

Commentary on Spatial Coverage Planning and Optimization for a Planetary Exploration Rover

Mark Giuliano

Space Telescope Science Institute
3700 San Martin Drive, Baltimore, MD 21218
giuliano@stsci.edu

Introduction

The paper explores spatial coverage issues for on-board planning and execution of autonomous planetary rovers. Techniques are presented for determining a planning preference that optimizes the spatial coverage of a set of experiments on a pre-imaged region of a planet. The overall goal is to improve the science return of the rover by enabling an on-board autonomous planner to adjust the uploaded plan in response to changes in the modeled resource profile. The uploaded plan would be slightly oversubscribed. If the plan requires less resources during execution than planned then, the on-board scheduler would use the spatial preference to select additional observations. If the plan requires more resources during execution than planned then, the on-board scheduler would use the spatial preference to select observations to remove from the plan.

I offer commentary on the paper in the following areas:

- Data structures for reasoning about spatial coverage.
- Closing the loop with the human scheduling process.
- Integrating spatial coverage preferences with other preferences during planning.

Data Structures

The application uses a terrain map to store elevation information for the spatial region under consideration. A terrain map is a matrix where each cell represents a region of the surface of the planet and the value of a cell gives the associated elevation. The authors state that “the resolution of the terrain map can be tuned to make a trade-off between accuracy of coverage quality predictions and computational complexity of the system.” The authors should consider the use of a hierarchical data structure (Samet & Hanan 1989) such as quad tree to represent the terrain map. A hierarchical data structure allows large

regions that have the same or similar elevations to be stored as a single value. Regions that have large fluctuations in elevations would have a deeper data structure allowing for more detailed elevation data. This approach may lessen the need to make trade-offs between accuracy and efficiency.

The approach given in the paper uses a matrix to track the quality of science coverage for spatial regions. As in the terrain map, each matrix cell represents a region of the surface area. Cell values in the quality matrix give the numeric quality of science observations that overlap the cell. When multiple observations cover the same region the value is the maximum of the values for the overlapping observations. Suppose a pending observation is removed from the plan due to a resource conflict. How can we quickly update the quality matrix for cells overlapping the observation? If the observation set the value for a cell then, computing the new value for the cell requires re-computing the overlapping observations and their quality values. One approach to avoiding this cost would be to add a hash table that takes the matrix position as a key and stores all the observations that overlap the matrix position. On removing an observation the hash table can be consulted to update the quality coverage matrix. If the surface area of the region is sparsely covered by observations, the space requirement for the hash table would not be burdensome.

Closing the Loop with Human Planners

The authors should address how human planners can utilize the spatial preference.

One question is how a spatial preference can be utilized in a mixed-initiative system where humans and computers interact to create schedules. While the terrain map and quality coverage matrix might be good for display, how would a human user understand the spatial tradeoffs between competing schedules? How would the system allow an expert user to modify specific spatial preferences while retaining the rest of the spatial model?

A related topic is using feedback from the autonomous scheduler to improve the ground based scheduling process.

The on-board scheduler could download status records when uploaded plans are modified. The downloaded records could be analyzed to help improve the ground based scheduling process. This feedback might be useful in several ways. First, the feedback can help determine the optimal level of over subscription for uploaded plans. The goal is to provide a level of over subscription that minimizes the re-planning of un-executed activities while maximizing rover utilization. Second, feedback can drive the types of science and spatial coverage to include in uploaded plans. If the planner typically utilizes a particular type of science or spatial region when re-planning then, ground planners can ensure that this type of observation is included in the over-subscription pool. Finally, this feedback could be used to reorder science in uploaded plans. Suppose that some types of otherwise high value observations are often removed during execution due to resource contention and poor spatial placement. In this case, uploaded plans might want to include this type of activity early in the schedule to ensure that they get executed before any resource problems occur.

Integration with other Preferences

As indicated by the authors a spatial coverage preference is one factor out of many to be used during planning. The last sentence of the paper states that the authors plan to investigate techniques for combining multiple preferences. I agree with the authors that this is a good next step for their work and offer the following questions for the authors to keep in mind:

- Can the spatial preference of an observation be determined independently of other factors? For example, within the specification of an observation there may be a trade-off between greater spatial coverage and an engineering resource.
- Should individual preferences be combined into a single numeric evaluation or should the evaluations be separately stored so that the system can determine the pareto optimal frontier (Johnston 2006)?
- Given that the system combines individual preferences, how should they be combined? Is the combination rule additive, gated, or multiplicative?

References

- Johnston, M., 2006. Multi-Objective Scheduling for NASA's Future Deep Space Network Array. In Fifth International Workshop on Planning and Scheduling for Space.
- Samet, Hanan, 1989. *Applications of Spatial Data Structures*. Addison-Wesley, Reading, MA.