS software architecture for these reasons, including its system engineering approach of state analysis.

3. SUMMARY: LOOKING FARTHER INTO THE FUTURE

The preceding on the emergence of autonomy included descriptions of the initial phases of development and deployment of autonomy capabilities for space platforms: Engineering-directed autonomy will place many traditional spacecraft functions such as mission planning, execution, resource management, guidance and control, navigation and fault protection onboard in closed-loop fashion. Science-directed autonomy will allow scientifically interesting phenomena to be detected via onboard capabilities, supported by onboard mission replanning.

Beyond these initial phases, we can project a phase where space platforms become the analogues of web nodes, with direct interaction enabled among space platforms, the science community, and the general public. Interested users may “register” with autonomous spacecraft to learn about breaking results.

The next phase may involve self-organizing constellations of space platforms consisting of heterogeneous assets performing joint, coordinated execution of mission objectives, with self-calibration and adaptation enabled.

A phase beyond that may be characterized by long-term survivability of ten years or more, even with zero ground support, achieved by onboard contingency handling and self-repair, including functional redundancy achieved via software, particularly planning capability.

We’re not done yet. The next phase may involve onboard reprogramming and discovery, with the spacecraft authorized to modify science objectives and the mission plan based on onboard learning and discovery capabilities.

And in a final phase, with spacecraft and mission evolution enabled, space platforms may not only repair themselves, but may also refine and improve their functionality, as a more efficient form of mission redesign and reprogramming.

As with any vision, the most remote projections get a bit harder to track and to immediately credit. But we can be confident that autonomy is here to stay as a central capability for achieving future NASA missions. Autonomy is multidisciplinary in nature and must be the product of the inputs of computer scientists, spacecraft engineers, mission designers, ground system engineers, mission operators, scientists, software engineers and systems engineers.

REFERENCES

- [1] Richard J. Doyle, "Spacecraft Autonomy and the Missions of Exploration," Guest Editor's Introduction, Special Issue on Autonomous Space Vehicles, *IEEE Intelligent Systems*, September/October 1998.
- [2] Daniel E. Cooke and Butler Hine, "Virtual Collaborations with the Real: NASA's New Era in Space Exploration," *IEEE Intelligent Systems*, March/April 2002.
- [3] Douglas Bernard, Gregory A. Dorais et al, "Design of the Remote Agent Experiment for Spacecraft Autonomy," in *Proceedings of the IEEE Aeronautics Conference*, Aspen, CO, March 1998.
- [4] Barney Pell, Erann Gat et al, "Plan Execution for Autonomous Spacecraft," *15th International Joint Conference on Artificial Intelligence*, Nagoya, Japan, August 1997.
- [5] Nicola Muscettola, Benjamin D. Smith et al, "On-board Planning for the New Millennium Deep Space One Spacecraft," *Proceedings of the 1997 IEEE Aerospace Conference*, Aspen, CO, February 1997.
- [6] Brian C. Williams and Pandu Nayak, "A Model-based Approach to Reactive Self-Configuring Systems," *13th National Conference on Artificial Intelligence*, Portland, OR, August 1996.
- [7] Douglas C. Bernard, Gregory A. Dorais et al, "Spacecraft Autonomy Flight Experience: The DS1 Remote Agent Experiment," *Proceedings of the AIAA Space Technology Conference and Exposition*, Albuquerque, NM. September 1999.
- [8] Nicola Muscettola, Gregory A. Dorais, et al, "IDEA: Unifying Planning and Execution for Real Time Autonomy," *Proceedings of the i-SAIRAS 2001*, Montreal, Canada, June 2001.
- [9] John L. Bresina, and Richard Washington, "Robustness via Run-time Adaptation of Contingent Plans" *Proceedings of the AAAI Spring Symposium on Robust Autonomy*, Stanford, CA, March 2001
- [10] Russell Knight, Gregg Rabideau, et al, "CASPER: Space Exploration Through Continuous Planning," *IEEE Intelligent Systems*, September/October 2001.
- [11] Paul Stolorz and Peter Cheeseman, "Onboard Science Data Analysis: Opportunities, Benefits, and Effects on Mission Design," Special Issue on Autonomous Space Vehicles, *IEEE Intelligent Systems*, September/October 1998.
- [12] Michael C. Burl, Tim Stough et al, "Automated Detection of Craters and Other Geological Features," *Proceedings of the i-SAIRAS 2001*, Montreal, Canada, June 2001.
- [13] Tara A. Estlin, Tobias P. Mann et al, "An Integrated System for Multi-Rover Scientific Exploration," *Proceedings of 16th National Conference on Artificial Intelligence*, Orlando, FL, July 1999.
- [14] Yang Cheng, Andrew Johnson, et al, "Passive Imaging-based Hazard Avoidance for Spacecraft Safe Landing," *Proceedings of the i-SAIRAS 2001*, Montreal, Canada, June 2001.
- [15] James Cutts, Samad Hayati et al, "The Mars Technology Program," *Proceedings of the i-SAIRAS 2001*, Montreal, Canada, June 2001.
- [16] A. Barrett, G. Rabideau et al, "Coordinated Continual Planning Methods for Cooperating Rovers," *Proceedings of the IEEE Aerospace Conference*, Big Sky, MT, March 2001.
- [17] Steve Chien, Rob Sherwood et al, "The Techsat-21 Autonomous Sciencecraft Constellation," *Proceedings of the i-SAIRAS 2001*, Montreal, Canada, June 2001.
- [18] Daniel Dvorak, Robert Rasmussen and Thomas Starbird, "State Knowledge Representation in the Mission Data System," *Proceedings of the IEEE Aerospace Conference*, Big Sky, MT, March 2002.