

Probabilistic Motion Planning of Balloons in Strong, Uncertain Wind Fields

Michael T. Wolf, Lars Blackmore, Yoshiaki Kuwata,
Nanaz Fathpour, and Alberto Elfes
NASA Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA USA

Claire Newman
Div. of Geological and Planetary Sciences
California Institute of Technology
Pasadena, CA USA

Abstract—This paper introduces a new algorithm for probabilistic motion planning in arbitrary, uncertain vector fields, with emphasis on high-level planning for Montgolfier balloons in the atmosphere of Titan. The goal of the algorithm is to determine what altitude—and what horizontal actuation, if any is available on the vehicle—to use to reach a goal location in the fastest expected time. The winds can vary greatly at different altitudes and are strong relative to any feasible horizontal actuation, so the incorporation of the winds is critical for guidance plans. This paper focuses on how to integrate the uncertainty of the wind field into the wind model and how to reach a goal location through the uncertain wind field, using a Markov decision process (MDP). The resulting probabilistic solutions enable more robust guidance plans and more thorough analysis of potential paths than existing methods.

I. INTRODUCTION

This report presents a new method for 3D motion planning of Montgolfieré (hot air) balloons using wind models for time-efficient paths. The method was developed for studying how exploration balloons (perhaps with some modest horizontal thrust capability) may be used in the atmosphere of Titan, a moon of Saturn, with an emphasis on goal-based traversals. Since wind velocities dominate over the buoyant vehicle’s own actuation in this scenario, using forecast models of the wind field surrounding the moon is critical to effective planning. In practical applications, however, the actual wind vector field is not known exactly and may deviate

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.

significantly from the wind velocities estimated by the model. To address this issue, our technique explicitly incorporates wind uncertainty into the planning algorithm.

Our algorithm is designed for high-level planning over large scale traversals, resulting in a policy that provides the desired actuation choice for every location. More fundamentally, this algorithm was also designed to enable mission trade studies, providing expected traversal times (and measures of the corresponding risk) for a range of vehicle actuation types and limits, insertion points, goal locations, etc. Because of these goals, we ignore balloon dynamics in the current implementation. The algorithm may be useful in other atmospheres, including Earth, Mars, Venus and the gas giants, and perhaps for any buoyant vehicles operating in strong vector fields, such as underwater vehicles in ocean currents.

Two types of vehicle actuation capability are considered: vertical-only and vertical-and-horizontal. Vertical actuation changes the vehicle’s buoyancy to exploit winds at different altitudes. Horizontal actuation is incorporated as additive velocity and is presumed weak compared to wind velocity (i.e., the wind fields are *strong*).

This report builds on earlier work by the authors [1]–[3], who previously reported deterministic, graph-based solutions to the same problem. Other previous approaches in path planning both for balloons [4], [5] and for the related problem of autonomous underwater vehicles (AUVs) in current fields [6]–[11] typically suffer from strong assumptions on the surrounding vector field. To our knowledge, ours is the first solution that meaning-

fully acknowledges the uncertainty of the vector field, thus providing a more robust solution and one useful in evaluating variability risk. Further, our method accepts arbitrary, 3D, even time-varying vector fields from the forecast model, rather than requiring analytical (even constant or linear) vector fields.

To prepare the motion planning problem, we first model the uncertainty in the wind field. Then we formulate the problem of reaching a particular goal location as a Markov decision process (MDP), using a discretized space approach. Solving the MDP provides a policy of what actuation option (how much buoyancy change and, if applicable, horizontal actuation) should be selected at any given location to minimize the expected time-to-goal. The results also provide expected time-to-goal values from any given location on the globe.

Below, we describe the problem setup in Sect. II and describe our approach in Sect. III. Finally, Sect. IV provides algorithm results for different goal locations and actuation models. The material in Sect. II and Sect. III-A follows closely from [2], with modifications for the probabilistic approach.

II. PROBLEM SETUP

Let the vehicle’s position be denoted by $\mathbf{r} \triangleq [x\ y\ z]^T$, where x is in degrees longitude, y is in degrees latitude, and z is altitude from the surface of the planet. The wind velocity is defined in a Cartesian coordinate frame fixed to the local surface tangent such that $\mathbf{w}(\mathbf{r}, t) \triangleq [w_x(\mathbf{r}, t)\ w_y(\mathbf{r}, t)\ w_z(\mathbf{r}, t)]^T$, where w_x , w_y , and w_z are the velocities in the easterly, northerly and vertically upwards directions, respectively.

We can simplify the problem using a partial decoupling of the Montgolfieré dynamics, based on the following assumptions: (1) The altitude of the balloon is fully controllable, subject to maximum rise and sink rates, denoted v_{rise} and v_{sink} , respectively. This assumption comes from the observation that the vertical control authority of the vehicle is large compared to the vertical winds predicted by the global circulation models of [12], and leads to ignoring the effects of vertical winds. (2) The unactuated horizontal velocity of the Montgolfieré is equal to the local wind velocity. This assumption comes originally from [4], and means that, in the horizontal plane, we need only consider the local

wind velocity and not any other thermal or dynamic state of the Montgolfieré.

We consider the case that the vehicle may have horizontal actuators that can generate additional velocity $\mathbf{u}(t)$ with respect to the air. Let $u_{h\ \text{max}}$ denote the maximum achievable horizontal actuation.

III. APPROACH

A. Spatial Discretization and Transitions

We choose to discretize space using a grid, where adjacent nodes are separated by Δx in longitude, Δy in latitude, and Δz in altitude. Consider for now the time-invariant problem, such that the wind provided by the model, denoted $\bar{\mathbf{w}}$, is fixed at some T for all t , i.e., $\bar{\mathbf{w}}(\mathbf{r}, t) = \bar{\mathbf{w}}(\mathbf{r}, T)$.

The decoupling described in Section II enables us to consider the discretization of the three-dimensional search space first in the horizontal plane, and then in the vertical plane. We also discretize the horizontal actuation into n_h vectors that are different in magnitude and/or direction. Transitions are modeled as follows. Let \mathbf{r}_i represent the vehicle position at state s_i , where each state designates a spatial cell in the grid. For every state s_i , the local wind $\mathbf{w}(\mathbf{r}_i, t)$ is selected at each visit from a probability distribution (the wind model and uncertainty is discussed in Sect. III-C). After selecting the desired horizontal actuation, the Montgolfieré velocity from state s_i can be written as:

$$\mathbf{v}_i = \frac{d\mathbf{r}_i}{dt} = \mathbf{w}(\mathbf{r}_i, t) + \mathbf{u}(s_i), \quad (1)$$

which, like \mathbf{w} , is defined on the tangent plane. The horizontal plane is considered first. We discretize the direction of the Montgolfieré velocity in the horizontal plane into one of eight segments, as shown in Figure 1. Which of these segments the resulting vector \mathbf{v}_i “points to” determines the cell to which the vehicle will transition, which we denote s_j . By assuming that the wind is constant in the interval until the next cell is reached, the time taken to travel from s_i to s_j is given by:

$$\delta t_i = \frac{\text{dist}(\mathbf{r}_i, \mathbf{r}_j)}{\|\mathbf{v}_i\|}. \quad (2)$$

Here, $\text{dist}(\cdot, \cdot)$ is a function that returns the cartesian distance between two points in a spherical coordinate frame.




Fig. 1. Discretization in the horizontal plane with horizontal transition boundaries.

By applying vertical actuation, cells above or below s_j may also be reached. By our earlier assumptions, the vertical range of the Montgolfieré depends only on the maximum rise and sink rates and the time available. For example, the maximum and minimum altitudes reachable in traveling from s_i to an adjacent cell are:

$$z_{\max}(s_i) = \frac{v_{\text{rise}}}{\delta t_i} \quad z_{\min}(s_i) = \frac{v_{\text{sink}}}{\delta t_i}. \quad (3)$$

The key idea is that the winds and horizontal actuation determine the next horizontal cell, and we allow the time of travel between states to be variable. Then, the vertical “layer” is selected by the vehicle from the allowable range defined by (3), resulting in the next state, s'_j .

B. Overview of the MDP

Because the wind velocity $\mathbf{w}(\mathbf{r}_i, t)$ is uncertain, the next horizontal state s_j may be considered a random variable, and a probability distribution can be constructed over all horizontally adjacent cells. Given these transition probabilities from all states, our motion planning problem is to select the actions (horizontal and vertical actuation of the Montgolfieré) that minimizes time-to-goal. This problem is thus naturally posed as a Markov Decision Process (MDP) [13] $(\mathcal{S}, \mathcal{A}, P, R)$, where: \mathcal{S} represents the set of possible states $\{s_i\}_i$ of the Montgolfieré; \mathcal{A} is the set of actions available from each state (horizontal/vertical actuation); P gives the transition probabilities $P_a(s_i, s'_j)$ from current state s_i to possible next states s'_j under action a ; R defines the expected immediate reward for each transition and each action a . In this problem, reward is negative travel time.

C. Modeling Wind Uncertainty

We choose to decompose the wind into direction and magnitude components, denoted respectively by θ_i and w_i , and assign independent probability distributions to each. The expected value of each component is equal to the field model value given by $\bar{\mathbf{w}}(\mathbf{r}_i, t)$. To model the uncertainty in θ_i , we employ a von Mises distribution, an analogue of a Gaussian distribution on the circle,

$$f_{\text{VM}}(\theta_i | \bar{\theta}_i, \kappa) = \frac{\exp(\kappa \cos(\theta_i - \bar{\theta}_i))}{2\pi I_0(\kappa)}, \quad (4)$$

where $\bar{\theta}_i = \angle \bar{\mathbf{w}}(\mathbf{r}_i, t)$ is the mean, κ is a concentration parameter, and $I_0(\kappa)$ is the modified Bessel function of the first kind of order 0.

The wind magnitude w_i uncertainty is modeled as Gaussian¹, with mean $\bar{w}_i = \|\bar{\mathbf{w}}(\mathbf{r}_i, t)\|$. The standard deviation is set as proportional to the magnitude, $\sigma_i = \rho w_i$ (ρ is chosen by the user).

D. Transition Probabilities

Next we define the transition probabilities $P_a(s_i, s'_j)$ that govern what state s'_j is entered after executing each action a from each state s_i . Note that all uncertainty is limited to the horizontal direction; thus, for the discussion below, only the projections onto the horizontal plane are relevant.

If there is no horizontal vehicle actuation, the wind direction wholly determines the next state’s horizontal position. In that case, the probability of the wind forcing the Montgolfieré from state s_i to a horizontally adjacent state s_j is

$$P_a(s_i, s_j) = \int_{\theta_{ij}}^{\theta_{ij} + \frac{\pi}{4}} f_{\text{VM}}(\theta | \bar{\theta}_i, \kappa) d\theta, \quad (5)$$

where θ_{ij} is the smaller bordering angle of the one-eighth circular sector pointing from s_i to s_j (see Fig. 1).

Introducing horizontal actuation as additive velocity on top of the wind velocity, we are interested in the transition governed by the resultant velocity, $\mathbf{v}_{i,a} = \mathbf{w}(\mathbf{r}_i, t) + \mathbf{u}_a$. A Monte Carlo method is used to determine the probabilities P and expected

¹More carefully, using a distribution with non-negative support, such as a log-normal distribution, may be preferable. However, because our choice of variance is generally a small percentage of the mean, the probabilities associated with negative magnitudes are negligible.

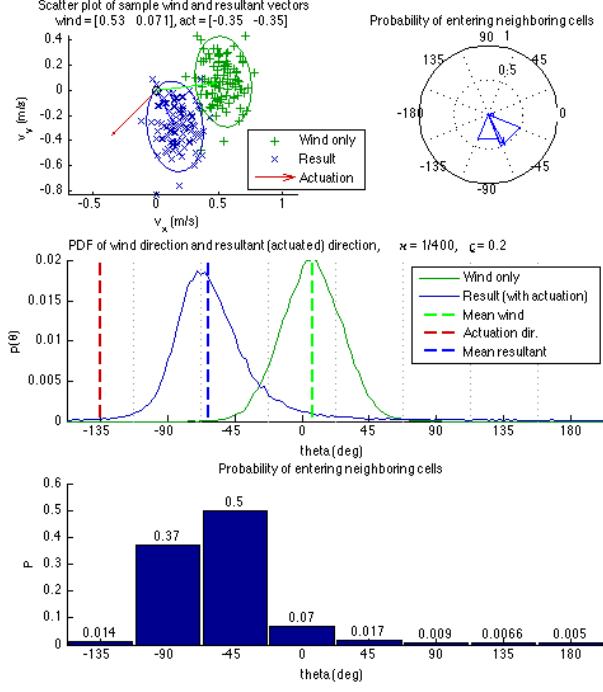


Fig. 2. Wind uncertainty model and horizontal transition probabilities. This example uses a mean wind direction $\bar{\theta}_i = 7^\circ$, with $\kappa = \frac{1}{400}$ (closely corresponding to Gaussian with $\sigma \approx \frac{1}{\sqrt{\kappa}} = 20^\circ$); mean wind magnitude $w_i = 0.54 \frac{m}{s}$, with $\rho = 0.2$ (so that $\sigma_i = \rho w_i = 0.11$). The figure includes a scatter plot of samples drawn from the wind distribution (top left); polar plot of probabilities of transitioning to each horizontally adjacent cell (top right); distribution of the wind direction (middle), and bar plot of probability of horizontal transitions (bottom).

rewards R from resultant vector for each action, described as follows and illustrated by example in Fig. 2. First, a set Θ_0 of N sample points are drawn from the von Mises distribution $f_{VM}(\theta|0, \kappa)$ and a set W_0 of N samples from the standard normal distribution. (This step is required only once, whereas the following steps must be done by iterating for each state s_i .) Second, we adjust these samples for the state s_i . The wind direction samples are rotated by $\bar{\theta}_i$, $\Theta_i = \Theta_0 + \bar{\theta}_i$, and the wind magnitude samples are adjusted by $W_i = \sigma_i W_0 + \bar{w}_i$ (pardoning the abuse of notation—each sample in Θ_0 and W_0 is adjusted individually). Third, we calculate the resultant velocity samples by converting the random wind samples to Cartesian coordinates and adding the horizontal actuation:

$$V_{i,a} = W_i \begin{bmatrix} \cos \Theta_i \\ \sin \Theta_i \end{bmatrix} + \mathbf{u}_a, \quad (6)$$

where again the operations involving W_i and Θ_i

are performed element-wise.

Once we have the resultant samples $V_{i,a} = \{\mathbf{v}_{i,a}^n\}_{n=1}^N$, we can calculate the transition probabilities P , which are simply made by counting the samples in each circular sector corresponding to s_j :

$$P_a(s_i, s_j) = \frac{1}{N} |V_{ij,a}|, \quad (7)$$

where $|\cdot|$ represents cardinality and:

$$V_{ij,a} = \left\{ \mathbf{v}_{i,a}^n \mid \angle \mathbf{v}_{i,a}^n \in \left[\theta_{ij}, \theta_{ij} + \frac{\pi}{4} \right] \right\}. \quad (8)$$

E. Action Space and Rewards

Travel time is used as a transition cost between states. Note that we do not know the resultant velocity exactly but that we seek the *expected* reward for the MDP. We estimate the expected value of the velocity magnitude as the population mean of the above Monte Carlo samples:

$$\langle v_{i,a} \rangle = \frac{1}{N} \sum_{n=1}^N \|\mathbf{v}_{i,a}^n\|. \quad (9)$$

The expected immediate reward for the transition from s_i to s_j under action a is then:

$$R_a(s_i, s_j) = -\frac{\text{dist}(\mathbf{r}_i, \mathbf{r}_j)}{\langle v_{i,a} \rangle}. \quad (10)$$

Note also that the travel time is identical for all actions a that have the same horizontal actuation (i.e., travel time is not dependent on altitude choice).

Finally, let us more formally define our set of available actions, \mathcal{A} . In the case that horizontal actuation is considered, this action are chosen from

$$\mathcal{A}^h = \{\mathbf{0}, u_{h \max} \hat{\mathbf{u}}_1, u_{h \max} \hat{\mathbf{u}}_2, \dots, u_{h \max} \hat{\mathbf{u}}_M\},$$

where $\hat{\mathbf{u}}_m$ is the unit vector pointing in each of the M directions allowed. M and the corresponding directions $\hat{\mathbf{u}}_m$ are chosen by the user as part of the problem setup; for example, we typically allow actuation in the direction of each horizontally adjacent cell and thus $M = 8$.

The vertical actuation options \mathcal{A}^v are governed by the time of travel to the next horizontal cell and the maximum rise and sink rates. However, the travel time is *a priori* unknown—not only is the wind velocity uncertain, but the distance traveled $\text{dist}(\mathbf{r}_i, \mathbf{r}_j)$ is also uncertain because s_j is unknown. Here we again rely on the expected value of the travel time, this time across all adjacent

horizontal cells. Let \mathcal{S}_i denote the set of cells horizontally adjacent to s_i . Then, given a starting state s_i and a horizontal action choice $a_h \in \mathcal{A}^h$, the expected travel time is

$$\begin{aligned} \widehat{\delta t}_{i,a_h} &\equiv E \left[\frac{d_{ij}}{\langle v_{i,a_h} \rangle} \mid s_i, a_h \in \mathcal{A}^h \right] \\ &= \sum_{j \in \mathcal{S}_i} P_{a_h}(s_i, s_j) \cdot R_{a_h}(s_i, s_j). \end{aligned} \quad (11)$$

Then $\widehat{\delta t}_{i,a_h}$ is used to determine the minimum and maximum altitude levels and a vertical actuation option assigned for each level in between. Defining the vertical “levels” $\{k_1, k_2, \dots, k_{n_z}\}$ at altitudes $\{z_1, z_2, \dots, z_{n_z}\}$, each separated in altitude by distance Δz , we can define \mathcal{A}^v for every starting cell and every horizontal action:

$$\mathcal{A}_{i,a_h}^v = \{k_j \mid z_i - v_{\text{rise}} \widehat{\delta t}_{i,a_h} \leq z_j \leq z_i + v_{\text{sink}} \widehat{\delta t}_{i,a_h}\}.$$

Finally, the action space at every cell is defined as the combination of horizontal and allowable vertical actions:

$$\mathcal{A}_i = \{a_h, a_v \mid a_h \in \mathcal{A}^h, a_v \in \mathcal{A}_{i,a_h}^v\}. \quad (12)$$

F. Solving for Minimum Time-to-Goal

To complete the problem setup, we define a goal location for the Montgolfieré and create a sink state s_g at this location, from which all transitions are set to have probability zero. Thus, the cumulative reward will decrease with every transition until the vehicle reaches the goal location s_g . Given this setup, the MDP solution will determine, for each given current state s_i , what is the optimal immediate action $a \in \mathcal{A}_i$ so that the expected cumulative time-to-goal is minimal. This collection of actions is referred to as the optimal policy π^* .

An undiscounted ($\gamma = 0$) MDP solution method is appropriate because we are interested in the cumulative time elapsed from start to goal. Then, the total expected reward of a state s_i (value of s_i under policy π^*) $V^*(s_i)$ indicates the expected value of the travel time from a given state s_i to the goal state. The value $V^*(s_i)$ thus provides the information needed to evaluate which starting states are most preferable across the globe. In our implementation, π^* and V^* were found by a standard value iteration approach [13].

G. Adjustments for Time-Varying Winds

Although the above description focuses on time-invariant wind models, extending the solution to time-varying wind models is relatively straightforward, with adjustments necessary in two areas: the definitions of the states \mathcal{S} (and their transitions) and the enforcement of terminal conditions. Below, we first describe our approach for the fully time-varying solution and then propose a hybrid approach that is more computationally manageable.

1) *Fully Time-Varying Wind Model*: For the fully time-varying case, we extend the set of states \mathcal{S} to include not only global locations but also time, so that s_i defines a particular (x_i, y_i, z_i, t_i) , and select a time step Δt and maximum number of time steps n_t . Thus, a transition will include an increase in the number of time steps equal to $\left\lceil \frac{R_a(s_i, s_j)}{\Delta t} \right\rceil$, where $\lceil \cdot \rceil$ indicates rounding to nearest integer.

This inclusion of time-variance complicates the terminal conditions. First, presuming we have no restriction on what time we must reach the goal location, we create multiple sink states, $\{s_i \mid (x_i, y_i, z_i) = (x_g, y_g, z_g)\}$, to represent the position of the goal over all time steps. Second, we must define what happens when a transition would exceed the maximum time $n_t \Delta t$ —We create an additional “penalty” terminal state s_p and send transitions that would exceed the time limit to s_p with an arbitrarily large cost (negative reward). Thus, the optimal policy will attempt to avoid exceeding the maximum time and drive the vehicle towards the goal location in time.

2) *Diurnally Cyclical Wind Model*: While the fully time-varying method provides a valuable framework, multiplying the number of position states by n_t may result in computationally prohibitive MDP for a desired spatiotemporal resolution. An effective simplification is possible by noticing that the wind models are often highly cyclical, with wind patterns similar every day for many consecutive days. For example, on Titan, this approach assumes that the largest effects of interest are due to solar heating and tidal influences, which are synchronized since Titan’s orbital period about Saturn is identical to its rotational period.

To set up the cyclical MDP, we simply wrap the states \mathcal{S} so that the time dimension now represents “time of day” (for the celestial body of interest)




Fig. 3. Expected time-to-goal plots using optimal MDP policy for a goal location at -15° N, -100° E (labeled ‘G’), with and without horizontal actuation. Color indicates the expected time to reach the goal from each cell (in Earth days; note different scales on each subplot). A cross section at an altitude of 1000 m is shown.

rather than absolute time. As a result, there is no maximum time limit. Terminal conditions are set to be all states whose position matches the goal location, as before. This hybrid approach can be the best compromise between computation / memory limitations and the desire to capture the time complexity of the winds.

IV. RESULTS

We applied the above approach to simulations of a Montgolfier balloon in the atmosphere of Titan. Nominal wind field values were taken from a wind field model [12], with uncertainty values $\kappa = 0.0025$ for direction and $\rho = 0.20$ for magnitude. The MDP was solved via value iteration until the maximum change in value over all states changes less than $\epsilon = 1$ [Earth day] between iterations.

A. Stationary Wind Model

Figure 3 displays the expected time-to-goal (AKA time-to-go) of the vehicle from anywhere on the Titan globe, equivalent to the “value” $V(s_i)$ of each state when the optimal policy π^* is employed.

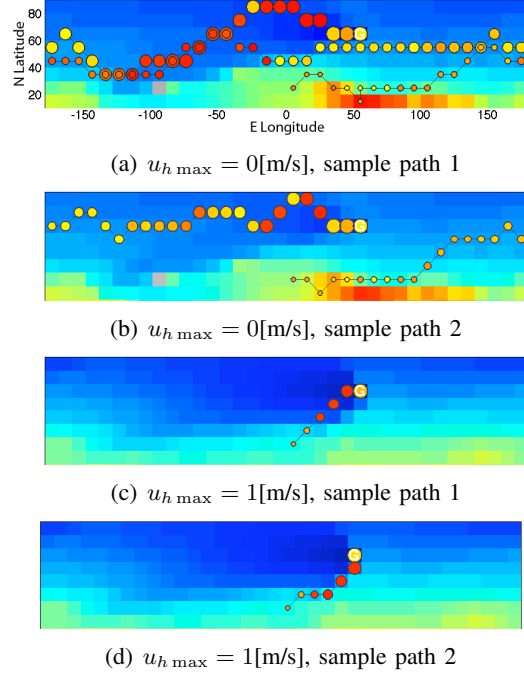


Fig. 4. Possible paths for goal location 68° N, 50° E and start location 20° N, 0° E, with and without horizontal actuation. The paths shown represent random samples of the transition probabilities generated along the optimal policy. The size of the circles along the path increases with time to show the progress of the vehicle, while the color inside the circle indicates vehicle altitude. Note that paths in subplots (a) and (b) involve wrapping the globe.

For each plot, the time-to-goal at only one horizontal slice (at $z = 1000m$) is shown. Using such plots, our method demonstrates its utility, for example, in quantifying how horizontal actuation changes the “reachability” of the goal location.

Figure 4 shows example vehicles paths superimposed on reachability plots for a different goal location. The path is randomly generated from a manually selected starting location, following the optimal policy π^* , and moving according to the state transition probabilities P . Note that since the transitions are uncertain, visiting the same location multiple times might result in different transitions, as happens in the unactuated case. The motion planning algorithm guides the vehicle to the goal, despite the uncertainty of the path and transitions.

B. Diurnally Cyclical Wind Model

Figure 5 provides the time-to-goal plots using a cyclical time wind field model for an actuated vehicle. When compared to the time-invariant versions, the time-varying results generally show a




Fig. 5. Time-to-goal plots with cyclically time-varying wind fields and $u_{h \max} = 1$ [m/s]. The superimposed white arrows represent the wind velocity vectors—there exist four vectors at each location, showing the winds at four evenly-spaced times throughout the Titan day.




Fig. 6. Histogram of time-to-goal durations, showing the distribution of sample path times under varying vehicle horizontal actuation capabilities.

shorter time-to-goal. Presumably, the vehicle can take advantage of shifting winds and be less likely to “get stuck” in areas that have small magnitude winds during part of the day.

As noted for Fig. 4, paths generated using the optimal policy π^* are subject to probabilities P and therefore uncertain. Figure 6 gives a sense of the distribution of the time taken over these paths. For this analysis, $N = 1000$ paths are generated from a selected starting location to the goal location under π^* , for each of four different horizontal actuation limits ($u_{h \max} = \{0, 0.25, 0.5, 1.0\}$ [m/s]). One can see that the times can vary dramatically, especially as the horizontal control authority of the vehicle decreases. Thus, our MDP approach is useful in assessing variability and risk across different scenarios.

V. CONCLUSION

We have presented a new method for 3D motion planning of Montgolfier balloons that incorporates the uncertainty of surrounding wind fields to plan robust, time-efficient paths. We have demonstrated

this method for planning vehicle paths in the atmosphere of Titan for varying horizontal actuation limits and both a stationary wind model and a time-varying (cyclical) model. Our stochastic approach can provide insights not accessible by deterministic methods; for example, one can evaluate variability and risk associated with different scenarios, rather than only viewing the expected outcome.

REFERENCES

- [1] L. Blackmore, “Graph-based path planning for titan balloons,” *NASA Tech Briefs*, vol. 46607, 2008.
- [2] L. Blackmore, Y. Kuwata, M. Wolf, C. Assad, N. Farthpour, C. Newman, and A. Elfes, “Path planning and global reachability for montgolfiere balloons,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2010.
- [3] Y. Kuwata, L. Blackmore, M. Wolf, C. Assad, N. Farthpour, C. Newman, and A. Elfes, “Decomposition algorithm for global reachability analysis on a time-varying graph with an application to planetary exploration,” in *IEEE/RSJ Int. Conf. on Intel. Robots and Syst. (IROS)*, 2009.
- [4] T. Das, R. Mukerjee, and J. Cameron, “Optimal trajectory planning for hot-air balloons in linear wind fields,” *AIAA Journal of Guidance, Control and Dynamics*, vol. 3, no. 26, pp. 416–424, 2003.
- [5] T. Kämpke and A. Elfes, “Optimal wind-assisted flight planning for planetary aerobots,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2004.
- [6] A. Alvarez, A. Caiti, and R. Onken, “Evolutionary path planning for autonomous underwater vehicles in a variable ocean,” *IEEE J. Ocean Engineering*, vol. 2, no. 29, pp. 418–429, 2004.
- [7] B. Garau, A. Alvarez, and G. Oliver, “Path planning of autonomous underwater vehicles in current fields with complex spatial variability: an A* approach,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2005.
- [8] C. Petres, Y. Pailhas, P. Patron, Y. Petillot, J. Evans, and D. Lane., “Path planning for autonomous underwater vehicles,” *IEEE Trans. Robotics*, vol. 2, no. 23, pp. 331–341, 2007.
- [9] D. Kruger, R. Stolkin, A. Blum, and J. Briganti, “Optimal AUV path planning for extended missions in complex, fast-flowing estuarine environments,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2007.
- [10] M. Soulignac, P. Taillibert, and M. Rueher, “Adapting the wavefront expansion in presence of strong currents,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2008.
- [11] J. Witt and M. Dunbabin, “Go with the flow: Optimal AUV path planning in coastal environments,” in *Australasian Conf. on Robotics and Automation*, 2008.
- [12] M. I. Richardson, A. D. Toigo, and C. E. Newman, “PlanetWRF: A general purpose, local to global numerical model for planetary atmospheric and climate dynamics,” *Journal of Geophysical Research*, vol. 112, 2007.
- [13] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 1998.