

Spatial Coverage Planning and Optimization for Planetary Exploration

Daniel M. Gaines Tara Estlin Caroline Chouinard
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109
{*firstname.lastname*}@jpl.nasa.gov

Abstract

We are developing onboard planning and scheduling technology to enable in situ robotic explorers, such as rovers and aerobots, to more effectively assist scientists in planetary exploration. In our current work, we are focusing on situations in which the robot is exploring large geographical features such as craters, channels or regional boundaries. In to develop valid and high quality plans, the robot must take into account a range of scientific and engineering constraints and preferences. We have developed a system that incorporates multiobjective optimization and planning allowing the robot to generate high quality mission operations plans that respect resource limitations and mission constraints while attempting to maximize science and engineering objectives. An important scientific objective for the exploration of geological features is selecting observations that spatially cover an area of interest. We have developed a metric to enable an in situ explorer to reason about and track the spatial coverage quality of a plan. We describe this technique and show how it is combined in the overall multiobjective optimization and planning algorithm.

1 Introduction

Onboard autonomy is becoming increasingly important for in situ planetary exploration. The success of the Mars Exploration Rovers mission, for example, can be attributed in part to the autonomous capabilities of the roves. Onboard science planning has been identified as a critical capability for a proposed Titan aerobot mission in order to deal with communication delays and blackouts.

Onboard activity planning for in situ spacecraft, such as rovers and aerobots, has the potential for increasing science return while, at the same time, increasing mission

duration and reliability. In current mission operations, the uplink team must make predictions about spacecraft resources in order to uplink an appropriate set of science observations. Due to the difficulty in predicting resources and the high-cost of over-subscribing the spacecraft, the team often under-utilizes spacecraft resources. The challenge is even greater for more distant exploration missions such as the exploration of Titan and Europa. Due to extended communication blackouts and limited bandwidth the robotic explorers would have to be capable of operating for much longer periods without human interaction. Making accurate predations about available resources over such a long period of time will be infeasible.

We are addressing this challenge through the use of onboard activity planning. With onboard planning, the operations team can uplink a large set of observations to the spacecraft and let the onboard planning system, using up-to-date knowledge of available resources and the state of the spacecraft, select from these observations. This increases mission reliability by ensuring that the spacecraft does not oversubscribe resources even if unexpected events result in a reduction of resources. This will also result in increased science productivity as the system can take advantage of extra resource availability.

While onboard planning for spacecraft control has been demonstrated on previous missions, additional challenges remain in extending its use in selecting science observations. In order to make appropriate selections and to be accepted with the science team, the onboard planner must take into account preferences from the scientists. Because there will be a variety of scientific and engineering objectives, the planner must be capable of reasoning about interacting preferences.

In our current work, we are focusing on situations in which the robot is exploring large geological features such as craters, channels or a boundary between two different regions. In these cases, an important factor in assessing

the quality of a plan is how the set of chosen observations spatially cover the area of interest. Thus, one of the considerations a robot should make when evaluating which observations should be included in a plan is how well the candidate observations will increase the spatial coverage of the plan. The overall goal of this technology is to enable the in situ explorer to generate and execute plans that makes an appropriate balance between detailed study and broad coverage of a region.

In this paper we describe a technique that allows a robot to evaluate the spatial coverage quality of a plan. We also describe our approach for incorporating the spatial coverage preference into a planning and optimization system that combines multiple science and engineering preferences while ensuring that generated plans respect mission and resource constraints. Our approach must also be efficient as the system may have to re-evaluate the coverage quality of the plan and the potential observations during plan execution.

2 Exploring Geological Features

We are developing onboard planning and scheduling technology to enable in situ explorers, such as rovers and aerobots, to more effectively assist scientists in exploring geological features. Figure 1 shows examples of geological features on Mars illustrating the types of features the robotic explorer may be directed to explore.

A scientific campaign for exploring a geological feature would employ a variety of science instruments for collecting data about the region. For example, each Mars Exploration Rover is equipped with remote sensing instruments including high-resolution panoramic stereo cameras with a variety of filters (Pancam), navigational (Navcam) and hazard avoidance (Hazcam) stereo cameras and a Mini Thermal Emissions Spectrometer (Mini-TES). Each rover also has an arm with a suite of instruments for close contact measurements: a microscopic imager (MI), two spectrometers and a rock abrasion tool (RAT) able to remove a few millimeters of a rock's surface.

When humans perform mission planning, there may be a variety of reasons why a particular observation is selected in a given plan. The observation may be selected for its scientific merit. As discussed previously, this may be further broken down into the observation's contribution to one or more of the scientific themes (e.g. atmospheric, geology, ...). An observation may provide benefit from an engineering perspective such as collecting data for long-term route planning or assessing atmospheric dust content to facilitate better energy modeling. In addition, an observation may be selected as it increases the area that has been covered by collected data. In many cases, a single observation may contribute to several of

these criteria in various degrees.

Of course, the mission planning team must also take into account the limited set of resources that the rovers have to perform observations. The rovers are constrained by limited energy, onboard data storage, downlink opportunities and bandwidth and time to complete observations. Each observation places a different set of demands on these resources. Some are very time consuming, such as long-term spectrometer integrations, while others are memory intensive, such as Pancam acquisitions. And some activities are constrained to occur at certain periods of the day due to sun angle or ambient temperature.

Each observation also varies in the amount of spatial coverage that it affords. For example, Navcams and Hazcams have a wide field of view while Pancams and Mini-TES have a narrow field of view. A further challenge is that terrain features, such as large rocks or hills, may occlude an observation thus limiting the area that it covers. The quality of coverage afforded by an observation also degrades as a function of distance from the rover.

Finally, for a given geological feature, scientists may be more interested in certain sub-regions of that feature than in others. Thus, observations should also be evaluated based on the relative importance of the area for which they provide coverage.

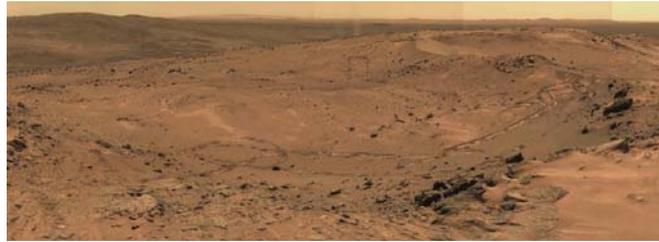
The current practice on planetary exploration missions is to have the science and engineering teams select a subset of observations that achieve an acceptable compromise among these considerations and respect the team's estimates of the available resources for accomplishing these tasks. The drawback to this approach is that the rover has extremely limited ability to alter the plan if things do not go as predicted. For example, if tasks take longer than expected, then later tasks may get dropped, even if they are deemed higher priority than earlier tasks. Or, if the rover has extra resources (e.g. more solar array energy than predicted) then it may have been able to accomplish more observations than were uplinked. Our objective is to enable rovers to reason about science quality onboard so that they can appropriately adjust the plan when state or environment information change.

3 Multiobjective Optimization and Planning Framework

Our objective is to enable an onboard planner to reason about multiple science and engineering objectives in assessing plan quality so that it can make more informed decisions about which observations to perform. This will enable the ground team to uplink a larger set of observations and let the planner dynamically select among them based on the scientific and engineering merit of the re-



(a) Channels (MGS MOC Image) (b) Craters (MGS MOC Image) (c) Boundary (MGS MOC Image)



(d) Large area characterization (Spirit Image)

Figure 1: Example geological features on Mars.

sulting plan and the planner’s assessment of available resources. During execution, the planner will modify the plan based on the current estimate of its resources.

Our approach is implemented within the CASPER system Estlin et al. [2002], Chien et al. [2000]. CASPER employs a continuous planning technique where the planner continually evaluates the current plan and modifies it when necessary based on new state and resource information. Rather than consider planning a batch process, where planning is performed once for a certain time period and set of goals, the planner has a current goal set, a current rover state, and state projections into the future for that plan. At any time an incremental update to the goals or current state may update the current plan. This update may be an unexpected event (such as a new science target) or a current reading for a particular resource level (such as battery charge). The planner is then responsible for maintaining a plan consistent with the most current information.

A plan consists of a set of grounded (i.e., time-tagged) activities that represent different rover actions and behaviors. Rover state in CASPER is modeled by a set of plan timelines, which contain information on states, such as rover position, and resources, such as energy. Timelines are calculated by reasoning about activity effects and represent the past, current and expected state of the rover over time. As time progresses, the actual state of the rover drifts from the state expected by the timelines, reflecting changes in the world. If an update results in a problem, such as an activity consuming more memory than expected and thereby over-subscribing RAM, CASPER re-plans, using iterative repair Zweben et al. [1994], to

address conflicts.

CASPER includes an optimization framework for reasoning about soft constraints such as reducing the distance traversed by the rover and increasing the value of science data collected. User-defined preferences are used to compute plan quality based on how well the plan satisfies these constraints. Optimization proceeds similar to iterative repair. For each preference, an optimization heuristic generates modifications that could potentially improve the plan score.

We explicitly represent a plan’s score relative to each science and engineering preference by defining a separate timeline for each objective. When the plan changes, associated functions are used to compute the value of each timeline. For this current work, we used three science objectives and one engineering objective:

increase_geology: collect observations that contribute toward geologic study

function: each observation specifies its geology contribution, overall plan score is the sum of each planned observation’s geology contribution.

increase_atmospheric: collect observations that contribute toward atmospheric study

function: each observation specifies its atmospheric contribution, overall plan score is the sum of each planned observation’s atmospheric contribution.

increase_spatial_coverage: collect observations that increase the area covered

function: section 4 describes how we track the spatial coverage quality of a plan

reduce_traverse_distance: minimize the overall distance traversed by the vehicle

function: our system uses a path planner to estimate the travel distance for each traverse activity in the plan. The overall traverse distance for a plan is the sum of these distances.

We use a weighted sum to combine these individual objectives into an overall plan score. This allows users to adjust the weights to direct the planner toward different objectives at different points of the mission.

CASPER provides a framework for generating, monitoring and repairing a plan and for managing user-defined preferences. However, an overall control algorithm must be created to enable integration of these capabilities in a manner appropriate for a given domain. Figure 2 shows the algorithm we developed for our work.

```
Input
  A set of prioritized science goals from Earth
  Time constraints
  Resource constraints
  Preferences

Repeat:
  Process any updates from Executive
  Save current plan
  For i = 1 to num.iterations
    If there are conflicts
      Select a conflict to work on
      Select a repair strategy for this conflict
      Apply repair strategy
    Else
      Select an optimization preference
      Select an optimization method for this preference
      Apply optimization method
      Compute plan score
      If current plan is best seen so far, save it
  Reload plan with highest score
```

Figure 2: CASPER control algorithm for rover domain.

The algorithm takes as input a set of goals with associated science priorities, a set of time and resource constraints and a set of user-defined preferences. The main loop of the algorithm interleaves iterative repair and iterative optimization to search for a conflict-free plan of high quality. The loop begins by processing any updates on state and resource timelines or on activity status. The current plan is saved. It then enters a loop in which it

attempts to improve the plan by repairing conflicts or performing optimization steps.

If the plan has a conflict, CASPER performs a repair iteration. Otherwise, CASPER attempts to improve the score by performing an iteration of optimize. If the score of the resulting plan is higher than the previously saved plan (based on a set of user-defined preferences), than the current plan is recorded.

The primary method used to enable the system to reason about spatial coverage of science observations was to develop an appropriate preference. The preference includes a means to score a plan based on the spatial coverage quality afforded by the plan. If the spatial coverage preference is selected to improve the plan score, the optimization method is to satisfy one of the requested observations that is not yet in the plan that will provide the biggest improvement to the coverage quality of the plan.

The next section provides details on how the spatial coverage quality of a plan is computed and how observations are selected to improve this score.

4 Spatial Coverage Preference

Figure 3) provides an example region of terrain that we want a rover to explore and Figure 4 shows an example set of observations that are under consideration for the plan.

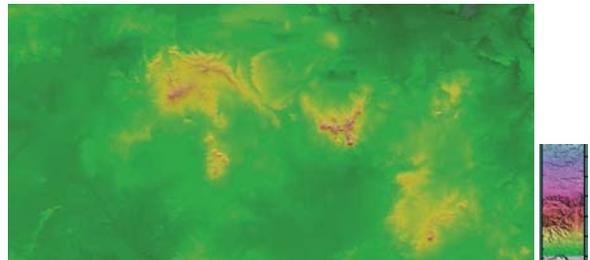


Figure 3: Digital elevation map of the an example terrain to be explored.

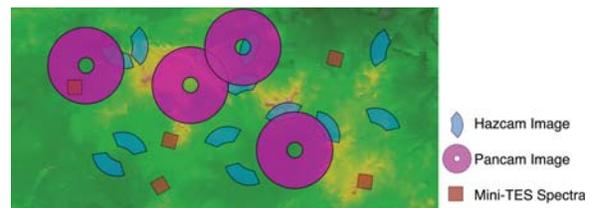


Figure 4: Set of observations.

With limited available resources, it is unlikely that the rover will be able to perform all of these observations. As discussed previously, there are many considerations for

determining which subset of observations should be included in a plan. The objective of this work is to develop a preference to encourage spatial coverage to be one of the considerations during plan generation and modification.

In this section we describe our approach to representing and reasoning about the spatial coverage quality of a plan. We begin by describing how we represent a priori information about the terrain to be explored along with scientists’ priorities indicating the relative importance of various sub-regions. We then describe how we model the coverage quality afforded by a given observation. These observation models are used to track the spatial coverage quality of plan, taking into account those observations that have already been executed and those that are scheduled to execute in the future. When resources and plan space is available, all of this information is then used to select which observations to add to the plan in an attempt to optimize the spatial coverage of the plan. Conversely, when resources are over-subscribed and observations must be shed, to select an observation that will make the smallest impact on the spatial coverage of the plan if it were to be removed.

4.1 Terrain and Terrain Priority Representation

Knowledge of terrain will enable the system to make better predictions about the coverage of observations as it will know about occlusions from terrain features such as rocks or hills.

Scientists typically have a variety of a priori information that is used to identify candidate observations that can contribute to the initial terrain map. Images from previous observations, such as NAVCAM and PANCAM observations, are the primary source of information for selecting new targets. In addition, images from orbiting spacecraft as well as images taken during the spacecraft’s descent, provide a coarse view of the geological features.

We represent a priori knowledge of the terrain to be explored as a digital elevation map where each pixel represents the height of the terrain at that point. Figure 3 shows an example terrain map.

The resolution of the map has a direct impact on the space and time complexity of the algorithm. It is not critical that we compute a highly accurate score for the amount of terrain covered by a given observation. Rather, it is important that the relative scores of different observations be correctly assessed. Thus, we convert the input terrain map into a coarser resolution such as the one in Figure 5. The resolution of the terrain map is a parameter that can be tuned to make a trade-off between accuracy of coverage quality predictions and computational complexity of the system.

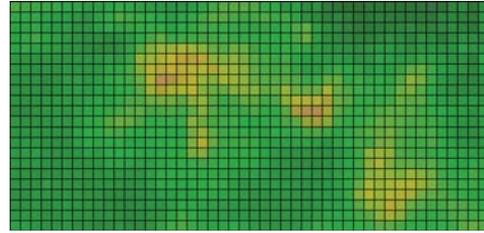


Figure 5: Lower resolution version of terrain in Figure 3.

It is also important to note that the approach does not require that a priori knowledge be complete or accurate. Missing or incorrect data in the terrain map will result in incorrect estimates of the spatial coverage that an observation will afford which, in turn, could result in lower quality plans. However, as observations are performed the terrain map will be updated and the coverage quality of upcoming observations will be re-assessed.

In addition to the terrain map, the system will take as input a matrix of weights that define the relative scientific importance of sub-regions of the terrain map. The matrix is the same size and dimensions as the input terrain matrix with each cell containing a value between 0 (least important) and 1 (most important).

4.2 Modeling Observations

In order to evaluate the coverage quality of a plan, it is necessary to compute the coverage afforded by a given observation. Figure 6 illustrates the key steps in this computation.

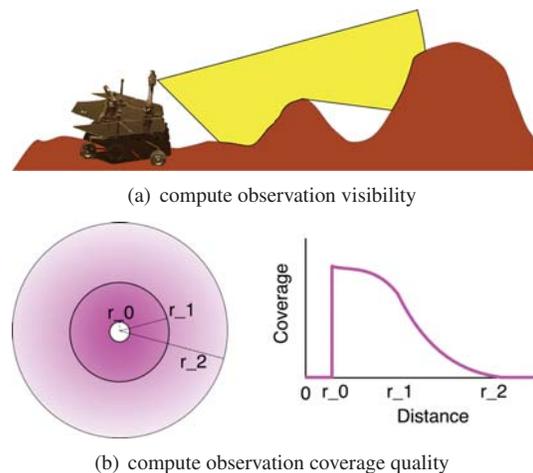


Figure 6: Modeling observation coverage quality.

The first step is to determine which cells of the terrain are visible from the location of the observation. The location of the observation, the range of the instrument and the

elevation of the terrain determine which cells. For each cell within range of the observation’s location (as determined by the range value for the instrument) we perform an intersection test between the terrain and a line from the location of the instrument to the cell using the intersection test Shapira [1990].

Once it has been determined that a cell is visible to an observation, we compute the coverage quality score which represents how well that cell is covered by the observation. A score of 0 indicates that the cell is not covered at all by this observation (e.g. it is occluded by the rover body). A score of 1 represents “perfect” coverage, in the sense that another observation of the cell would not improve upon the cell’s coverage.

This computation of the coverage quality score will vary based on the type of instrument used. In general, cells further from the origin of the observation are not covered as well as those closer in. Figure 6 (b) illustrates the coverage quality model for a panoramic image observation. Let d be the distance of the visible cell from the origin of the observation. The radius r_0 represents the diameter of the rover body. If $d \leq r_0$ then the cell is occluded by the rover body and not covered by this observation. Cells that lie between r_0 and r_1 are within the primary range of this instrument. Coverage is high for cells close to r_0 but drops off moving out toward r_1 .

The range between r_1 and r_2 is used to encourage “spreading out” of observations. The idea is that, all else being equal, one might prefer to have observations spread out across the terrain rather than clustered in a small region. The extend to which spreading out is encouraged can be tuned by increasing or decreasing r_2 .

Finally, the coverage of the cell is multiplied by the weight in the scientific priority matrix resulting in the cell’s spatial coverage quality provided by this observation.

4.3 Tracking Coverage Quality

Now that we have a way of computing the spatial coverage for a given observation, the next step is to keep track of the spatial coverage provided by a set of observations. We do this by recording the spatial coverage quality afforded by the observations into a coverage quality matrix. A coverage quality matrix is the same dimension and resolution as our terrain matrix with each cell containing a coverage quality value. If multiple observations cover the same cell in the coverage quality matrix we record the max coverage quality score afforded by these observations. Figure 7 shows an example coverage quality matrix reflecting the coverage quality for a set of observations.

We maintain two separate coverage qualities matrices. We keep track of an Executed Coverage Quality Matrix

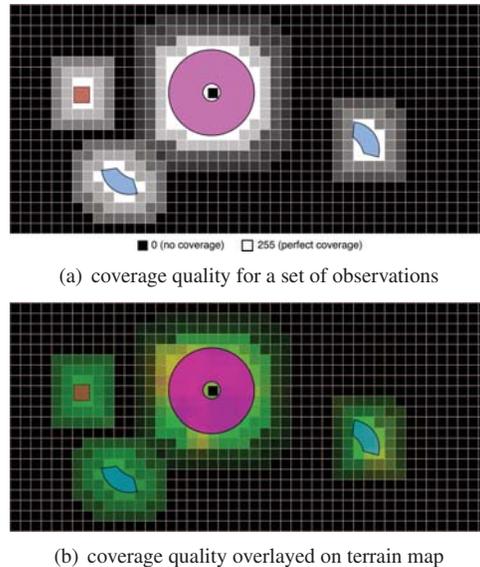


Figure 7: Example coverage quality for a set of observations.

that keeps track of the coverage quality afforded by the observations that have already executed. The second matrix is the Pending Coverage Quality Matrix and includes the coverage quality from the executed observations and the predicted coverage quality that will be obtained after the pending observations in the plan have been executed. As will be seen, maintaining these matrices will improve efficiency when the terrain map is updated and when selecting observations to add to or remove from the plan.

Each coverage quality matrix has a score which is equal to the sum of the coverage quality of each of its cells.

4.4 Ranking Observations

We rank observations with respect to how well they are expected to improve the coverage quality of the plan. We maintain two rankings, one for the requested observations, those that are not yet in the plan, and the pending observations, those that are in the plan but have not yet executed. When selecting a requested observation to add to the plan we select the highest ranked observation from the requested observations ranking. If we must shed a pending observation to resolve a conflict, we select the lowest observation from the pending observations rankings.

In actuality, rather than always selecting the highest observation to add (or lowest when deleting) we perform a probabilistic selection from the ranked list of observations with a probability of selecting a particular observation proportional to the coverage quality it is expected to contribute to the plan. This enables the system to avoid getting stuck trying to satisfy an observation for which

there are insufficient resources to perform.

Because observations may overlap, the coverage quality an observation will contribute to a plan depends in part on the other observations already in the plan. Thus, when we rank the requested observations we do so relative to the pending coverage quality matrix. Similarly, the coverage quality contribution of a pending observation depends on the observations that have already been executed. Therefore, the pending observations are ranked relative to the executed coverage quality matrix.

Figure 8 shows the algorithm used to rank a set of observation relative to a coverage quality matrix. A contribution score is computed for each observation which is equal to how much the score of the coverage quality matrix would increase if this observation were added.

```
Input
  Unranked Observations
  Initial Coverage Quality Matrix

For each observation in Unranked Observations
  Observation's contribution score = how
  much the score of the coverage quality
  matrix would improve if this observa-
  tion were included
Sort observations based on contribution
score
```

Figure 8: Ranking a set of observations relative to a coverage quality matrix.

Note that the algorithm in Figure 8 ranks only the single next observation to add (or remove) and does not indicate which order the remaining observations should be added (or removed). Instead, we use an iterative approach since we will add or remove activities one at a time.

This iterative approach represents a greedy algorithm for selecting observations to add and remove and does not guarantee an optimal solution except in the case where the observations do not overlap. We have chosen not to attempt an optimal solution for three main reasons. First, observations are selected for a variety of reasons, not just spatial coverage. Thus, we cannot count on the spatial coverage ranking being honored when observations are added and removed. Second, because we will have to re-rank observations during execution (when observations are added and removed or when the terrain map is updated) we want a fast computation. Finally, we expect that observations will not overlap significantly and thus the greedy approach will not be far from optimal.

4.5 Updating Spatial Coverage During Execution

During the course of executing the plan, the system will need to update its rankings. Through the collection of observations, we will be collecting new information about the terrain being explored. We can update the terrain map when this happens. Doing so will improve the accuracy of the coverage quality predictions. However, when the terrain is updated we will need to re-compute the coverage quality afforded by each of the observations and re-compute our rankings.

We must also re-compute rankings when observations are added to or removed from the plan since the contribution score of an observation depends on the order in which it is added to the plan.

5 Related Work

The spatial coverage problem we are solving is similar to a classic computational geometry problem called the Art Gallery Problem [O'Rourke, 1987]. Given a polygon, representing the floor plan of an art gallery, the problem is to select the minimum number of locations to place guards such that every point in the polygon is in view of at least one guard. It is assumed that guards can see in all directions and can see out to infinity unless an edge of the polygon obstructs its view. This has been shown to be NP-Hard, but a popular approximation runs in time $\mathcal{O}(n \log n)$. The approximation triangulates the polygon and then performs three-coloring on the resulting vertices. Guards are posted at the vertices that were colored by the minimum color class (the color class that was used the least amount of times to color vertices).

While similar, there are significant differences between the Art Gallery Problem and the problem we wish to solve. Most significant is that the Art Gallery Problem is restricted to 2D while we are modeling and selecting observations in 3D. The triangulation approximation described above does not scale to the 3D case. Furthermore, it is unrealistic to model observations the way guards are modeled. Some observations cannot see in all directions and the quality of the observation is not constant with respect to the distance of a point from the origin of the observation. Finally, the Art Gallery Problem does not consider the costs of posting guards.

The ROPE (Rank and Overlap Elimination) system selects locations for video cameras for visual surveillance of large 3D open spaces [Rana, 2005]. ROPE discretizes the area into cells and then uses a greedy algorithm to select camera locations by first placing a camera in the cell from which the largest number of other cells are visible. The selected cell and the visible cells are removed from

the area and the process is repeated until no cells remain. This greedy approach is similar to the approach we take in ranking observations. However, like the Art Gallery Problem, ROPE does not model the quality of coverage and it does not consider cost.

The swath coverage problem for orbital satellites is similar to the spatial coverage problem addressed in this paper. A satellite collects images as it orbits and may acquire images of the same location on subsequent orbits. The problem is to select which images to downlink to maximize science value while respecting onboard storage and downlink capacities. Knight presents an optimal solution using depth first branch and bound and a network flow formulation for a heuristic reward estimator [Knight and Smith, 2005]. The ASTER system uses a greedy approach for selecting images similar to ROPE [Muraoka et al., 1998]. Both approaches take into account resource cost of observations and can take into account quality of coverage. The heuristic used in Knight assumes that the reward for an observation scales with its cost. This is not necessarily the case for rovers where the cost of an observation may vary independent of the reward due to many factors such as distance to the observation and terrain conditions. Our greedy algorithm is essentially the same as that used in ASTER. The main difference in our work is the modeling of instrument observations and the incorporation of this work into surface operations and other optimization preferences.

Leonard et al. [2007] present an optimization method for coordinating the search patterns of multiple vehicles for optimal data collection. Their system uses a global optimization function that combines all constraints and preferences. In contrast, our approach separates these factors. The trade-off is that a combined approach may result in more optimized results while the factored approach is typically easier to incorporate new or changing capabilities, constraints and preferences

6 Conclusions

We have presented a set of algorithms that enable a rover to compute the spatial coverage quality of a plan and to rank candidate observations by how well they are expected to improve coverage quality. Using this technique, a rover is better able to assist in the exploration of geological features by generating high quality operations sequences that take into account spatial coverage along with other science considerations. We have currently implemented and tested these algorithms and have tested them in a stand-alone mode. We are in the process of integrating the preferences into our rover planning and execution system. In future work, we will focus on techniques for combining multiple preferences functions so that the sys-

tem can more effectively trade-off science and engineering objectives when generating and executing plans.

Acknowledgments

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This work was funded by the Jet Propulsion Laboratory Research and Technology Development program.

References

- Steve Chien, Russell Knight, Andre Stechert, Rob Sherwood, and Gregg Rabideau. Using iterative repair to improve the responsiveness of planning and scheduling. In *Fifth International Conference on Artificial Intelligence Planning and Scheduling*, Breckenridge, CO, April 2000.
- Tara Estlin, Forest Fisher, Daniel Gaines, Caroline Chouinard, Steve Schaffer, and Issa Nesnas. Continuous planning and execution for a Mars rover. In *Third International NASA Workshop on Planning and Scheduling for Space*, Houston, TX, October 2002.
- R. Knight and B. Smith. Optimally solving nadir observation scheduling problems. In *International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2005)*, Germany, September 2005.
- Naomi Ehrich Leonard, Derek Paley, Francois Lekien, Rodolphe Sepulchre, David M. Fratantoni, and Russ Davis. Collective motion, sensor networks and ocean sampling. *Proceedings of the IEEE*, 95(1):48–74, 2007.
- H. Muraoka, R. H. Cohen, T. Ohno, and N. Doi. Aster observation scheduling algorithm. In *SpaceOps 98*, Tokyo, Japan, June 1998.
- J. O’Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, New York, 1987.
- S. Rana. Use of gis for planning visual surveillance installations. In *ESRI Homeland Security GIS Summit*, Denver, 2005.
- Andrew Shapira. *Fast line-edge intersections on a uniform grid*, chapter 6, pages 29–37. Morgan Kaufmann, 1990.
- M. Zweben, B. Daun, E. Davis, and M. Deale. Scheduling and rescheduling with iterative repair. In M. Fox and M. Zweben, editors, *Intelligent Scheduling*, pages 241–256. Morgan Kaufmann Publishers Inc., 1994.