Combined EDL-Mobility Planning for Planetary Missions

Yoshiaki Kuwata and Bob Balaram*

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA †

This paper presents an analysis framework for planetary missions that have coupled mobility and EDL (Entry-Descent-Landing) systems. Traditional systems engineering approaches to mobility missions such as MERs (Mars Exploration Rovers) and MSL (Mars Science Laboratory) independently study the EDL system and the mobility system, and does not perform explicit trade-off between them or risk minimization of the overall system. A major challenge is that EDL operation is inherently uncertain and its analysis results such as landing footprint are described using PDF (Probability Density Function). The proposed approach first builds a mobility cost-to-go map that encodes the driving cost of any point on the map to a science target location. The cost could include variety of metrics such as traverse distance, time, wheel rotation on soft soil, and closeness to hazards. It then convolves the mobility cost-to-go map with the landing PDF given by the EDL system, which provides a histogram of driving cost, which can be used to evaluate the overall risk of the mission. By capturing the coupling between EDL and mobility explicitly, this analysis framework enables quantitative tradeoff between EDL and mobility system performance, as well as the characterization of risks in a statistical way. The simulation results are presented with a realistic Mars terrain data.

I. INTRODUCTION

Future potential planetary missions, such as Mars Sample Return, are expected to use rovers to reach scientifically interesting sites after landing. In order to minimize the risk of mobility failure, it is critical to land the rover in a safe place, close to the science targets, and without the risk of getting stuck. Landing a rover on a planet or a moon, however, is not a deterministic operation due to various sources of uncertainties such as entry state, vehicle dynamics, atmospheric model, and wind. When performing a mission analysis, the landing position is described only in a probabilistic sense, typically with a landing error ellipse. Motivated by such mission sequence, this paper addresses the following problem: "Given the EDL model, the landing hazards, the mobility hazards, and science targets, determine the best landing ellipse location that minimizes the drive cost to the nearest target, and find the path to get there once landed."

The proposed approach is to build a mobility cost-to-go map and perform a 2D search of the landing ellipse on the map. On a discretized 2D map, it first computes a cost of driving from each cell to the nearest target while avoiding various types of hazards, such as rocks, craters, slopes, and roughness. A standard Dijkstra's algorithm is used for this computation. It then sweeps the region by placing a landing PDF (Probability Density Function) at each cell center. By convolving the landing PDF with the landing hazards, it computes the probability of landing failure; by convolving it with the mobility cost-to-go map, a histogram of driving cost to the nearest target is obtained. This histogram contains the distribution of drive distance and can produce not only the expected drive distance but the statistical values such as standard deviation, which can be used to evaluate the overall risk of the mission. The final step is to find the best location to place the center of the landing PDF from the histogram.

The rest of the paper is organized as follows. Section II presents the problem statement. Section III presents the generation of the mobility cost-to-go map and how to combine it with EDL analysis. Finally, Section IV shows results from the numerical simulations.

*Robotics and Mobility Section, 4800 Oak Grove Drive, Pasadena, CA, USA Yoshiaki.Kuwata@jpl.nasa.gov

†©2010 California Institute of Technology. Government sponsorship acknowledged.

II. PROBLEM SETUP

The problem of interest is to find an optimal location to place a landing footprint. There could be multiple cost metrics to measure the optimality, including distance to science targets, closeness to landing hazards and/or driving hazards, uncertainties in the terrain properties. This paper considers the subsequent driving cost to a target as the most important objective to minimize. The driving cost is a generic notion that could include several metrics such as distance to travel, time, wheel slippage, slope, and distance to hazards.

II.A. Landing PDF

The EDL aims to a nominal landing point, but because of the inherent uncertainties of the EDL process, the predicted landing position is described only probabilistically. Let $p(\boldsymbol{x}_{\text{land}})$ denote the landing PDF (Probability Distribution Function) over the two-dimensional plane. The vector $\boldsymbol{x} \in \mathbb{R}^2$ represents the position of the vehicle.

$$\int p(\boldsymbol{x}_{\text{land}}) d\boldsymbol{x} = 1$$

The accurate PDF can be computed using a tool such as DSENDS,[1] which runs numerous Monte-Carlo simulations with a detailed model of the spacecraft, atmosphere, wind uncertainties, and set of initial conditions of the entry phase. Assuming that there is no online guidance correction in the EDL phase, as in JPL's Mars surface missions such as MERs, Phoenix, and MSL, the PDF is well approximated as a Gaussian distribution,[2] although the approach presented in the paper is not limited to Gaussian landing PDFs.

$$p(\boldsymbol{x}_{\text{land}}, \boldsymbol{x}_{\text{c}}, \Sigma) = \frac{1}{2\pi|\Sigma|} \exp\left(-\frac{1}{2}(\boldsymbol{x} - \boldsymbol{x}_{\text{c}})'\Sigma^{-1}(\boldsymbol{x} - \boldsymbol{x}_{\text{c}})\right)$$

where $\boldsymbol{x}_{\text{c}} \in \mathbb{R}^2$ is the mean landing position and a 2-by-2 matrix Σ specifies the size and the orientation of Gaussian. The mean $\boldsymbol{x}_{\text{c}}$ is also a center of the 3-σ ellipse of the distribution, and is referred to as "center" of the landing position PDF. Once the EDL system is designed and the launch date of the spacecraft is determined, the shape and the orientation of the PDF cannot be changed. Even after launch, however, the latitude and the longitude of the center $\boldsymbol{x}_{\text{c}}$ of the PDF can be shifted translationally through orbital maneuvers during the cruise phase. Thus, Σ is fixed and the center $\boldsymbol{x}_{\text{c}}$ is what we solve for in order to minimize the appropriate cost metric.

II.B. Terrain

In order to characterize the terrain hazards, there are several types of data sources available. From the elevation data, a slope map and a roughness map could be obtained. From the image with shadow, rocks of various shapes can be detected.[?] From the thermal inertial map, cohesiveness of the terrain surface could be inferred.[?] These different source of hazards are lumped into a single hazard map by taking the worst value of different hazards types.

The threshold to declare something hazardous could be different among vehicles. For example, a larger vehicle with larger wheels (for rovers) or longer legs (for landers) can handle areas with larger rocks. On the other hand, a lander could have a higher center-of-gravity than a rover and might not be able to handle steep slopes due to tip-over constraints. Let $\mathcal{X}_{\text{lander}} \subseteq \mathbb{R}^2$ denote a set of region that is feasible (i.e., non-hazardous) to the lander, and $\mathcal{X}_{\text{rover}} \subseteq \mathbb{R}^2$ to the rover. In this paper, the vehicle design is assumed to be given, together with the thresholds for the maximum slope level, rock size, and roughness. Therefore, $\mathcal{X}_{\text{lander}}$ and $\mathcal{X}_{\text{rover}}$ are assume to be fixed.

II.C. Science Targets

The goal of the mission is to drive the rover to scientifically interesting sites to collect useful science data with its on-board instruments. There can be different types science sites, such as point targets (e.g., small craters – going to a point is sufficient), linear features (e.g., ravines – need to take science samples along a line), correlated targets (e.g., if the rover is to take a sample in one location, another specific location needs to be sampled in order for the data to be useful), and sized areas (reaching any point in the area is sufficient). For the purpose of this paper, the science targets are assumed to be represented as a set of 2D points $\boldsymbol{x}_{\text{target},i} \in \mathbb{R}^2$, $(i = 1, \ldots, n_{\text{target}})$, where n_{target} is the number of targets, and the rover's goal is to visit any of them.

II.D. Mission Risks

There are two major sources of mission failure risks, the landing failure and mobility failure.

Because the EDL is a probabilistic event with large uncertainties (e.g., the size of the landing error ellipse of MERs was 63km-by-9km[2]), the probability of landing failure cannot be made 0. The mission designer must set an acceptable probability of failure (e.g., 0.01% chance of landing on a hazard), which is denoted by $p_{\text{land,max}}$.

Another mission failure case for combined EDL-Mobility missions is that the EDL lands the vehicle at a location from which there is no feasible path to reach any target. One such example is that the landed location is in the middle of a crater whose slope at the perimeter is too steep for the rover to climb over. Even if it is safe for the lander to land, this is a failure case for the rover mission. Let $p_{\text{trap,max}}$ denote the maximum probability that the user specifies to tolerate such entrapment. These risk bounds are used as a constraint when choosing the center of the landing PDF.

The mobility specific failure could be caused by several sources of uncertainties such as hardware fatigue, high centering due to unforseen obstacle, and entrapment in the soil that is softer than expected.[?] The probability of such uncertain event is a function of path integral of distance traveled and closeness to hazards. Let $p_{\text{mob,max}}$ denote the maximum acceptable probability of such event happening.

II.E. Problem Statement

Finally, the problem is stated as follows. Given

- the shape of the landing position PDF Σ,
- a set of science targets $\boldsymbol{x}_{\text{target},i}$, $(i = 1, \ldots, n_{\text{target}})$,
- a list of feasible regions $\mathcal{X}_{\text{lander}}$ and $\mathcal{X}_{\text{rover}}$ for the lander and the rover,
- the size of the lander and the rover, and
- thresholds for mission failure probability $p_{\text{land,max}}$, $p_{\text{trap,max}}$, $p_{\text{mob,max}}$

find the best location for the center $\boldsymbol{x}_c$ of the landing footprint, minimizing the expected cost-to-go of the rover driving, subject to the mission failure probability less than the user-specified bounds.

III. ALGORITHMS

The approach consists of two phases. Although the actual EDL happens before the rover driving in the mission sequence, the analysis goes backwards by first performing a mobility cost-to-go analysis, then selecting the best center of the landing PDF. The following subsections gives a more detailed procedure.

III.A. Mobility Cost-To-Go Map

In the first phase, a mobility cost-to-go map is computed. This is a 2D map that stores the cost of driving from an arbitrary location $\boldsymbol{x}$ to the nearest target along an optimal path, and is denoted by $\text{CTG}(\boldsymbol{x})$. This cost-to-go map represents the results of the mobility analysis and serves as an interface to the EDL analysis presented in the next subsection.

The cost-to-go map is computed by constructing a graph and running a shortest path algorithm with each target $\boldsymbol{x}_{\text{target},i}$ as a goal node. Figure 1 shows the information flow of this process.

Several types of terrain hazards, as discussed in Section II.B, are combined into a rasterized hazard map, as shown in the upper left blocks of the figure. The hazards are dilated by the size of the vehicle radius, so that the vehicle can be treated as a point.

The trajectory search space is defined on a discrete grid. This paper uses a simple 8-connected grid, but the presented approach extends to more sophisticated grid such as 16-connected grid or state lattice[3] if wider range of maneuvers need to be captured. The edge cost $c(\boldsymbol{x}_i, \boldsymbol{x}_j)$ encodes the cost of traversing from one node $\boldsymbol{x}_i$ to its neighbor node $\boldsymbol{x}_j$, and is computed by looking up the hazard level along the edge. The

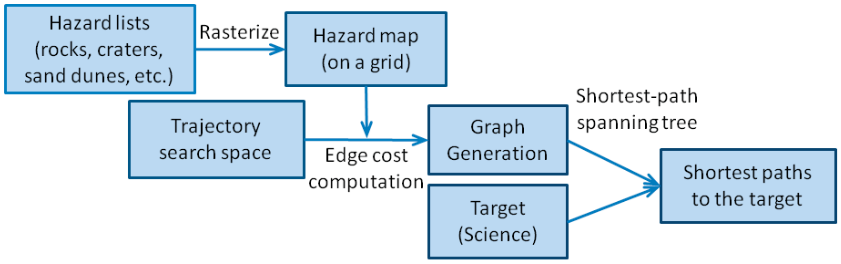

Figure 1. Information flow for computing shortest paths from every location on the map to the target

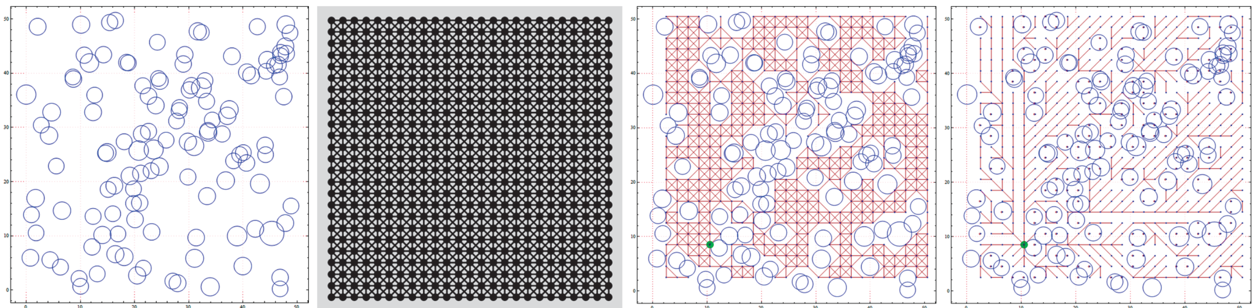

(a) Circular rocks (b) 8-connected graph (c) Graph with finite edge costs (d) Minimum spanning tree

Figure 2. A simple example to illustrate the cost map generation procedure

edge cost is infinity if the hazard (e.g., slope, roughness, and rock height) is beyond the threshold of the rover. Otherwise, Euclidean distance with some hazard penalty is used, so that

$$c(\boldsymbol{x}_i,\, \boldsymbol{x}_j) = \|\boldsymbol{x}_i - \boldsymbol{x}_j\| + wh(\boldsymbol{x}_i, \boldsymbol{x}_j)\|\boldsymbol{x}_i - \boldsymbol{x}_j\| \tag{1}$$

where w is a weight on the hazard penalty, $h(\boldsymbol{x}_i, \boldsymbol{x}_j)$ is a hazard value along the (i, j) edge. The first term in (1) represents the distance to traverse, and the second term represents the path integral of the hazard risk.

Once the edge costs are defined, a set of science targets $\boldsymbol{x}_{\text{target},i}$ are used as the source nodes of the graph, and the shortest-path spanning tree is computed. The shortest-path spanning tree also gives a cost of moving from each node to the target. Let $\text{PP}(\boldsymbol{x}_1,\, \boldsymbol{x}_2)$ denote a function that gives a cost associated with the minimum-cost path from $\boldsymbol{x}_1$ to $\boldsymbol{x}_2$. Then, the cost-to-go value defined at each node location $\boldsymbol{x}$ is given by

$$\text{CTG}(\boldsymbol{x}) = \min_{i=1,\dots,n_{\text{target}}} \text{PP}\big(\boldsymbol{x},\, \boldsymbol{x}_{\text{target},i}\big)\,. \tag{2}$$

Figure 2 shows a simple example that illustrates this procedure using a single target and circular obstacles. Figure 2(a) has a set of circular hazards representing rocks. Figure 2(b) is the 8-connected grid used as a search space for rover's path. By combining these hazard map and a search space, we can obtain a graph with a cost, as shown in Figure 2(c). The edge connections with ∞ cost are infeasible and not drawn. Finally, by running a shortest path algorithm from a given goal (marked with a green circle in the lower left) produces a minimum spanning tree, which shows which neighboring node to transition from every node on the map in order to reach the goal with the minimum cost.

III.B. Landing Ellipse Placement

In the second phase, the optimal landing ellipse is placed using the cost-to-go map obtained in the first phase.

As discussed in Section II.A, the result of the EDL operation is represented as a PDF of the landing location. Therefore, minimization of the drive cost must be done using expected value. The cost-to-go map provides the drive cost from any point $\boldsymbol{x}$ to the target, so that the problem of finding the optimal landing ellipse is formulated as

$$\min_{\boldsymbol{x}_c} \mathop{\mathbb{E}}_{\boldsymbol{x}_{\text{land}}} \left[\text{CTG}(\boldsymbol{x}_{\text{land}})\right] = \min_{\boldsymbol{x}_c} \int \text{CTG}(\boldsymbol{x}_{\text{land}})\, p(\boldsymbol{x}_{\text{land}}, \boldsymbol{x}_c, \Sigma) d\boldsymbol{x}_{\text{land}} \tag{3}$$

where the integral is taken over all the 2D region that is covered by the landing PDF centered at $\boldsymbol{x}_c$. This integral is implemented by defining the cost-to-go function CTG on a regular grid, and also rasterizing the PDF $p(\cdot)$ on the same grid. This has the benefit of being able to incorporate any types of landing PDF (e.g., a powered-descent guidance avoids landing hazards and/or steers the lander to a more preferable landing location[4]) and also allowing us to easily evaluate the term minimized in (3) numerically.

IV. SIMULATION RESULTS

IV.A. Setup

The simulation results presented in this section uses the actual Mars terrain data obtained by MRO (Mars Reconnaissance Orbiter). The elevation map was downloaded from a publicly available web site.[5]

The size of the grid in the following subsections is 500-by-500, where the physical size of each pixel is 1m-by-1m.

IV.B. Results

IV.B.1. Mobility cost-to-go map

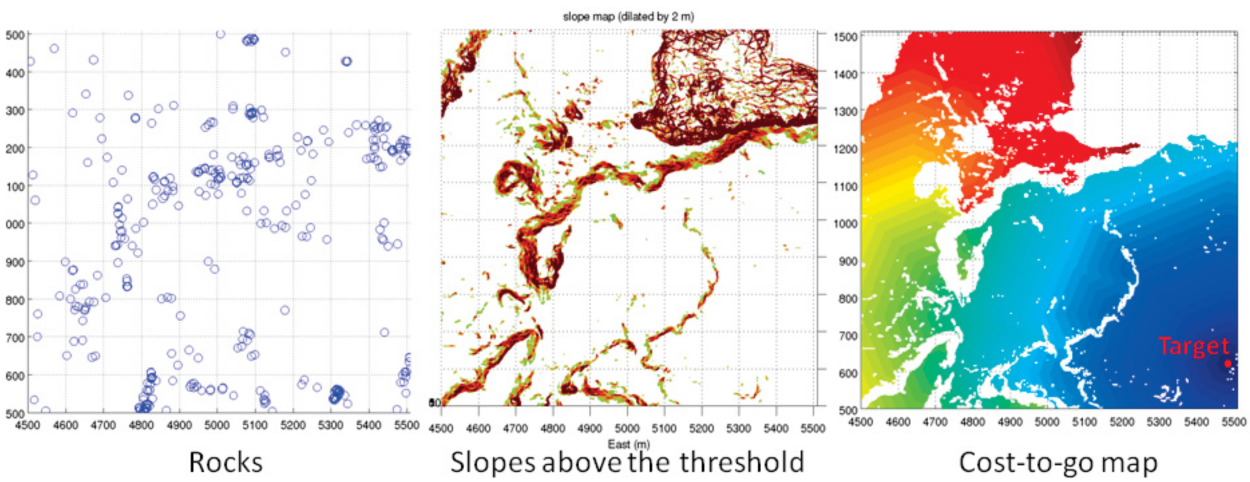

Figure 3. Cost-to-go map computed with rocks and slopes as hazards

The left-most and middle figures of Figure 3 show hazard maps for rocks and slopes respectively. The target is at the lower right, as shown in the right-most figure. The right-most figure shows the resulting mobility cost-to-go map based on these two maps. The color represent the cost-to-go from each grid cell to the target. For the locations that are above the hazard thresholds or that is not possible to reach the target from, a blank color is used. Note that the top part of the map has a higher cost-to-go compared to lower left of the map. This is because the linear terrain feature with steep slope (from the upper right to the middle left of the figure) prevents the rover from driving straight towards the target.

IV.B.2. Optimal landing ellipse placement

Figure 4 shows how the failure probability of EDL is computed. The plot on the left shows the landing hazard map. The color white indicates that the hazard level at that location is above the threshold of the lander.

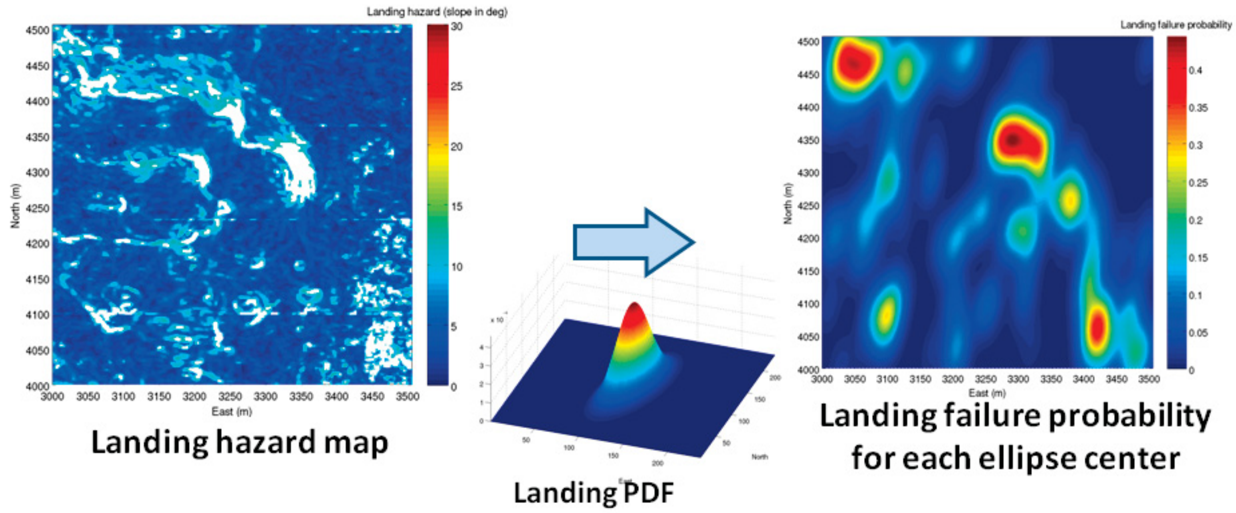

Figure 4. EDL failure probability map

The plot in the middle represents the landing PDF. By placing the landing center $\boldsymbol{x}_c$ at each cell of the 500-by-500 grid, and convolving the landing hazard map and the landing PDF using the integral in (3), we can compute the landing failure probability. The result is shown in the right figure. The color of each cell represents the landing failure probability if the *center* $\boldsymbol{x}_c$ of the landing PDF is placed at that cell.

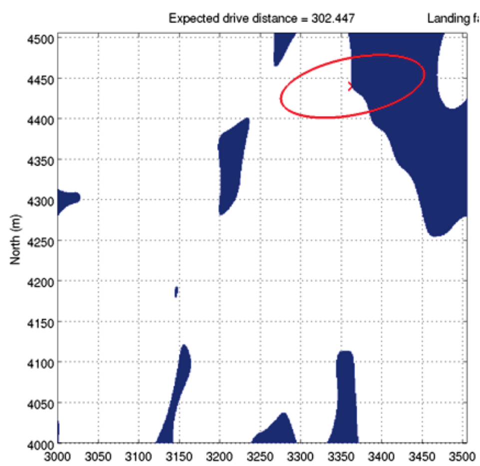

Figure 5. The optimal landing ellipse that minimizes the expected rover driving

By retaining only the region where the landing failure probability is within the acceptable threshold $p_{\text{land},\text{max}}$, we obtain a region where placing the landing PDF center leads to successful landing. Then, for each cell in this region, we compute the expected drive distance via (3) and select the best center of the landing PDF $\boldsymbol{x}_c^*$. The final result is shown in Figure 5.

V. CONCLUSION

This paper presented an analysis framework for combined EDL and mobility system. The landing position after EDL is described with a PDF, and in order to maintain the probabilistic framework through the mobility analysis, a cost-to-go map is presented. By convolving the cost-to-go map with the landing PDF, the approach finds the optimal landing location that minimizes the subsequent driving cost in an expected sense. Future work will investigate how to extend this approach to a larger map, and how to extend this framework to more general landing PDFs such as the one obtained by pin-point landing.

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.

References

[1]Balaram, J., Austin, R., Banerjee, P., Bentley, T., Henriquez, D., Martin, B., McMahon, E., and Sohl, G., "DSENDS – A High-Fidelity Dynamics and Spacecraft Simulator for Entry, Descent and Surface Landing," *Proceedings of IEEE Aerospace Conference*, Big Sky, Montana, March 2002.

[2]Knocke, P., Wawrzyniak, G., Kennedy, B., Desai, P., Parker, T., Golombek, M., Duxbury, T., and Kass, D., "Mars Exploration Rovers Landing Dispersion Analysis," *Proceedings of the AIAA/AAS Astrodynamics Specialist Conference*, No. AIAA 2004-5093, 2004.

[3]Pivtoraiko, M., Knepper, R. A., and Kelly, A., "Differentially Constrained Mobile Robot Motion Planning in State Lattices," *Journal of Field Robotics*, Vol. 26, No. 3, 2009, pp. 308–333.

[4]Acikmese, B. and Ploen, S., "Convex Programming Approach to Powered Descent Guidance for Mars Landing," *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 5, 2007, pp. 1353–1366.

[5]"HiRISE — High Resolution Image Science Experiment," http://hirise.lpl.arizona.edu/ as of 2011.