28

Optimal Path Planning of Multiple Mobile Robots for Sample Collection on a Planetary Surface[*]

J.C. Cardema and P.K.C. Wang
University of California
Los Angeles, California
U.S.A.

1. Introduction

In the exploration of a planetary surface such as that of Mars using mobile robots, rock and soil-sample collection and analysis are essential in determining the terrain composition and in searching for traces of ancient life (Malin & Edgett, 2000). Several missions to Mars have already been sent. In the 1997 Mars Pathfinder mission (Mars Pathfinder Homepage), the Sojourner rover used an alpha-proton-X-ray spectrometer (APXS) to analyze rock and soil sample compositions. It also had a simple onboard control system for hazard avoidance, although the rover was operated remotely from Earth most of the time. The method for rock and soil-sample collection is as follows. After landing, the rover used its black-and-white and color imaging systems to survey the surrounding terrain. The images were sent back to Earth, and analyzed by a team of geologists to determine where interesting samples might be found. Based on that information, the next destination for the rover was selected and the commands to get there were sent to the rover via radio with transmission delays ranging from 10 to 15 minutes (depending on the relative orbital positions of Earth and Mars). The set of commands were sent out over a day with the rover moving only a small distance each time. This was done to allow the mission control to constantly verify the position, with time to react to unforeseen problems. When the rover finally reached its destination and analyzed the sample, it spent another day transmitting the information back to Earth. The cycle was repeated as soon as the geologists had decided on the next destination for the rover. Clearly, an automated system for rock and soil sample collection would expedite the process. In the 2004 Mars Exploration Rover (MER) mission, the Spirit and Opportunity rovers (Mars Spirit & Opportunity Rovers Homepage) featured an upgraded navigation system. Imagery from a stereo camera pair was used to create a 3-D model of the surrounding terrain, from which a traversability map could be generated. This feature gave the mission controllers the option of either directly commanding the rovers or

allowing them to autonomously navigate over short distances. Consequently, the rovers were often able to traverse over 100 meters a day (Biesiadecki & Maimone, 2006). The rovers were also programmed to autonomously select interesting soil samples, but this feature was seldom used. Nevertheless, this was a significant first step toward fully automating the soil-sample collection process.

In this study, an attempt is made to formulate the path planning problem for single and multiple mobile robots (referred to hereafter as "rovers" for brevity) for sample collection as a mathematical optimization problem. The objective is to maximize the value of the mission, which is expressed in the form of a mission return function. This function contains the performance metric for evaluating the effectiveness of different mission setups. To the best of our knowledge, the problem of sample collection has not yet been studied in conjunction with optimal path planning. There are many considerations in the mathematical formulation of this problem. These include planetary terrain surface modeling, rover properties, number of rovers, initial starting positions, and the selection of a meaningful performance metric for rovers so that the performance of single versus multiple rovers in representative scenarios can be compared. The basic problem is to find a sample-collection path based on this performance metric. The main objective is to develop useful algorithms for path planning of single or multiple planetary rovers for sample collection. Another objective is to determine quantitatively whether multiple rovers cooperating in sample collection can produce better performance than rovers operating independently. In particular, the dependence of the overall performance on the number of rovers is studied. To clarify the basic ideas, we make use of the Mars rover rock and soil-sample collection scenario in the problem formulation and in the numerical study.

2. Problem Description

We begin with a discussion of the problem of planetary surface modeling, followed by various operational considerations of the rovers. To facilitate the mathematical formulation of the optimal path-planning problems, a few basic definitions will be introduced. Then, precise mathematical statements of the optimal path-planning problems will be presented for both single and multiple rover cases.

2.1 Planetary Surface Modeling

Assuming that a region on the planetary surface has been selected for detailed scientific study, the main task is to develop a suitable terrain surface model for rover path-planning. Initially, a crude surface model for the selected spatial region may be constructed from the aerial planetary survey data obtained by fly-by spacecraft or observation satellites such as the Mars Orbiter. Once the rovers are on the planetary surface, more refined models (usually localized model) may be constructed from the image-data generated by on-board cameras. Although the refined models may be useful for scientific studies, they may not be useful for practical optimal path planning. Therefore we resort to approximate models that simplify the mathematical formulation and numerical solution of the optimal path-planning problems. In our model, we assume that the area of the spatial domain for exploration is sufficiently small so that the curvature of the planetary surface can be neglected. Moreover, the surface is sufficiently smooth for rover maneuvers.

2.1.1 Approximate Surface Model

Let Ω be a bounded spatial domain of the two-dimensional real Euclidean space R^2 and the representation of a point in R^2 with respect a given orthonormal basis be denoted by x. Let $f = f(x)$ be a real-valued continuous function defined on Ω. Let $G_f \stackrel{\text{def}}{=} \{(x, f(x)) \in R^3 : x \in \Omega\}$ denote the *graph of* f, which represents the planetary surface under consideration. In this work, we use a polygonal approximation for the planetary surface G_f via triangulation that partitions G_f into adjacent, non-overlapping triangular patches, where each edge of a triangular patch is shared by exactly two triangular patches except on the boundaries of G_f. It has been proved that every C_1-surface defined on Ω with a sufficiently smooth boundary has a triangulation, although an infinite number of triangular patches may be required (Weisstein). Here we make use of the Delaunay triangulation, which produces a set of lines connecting each point in a given finite point set to its neighbors. Furthermore, it has the property that the triangles created by these lines have empty circumcircles (i.e. the circumcircles corresponding to each triangle contains no other data points). The Delaunay triangulation of G_f is a polygonal approximation of the original planetary surface. It can also be thought of as the projection of the planetary surface onto a mesh space. The domain of the triangulation is a mesh space denoted by $\hat{\Omega} \subset \Omega$, where $\hat{\Omega}$ is the discrete version of Ω. The resulting polygonal approximation of the planetary surface G_f will be denoted by $\hat{\mathsf{G}}_f$. This approximate surface model will be used in formulating the optimal path-planning problem. Although one might use other forms of approximation for G_f that lead to smoother approximate surfaces, our choice provides significant simplification of the optimal path-planning problem, since the paths are restricted to lie on the edges of the triangular patches.

2.1.2. Rock and Soil-sample Properties

Rock and soil samples have different values to geologists based on the questions they are trying to answer. For example, in Mars exploration, sedimentary rocks are important since two of the primary questions about early Martian geological history are whether liquid water could exist on its surface and, if so, whether liquid water ever took the form of lakes or seas (Malin & Edgett, 2000). According to Malin and Edgett, outcrop materials are interpreted as Martian sedimentary rock, and they are of particular interest to geologists for answering these questions. The outcrop materials occur in three types: layered, massive, and thin mesas, which differ in visual tone, thickness, texture, and configuration. The locations of these outcrops are limited to specific regions mostly between ± 30 degrees latitude. One of the regions with large outcrop occurrence is in Valles Marineris. The terrain in a portion of this region is used for our case study. The three types of outcrops are speculated to come from different Martian ages. The rover should have the capability of identifying and distinguishing these different types. To model a portion of Valles Marineris, we assume that the rock and soil samples are randomly distributed over various sub-regions. An appropriate model should allow for the specification of any rock and soil-sample distribution on the given Mars terrain. A simple way to do this is to divide the terrain into sub-regions and assign weights to determine how many samples to uniformly

distribute within each sub-region. Moreover, we assume there are a finite number of samples, each with an assigned value in a prescribed range. A high sample value implies high scientific value. In practical situations, this task may be accomplished by a careful study of the aerial survey data.

2.2. Single Rover Case

2.2.1 Operational Considerations

In what follows, we consider the main factors that are relevant to the formulation of the path-planning problem.

2.2.1.1 Dynamic properties

The rover is modeled simply. We only consider mass m, maximum traversable slope or tilt angle θ, maximum speed v_{max}, and maximum power P_{max}. They are used to calculate the rover's traveling time on terrains with varying slopes. Higher-order dynamics involving the acceleration of the rover and the damping effects of the suspension system are not included, since the actual motion of the rover is relatively slow compared to that of mobile robots in a laboratory environment. In what follows, the term "sample collection" is used interchangeably with "sample analysis", although they do not necessarily have the same connotation. (i.e. sample collection can be thought of as the collection of sample data.) In the case where the rover physically picks up the sample, the mass of each collected sample is added to the overall mass of the rover. There is also a loading constraint that limits the number of samples that the rover can physically hold in its storage compartment. In this study, we do not consider the situation where the rover can only use its imaging system to identify and detect samples. This leads to the problem of determining a path that maximizes the visual coverage of the terrain (Wang, 2003, 2004).

2.2.1.2 Mission time limit

The mission length is an important consideration in path planning. For the 1997 Pathfinder mission, the planned mission duration was 30 days. The algorithm for path planning should verify that the time duration for sample collection is within the prescribed mission time limit (denoted by τ_{max}), which in turn determines the maximum terrain coverage.

There should be a clarification about the distinction between the overall mission time and the time it takes to execute a planned path. In practical situations, it would be difficult to plan a path for the entire mission duration. It would be more reasonable to plan paths of shorter duration that can be executed at specific intervals during the mission. However, to simplify our formulation, we do not make this distinction and assume that we can plan a path for the entire mission duration.

2.2.1.3 Sample analysis time

As mentioned earlier, the Sojourner rover in the Pathfinder mission was equipped with a spectrometer (APXS) for rock and soil-sample analysis. The sensor head of the APXS was placed on the sample for 10 hours during the analysis. To account for this, the sample analysis time τ_{wait} is introduced into our model. It represents the amount of time required to analyze the sample. To simplify the model, we assume that τ_{max} is the same for every sample, regardless of its type. With the inclusion of the sample analysis time, the rover is

forced to consider more carefully which rock and soil samples to collect while still operating within the time limit.

2.2.1.3 Mission return function

To characterize the performance of a rover in rock and soil sample collection, we need to choose an appropriate mission return function to quantify the rover's performance throughout the mission. The mission return function used in this study is simply the sum of the collected sample values.

2.2.1.4 Terrain risk

Let $\theta_{\max}$ denote the angle of the maximum traversable slope corresponding to the maximum allowable tilt angle of the rover before it topples over. To determine if a point on the surface is too risky to traverse, the terrain slopes at that point in all directions are computed. If the magnitude of the slope angle in any direction exceeds $\theta_{\max}$, that point is deemed un-traversable. Usually, the rover is more susceptible to tipping over sideways than forwards or backwards, although the dimensions of the rover are not considered in this study.

2.2.1.5 Terrain properties

Ideally, the terrain surface in the spatial region chosen for exploration should be sufficiently smooth to facilitate rover maneuverability but also has features to indicate the possible presence of interesting samples. For rover traversability, the terrain texture and hardness are also important (Seraji, 2000). Terrains that are rough and rocky are avoided in favor of smoother ones. Terrain risk depends on both the terrain texture and hardness.

2.2.2. Definitions

Having specified an approximate surface in the form described in Sec. 2.1.1, a path can be specified on this surface. First, we introduce the notion of an admissible segment.

2.2.2.1 Definition 1

(*Admissible segment*): A segment γ connecting a point $(x_a, f(x_a)) \in \hat{G}_f$ with an adjacent point $(x_b, f(x_b)) \in \hat{G}_f$ along an edge of a triangular patch formed by the Delaunay triangulation of G_f is said to be admissible if it satisfies the following constraints induced by the terrain risk:

$$\theta_a = \arcsin \left| \frac{f(x_a) - f(x_{ai})}{d} \right| \leq \theta_{\max} \quad \text{for all } x_{ai},$$

and

$$\theta_b = \arcsin \left| \frac{f(x_b) - f(x_{bi})}{d} \right| \leq \theta_{\max} \quad \text{for all } x_{bi},$$

where x_{ai} and x_{bi} are the adjacent points of x_a and x_b, respectively, and $d = \left\| (x_i, f(x_i)) - (x_{i+1}, f(x_{i+1})) \right\|$ is the Euclidean distance between points $(x_i, f(x_i))$ and $(x_{i+1}, f(x_{i+1}))$. (i.e. The slopes of $\hat{G}_f$ from x_a and x_b to all their neighboring points satisfy the maximum traversable slope angle constraint $\theta_{\max}$).

2.2.2.2 Definition 2

(*Admissible path*): A path Γ composed of connected segments in $\hat{\mathsf{G}}_f$ is said to be admissible if each segment is admissible.

Each admissible path can be represented by an ordered string of points that will be denoted by $\mathsf{S} \subset \hat{\Omega}$. The string of points S may include repeated points since partial backtracking along the path is allowed. To account for different sample values in the model, the samples are individually indexed. Each sample σ_k has a corresponding value λ_k, where k is the index number. We assume that the sample values as well as the sample distribution $D_{ss}(x)$ on the terrain defined below are known *a priori*.

2.2.2.3 Definition 3

(*Sample distribution*): The sample distribution $D_{ss} = D_{ss}(x)$ is a set-valued function defined as follows: If at a point $x \in \hat{\Omega}$, there are m samples with indices $\mathsf{J}_x = \{k_{x,1}, k_{x,2}, ..., k_{x,m}\}$, then $D_{ss}(x) = \mathsf{J}_x$. If there are no samples at $x \in \hat{\Omega}$, then $D_{ss}(x)$ is an empty set.

The entire set of all samples indices is denoted by E_{ss}. Along each admissible path Γ with the corresponding string S, there is a set of *collectable samples* C_{ss} that includes all the samples contained in S as defined below:

2.2.2.4 Definition 4

(*Collectible sample set*): If S_Γ is the string of points associated with an admissible path Γ, and $\mathsf{J}_\Gamma = \{k_1, k_2, ..., k_N\}$ is the set of sample indices such that

$$\mathsf{J}_\Gamma = \bigcup_{x \in \mathsf{S}_\Gamma} D_{ss}(x),$$

where N is the number of soil samples along the path, then the collectable sample set is $\mathsf{C}_{ss} = \mathsf{J}_\Gamma$.

Next, we define the attainable set associated with an initial point x_0 at time t_0.

2.2.2.5 Definition 5

(*Attainable set*): The attainable set at time t starting from $x_0 \in \hat{\Omega}$ at time t_0 (denoted by $\mathsf{A}(t; x_0, t_0) \subset \hat{\Omega}$) is the set of all points in $\hat{\Omega}$ that can be reached at time t via admissible paths initiating from x_0 at time t_0.

Since $\hat{\mathsf{G}}_f$ is time-invariant, we can set $t_0 = 0$. Successive points along the admissible path Γ will be restricted to this attainable set. An attainable set associated with a maximum mission time duration τ_{max} will be denoted by $\mathsf{A}_{max}(x_0) = \mathsf{A}(\tau_{max}; x_0, 0)$, which will be referred to as the *maximal attainable set from x_0*.

Evidently, we can find the admissible paths and their collectible sample sets associated with a given $\mathsf{A}(t; x_0, 0)$. Once we determine the samples to collect along the path, we can introduce the notion of an *admissible tour*.

2.2.2.6 Definition 6

(*Admissible tour*): An admissible tour T_a is a pair $(\Gamma_a, \mathsf{L}_{ssa})$, where Γ_a is an admissible path with point set S_a and collectible sample set C_{ssa}, and L_{ssa} is the list of soil samples collected along the path Γ_a.

In an admissible tour, the time for traversing the path (including the sample analysis time) is within the mission time limit. If n is the number of samples collected, then the mission time τ_m is given by

$$\tau_m = \sum_i \tau_i + n \cdot \tau_{wait} \leq \tau_{max}, \tag{1}$$

where τ_{wait} is the time required to analyze each sample; τ_{max} is the mission time limit; τ_i is the traveling time between the successive points x_i and x_{i+1} in S_a and is given by

$$\tau_i = \max \left\{ \frac{m g_m h}{P_{max}}, \frac{d}{v_{max}} \right\}, \tag{2}$$

where g_m is the acceleration due to gravity of the planet, $h = f(x_i) - f(x_{i+1})$ is the difference in terrain elevation at the points x_i and x_{i+1}, and $d = \|(x_i, f(x_i)) - (x_{i+1}, f(x_{i+1}))\|$. Thus, τ_i corresponds to the maximum of the traveling times under the power and speed constraints. Physically, when climbing up a slope, the power constraint is used to compute τ_i. When the terrain is flat or sloping downward, the speed constraint is used instead.

For an admissible tour T_a, the sample list L_{ss} contains the index of each sample in the order of collection. Each sample σ_k in the sample list has value λ_k. Let $\mathsf{I} = \{k_1, k_2, ..., k_j, ..., k_n\}$ be the set of sample indices in the order of collection, where $\mathsf{I} \subset \mathsf{C}_{ssa}$ and k_j is the index of the j th sample collected along the path. Then, $\mathsf{L}_{ss} = \mathsf{I}$.

2.2.3. Problem Formulation

Assume that an approximate planetary surface $\hat{\mathsf{G}}_f$, initial starting point x_0, mission time limit τ_{max}, sample analysis time τ_{wait}, sample index set E_{ss}, samples σ_k's with values λ_k, $k \in \mathsf{E}_{ss}$, and sample distribution $D_{ss} = D_{ss}(x)$ are given.

2.2.3.1 Problem P1

Find an optimal tour T^o (with admissible path Γ_a^o and the corresponding sample list L_{ssa}^o) that maximizes the mission return function

$$V(\mathsf{L}_{ssa}) = \sum_{k \in \mathsf{L}_{ssa}} \lambda_k, \tag{3}$$

where λ_k is the value of sample σ_k, i.e.

$$V(\mathsf{L}_{ssa}^o) \geq V(\mathsf{L}_{ssa}) \tag{4}$$

for all L_{ssa} associated with admissible paths Γ_a.

To find an optimal tour, we must first consider all possible admissible paths from the given initial starting point x_0. The total traveling time τ_{at} along each admissible path Γ_a (with path point set S_a and collectible sample set C_{ssa}) satisfies the constraint:

$$\tau_{at} = \sum_i \tau_i \le \tau_{\max},$$ (5)

where τ_i is the traveling time between successive points x_i and x_{i+1} in S_a.

Along each admissible path, we have a set C_{ssa} of collectable samples. Let $\mathsf{I} \subset \mathsf{C}_{ssa}$ be a set of sample indices in the order of collection such that

$$\tau_{am} = \sum_i \tau_i + n \cdot \tau_{wait} \le \tau_{\max},$$ (6)

where τ_{am} is the total mission time, n is the number of samples in I, and each τ_i is the traveling time between successive points x_i and x_{i+1} in S_a along the path. All possible sets of sample indices I are considered.

For each admissible path Γ_a (with path point set S_a and collectible sample set C_{ssa}), we search through all possible sets of sample indices $\mathsf{I} \subset \mathsf{C}_{ssa}$, and find a $\mathsf{I}^* \subset \mathsf{C}_{ssa}$ that maximizes the mission return function

$$V(\mathsf{I}) = \sum_{k \in \mathsf{I}} \lambda_k,$$ (7)

i.e.

$$V(\mathsf{I}^*) = \max\{V(\mathsf{I}) : \mathsf{I} \subset C_{ssa}\} \text{ and } \mathsf{L}_{ssa} = \mathsf{I}^*.$$ (8)

Let Γ_A denote the set of all admissible paths, and Λ_A the corresponding set of sample lists. Once we have performed the maximization for each admissible path, the optimal sample list L_{ss}^o is found by taking the path and sample list that maximizes the mission return function, i.e.

$$V(\mathsf{L}_{ss}^o) = \max\{V(\mathsf{L}_{ss}) : \mathsf{L}_{ss} \in \mathsf{L}_A\}.$$ (9)

The optimal path Γ^o is the path associated with L_{ss}^o, and the optimal tour is $\mathsf{T}^o = (\Gamma^o, \mathsf{L}_{ss}^o)$.

The optimal sample collection path generally depends on the initial starting position of the rover. Intuitively, we would like to place the rover as close as possible to the highest-valued samples. Since the distributions of the samples are known, we could also maximize the mission return function with respect to the starting position. But in practical situations, the starting position cannot be specified precisely. Moreover, the sample distributions are not known beforehand.

2.3. Multiple Rover Case
In the multiple rover case, it is necessary to introduce a few additional notions.

2.3.1. Operational Considerations

2.3.1.1 Starting positions
First, let us consider the attainable set of each rover corresponding to a given initial starting position at time $t = 0$. If the maximal attainable sets of two or more rovers overlap, then a decision has to be made on the assignment of rovers to cover the overlapping region. This decision can be made either by the mission planner (centralized operation) or by the rovers themselves (autonomous operation).

As in the single rover case, the choice of starting positions is an important issue. If the rovers are placed too close together, they could interfere with each other's collection tasks. If the

rovers are placed too far apart, then there is little cooperation among them since the overlapping region is small. Therefore, it is desirable to place the rovers such that a balance between cooperation and interference can be reached. The problem of finding the optimum starting positions for m rovers, with or without *a priori* knowledge of the sample distributions, is an interesting one. This problem will not be considered here. In what follows, we assume that the starting positions of the rovers are pre-assigned.

2.3.1.2 Interaction

The interaction between the rovers can lead to either cooperation or interference. In order to promote cooperation, the rovers can actively communicate with each other and decide among themselves on how to split up the terrain to make the collection process most efficient. Alternatively, a central supervisor or mission planner can make all these decisions beforehand and predetermine the rover's sample collection path. We expect that the performance of multiple rovers with interaction and cooperation is better than that of multiple rovers operating independently.

2.3.1.3 Centralized vs. Autonomous Operation

The distinction between the centralized and autonomous operations depends on information utilization. In the centralized operation, all the information about the terrain including sample locations is known beforehand. This information is analyzed and the optimum paths for the rovers are predetermined. This is a simplified version of the real-world scenario. In practical situations, the terrain details as well as the rock and soil-sample locations are not completely known, hence autonomous operation would be more desirable. Here, each rover must rely on its vision system to examine the surrounding terrain and to detect samples. Since the range of view of the rover is limited, cooperation between the rovers is more desirable. Using the data from its vision system, each rover would then plan its own optimum path while keeping in communication with the other rovers for promoting cooperation and avoiding interference. The autonomous operation could account for inaccuracies or uncertainties in the rover's terrain information. In what follows, we consider only the centralized operation.

2.3.2. Centralized Operation

The multiple rover case is similar to the single rover case when the rovers are placed far apart. Since there is little or no interaction, the problem can be reduced to one involving separate single rover cases. When the rovers are placed close enough such that interaction occurs, a new approach to the problem must be developed. We shall consider several different sample collection strategies for m rovers with given initial starting positions and sample distributions.

2.3.2.1 Best Route First

Here, we solve the single rover case for each rover and compute the optimal tours. We search for the rover whose tour has the highest mission return function value and keep only that tour. The tour is assigned to that rover and its collected samples are removed from the terrain. The process is repeated for the remaining rovers until a tour has been assigned to each rover. The drawbacks with this strategy include long computation time, especially for a large number of rovers and samples. Moreover, the tour value of each successive rover is less than that of the preceding rover in the iteration process.

2.3.2.2 Partition of Overlapping Attainable Set: Closest Rover

The maximal attainable sets of the m rovers are examined and the terrain is divided into overlapping and non-overlapping regions. The overlapping region includes all points on the terrain that can be reached by more than one rover within the given mission time limit. Since each sample in the overlapping region can be assigned to only one rover, the set of samples in the overlapping region must be partitioned according to some given criteria. Here, the samples in the overlapping region are each assigned to one of the rovers based on the distance from the sample's position to each rover's starting position. The rover whose starting position is closest to the sample's position is assigned to collect that sample. Once the partitioning of the samples in the overlapping region is completed, the problem reduces to m single rover problems with each rover limited to a subset of the collectable samples. However, a rover may be assigned more samples than it can collect within the prescribed mission time limit. In that case, after determining the rover's tour, the uncollected samples are passed on to the other rovers for consideration.

2.3.2.3 Partition of Overlapping Attainable Set: Closest Path

This strategy is similar to the previous one except that the criterion for partitioning samples in the overlapping region of the maximal attainable sets is different. The samples in the overlapping region are assigned to one of the rovers based on the distance from the sample's position to each rover's preliminary path (a path planned before the overlapping region is taken into account). The rover whose preliminary path comes the closest to the sample's position is assigned to collect that sample. This criterion makes the sample collection task easier, since it involves only a slight deviation from the preliminary path. Again, if a rover cannot collect all of its assigned samples, the uncollected samples are passed on to the other rovers.

One possible modification is to insert the sample into the preliminary path when considering the subsequent samples in the overlapping region. This may result in a better partitioning of the overlapping region. For simplicity, this is not done in this work. Rather, when considering other samples, we use the original preliminary path. In what follows, only the "Partition of Overlapping Attainable Set: Closest Path" strategy will be considered.

2.3.3. Definitions

The sample collection path for rover j is Γ_j and is associated with a collectible sample set $C_{ss,j}$. Each path Γ_j is represented by an ordered set of points that will be denoted by $S_j \subset \hat{\Omega}$. This set may include repeated points. Let the sample distribution be denoted by $D_{ss} = D_{ss}(x)$, and the entire set of sample indices by E_{ss}.

The attainable set of rover j at a particular time t, starting from an initial point $x_{0,j} \in \hat{\Omega}$ at t_0, is denoted by $A_j(t; x_0, t_0) \subset \hat{\Gamma}$. Successive points along the admissible path Γ_j will be restricted to this attainable set. The attainable set of rover j associated with the maximum mission time τ_{max} is denoted by $A_{max,j} = A_j(\tau_{max}; x_0, t_0)$. This attainable set $A_{max,j}$ is the maximum attainable set of rover j.

The spatial domain $\hat{\Omega}$ is partitioned into sub-regions based on the maximal attainable sets $A_{max,1}, ..., A_{max,m}$. Let $\hat{Y} \in \hat{S}$ denote the set of points in the overlapping region of the maximal

attainable sets. (i.e. All points in $\hat{Y}$ can be reached by more than one rover.) Let $M_j =$ $A_{max,j} - (A_{max,j} \cap \hat{Y})$ be the sub-region of the spatial domain that can only be reached by rover j. The set of sample indices in the overlapping region will be denoted by $I_{ss} \in E_{ss}$. The sample indices in each sub-region M_j are denoted by $F_{ss,j} \subset E_{ss}$.

2.3.4 Problem Formulation
There are many possible formulations of the optimal path-planning problem for sample collection involving multiple rovers. We give one such formulation.

2.3.4.1 Problem P2
Given the approximate planetary surface $\hat{G}_f$; m rovers with initial starting points $x_{0,1},...,x_{0,m}$; mission time limit τ_{max}; sample analysis time τ_{wait}; sample index set E_{ss}; samples σ_k with values λ_k, $k \in E_{ss}$; and sample distribution $D_{ss} = D_{ss}(x)$, find optimal tours $T_1^o,...,T_m^o$ (each with path Γ_j^o and sample list $L_{ssopt,j}$) that maximize the mission return function

$$V(L_{ss,1},...,L_{ss,m}) = \sum_{j=1}^{m} \sum_{k \in L_{ss,j}} \lambda_k. \tag{10}$$

Consider the "Partition of Overlapping Attainable Set: Closest Path" strategy for finding the optimal tours. We need to first partition the terrain into sub-regions based on the maximal attainable sets. By considering the terrain without including the overlapping region $\hat{Y}$, we can solve m single rover problems: "Rover 1 starting at $x_{0,1}$ with sub-domain M_1 " to "Rover m starting at $x_{0,m}$ with sub-domain M_m." For each rover j, we maximize the mission return function

$$V(L_{ss,j}) = \sum_{k \in L_{ss,j}} \lambda_k \tag{11}$$

over the sub-domain M_j and the corresponding set of sample indices $F_{ss,j}$. From these m single rover problems, we obtain m tours $T_i = (\Gamma_i, L_{ss,i})$, $i = 1,...,m$. These tours are not optimal since we have not yet considered the samples in the overlapping region.

Next, we consider each sample in the overlapping region $\hat{Y}$. For each sample with index $k \in I_{ss}$, we find the rover j whose path Γ_j (with path point set S_j) is closest to the sample. We minimize

$$d_p(x_i) = \min_j \min_{x \in S_j} \| (x_i, f(x_i)) - (x, f(x)) \| \tag{12}$$

ith respect to x_i, the spatial coordinate of the ith sample in the overlapping region. The rover j that minimizes this function is assigned to collect the sample at x_i. This is repeated for all the samples in the overlapping region. After the assigning of the samples in the overlapping region, each rover j has a new set of samples to collect from. This new set of sample indices is denoted by $F'_{ss,j}$.

Now, we repeat solving m single rover problems with each rover j limited to the sample set $F'_{ss,j}$. The resulting tours are near-optimal.

2.4 Sample Collection Problem

The Sample Collection Problem (SCP) is an instance of the well-known Traveling Salesman Problem (TSP). A brief discussion of the TSP can be found in Appendix A. In the TSP, the problem is to find a path that visits all the nodes with minimum total traveling distance. In the SCP, the problem is to find a path that maximizes the value of the nodes visited within a specified maximum mission time limit. The differences between the TSP and the SCP are: (i) the TSP begins and ends at the same node while the SCP can end anywhere, (ii) the SCP has a waiting time associated with picking up a sample to account for the sample analysis time, (iii) the samples have different values, so different payoffs are associated with different nodes, (iv) instead of finding the minimum total distance, the SCP tries to maximize the value of all the collected samples, and (v) not all the nodes need to be visited. These modifications make the SCP a much more complex problem to solve than the original TSP, which is known to be NP-hard.

The heuristic used in solving the SCP is the maximum-value heuristic, which is similar to the minimum-distance heuristic used in solving the TSP. Instead of minimizing the total distance traveled, the maximum-value heuristic calls for maximizing the value of a weighting function that takes into account the value of each sample as well as the distance. At a given position x_i, the weighting function is used to decide on the next sample to collect.

3. Algorithms

As mentioned in Sec. 2.1.1, we only consider the rover's mass m, maximum traversable slope (or tilt) θ_{max}, maximum velocity v_{max}, and maximum power P_{max}. The samples consist of three different types, each type with a corresponding value (1, 3, or 9) representing its relative worth to geologists studying the planetary surface. The samples are randomly distributed on the terrain in the following way. The terrain is first divided into sub-regions and each sub-region j is assigned a weight $W_{i,j}$ for each sample type i such that

$$\sum_j W_{i,j} = 1 \quad \text{for all } i. \tag{13}$$

The total number of samples for each type $n_{i,tot}$ is given beforehand, and the numbers for each sub-region j, $n_{i,j}$, depend on the weights according to the following:

$$n_{i,j} = C\left(\frac{W_{i,j}}{\sum_j W_{i,j}} n_{i,tot} \right), \tag{14}$$

where the ceiling function $C(v)$ rounds v to the integer $\geq v$. Once the number of each sample type for each sub-region is given, the samples are uniformly distributed within that sub-region. Multiple samples are allowed to occupy the same spatial point. If a certain type of sample is collected, the values of the other samples of that type are not affected. This may be different from the real scenario where once one sample type has been collected; there is no need to find other samples of that type.

Starting from the rovers' initial positions, the maximal attainable sets corresponding to the mission time limit τ_{max} are computed. These maximal attainable sets represent the maximum range of each rover when no samples are collected. As the maximal attainable

sets are computed, the best path to each point on the terrain is stored. All possible paths from the starting point are explored. We examine one path at a time and follow it until time expires. As the path reaches each point, we do the following: If this is the first time a path has reached this point, the traveling time to the point is stored. If this is not the first time, the traveling time to the point by the current path is compared to the previously stored time. If the current time is less than the stored time, the current time replaces the stored time. The point may be reached many more times as other paths are tried and the best time is always kept. In this way, the best times to each point are stored as the computation progresses.

In order to retrace a path given the starting and terminal points, we introduce the "previous path point" variable. Whenever the best time for a path point is stored, the path point the rover came from (or previous path point) is also stored. Therefore, by going from a previous path point to the previous path point, the path can be retraced. This method saves memory space, since only one previous path point variable has to be stored for each attainable point on the terrain instead of an entire list of path points.
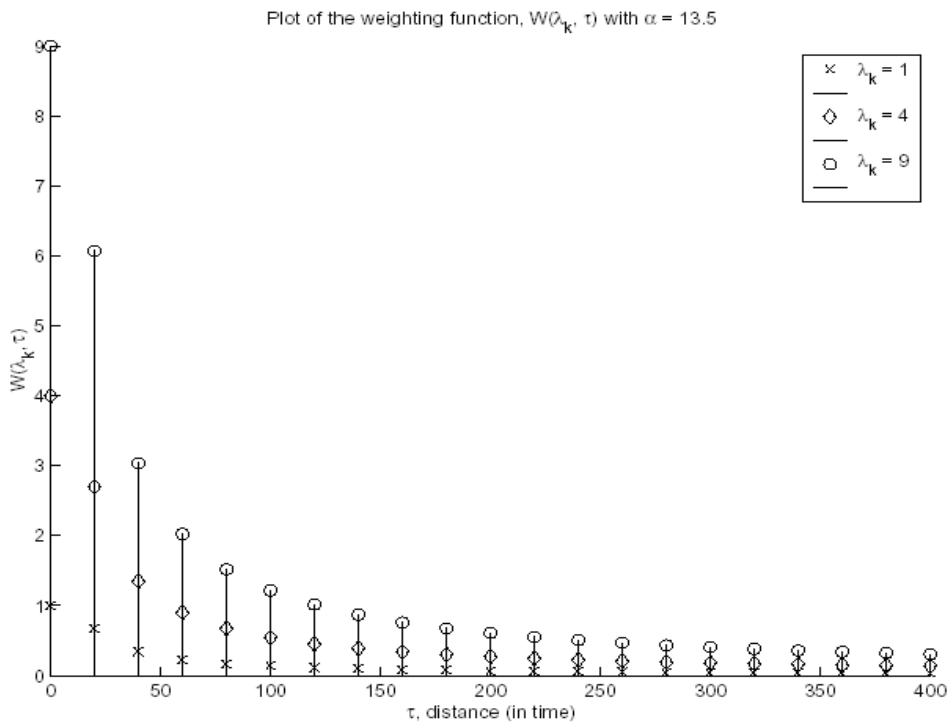

Fig. 1. Plot of the weighting function.

The set of best times (TimeMatrix), previous path points (PMatrix), and maximal attainable set (AMatrix) for each starting point are saved as Path Planning Data Sets (PPDS), which depends on the starting point and the maximum mission time. This

pre-computation saves time later when solving the SCP. The PPDS starting at each sample are also computed, which is the most time-consuming operation in the program.

Since considering all possible admissible paths is time consuming and memory intensive, an approximate solution is obtained by applying the "maximum-value heuristic" and "3-opt switching" to solve the SCP. The "maximum-value heuristic" is almost a "greedy" heuristic, which ignores all the samples except for the highest-valued ones.

The "maximum-value heuristic" is based on a weighting function that weights each sample based on its value and its distance from the starting point. Sample σ_k with value λ_k from a point x_0 is given and $\tau(x_0, \sigma_k)$ (the time it takes to get from point x_0 to sample σ_k). We define the weighting function as

$$W(\lambda_k, \tau(x_0, \sigma_k)) = \begin{cases} \lambda_k, & \text{if } \tau(x_0, \sigma_k) = 0; \\ \alpha \lambda_k / \tau(x_0, \sigma_k), & \text{otherwise.} \end{cases} \tag{15}$$

We collect the sample that maximizes this weighting function. The value for α in (15) is determined by setting the value of a sample 1 meter away from x_0 to be 3/4 of the λ_k value. In Figure 1, the weighting function for α = 13.5 and different values of λ_k are plotted with $\tau(x_0, \sigma_k)$ as a variable.

The algorithm also looks ahead two steps and maximizes the weighting function for the next two samples to collect, as well as taking into account the remaining mission time. Looking ahead two steps helps steer the rover toward the higher concentrations of samples. Let $W_i(\lambda_{k,i}, \tau_{i,i-1})$ be the weighting function value for the ith step, with sample k_i and where $\tau_{i,i-1}$ is the time required to traverse from step $i - 1$ to step i. The algorithm weights the second step a fraction $1/\beta$ of the value of the first step. The algorithm maximizes the function

$$W_{tot} = W_1(\lambda_{k,1}, \tau_{1,0}) + \frac{1}{\beta} W_2(\lambda_{k,2}, \tau_{2,1}) \tag{16}$$

with respect to $\tau_{1,0}$ and $\tau_{2,1}$ satisfying

$$(\tau_{1,0} + \tau_{wait}) + (\tau_{2,1} + \tau_{wait}) \le \tau_{max}, \tag{17}$$

where $\tau_{1,0}$ is the time to reach sample k_1 from the starting position x_0, and $\tau_{2,1}$ is the time to reach sample k_2 from sample k_1. In our case studies, the value for β is found by setting $W_1(9, 250) > W_1(1, 500) + (1/\beta)W_2(9, 0)$. The value of 500 for the time length was chosen since it is equal to the maximum mission time used for the flat terrain experimental case. The calculated value is $\beta \approx 20$.

After the list of collected samples has been determined, we apply "3-opt switching" (see Appendix A3) to obtain a more time-efficient solution. Then, we determine whether more samples can be collected. If so, the "3-opt switching" is performed again, repeatedly as necessary. The resulting tour is locally optimized, but is not necessarily optimal.
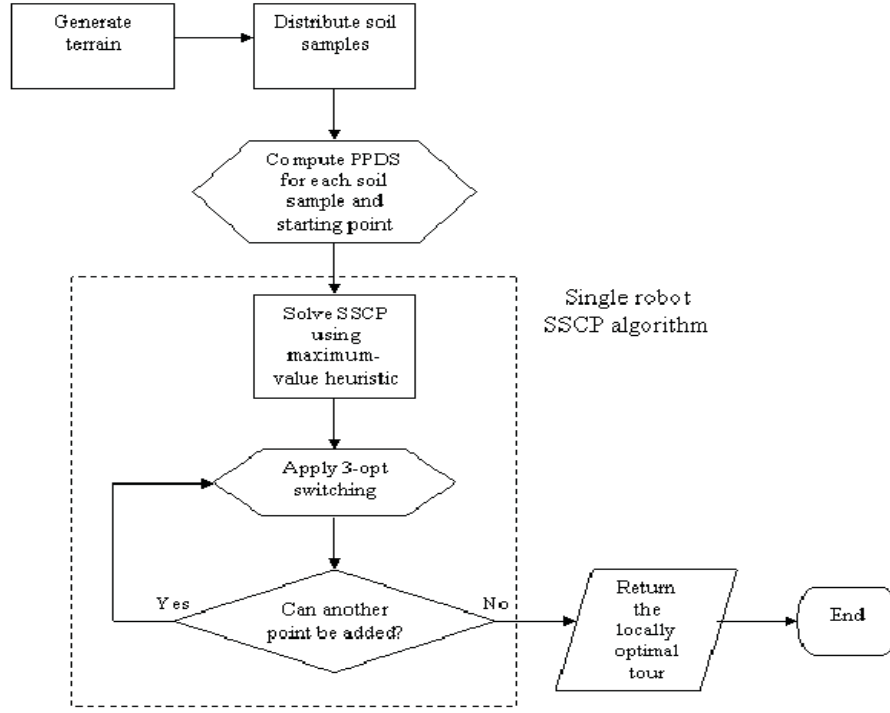
Fig. 2. Flow chart for the single rover case.

3.1 Single Rover Case

Starting from an initial point x_0 and a list of sample indices E_{ss}, an admissible tour T_a (with path Γ_a and sample list L_{ssa}) is composed by using the "maximum value heuristic" that looks ahead two steps. This tour is subject to the time constraint (6), i.e.

$$\tau_{am} = \sum_i \tau_i + n \cdot \tau_{wait} \le \tau_{max}, \tag{18}$$

where τ_{am} is the elapsed mission time for the admissible tour and n is the number of samples in L_{ssa}.

Next, "3-opt switching" is applied to the sample list L_{ssa} to determine if the same samples can be collected in less time. Each remaining uncollected sample is then considered to determine if it can be collected and analyzed within the remaining mission time. We collect the remaining sample with the highest weighting function value that still satisfies the mission time constraint. Every time an additional sample is added to the sample list L_{ssa}, "3-opt switching" is applied. This process is repeated until the list is unchanged through one iteration. The resulting tour is locally optimized, but not necessarily optimal.
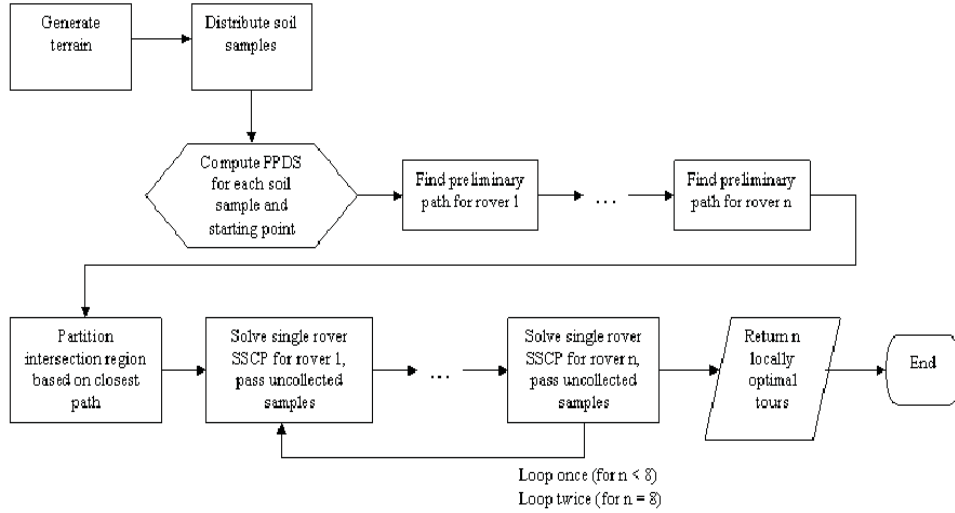
Fig. 3. Flow chart for the multiple rover case.

3.2 Multiple Rover Case

In the multiple rover cases, we assign the samples in the overlapping region to the rover whose preliminary path is closest to the sample as discussed in Sec. 2.2.1. Let $m \geq 2$ be the number of rovers. Each rover's preliminary path is found by removing the samples in the overlapping region and solving the SCP for each rover. To reduce the computation time, only the collected samples along the path are used instead of every point along the path when computing the distance to the samples in the overlapping region. If σ_j is the sample in the overlapping region in question, the distance from σ_j to rover i's path, denoted by $\delta_{i,j}$, is found according to:

$$\delta_{i,j} = \min_{x_k} \left\| (x_j, f(x_j)) - (x_k, f(x_k)) \right\|,$$

where x_k is the position of sample $\sigma_k \in \mathsf{L}_{ssa}$, and x_j is the position of soil sample σ_j in the overlapping region $\hat{Y}$. After partitioning the overlapping region, each rover has its own assigned set of samples to collect, and each sample can only be assigned to one rover at a time.

After partitioning the samples in the overlapping region among the rovers, the multiple rover case reduces to solving m single rover cases. After the first rover's tour is determined, we assign its uncollected samples to the next rover, since considering this extra set of samples may result in a better tour for the second rover. Similarly, after each rover's tour is computed, the uncollected samples are always passed on for the next rover to consider. After the m th rover's tour is determined, its uncollected samples are passed back to the first rover and all the tours are computed again. For the one-, two-, and four-rover cases, this loop-back only occurs once. In the eight-rover case, the loop-back occurs twice. This ensures that each rover has a chance to consider all the uncollected samples.

One additional note is that collision avoidance was not implemented. Adding this consideration would further complicate the path-planning algorithm and may not be useful at this point, since each rover is assumed to be a point mass in this study. In the real-life scenario, collision avoidance must be included and the dimensions of the rover must be taken into account.

4. Case Study

In this study, real Mars terrain data obtained from the Mars Orbiter Laser Altimeter (MOLA) Science Investigation are used. The terrain area chosen is relatively flat with a high plateau. The terrain is assumed to be smooth. Delaunay triangulation is used to create a 24×24 m^2 mesh for approximating the terrain. From this point on, *terrain* is used to refer to the approximated version of the Mars data.

4.1 Flat Terrain
First, we consider the case of a flat 24×24 m^2 terrain after Delaunay triangulation, shown in Fig. 4. This case provides a test of the path-planning algorithm under idealized conditions. Here, each node on the terrain is labeled from 1 to 576, starting from the bottom left corner and moving up the vertical axis. This label will be referred to as the *x-reference*.
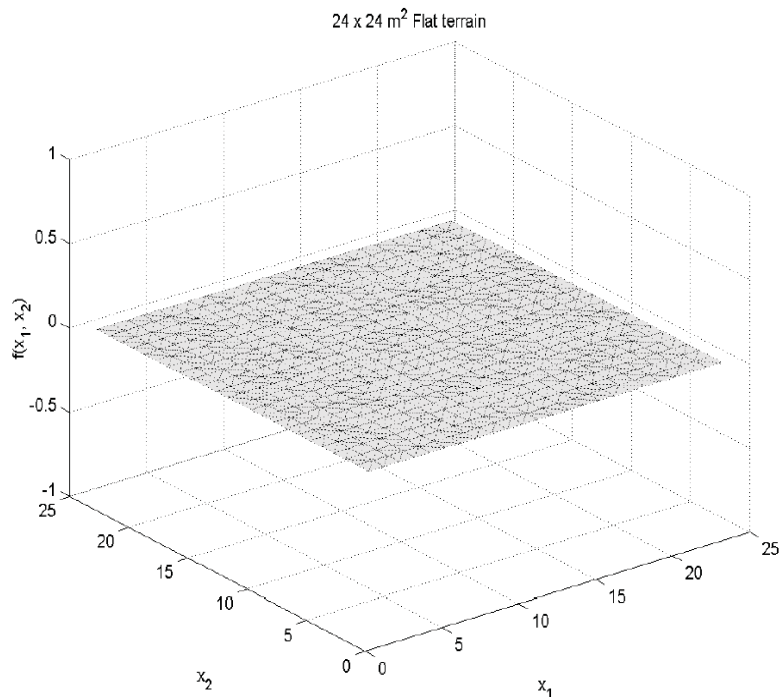

Fig. 4. Flat terrain after Delaunay triangulation.

The parameter values for the 24×24 m^2 flat terrain case are given in Table 1. They include the rover vehicle and mission specifications based on the NASA/JPL FIDO rover. Based on these data, the rover is capable of traveling a distance of 1 m in 17.99 seconds. These values

Thank You for previewing this eBook

You can read the full version of this eBook in different formats:

 ➢ HTML (Free /Available to everyone)

 ➢ PDF / TXT (Available to V.I.P. members. Free Standard members can access up to 5 PDF/TXT eBooks per month each month)

 ➢ Epub & Mobipocket (Exclusive to V.I.P. members)

To download this full book, simply select the format you desire below