On-board Plan Modification for Opportunistic Science

Derek Long University of Strathclyde, UK
Mark Woods and **Andy Shaw** SciSys UK Ltd, Bristol, UK
Derek Pullan University of Leicester, UK
Dave Barnes and **Dave Price** University of Aberystwyth, UK

Abstract

The pressure of demands for greater cost-effectiveness is leading space agencies to seek more efficient ways to exploit the science platforms they launch. It is clear that one route to greater efficiency is to migrate some of the decisions about operations from the ground to the platforms themselves. EO-1 has been an effective demonstration of the potential benefits that this migration can offer, with rewards in the increased density of science value in communicated data and the improved responsiveness of the craft in its task of collecting relevant data. In this paper we present some findings from a project conducted to explore the possibility of performing *opportunistic science gathering* in a planetary rover mission, focussing on geological science data. This problem raises several distinct challenges, including the automatic identification and evaluation of possible science targets, estimation of the impact of attempting to collect the data these opportunities might offer, making decisions about the operations required to collect the data and, finally, the task of enacting those operations. We outline our solutions to all of these elements, but concentrate, in this paper, on the decision-making elements.[1]

1 Introduction

It is well understood that current procedures for operating remote space missions are a compromise between the conservative use of a vulnerable system and the ambitious demands of scientists seeking scientific data. This compromise is significantly impacted by communication constraints which mean that operations are typically planned on the ground for a window of hours or days ahead and then downlinked to the craft for subsequent unsupervised execution. There are several consequences of this process: firstly, with very limited capacity on-board to react to unexpected results of actions, the standard response to significantly off-nominal behaviour is for the craft to enter safe mode, and secondly, the science gathering operations are limited to targets that are already known to the scientists at the point of planning. Both of these limitations have motivated work exploring the introduction of intelligent on-board systems, offering on-board diagnosis and replanning or plan repair (Jönsson *et al*. 2000; Muscettola 1993; Fox *et al*. 2006) and on-board science evaluation and response (Chien *et al*. 2004; Estlin *et al*. 1999; Castaño *et al*. 2007). The former work is focussed on attempting to respond to problems during plan execution, recovering from operational failures by replanning in an attempt to avoid losing operations that might be unaffected by minor execution problems. The latter work, in contrast, is concerned with attempting to exploit opportunities to acquire scientific data that might otherwise be lost. The OASIS system, in particular (Castaño *et al*. 2007), has very similar objectives to the work described here, including practical demonstration on physical hardware. In OASIS the planning function is supported by CASPER, with execution achieved via a mapping to Task Description Language (TDL) (Simmons & Apfelbaum 1998).

In this paper, we report on our experience in building and testing a system demonstrating the on-board recognition of science targets, the response to target identification and the management of the replanning necessary to integrate operations on the targets within an existing plan.

2 The Context

Remote planetary rovers, such as the current MER mission (Squyres 2005) or the planned ExoMars mission (van Winnendael, Baglioni, & Vago 2005), all face the problem that limited communications impose significant delays in operations and dramatically reduce responsiveness of the missions. These missions are essentially mobile platforms for supporting geological investigations and, in the case of the Pasteur payload planned for ExoMars, organic molecular analysis. In order to perform experiments, in most cases the instruments, which are mounted on an arm, must be placed close to, or touching, the science target. Therefore, the operation of one of these platforms consists of a fairly standard pattern: an initial imaging survey of the landscape around the rover is used by scientists to select science targets. Routes are then planned by which the rover can navigate to the first of these targets, usually in a series of waypoint-to-waypoint legs. On arrival, further imaging surveys are conducted on the target and its surroundings. A series of short manoeuvres (micro-traverses) is often required to approach the target and then science operations can begin. Science operations on a target of interest typically con-

[1] Several other papers have been written based on this work, including (Woods *et al*. 2009) which covers the project in considerable detail. Nevertheless, this paper includes new material describing the function of the replanning subsystem.

sist of arm deployments for close-up and microscopic imaging, contact analysis with spectrometers, surface preparation (using an abrasion tool), further imaging and analysis of the prepared surface and, finally, perhaps sample acquisition for analysis inside the rover.

The individual steps in this process are often short — navigation to a target might take tens of minutes and most science operations, including deployment of a robotic arm, have similar orders of magnitude. The most time and power-consuming activity is the use of power tools such as rock abrasion or drill tools, which might consume as much as a few hours of operation. In principle, the whole survey might be completed well within a single Sol. In practice, the need to monitor progress, to respond to unfolding events and to make intelligent choices at various stages in the process, coupled with constraints on communications, all conspire to make these operations extend over several Sols. During this time, the rover will spend significant periods idle, awaiting new instructions.

A similarly frustrating sluggish rate of science acquisition by Earth monitoring satellites motivated some of the design of EO-1 (Chien *et al*. 2004). This craft carries software that is able to perform initial analysis of image data to determine whether cloud cover has obscured the target. In cases where the data is worthless, the on-board system can choose to discard it and reschedule a new observation. This reactive response to on-board data evaluation, with replanned activity also generated on-board, is judged to be responsible for vast savings in the operations management for this satellite.

A further consequence of the current procedures for rover operations is that transitory phenomena, such as dust devils (Bornstein *et al*. 2007), are hard to explore. The difficulty is that these phenomena cannot be identified in advance, with operations planned to gather data. The MER team have been fortunate to capture images of dust devils by chance, but there is no possibility of reactively tracking these phenomena, or performing any planned study. This problem extends to less transitory science targets, where the opportunity to carry out investigations is limited not by the movement of the target, but by the movement of the rover. During extended traverses, a rover might pass by interesting targets that were occluded or obscured in the original surveys, so not selected as mission objectives during planning. A traverse might pass close by such a missed target and gather tantalising images of the lost opportunity, transmitted back to Earth long after the traverse is complete.

The authors undertook the construction of a demonstrator to explore the use of on-board science evaluation software as a project funded by the UK Science and Technology Facilities Council CREST (Collaborative Research in Exploratory Systems Technology) program. The objective of the project was to address the problems outlined above. In particular, we sought:

1. to address the problem of single-cycle instrument placement (SCIP), also sometimes referred to as single command approach and instrument placement (SCAIP) (Backes *et al*. 2005; Hunstberger *et al*. 2005; Pederson *et al*. 2005), which alleviates the first problem outlined above, the extended delays during operations;
2. to perform on-board science target identification and evaluation;
3. and to support opportunistic science data acquisition by combining the first two functions with replanning.

Our objective was to demonstrate this combined functionality in a 'Mars Yard' environment.

3 The Architecture

The architecture of our demonstration system is based around three key components: a science assessment and response agent (SARA), the arm agent and perception interface (AAPI) and the timeline validation, control and replanning system (TVCR). The broad structure of the architecture is shown in figure 1. The executive receives an initial plan from mission planners on Earth. During execution, TVCR is used to monitor execution, as described in (Fox *et al*. 2006), responding to plan failure by proposing plan modifications to the executive. The executive can choose to accept the modifications or to behave more cautiously, perhaps entering safe mode. During normal operation, the executive will also request, at intervals, a science assessment to be performed on image data acquired from the navigation cameras. This analysis is performed by SARA. If SARA determines that there is a possible target, further evaluation will be requested and the executive can then request that the necessary resources for this evaluation can be scheduled by TVCR. Assuming further evaluation is performed and determines that a target of appropriate priority has been found, further replanning requests will be issued to TVCR to plan science operations. If a plan can be constructed to include the operations, then requests will be issued by the executive, according to the plan, to deploy instruments. These requests are handled by the AAPI, which solves the arm placement problem and deploys the appropriate instruments. The AAPI uses closed-loop control of the arm to resolve the approach and placement problem.

In order to perform useful close-up and contact-based experiments, the targets selected for evaluation must be situated in locations that can be approached. Since the navigation of the rover is a complex and hazardous task, this effectively limits operations to targets that are within reach of the rover as it carries out planned traverses. However, it is generally the case that targets will not be positioned ideally for direct interactions using the arm, so we assume, in the design of the AAPI subsystem and in the replanning, that micro-traverses are part of the repertoire of operations permitted to the on-board operations planning and execution. This is an ambitious assumption given the current state of missions planning: it would already require a significant new level of trust of on-board control systems to allow payload operations to be planned on-board (the operation of the arm and instruments), but to support rover operations, even if restricted to a very constrained subset, would be a dramatic advance. Nevertheless, micro-traversals are a grey area at the intersection of payload and rover operations and this makes it plausible to anticipate that future missions might support on-board control of this function.

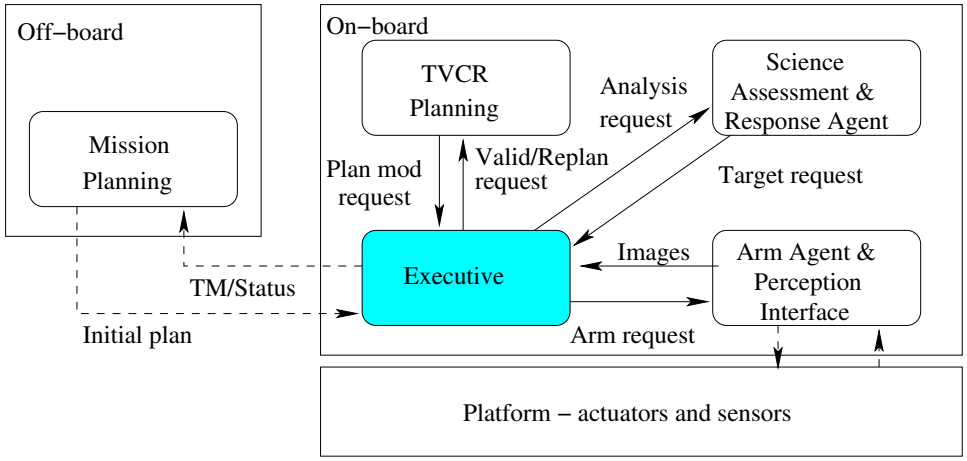

Figure 1: Architecture of the Autonomous Science Evaluation and Acquisition System

We now consider the elements of the system in more detail.

4 The Science Assessment and Response Agent

Geologists select the targets of their investigations based on experience, knowledge and visual assessments. Visual assessments are best performed at a variety of scales from remote to close-up, since the key features that influence selection of targets, including texture, structure and composition, are hard to identify at long range. This is one of the reasons that opportunistic selection of science targets when the rover passes close by possible candidates could lead to discovery of important new scientific data.

SARA uses image analysis to extract clues that indicate the possible value of candidates. A series of features characterising geological aspects of candidate targets was developed, representing structure (eg shape, scale, orientation), texture (eg luster, relief, grain) and composition (eg albedo, colour, mineralogy). Candidate targets are evaluated by summing weighted scores associated with these elements to arrive at a science value score (SVS). Weights are dependent on the regional geology of the area being explored and the expectations and interests of the scientists. The weighted sum promotes feature-rich targets, since each new feature accumulates with others to increase the SVS for the target.

Image data can be acquired at various ranges from a possible target, and the range impacts on the way that the target is evaluated and the response that might be expected. For this work we defined three ranges: proximal ($\sim$100cm), macroscopic ($\sim$10cm) and microscopic ($\sim$1cm). Initial target selection, on the ground, is performed at even longer ranges, but the function of SARA is intended to be for evaluation of potential targets close to the rover during traverse operations. As noted above, long traverses are hazardous to perform without an initial survey and it is unlikely that such operations will be devolved to on-board control in the near term. Therefore, the range of our proximal observation area is deliberately chosen to be within the range of microtraverse and arm placement. In our study, we only considered assessment of image data at proximal and macroscopic ranges. In part, this was a consequence of constraints on time and resource within the project, but it is also the case that microscopic imaging depends on the deployment of instruments, so already depends on a decision having been made to commit the necessary resources to closer inspection of a target, possibly coupled with additional experiments on the target.

Image interpretation includes an initial segmentation phase, to identify individual candidates within the image, followed by a feature extraction process (Shaw & Barnes 2003). In this work we focussed on successful extraction of features indicating *bedding*, which is potentially associated with sedimentary rock formation by deposition of layers and possible water erosion. However, the techniques are designed to be extensible to include other features.

5 Timeline Validation, Control and Replanning

In earlier work (Fox *et al*. 2006; Howey, Long, & Fox 2005), we have explored the development of an on-board plan management system, (Timeline Validation, Control and Replanning, TVCR) that can intelligently monitor plan execution and repair plans that have suffered from execution failure. In the development of this system we have maintained a careful sensitivity to the need for control to be managed, as far as possible, within parameters determined by mission planners. It is certain that mission planners will want to retain control of planned operations while their is sufficient time and opportunity for them to do so, while ceding control when the situation the rover finds itself in is different to that predicted on the ground. To achieve this balance, the design we have adopted is one in which the initial plan is constructed on the ground, by missions planners. The plan is composed of primitive actions organised into coherent connected collections, called *plan fragments*.

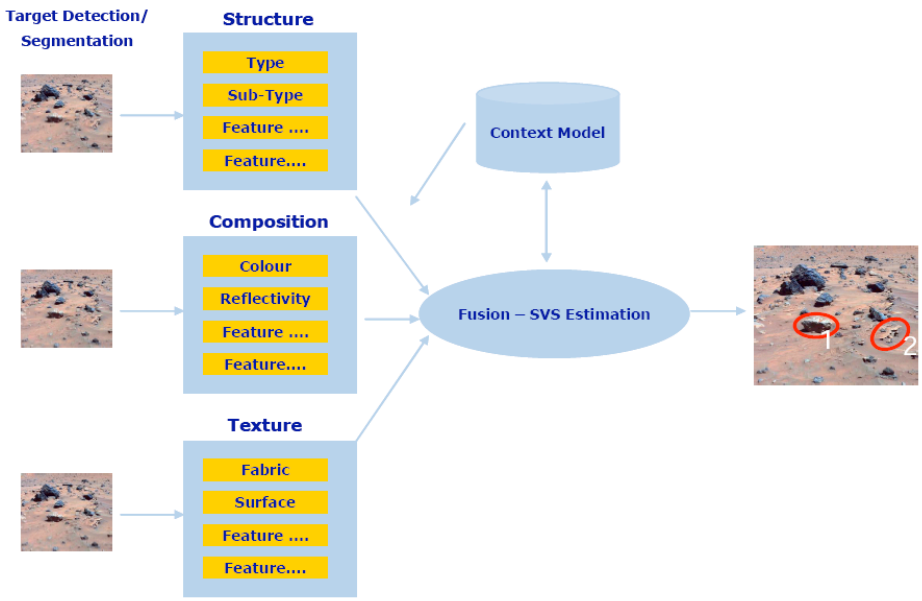

Figure 2: SARA: Target assessment is built up through the analysis of individual attributes. Context is included to qualify the final score.

5.1 Plan Fragments

Plan fragments contain not only the primitive actions that make up the activity they describe, but also constraints between the actions that must hold in order for the fragment to be executed usefully. It is worth emphasising that the causal relationships between the primitive actions are not explicitly recorded in the fragments, since the primitive action model is supplied, complete with pre- and post-condition descriptions, as a separate initial information source, written in PDDL+ (Fox & Long 2006). The purpose of the constraints within a plan fragment is to capture methodological constraints on the use of actions that are not implied by the logical constraints enforced by their pre- and post-conditions. It also allows constraints to be imposed that reflect elements of the real world interactions between actions that might prove hard to model in the action structures we define. For example, the gap between warming up an instrument and using it is typically constrained by the rate at which heat might be lost by the instrument if it is turned off. Modelling this accurately is very difficult and not very useful, since, in practice, standard operations will require an instrument to be used within a short period of being warmed up. This constraint can be captured in the PDDL model of the actions themselves, but only with cumbersome artificial actions that capture the temporal relations involved; a simpler solution is to add a constraint to the plan fragment in which these actions appear. PDDL3 (Gerevini & Long 2006) offers a more convenient way to capture constraints of this sort than PDDL2 (Fox & Long 2003), using trajectory constraints, but the language we have used for describing plan fragments is less general, having developed by incremental extension as examples have been encountered.

Further constraints can be imposed between plan fragments themselves. These include ordering constraints, which require that plan fragments be executed in a particular order if they both appear in a plan, dependencies where one plan fragment can only be included in a plan if its dependencies are met by other plan fragments being included in the plan, and mutual exclusion constraints that prevent certain pairs of fragments from both appearing in the same plan. Temporal constraints can also be imposed, so that fragments can only be executed during the day, or during the night, or at a particular absolute time. These constraints are used to ensure that modifications to the plan maintain properties that do not necessarily follow from the logical structures of the actions. For example, a diagnostic operation for an instrument can be made mutually exclusive with experiments using the instrument, since there will never be any reason to perform a diagnostic if the instrument is assumed to be operational, even though there is no logical constraint on the operation itself to prevent it.

5.2 Plan Repair

The original use of TVCR was in the monitoring of plan execution, using plan validation (Howey, Long, & Fox 2004; Fox, Howey, & Long 2005) to confirm that the remaining part of the planned operations remain executable (according to the model). If a plan fails to validate, then repair strategies are used to modify the plan to achieve a new executable plan. The repair strategy consists of two phases: a reduction phase in which fragments are removed from a plan until the plan is executable, followed by an extension phase in which the released resources are allocated to new plan fragments that can exploit them.

Each plan fragment has an associated priority value, indicating the importance attached to the execution of the frag-

ment by the missions planners (and the scientists who might have requested the experiments they represent). The priorities are used to determine which plan fragments should be dropped first in order to achieve an executable plan. We assume that there is a simple minimal plan that is guaranteed to be executable, which is the plan in which all fragments are removed and the single action instructing entry into safe mode is placed in the plan, together with the action committing to interaction in the next communication opportunity. In general, we would also expect that a plan containing only the communication commitment should be executable, provided that the current state satisfies any constraints on energy and other safety thresholds that govern operations. Reduction to this plan is preferred since it avoids the need to switch the rover back out of safe mode, which can be a time consuming process. Reduction is performed by iterative greedy elimination of plan fragments, using priority as the metric. Note that priorities simply provide an ordering on fragments — their absolute values are not intended to be interpreted as an indication of utility, so there is assumed never to be a situation in which a combination of lower priority fragments might be preferred to a higher priory fragment. This greatly simplifies the problem of deciding which elements to maintain and which to drop, which is a combinatorial problem sharing structure with Knapsack and Orienteering problems. Indeed, one way to tackle the problem of intelligent selection of goals when faced with oversubscription is to view the problem as an orienteering problem (Smith 2004).

Repairs can remove fragments, move fragments (in order to exploit availability of temporally dependent resources) and reorder fragments. In general, the removal or reordering of fragments can leave breaks in the causal linkage of the activities that form a plan, where preconditions of a fragment are no longer satisfied in the state in which the system is left by the execution of the preceding fragments. In our application, there are key activities that 'glue' fragments together and which we do not see as part of the fragments themselves. These activities move instruments to specific locations (arm deployments or paw rotations) or place instruments in particular modes. Selection of the correct sequence of activities to link together fragments is a planning problem, but it is highly specialised and constrained, so can be solved by a very simple means-ends analysis, apart from the planning of arm movements between specific joint-configuration waypoints. The latter is treated as a symbolic path-planning problem (the actual kinematics required to move the arm between waypoints are solved separately, as discussed below) and solved by a simple shortest path algorithm.

Once an executable plan is achieved, resources might have been released by the removal of activities from the original plan. Along with the fragments that form the original plan, missions planners have the opportunity to add a library of plan fragments we call *opportunities* and these may be inserted into the plan if the resources become available (subject to the constraints that might hold between these fragments and those in the plan). To simplify the problem of determining whether an insertion might be possible, each fragment has an associated estimated cost in terms of power, time and other resources (in our case, data store). They also list the instruments they rely on. This information allows a rapid initial pass to eliminate candidate fragments for which there is insufficient resource budget or for which instrument requirements cannot be met. In our previous experiments (Howey, Long, & Fox 2005), the numbers of opportunities have been small and it is a lightweight task to consider insertion using a greedy iterative incremental selection process.

The iterative plan repair strategy we use shares a great deal in terms of overall structure and motivations with work by Steve Chien and his team on iterative plan repair for space operations (Chien *et al*. 1999; Rabideau *et al*. 1999; Chien *et al*. 2000).

5.3 Handling Opportunistic Science Targets

In one sense, the need to address the existence of opportunistic science targets has little impact on the TVCR behaviour described above: exploiting an opportunity amounts to inserting a new plan fragment into the current plan, if possible. This is a slight change to the way in which TVCR is expected to respond to events, since the opportunity is not a failure in an existing plan, but rather a possibility to enhance the current plan. This can be handled by invoking the extension phase described above, restricting attention to the new opportunity that has been recognised. There are, however, several reasons why this process is complicated by the discovery of these targets during execution of a plan.

1. The location of newly discovered potential targets will typically not be a waypoint, but will lie between waypoints on a traverse. This means that the location has no identifier and does not appear in the problem description for the current plan.
2. The target itself is not named in the problem description.
3. No parameters that describe the characteristics of the tasks involved in interacting with the new target are given in the problem description, so it is impossible to check whether actions that depend on availability of resources can be applied, since their preconditions are not instantiated.
4. Obviously no fragments referring to the target will appear in the opportunities library.

In the original version of TVCR there is a component responsible for forming a symbolic representation of the current state to act as the initial state from which plan validation proceeds. This is constructed incrementally, by updating the model from the original initial state as actions are dispatched and observed to have successfully completed. Aspects of the current state that can be observed via sensors are entered into the initial state directly, so, for example, the power levels are taken from measured values rather than from predicted levels calculated by forward simulation of the actions so far executed from the original initial state. The expected costs of actions that involve interaction with waypoints or targets in the original plan are available in the original initial state model. Problem 1 listed above can be solved by making a change to this state-construction component, allowing new waypoints to be introduced for navigation and new targets to

be introduced for science operations. In the case of navigation waypoints, we only have to be able to introduce a new name for a location that lies part way along the current traverse, corresponding to the point at which a potential target has been identified and the current location of the rover. It is straightforward to split the expected costs of the current traverse into that part that has already been spent on getting to the new waypoint and the part that remains unspent, modifying the plan to replace the current traverse with the linked pair of traverses passing via the current location. A further new location must be added, corresponding to the operations site for the new target. This location can only be reached by micro-traverse from the current location and, if it is visited, a similar micro-traverse is required to return to the current location in order to resume the interrupted traverse.

A further extension of the current state model is required to capture the properties of the new target, in order to solve problem 3. Firstly, solving problem 2, the target must be named and its name added to the collection of objects in the problem description. This is easy, but it is important to observe that the addition of new objects to the problem description (both waypoints and targets) has an impact on any use of grounding. In our previous implementation, we used off-line grounding of actions to allow telecommands to use indices to refer to action instances. To retain validity of these indices when the problems are extensible requires care in the process by which they are generated on-board. Furthermore, in order to communicate a history of events from on-board back to the missions team on the ground, it is necessary to ensure that the generation technique can be repeated on the ground to obtain a consistent set of indices for action instances. This instantiation process is important as a way of solving problem 4: mission operations personnel can define standard template plan fragments for basic science operations that can be conducted on opportunistic science targets. These are then instantiated using the objects that are created to name the new targets. The priority of these instantiated fragments is derived from a template function applied to the SVS returned by SARA. This allows missions personnel to decide the relative priority to be given to opportunistic science gathering, while also accounting for the anticipated value of the opportunities themselves.

The properties associated with the new targets are understood in stages, so the initial state will be extended as a new target is explored. As a consequence, TVCR will be called on several times during the evaluation process, each time requiring a new initial state including the emerging properties of the target being considered.

The sequence of events is as follows:

1. A candidate target is spotted in an image at proximal range.
2. TVCR is called to determine whether a further science evaluation activity is possible. This can be determined by requesting insertion of a new fragment corresponding to the instantiated relevant library fragment template. At this stage, the only information that is known about the target is its position relative to the rover and the only activity that is sought is capture of a zoom pancam image. The impact of this activity will be to cause the rover to halt while the pancam is targeted and the image captured. The parameters that govern the use of this fragment are the power, time and data store demands. These values are approximated, conservatively, by the Arm Agent, AAPI, and used by TVCR to determine whether the fragment can be inserted without preventing completion of any higher priority activities. In contrast to the plan repair function considered in our previous work, in this case it is no longer reasonable to assume that the opportunities are lower priority than anything already in the plan. This means that TVCR might be faced with the need to consider reducing the plan before extending it, by removing lower priority fragments than the new opportunity.

   In principle, many lower priority fragments might need to be removed, incrementally, before the new fragment might be inserted. Recall that the priorities are not considered to be a measure of utility, so the problem is not a knapsack packing problem. In practice, it is unlikely that more than one or two fragments might be threatened in this process and the cost of considering each reduction in turn is not high. Furthermore, a quick filter can be applied to use the estimated costs of the new fragment and the resources released by lower priority fragments to identify a minimal subset (in priority order) that must be removed before the new fragment might be inserted. The original plan is backed up during this process, to restore if the new fragment cannot be inserted.

   TVCR responds at this point with a 'go/no go' signal and modifies the plan accordingly.
3. Assuming that the response is 'go', the Arm Agent completes deployment of the pancam to obtain a zoom image. SARA evaluates the zoom image to decide whether macroscopic (close-up) investigation should be considered. If the target generates a high SVS at this stage, a new fragment will be instantiated, corresponding to an approach using micro-traverse, instrument deployment and data capture.
4. TVCR is again called to determine whether this new fragment can be inserted. A similar process is used as before. In this case, the estimate of resource demands is required for each of the individual elements of the fragment: the micro-traverses to and from the target operations site, the arm deployment and the instrument use. The most difficult of these to assess is the arm deployment costs and an estimate for this is provided by the Arm Agent using an approximation of the kinematics it uses to solve the actual deployment problem. It is important that this should be a lightweight estimate, in case the cost of assessment of the opportunity threaten the availability of resources to complete other activities in the plan.

   It might be considered that the original request to TVCR should consider whether there is sufficient resource to follow through with the instrument deployment, since it might seem pointless to perform the evaluation if the subsequent response will be to reject the opportunity. Although there are some circumstances in which the follow up response could be determined in advance, there are two

reasons why we chose to delay consideration of the experiment fragment pending the second science evaluation. The first is that the zoom image and its analysis already has a scientific value, reflected in the SVS attached to the first image. The second reason is that the precise priority of the second fragment depends on the SVS determined by the SARA evaluation of the zoom image, making it harder to anticipate which fragments might be dropped from the current plan to make resource available for the new opportunity.

If TVCR determines that the new fragment can be successfully added to the plan, a 'go' signal is returned, together with the modified plan, and the executive instructs the AAPI to proceed to approach and perform instrument placement.

5. At completion of the experiment, the plan continues with the micro-traverse back onto the original path and the continuation of the interrupted traverse. In principle, this fragment of the plan might have a low enough priority that, if resources are limited, it is displaced from the plan. This, of course, can only happen as a consequence of a deliberate selection of an appropriately low priority by the mission planners on the ground, but it offers the team a way to pause the rover at a site of particular scientific interest through the next communication window, so that a ground-based assessment can be completed and an appropriate follow through decided.

The preliminary identification and evaluation of science targets requires resources: power, computation and data storage. This means that it is actually necessary to plan when to perform this task, taking into account the parallel demands on these resources from other activities. Therefore, the science evaluation is not performed continually on the images from the navigation cameras, but is performed on pancam images that are captured as part of the original planned data acquisition, always performed at identified waypoints. Thus, although the TVCR can support the addition of new waypoints at arbitrary points along a traverse, in fact the science assessments are always at known locations. Nevertheless, the waypoint required for the base of operations at a target is still required in order to model the required micro-traverses. The ability to split traverses is functionality that will support a more ambitious opportunistic science assessment in the future.

6 Arm Agent: Approach and Placement

Control of the arm is achieved using closed-loop control with the cameras acting as the key sensors in the loop. The cameras include panoramic cameras and a zoom, high-resolution camera (pancam) with pan and tilt control. The platform we used in our experiments, which is an approximately half-sized chassis model of the ExoMars rover design, is shown in figure 3. The arm was modelled as a lightweight and low-powered assembly of shafts, offering three degrees of freedom and using potentiometers to sense joint positions.

The first task of the AAPI is to respond to a request initiated by SARA (mediated by the executive) to capture a close-up image of a potential target. To achieve this requires calculation of appropriate pan, tilt and zoom settings for the pancam. Before the activity is actually performed, the calculation is used to estimate power and time demands and this information is used by TVCR to determine a 'go/no go' response, based on the impact that this estimated cost will have on the plan.

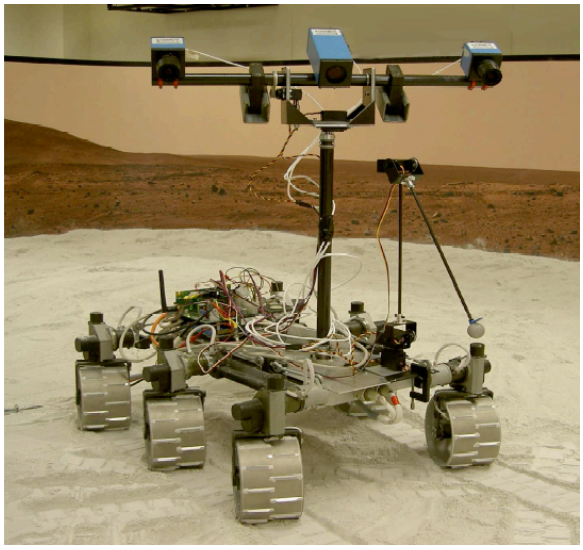

Figure 3: The rover chassis used in the experiments, with stereo panoramic cameras and pancam and the simple arm model.

When instructed to proceed to the instrument placement step, the AAPI uses a stereo wide-angle camera image pair to construct a digital elevation map of the target. This is combined with arm kinematics to determine candidate placement sites and arm trajectories. This calculation checks for arm collision risks along the trajectory. A key problem in this context is handling occlusion, which is a common phenomenon at the macroscopic level when considering arm placement on rock targets. Arm placement trajectories are computed using an inverse kinematics algorithm.

Further details concerning the AAPI, which is not the focus of this paper, can be found in (Woods *et al*. 2009).

7 Field Trial

Figure 4 shows the set-up used in testing. The tests were carried out on a Mars analogue surface, using a half-sized scale model of the ExoMars chassis, with purpose-built camera mounts and a simple dummy arm. The area of the test site is clearly rather limited, which has the beneficial effect of making trials of reasonable duration, since traverse times are short, but it limits the size of the plans we could use with TVCR. In the tests on the physical device, TVCR is presented with a very simple problem involving one iteration of the complete cycle described in section 5.3. By altering the placement of a planned communications action, which has highest priority, we confirmed that TVCR could accept or reject attempts to perform opportunistic science operations,

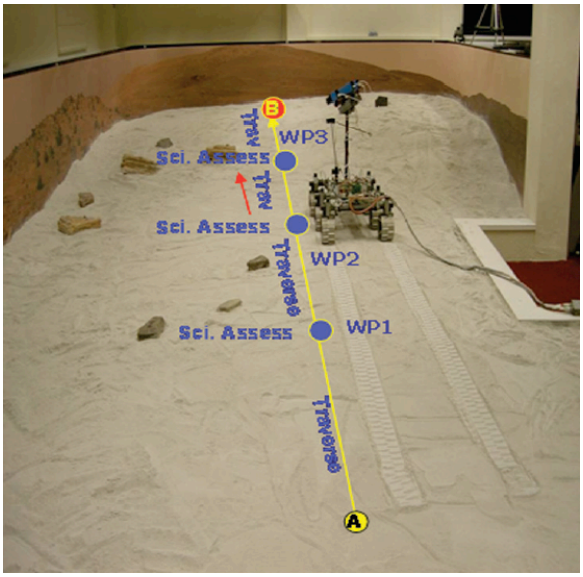

Figure 4: Field trial experimental set up.

with both 'go' and 'no go' cases being possible at each decision point.

Although the scenario we envisage includes possible micro-traverse operations to approach the target, we did not have an autonomous navigation component implemented in this field test. Therefore, the micro-traverse operation was performed manually, with appropriate time allowance in the plan to allow for manual operation.

In stand-alone tests, TVCR is able to handle more complex scenarios involving combinations of plan repair and opportunistic science insertion, including in cases where opportunities are identified during a traverse rather than at a waypoint. By simulating various operating conditions, it is also possible to create scenarios in which early parts of a plan execute more efficiently than predicted, releasing resources that can then be used to exploit opportunities that would be excluded by constraints on future operations were these unexpected efficiencies not to have materialised.

TVCR is built on the VAL validator software (Howey, Long, & Fox 2004), which is not written for deployment on-board, but has been developed incrementally as an academic research software tool. Coupling it with the systems in the rest of the trial was achieved using CORBA, making communication relatively expensive. Fortunately, the trials were not hampered by the limited computational resources likely to be available on real mission hardware.

8 Conclusion

Our experiments have generated valuable data for the design of ExoMars, in all aspects. Development of the TVCR subsystem to support on-board plan repair and modification continues to be a strategic priority and reimplementation of the system as a verified and deployable software component is planned for the next 12 months.

The translation of progress in automated planning into the control of autonomous space systems is challenging. The challenges are not only technical, but also operational: it is critical that the process achieves growing levels of trust and experience has to be accumulated by both researchers in planning and operations staff who are faced with the prospect of interacting with systems with increasing levels of on-board autonomy. It is clear that the pressure for improving efficiency will lead to a steady adoption of technologies such as these, but it is equally clear that a leap from manual planning to fully goal-based on-board autonomous planning is impossible. TVCR represents an acceptable compromise, offering the means to perform constrained replanning and plan repair on-board, but within parameters that are tightly defined by missions planners on the ground.

References

Backes, P.; Diaz-Calderon, A.; Robinson, M.; Bajracharya, M.; and Helmick, D. 2005. Automated rover positioning and instrument placement. In *Proceedings of IEEE Aerospace Conference*.

Bornstein, B.; Fukunaga, A.; Castano, A.; Biesiadecki, J.; Castano, R.; Chien, S.; Greeley, R.; Whelley, P.; Neakrase, L.; and Lemmon, M. 2007. Onboard science on the mars exploration rovers: Cloud and dust devil detection. In *38th Lunar and Planetary Science Conference, (Lunar and Planetary Science XXXVIII)*.

Castaño, R.; Estlin, T. A.; Anderson, R. C.; Gaines, D. M.; Castano, A.; Bornstein, B.; Chouinard, C.; and Judd, M. 2007. Oasis: Onboard autonomous science investigation system for opportunistic rover science. *J. Field Robotics* 24(5):379–397.

Chien, S.; Knight, R.; Stechert, A.; Sherwood, R.; and Rabideau, G. 1999. Integrated planning and execution for autonomous spacecraft. In *Proceedings of the IEEE Aerospace Conference (IAC)*.

Chien, S.; Knight, R.; Stechert, A.; Sherwood, R.; and Rabideau, G. 2000. Using iterative repair to improve responsiveness of planning and scheduling. In *International Conference on Artificial Intelligence Planning Systems (AIPS 2000)*.

Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Rabideau, G.; Castano, R.; Davies, A.; Lee, R.; Mandl, D.; Frye, S.; Trout, B.; Hengemihle, H.; D'Agostino, J.; Shulman, S.; Ungar, S.; Brakke, T.; Boyer, D.; Gaasbeck, J. V.; Greeley, R.; Doggett, T.; Baker, V.; Dohm, J.; and Ip, F. 2004. The EO-1 autonomous science agent. In *Proceedings of AAMAS-04*.

Estlin, T.; Gray, A.; Mann, T.; Rabideau, G.; no, R. C.; Chien, S.; and Mjolsness, E. 1999. An integrated system for multi-rover scientific exploration. In *Proceedings of AAAI*.

Fox, M., and Long, D. 2003. PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains. *Journal of AI Research* 20:61–124.

Fox, M., and Long, D. 2006. Modelling mixed discrete-continuous domains for planning. *Journal of AI Research* forthcoming.

Fox, M.; Long, D.; Baldwin, L.; Wilson, G.; Woods, M.; Jameux, D.; and Aylett, R. 2006. On-board timeline validation and repair: A feasibility study. In Giuliano, M., ed., *Proceedings of 5th International Workshop on Planning and Scheduling in Space*.

Fox, M.; Howey, R.; and Long, D. 2005. Validating plans in the context of processes and exogenous events. In *Proceedings of The 20th National Conference on Artificial Intelligence (AAAI-05)*.

Gerevini, A., and Long, D. 2006. Plan constraints and preferences on PDDL. In *ICAPS Workshop on Soft Constraints and Preferences in Planning*.

Howey, R.; Long, D.; and Fox, M. 2004. VAL: Automatic plan validation, continuous effects and mixed initiative planning using PDDL. In *Proceedings of 16th IEEE International Conference on Tools with Artificial Intelligence*.

Howey, R.; Long, D.; and Fox, M. 2005. Plan validation and mixed-initiative planning in space operations. In Castillo, L.; Borrajo, D.; Salido, M. A.; and Oddi, A., eds., *Planning, Scheduling and Constraint Satisfaction: From Theory to Practice*, Frontiers in Artificial Intelligence and Applications. vol. 117. IOS Press. 60–70.

Hunstberger, T.; Cheng, Y.; Stroupe, A.; and Aghazarian, H. 2005. Closed loop control for autonomous approach and placement of science instruments by planetary rovers. In *Proceedings of IEEE/RSJ Int Cong. on Intelligent Robots and Systems*.

Jönsson, A.; Morris, P.; Muscettola, N.; and Rajan, K. 2000. Planning in Interplanetary Space: Theory and Practice. In *Proceedings of International Conference on Artificial Intelligent Planning and Scheduling (AIPS)*.

Muscettola, N. 1993. HSTS: Integrating Planning and Scheduling. In Zweben, M., and Fox, M., eds., *Intelligent Scheduling*. San Mateo, CA: Morgan Kaufmann. 169–212.

Pederson, L.; Deans, M.; Lees, D.; Rajagoplan, S.; and Smith, D. 2005. Multiple target single cycle instrument placement. In *Proceedings of 8th Int. Symp. on Artificial Intelligence, Robotics and Automation in Space*.

Rabideau, G.; Knight, R.; Chien, S.; Fukunaga, A.; and Govindjee, A. 1999. Iterative repair planning for spacecraft operations in the ASPEN system. In *International Symposium on Artificial Intelligence Robotics and Automation in Space (ISAIRAS)*.

Shaw, A., and Barnes, D. 2003. Landmark recognition for localisation and navigation of aerial vehicles. In *Proceedings of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*.

Simmons, R., and Apfelbaum, D. 1998. A task description language for robot control. In *Proceedings of Conference on Intelligent Robotics and Systems*.

Smith, D. 2004. Choosing objectives in over-subscription planning. In *Proc. 14th Intl. Conf. on Automated Planning and Scheduling (ICAPS-04)*.

Squyres, S. 2005. *Roving Mars: Spirit, Opportunity, and the Exploration of the Red Planet*. Hyperion.

van Winnendael, M.; Baglioni, P.; and Vago, J. 2005. Development of the ESA ExoMars rover. In *Proceedings of 9th Int. Conf. on Artificial Intelligence, Robotics and Automation in Space*.

Woods, M.; Shaw, A.; Barnes, D.; Price, D.; Long, D.; and Pullen, D. 2009. Autonomous science for an exomars rover-like mission. *Journal of Field Robotics* 26(4):358–390.