

Feasibility Studies on Guidance and Global Path Planning for Wind-Assisted Montgolfière in Titan

Nanaz Fathpour, *Senior Member, IEEE*, Lars Blackmore, Yoshiaki Kuwata, Chris Assad, Michael Wolf, Claire Newman, Alberto Elfes, and Kim Reh

Abstract—Recent studies have proposed the use of a hot-air (Montgolfière) balloon for possible exploration of Titan, Mars and Venus. One of NASA’s Outer Planet Flagship mission concepts is the Titan Saturn System Mission (TSSM), which would be a joint NASA-ESA partnership that plans to employ a Montgolfière along with a lake lander and an orbiter. This Montgolfière would circle Titan, investigating how Titan and Saturn operate as a system and determining how far prebiotic chemistry has developed. This paper provides a new method to analyze global path planning with balloons on Titan. The main objective of this study is to determine whether the balloon could reach a particular location of interest from a given initial position at its insertion point in the atmosphere using the wind fields on Titan. This study is the first comprehensive analysis and quantitative assessment of balloon guidance in Titan that proactively uses wind for global path planning. The paper will investigate and characterize the guidance and path-planning performance of Montgolfière balloons in Titan’s atmosphere for lower-atmosphere and surface exploration in the presence of variable wind fields using TitanWRF (Titan Weather Research and Forecasting) model. The study focuses on determining the altitude profile that a balloon could follow, using variable wind fields, in order to reach its target most quickly. Our results show that a simple unpropelled Montgolfière without horizontal actuation would be able to reach a broad array of science targets within the constraints of the wind field. The study also indicates that even a small amount of horizontal thrust allows the balloon to reach any area of interest on Titan, in a fraction of the time needed by the unpropelled balloon. The results show that using the Titan wind field allows a balloon to significantly extend its scientific reach, and that a Montgolfière (unpropelled or propelled) is a highly desirable architecture that can significantly enhance the scientific return of a future Titan mission.

I. INTRODUCTION

Titan is ideally suited for aerial vehicle exploration due to its high atmospheric density, its low gravity, and its widely spaced sites of interest. In the case of planets and moons with atmospheres, such as Titan and Venus, a number of authors have proposed the use of aerial systems that can combine extensive coverage with high-resolution data collection and in-situ science capabilities. [1], [2], [3], [4]. Recent studies have proposed the use of a Montgolfière balloon for possible exploration of Titan, Mars and Venus [2], [3], [4]. A Montgolfière, also known as a hot air balloon, maintains buoyancy by heating the atmospheric gas inside the balloon. Such a balloon can control its altitude by changing the heating rate or venting gas, but has no actuation capability in the horizontal plane. The motion of the Montgolfière in the horizontal plane is driven by the local winds. However, it may be possible for the balloon to use the difference in winds at different altitudes to guide itself to a desired location. This approach relies on predictive models of the winds on a planet or moon, such as Titan, which vary both spatially and temporally. These models are known as General Circulation Models, and in recent years much attention has been devoted to their development, for example [5], [6], [7], [8], [9], [10], [11], [12]. In our research we

have chosen to make use of the work of [5], who have developed a general purpose numerical model of planetary atmospheric and climate dynamics known as TitanWRF. This general model has been specialized to generate global three-dimensional wind models for Titan. In planning an aerial mission in Titan, it is extremely important to assess how the moon-wide wind field can be used to extend the navigation capabilities of a balloon, so that widely dispersed science targets can be visited during the mission life time, thereby enhancing the scientific return of the mission.

One of the important questions that needs to be addressed is: “What surface locations could the balloon reach from an initial location, and if so how long would it take?” One of the challenging problems in balloon guidance is its limited control authority. In this study we will develop new path planning methods for exploration balloons that proactively use the wind field for guidance purposes. The primary objective of this study is to demonstrate the advantages of this new approach over passive floating balloons and to answer mission-level questions such as how long it takes to reach potential science targets, how much fuel/energy it consumes, and how much uncertainty the planned path would involve. The resultant product will be a general purpose guidance algorithm that can be applied to exploration balloons on any moon/planet with atmosphere, including Titan, Mars, Venus, as well as the gas giants (Jupiter, Saturn, Uranus and Neptune). This study will be the first comprehensive analysis and quantitative assessment of balloon guidance that proactively uses wind for global path planning. Unlike existing studies, which are limited to linear wind fields and cannot incorporate more realistic global wind fields with sufficient resolution, our approach will address stochastic, non-linear, and time-varying wind fields. Through this wind-assisted guidance, we expect to significantly extend the range of the balloon and to reach science targets more quickly, compared to the simple floater. This approach would also enable repeated visits to a site of interest (“loitering”) that would require significantly less fuel/energy compared to the simple guidance strategy. Results could provide a breakthrough approach for exploration of any planet or moon with significant atmosphere.

It is important to note that the study did not focus on control strategies for the balloon to track trajectories. This work primarily focus on high level guidance architecture with simplified dynamics. A future work can be dedicated on studying the full dynamics of the system and control strategies that drive the balloon to desired guidance trajectories in presence of all kinds of disturbances.

In the analysis, these options have been explored: a free floater, an unpropelled balloon with altitude control, and a propelled balloon. For a given fixed Titan wind field, we have investigated what regions on the surface of Titan could be reached by the above balloon options, and how much time would be required to reach these regions. Results to date have shown that the regions reachable by an unpropelled Montgolfière depend critically on the latitude/longitude insertion point of the vehicle in Titan’s atmosphere, while a propelled balloon can essentially reach any point on the surface of Titan regardless of the initial insertion state.

The rest of the paper is organized as follows: We will first describe

Copyright 2011. All rights reserved.

N. Fathpour, L. Blackmore, Y. Kuwata, C. Assad, M. Wolf, A. Elfes, and K. Reh are with Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA. Nanaz.Fathpour@jpl.nasa.gov.

C. Newman is with Division of Geological and Planetary Sciences, California Institute of Technology, Pasadena, CA, USA.

Titan and its wind field model briefly. We will then present the problem statement and give an overview of Graph Search for Path Planning and Global Reachability. In Section IV we describe our graph generation approach for deterministic, static (time-invariant) wind fields only. In Section V we generalize the time-invariant approach to deterministic, but time varying wind fields. In Section VI we describe an approach to solving path planning problem in the presence of stochastic uncertainty in the wind field. Finally, Section VII deals with “loitering” problem where repeat visits of an interesting site is required.

II. TITAN AND ITS WIND MODEL

Titan is the largest moon of Saturn and the only moon in our solar system to have a substantial atmosphere. The atmosphere is poorly understood and obscures the surface, leading to intense speculation about Titan’s nature. The successful entry of the Huygens probe into Titan’s thick nitrogen-methane atmosphere has revealed a new world, strangely Earth-like, with methane playing the role of water, low temperature ice substituting for rock, and organic aerosols precipitated from the atmosphere taking the place of soil. Streams of liquid methane flow over the icy bedrock of a world nearly frozen in time shortly after its formation. This complex world is very similar to Earth in many ways. Studying Titan and its prebiotic chemistry can give clues to the origin of life.

The work presented in this paper is based on aerodynamic models that characterize balloon performance in Titan using TitanWRF (Titan Weather Research and Forecasting) model.

TitanWRF is a global model of Titan’s atmosphere, extending from the surface to ~ 400 km. Based on a global version of the terrestrial WRF (Weather Research and Forecasting) model, TitanWRF basically solves the primitive equations of atmospheric physics (“ $F = ma$ ” in a rotating frame plus conservation of mass and energy) discretized onto a three-dimensional grid. As described in Ref. [5], TitanWRF has been fully adapted to Titan conditions (low gravity, slow rotation rate, atmospheric composition etc.) and is typically run with ~ 5 degrees between horizontal grid points and with 55 vertical levels (spaced more closely within Titan’s troposphere, below ~ 40 km). The model includes parameterizations of turbulent mixing, sub-surface heat diffusion, surface energy and momentum exchange, and a radiative transfer scheme for Titan’s thick, hazy N_2 - CH_4 atmosphere. TitanWRF includes the seasonal and diurnal variation in incident solar energy at each location. Finally, TitanWRF includes the gravitational accelerations due to Titan’s eccentric orbit around Saturn. This produces a time-dependent change (“tide”) in the gravitational forcing which modifies horizontal wind directions and speeds, particularly in the lower atmosphere where background wind speeds are quite low. We define seasons using Titan (Saturn’s) angular position in its orbit around the Sun, i.e. its planetocentric solar longitude [L_s]. $L_s = 0$ [deg] is northern spring equinox, $L_s = 90$ [deg] northern summer solstice, and so on. Perihelion (the time at which Saturn and thus Titan are closest to the Sun) occurs at $L_s = 278$ [deg] thus during northern winter / southern summer.

III. PATH PLANNING APPROACH

In this study we present a new method for global path planning with Montgolfières that solves both the problems of path planning and reachability [13]. The new method first performs a principled simplification and decoupling of the dynamics of the Montgolfière. This enables us to perform an efficient discretization of the search space, converting the planning problem into a graph search problem. We then use Dijkstra’s algorithm [14] to calculate the minimum-time path from the start location to every possible location in the graph.

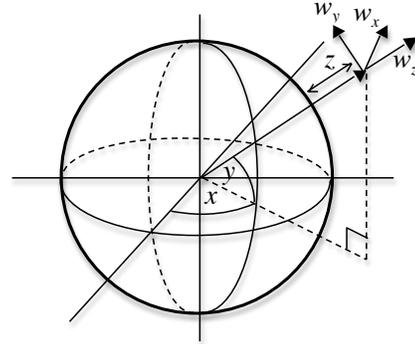


Fig. 1. Coordinate frame definitions. Degrees longitude is denoted x , degrees latitude is denoted y and z is the altitude from the surface of the planet. The velocities w_x , w_y and w_z are defined in a Cartesian coordinate frame fixed to the local surface tangent.

This solves the reachability problem, and then, for a given desired location, we can extract the minimum-time altitude profile to the goal, thereby solving the path planning problem. In this study we show planned paths and reachability maps for a number of scenarios on Titan. The proposed path planning method is a general purpose guidance algorithm that can be applied to exploration balloons on any planet or moon with atmosphere.

We assume that we have a general set of dynamic equations for the Montgolfière in a time-varying wind field of the form:

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t), t), \quad (1)$$

where $\mathbf{u}(t)$ are the heating, venting and horizontal actuation control inputs applied at time t , and $\mathbf{x}(t)$ is the state of the Montgolfière at time t . Equations in this form are derived by [15] using the thermal and dynamics balloon models of [16]. In this case, the state includes the temperature and volume of the balloon, as well as the three-dimensional position and velocity of the Montgolfière, which we denote by \mathbf{r} and $\dot{\mathbf{r}}$ respectively. The dynamics in (1) rely on a predictive model of the winds at any location \mathbf{r} and time t . We assume we have such a model in the general form $\mathbf{w}(\mathbf{r}, t)$. Here, position is defined in a spherical coordinate frame such that:

$$\mathbf{r} \triangleq \begin{bmatrix} x \\ y \\ z \end{bmatrix}, \quad (2)$$

where x is in degrees longitude, y is in degrees latitude, and z is altitude from the surface of the planet. Wind velocity is defined in a Cartesian coordinate frame fixed to the local surface tangent such that:

$$\mathbf{w}(\mathbf{r}, t) \triangleq \begin{bmatrix} w_x(\mathbf{r}, t) \\ w_y(\mathbf{r}, t) \\ w_z(\mathbf{r}, t) \end{bmatrix}, \quad (3)$$

where w_x , w_y and w_z are the velocities in the easterly, northerly and vertically upwards directions respectively. The coordinate frames are shown in Figure 1.

The path planning and reachability problem may now be stated as follows:

Problem 1: Given a Montgolfière with dynamics $f(\cdot)$ initially at location \mathbf{r}_0 , and a wind model $\mathbf{w}(\cdot)$, determine, for every possible end location \mathbf{r}_f , the minimum time to reach \mathbf{r}_f and the sequence of control inputs $\mathbf{u}(\cdot)$ that achieves this minimum.

A. Related Work

The problem of path planning for Montgolfières was studied by [15], [17]. In [15], the authors assume that wind fields vary linearly in space and are fixed in time. Given this assumption, they solve the optimal control problem to find the sequence of heating inputs that takes the balloon from its initial state to the goal. The optimal control approach is inherently limited to linear wind fields, whereas the fields predicted by global circulation models are highly nonlinear. In our previous work we extended this approach to the case where the wind field consists of a discrete, finite set of layers, each of which has a constant wind direction and magnitude [17]. In this case the wind varies with altitude, but not with horizontal location. In the present paper we provide a new method that can handle arbitrary wind fields that vary nonlinearly in the horizontal and vertical directions, as well as in time. Such a capability is necessary given that the available General Circulation Models predict a strong dependence of the winds on all three of these parameters.

A problem that has received considerable attention is that of path planning for autonomous underwater vehicles (AUVs) in current fields [18], [19], [20], [21], [22], [23]. While the AUVs are assumed to have significant horizontal actuation capabilities, the currents have a significant effect on the motion of the AUV, meaning that this problem shares some aspects of the Montgolfière path planning problem. The approach of [23] poses the path planning problem as a nonlinear optimization problem, and uses a ‘swarm’ of feasible paths to provide multiple initial guesses to an optimizer such as local random search, or simulated annealing. The performance of such optimizers is highly dependent on the quality of the initial guess, and in the Montgolfière case finding feasible paths to use as guesses is very challenging. Alternative approaches use a spatial discretization approach combined with a variety of search techniques to solve the planning problem. [18] uses a genetic algorithm to search for the optimal path in a two-dimensional field; however, this search algorithm does not guarantee convergence to a global (or even feasible) solution, and does not solve the reachability problem. [19] instead use A* graph search to guarantee that the optimal path is found; however, this work is restricted to a two-dimensional, time-invariant current field and assumes that the AUV has significant actuation capabilities in the horizontal plane. [20] use a fast-marching search technique to find optimal paths; however, these approaches can return infeasible plans if the current is stronger than the actuation capability of the AUV [22]. This is clearly the case with a Montgolfière, which has no horizontal actuation capability. [22] extends fast-marching techniques to the case where currents are stronger than actuators, however it is not clear that this extends to the case of a Montgolfière balloon; in addition [22] considers only the path planning problem and not the reachability problem.

B. Simplification of Dynamics

The Montgolfière planning problem is complex, because it takes place in a three-dimensional environment. We can, however, simplify the problem using a partial decoupling of the Montgolfière dynamics. This decoupling is based on the following assumptions:

Assumption 1: The altitude of the Montgolfière is fully controllable, subject to maximum rise and sink rates, denoted v_{rise} and v_{sink} respectively.

Assumption 2: The horizontal velocity of the Montgolfière is proportional to the local horizontal wind velocity at all times.

Assumption 1 comes from the observation that the vertical control authority of the Montgolfière is large compared to the vertical winds predicted by the global circulation models of [5]. This means we can assume that a separate altitude controller exists that issues heating and

venting commands to reject wind disturbances and to drive the Montgolfière to a desired altitude setpoint. This allows us to ignore the effects of vertical winds and the complicated thermodynamic model used by [15] to describe the vertical motion of the Montgolfière.

Assumption 2 comes from [15], who use the following relationship for the horizontal dynamics of the Montgolfière:

$$\dot{\mathbf{r}} = \gamma \cdot \mathbf{w}(\mathbf{r}, t). \quad (4)$$

where γ is a measure of the drag of the Montgolfière in the horizontal plane. These dynamics mean that, in the horizontal plane, we need *only* consider the local wind velocity. We do not need to consider any other thermal or dynamic state of the Montgolfière. This simplification, along with that given by Assumption 1 is critical in efficiently generating a discretized graph, as described in the following sections.

In addition to the model of [15], we assume that the vehicle has horizontal actuators that can generate additional velocity $\mathbf{u}(t)$ with respect to the air. Using Assumption 1 and setting $\gamma = 1$, we have:

$$\dot{\mathbf{r}}(t) = \mathbf{w}(\mathbf{r}, t) + \mathbf{u}(t), \quad (5)$$

and

$$\mathbf{u}(t) = \begin{bmatrix} u_x(t) & u_y(t) & u_z(t) \end{bmatrix}^T \quad (6)$$

$$-v_{\text{sink}} \leq u_z(t) \leq v_{\text{rise}} \quad (6)$$

$$\sqrt{u_x(t)^2 + u_y(t)^2} \leq u_{h \max}, \quad (7)$$

where $u_{h \max}$ is the maximum achievable horizontal actuation.

The key idea to solving a discretized approximation of Problem 1 is to perform an efficient discretization of the search space, converting the planning problem into a graph search problem. We then use Dijkstra’s algorithm to calculate the minimum-time path from the start location to every possible location in the graph.

The graph generation problem may be stated as follows:

Problem 2: Generate a graph \mathcal{G} consisting of a set of nodes \mathcal{S} , where each node $s_i \in \mathcal{S}$ consists of an index i and a position $\mathbf{r}(s_i)$, a set of arcs between nodes, and a weighted adjacency matrix A defined such that $A(i, j)$ is the cost to traverse the arc from s_i to s_j . $A(i, j) = \infty$ implies that no arc exists between s_i and s_j .

Given a graph \mathcal{G} , Dijkstra’s algorithm finds the minimum cost path from a start node to all other nodes in the graph. The running time of Dijkstra’s algorithm scales with the square of the number of nodes, and is hence an appealing algorithm even for large graphs. In this paper we use Dijkstra’s algorithm to find the minimum-time path between the node nearest to the initial location of the Montgolfière, denoted s_i , and all other nodes. This solves a discretized approximation to the reachability problem. Given the reachability results, the path planning problem can be solved simply, by extracting the path $p \in \mathbb{Z}^N$ corresponding to the particular target node s_l . This path consists of a sequence of node indices such that $p(1) = i$ and $p(N) = l$.

The above approach was considered for the following cases:

- Deterministic, time-invariant (static) wind fields,
- Deterministic, time-varying wind fields,
- Stochastic, time-varying wind fields.

IV. TIME INVARIANT WIND FIELDS

In this section we consider time-invariant wind fields, such that $\mathbf{w}(\mathbf{r}, t) = \mathbf{w}(\mathbf{r}, T)$ for some fixed T and all t . In generating the nodes of the graph, we choose to discretize space using a uniform grid, where adjacent nodes are separated by Δx in longitude, Δy in latitude, and Δz in altitude. The decoupling described in

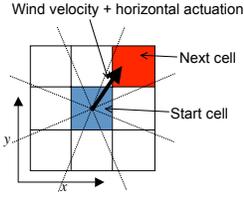


Fig. 2. Discretization in the horizontal plane.

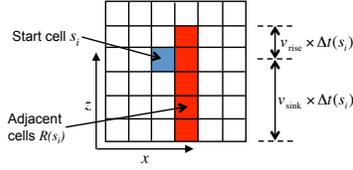


Fig. 3. Discretization in the vertical plane.

Section III-B 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. Thus each node s_i , known as the source node, corresponds to a particular location \mathbf{r} and actuation vector $\mathbf{u}(s_i)$.

The weighted adjacency matrix A is populated as follows. For every node s_i , we use the wind model to determine the local wind $\mathbf{w}(\mathbf{r}(s_i), T)$. Considering the horizontal plane first, from (4), the horizontal velocity of the Montgolfière can be written as:

$$\dot{\mathbf{r}}(s_i) = \mathbf{w}(\mathbf{r}(s_i), T) + \mathbf{u}(s_i). \quad (8)$$

We discretize the direction of the Montgolfière velocity in the horizontal plane into one of eight segments, as shown in Figure 2. Which of these segments the local wind velocity falls into determines the cell that the Montgolfière will transition to if no vertical actuation is applied, which we denote s'_i . 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'_i is given by:

$$\Delta t(s_i) = \frac{\text{dist}(\mathbf{r}(s_i), \mathbf{r}(s'_i))}{\|\mathbf{w}(\mathbf{r}(s_i), T) + \mathbf{u}(s_i)\|}. \quad (9)$$

Here, $\text{dist}(\cdot, \cdot)$ is a function that returns the cartesian distance between two points in a spherical coordinate frame, while $\|\cdot\|$ is the standard vector 2-norm.

By applying vertical actuation, cells above or below s'_i may also be reached. By Assumption 1, the vertical range of the Montgolfière depends only on the maximum rise and sink rates and the time available. Hence the maximum altitude increase and decrease possible in traveling from s_i to an adjacent cell are:

$$z_{\text{rise}}(s_i) = \frac{v_{\text{rise}}}{\Delta t(s_i)} \quad z_{\text{sink}}(s_i) = \frac{v_{\text{sink}}}{\Delta t(s_i)}. \quad (10)$$

The set of cells $\mathcal{R}(s_i)$ for which an arc exists between s_i and $s_j \in \mathcal{R}(s_i)$ is therefore given by:

$$s_j \in \mathcal{R}(s_i) \iff \begin{cases} x(s_j) = x(s'_i) \\ y(s_j) = y(s'_i) \\ z(s'_i) - z_{\text{sink}} \leq z(s_j) \leq z(s'_i) + z_{\text{rise}}(s_i). \end{cases} \quad (11)$$

We can now populate the weighted adjacency matrix A . Since Problem 1 requires us to find the minimum time to get from every node to every other node, the cost on an arc must be the time taken to traverse that arc. The matrix A is therefore populated as follows:

$$A(i, j) = \begin{cases} \Delta t(s_i) & s_j \in \mathcal{R}(s_i) \\ \infty & s_j \notin \mathcal{R}(s_i). \end{cases} \quad (12)$$

Notice that, even though we have discretized space and wind direction, we have retained time and wind magnitude as continuous variables. This ensures that arbitrarily large variations in wind magnitude can be captured without requiring an intractably large number of grid

cells. This is essential, since the PlanetWRF wind model predicts winds that vary in magnitude over several orders of magnitude.

In what follows, we define a Titan [solar] day as the time taken for Titan to rotate once (roughly 16 Earth days). Titan is tidally locked which means that it keeps roughly the same face to Saturn at all times, so one Titan day is also the time taken for it to orbit Saturn. Titan is in orbit around Saturn so a Titan year [~ 673 Titan days] is the same as a Saturn year and lasts roughly 30 Earth years. We also define a Titan hour as $1/24$ th of a Titan day.

We describe seasons using Titan (Saturn's) angular position in its orbit around the Sun, i.e. using its planetocentric solar longitude L_s . $L_s = 0$ is the northern spring equinox, $L_s = 90$ the northern summer solstice, and so on. Perihelion (the time at which Saturn and thus Titan are closest to the Sun) occurs at $L_s = 278$, during the northern winter / southern summer.

a Titan [solar] day as the time taken for Titan to rotate once (roughly 16 Earth days). Titan is tidally locked which means that it keeps roughly the same face to Saturn at all times, so one Titan day is also the time taken for it to orbit Saturn. Titan is in orbit around Saturn so a Titan year [~ 673 Titan days] is the same as a Saturn year and lasts roughly 30 Earth years. We also define a Titan hour as $1/24$ th of a Titan day.

A. Results

In this section, we give results demonstrating the graph search algorithm in the time-invariant wind field case. For the following results, winds at time $L_s = 90$ [deg] were used, corresponding to summer solstice. The following parameters were used: Figures 4(a), 4(b), 4(c), and 4(d) show reachability maps for Titan with no horizontal actuation (only vertical control) at 4 different starting locations: the north pole, 25°N , the equator and 45°S . The color in a cell shows how much time it takes for the optimal path from the starting location 'S' to all cells across the moon. The color white on the figure shows that the cell is not reachable for that particular starting location. When there is no horizontal actuation, with only the ambient winds and altitude control available to move the balloon, it may take many days to get to a destination of interest, sometimes up to 200 days, and sometimes the location is simply not reachable. Figures 5(a), 5(b), 5(c), and 5(d) show global reachability maps for static wind fields with a horizontal actuation of 1 m/s. As seen in Fig. 5, unlike the case where there is no horizontal actuation, the entire moon is reachable from all 4 starting locations when there is horizontal actuation of 1 m/s. It is also interesting to note that, for this actuation level, the time taken to reach a given cell is roughly a factor of 6 shorter.

V. TIME VARYING WIND FIELDS

This section presents the approach and results with a time varying wind field. Because the wind at a given location differs depending on when the balloon reaches it, a temporal as well as spatial discretization is performed. One challenge is that the size of the graph becomes excessive when considering the global reachability problem for planetary exploration. This is because for each dimension (e.g., latitude, longitude, altitude, and time) the size of the graph increases exponentially, and also because of the resolution required to capture the various time scales of the wind field while planning over the long mission duration. In such cases, the memory requirement for graph construction and the search algorithm becomes fairly significant.

To address this issue, a decomposition algorithm for reachability analysis of a time-varying graph has been developed. Because the balloon only moves in the positive direction in time, the adjacency matrix of the graph can be represented with an upper block-triangular

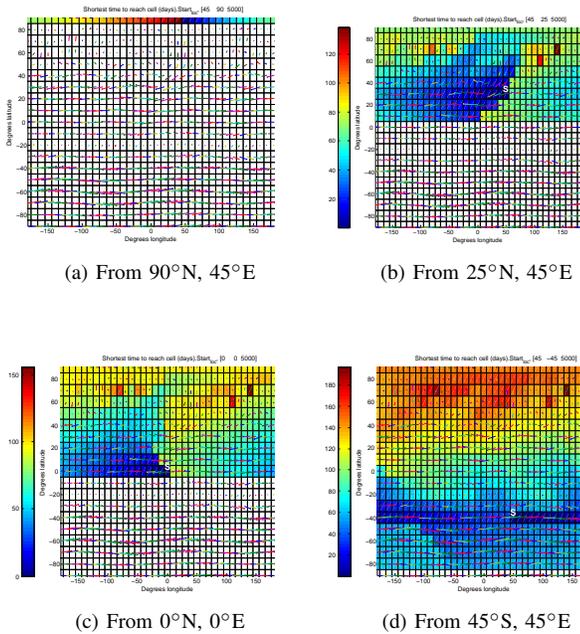


Fig. 4. Global reachability map with no horizontal actuation and static wind field

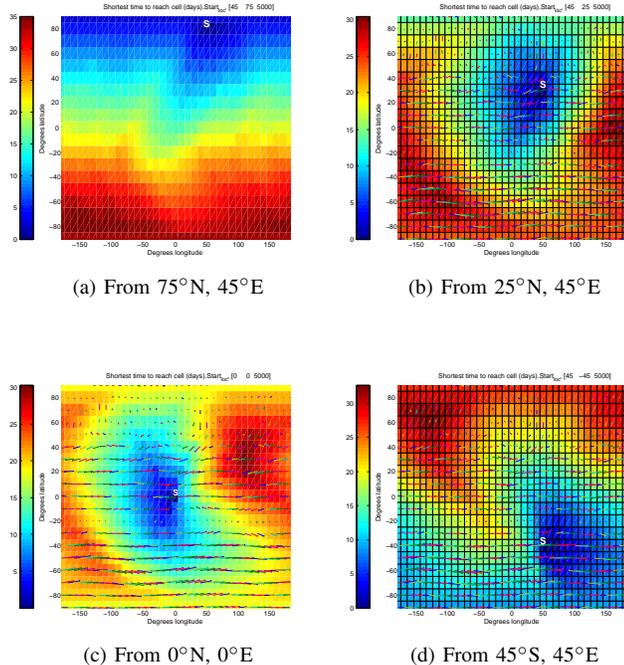


Fig. 5. Global reachability map with with 1 m/s horizontal actuation and static wind field

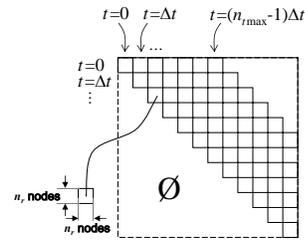


Fig. 6. Adjacency matrix.

matrix, and this upper block-triangular structure can be exploited to decompose a graph search problem. Instead of solving a single large problem, our algorithm solves subproblems sequentially whose size is much smaller than the original problem. The new approach therefore consumes much smaller amounts of memory, which also helps speed up the overall computation when the computing resource has a limited physical memory compared to the problem size.

A. Problem Statement – Decomposition Algorithm

Let n_s denote the number of starting nodes. Each node s_i is a function of the position (x_i, y_i, z_i) and the time t_i . A uniform grid is used to represent the world, and we let n_x , n_y , n_z , and n_t respectively denote the number of cells in x , y , z , and the time axes. We define $n_r \triangleq n_x n_y n_z$ as the number of cell positions in the environment, and $N \triangleq n_r n_t$ as the total number of cells. The problem statement for the graph search is to find the shortest paths from n_s starting nodes s_i to all n_r 3D locations. Once the graph is constructed, it is straightforward to apply Dijkstra’s algorithm to find the shortest paths. A matrix of size $n_s \times n_r$ can represent the minimum time of arrival at each 3D location from each starting node. Let this matrix be denoted by C^* . Its (p, q) element C_{pq}^* stores the time of travel from the p^{th} starting location to the q^{th} 3D location, and is set to be ∞ if there exists no such trajectory.

The primary components of the time-varying wind field of Titan include (a) seasonal changes in the global atmospheric circulation driven by the changing solar forcing as Saturn orbits the Sun (period of 1 Titan year, ~ 30 Earth years), (b) tidal effects driven by Titan’s eccentric orbit around Saturn (period of 1 Titan day, ~ 16 Earth days), and (c) small- and large-scale waves occurring naturally in Titan’s atmosphere on a range of time-scales. A potential future mission to explore Titan would have a mission duration of 6–12 Earth months, and capturing the time-varying wind field of various time scales would require a large n_t to be used in the graph.

Figure 6 shows the adjacency matrix of the graph, which is of size $N \times N$. Because the balloon only moves in the positive direction in time, it can be represented with an upper block-triangular matrix, by ordering the nodes in terms of their time t_i . Each block contains a snapshot of the 3D world, whose size is $n_r \times n_r$. With a large n_t , the memory requirement for graph construction and Dijkstra’s algorithm becomes significant. However, this upper block-triangular structure of the adjacency matrix can be exploited to decompose the problem into several smaller subproblems that use far less memory.

We assume the balloon starts at $t = 0$, so that all the starting nodes are in the $(1, 1)$ block. This approach splits the weighted adjacency matrix A into several submatrices M_k , $k = 1, \dots, k_{\text{max}}$, and repeatedly applies Dijkstra’s algorithm to each submatrix. The result of each subproblem can be represented by a matrix whose size is much smaller than the submatrix used in the subproblem. The next subproblem is formed by appending to its submatrix the small matrix obtained in the previous subproblem. This process is repeated until all the submatrices are processed or all the shortest paths from the

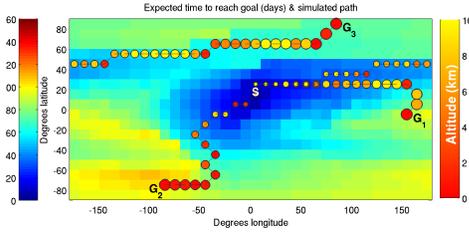


Fig. 7. The trajectories from a start to 3 different goals (G_1 , G_2 , G_3). Trajectories are marked with circles whose color represents the altitude. The corresponding color bar is given at the right. The left color bar shows the time to reach each cell at an altitude of 250 m (in Earth days).

starting nodes to the 3D locations are found. A detailed discussion of the algorithm is available in [24].

B. Simulation Results

Figure 7 shows several trajectories from a start location of (5°S , 5°E) to 3 different goals $G_1(5^\circ\text{S}, 155^\circ\text{E})$, $G_2(75^\circ\text{S}, 85^\circ\text{W})$, and $G_3(85^\circ\text{N}, 85^\circ\text{E})$. The size of the circles represents the elapsed time from start, and the color of the circles represents the altitude of the trajectory. Note that because of the nonlinear and time-varying wind field, the minimum-time trajectories involve several vertical actuation steps and are far from straight lines.

Figure 8 shows the percentage of the areas of Titan's surface that would be reachable in a given time with 4 different horizontal actuation levels (0.0 m/s, 0.25 m/s, 0.50 m/s, and 1.0 m/s). For example, 50% on the y axis means that 50% of the points on Titan's surface could be reached in a given time if set as a destination; it does not mean that the balloon would sweep over 50% of the Titan surface in this time. The plots have 24 different starting locations with the following combinations of longitude and latitude.

- 4 longitudes (175°W , 85°W , 5°E , 95°E)
- 6 latitudes (85°S , 45°S , 15°S , 15°N , 45°N , 85°N)

Because the wind field varies far more with latitude than with longitude, the lines corresponding to the same latitude are plotted with the same color and show similar trends: for example, a balloon starting at 15°S (shown in red) can initially reach only limited areas, but the reachable area grows rapidly after a few months, whereas for a balloon starting near the southern pole at 85°S (shown in blue) the reachable area does not grow as fast with time.

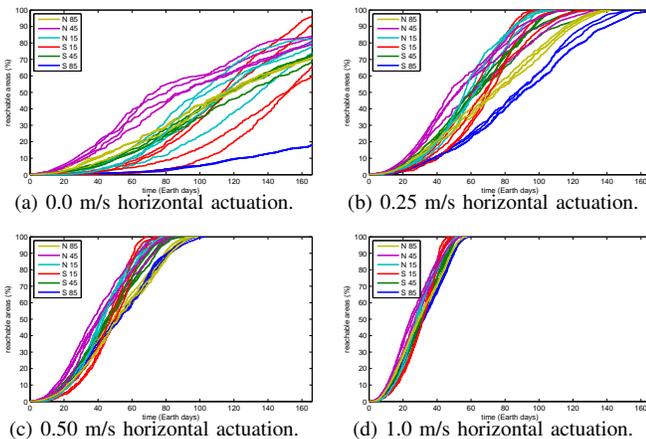


Fig. 8. Percentage of reachable area given a starting location. $L_s = 0$ [deg].

It is clear that both the reachable area and the time of arrival improves dramatically by allowing more horizontal actuation. Also

note that with no horizontal actuation (e.g., 0.0 m/s), some of the area is not reachable.

VI. UNCERTAIN WIND FIELDS

In practical applications, the wind vector field is not known exactly and may deviate significantly from the wind velocities estimated by the TitanWRF model. This section describes how the problem framework may also be adapted for *uncertain* wind fields for probabilistic solution methods. Our primary contribution in this scenario is in appropriately formulating the problem as a Markov Decision Process (MDP), for which a number of solution methods may be applied.

A. Approach

With uncertain wind fields, the transition from each state is no longer deterministically specified by the wind model as it was in the graph solution above. From a given state s_i , the next state, dubbed s'_i , may be considered a random variable and a corresponding probability distribution for s'_i can be constructed over all horizontally adjacent cells. Given these transition probabilities from all states, we wish to select the actions (horizontal and vertical actuation of the Montgolfière) that minimizes time-to-goal. This optimality problem is thus naturally posed as a Markov Decision Process (MDP) $(\mathcal{S}, \mathcal{A}, P, R)$, where: \mathcal{S} represents the set of possible states, i.e., (x, y, z) locations, of the Montgolfière; these are the same discretized longitude / latitude / altitude cell positions as defined for the graph \mathcal{G} in Section IV. \mathcal{A} is the set of actions available from each state, namely the combination of horizontal and vertical actuation options. P gives the transition probabilities between states under a given action: given current state s_i and action a , $P_a(s_i, s_j) = \Pr[s'_i = s_j | s_i, a]$. R defines the expected immediate reward for each transition and each action a . Since we wish to optimize for least travel time, we use a reward of negative travel time (or, equivalently, a cost of positive travel time).

In order to model the wind uncertainty, we choose to decompose it into direction and magnitude components and assign independent distributions, each with an expected value equal to the field model value. Let $\tilde{w}(r_i, t)$ denote the (random) wind velocity at position r_i at time t and $w(r_i, t)$ denote the velocity given by the wind field (TitanWRF) model. The direction and magnitude of $\tilde{w}(r_i, t)$ are represented respectively by θ_i and w_i . The wind direction θ_i is of primary importance as it principally determines the next horizontal cell of the vehicle. To model the uncertainty in θ_i , we employ a von Mises distribution, an analogue of a Gaussian distribution on the circle. The von Mises distribution is defined by

$$f_{\text{VM}}(\theta_i | \phi_i, \kappa) = \frac{\exp(\kappa \cos(\theta_i - \phi_i))}{2\pi I_0(\kappa)}, \quad (13)$$

where ϕ_i is the mean, κ is a concentration parameter, and $I_0(\kappa)$ is the modified Bessel function of the first kind of order 0. We set ϕ_i to the angle given by the wind field model at position r_i , $\phi_i = \angle w(r_i, t)$, and κ to a constant chosen by the user. Note that κ could also be chosen to vary by position if desired. The wind magnitude w_i uncertainty is modeled as Gaussian, with distribution

$$f_{\text{N}}(w_i | \mu_i, \sigma_i^2) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(w_i - \mu_i)^2}{2\sigma_i^2}\right). \quad (14)$$

Here, we set the mean to the value given by the wind field model, $\mu_i = \|w(r_i, t)\|$, and the standard deviation set as proportional to the magnitude, $\sigma_i = \rho \|w(r_i, t)\|$. The value of ρ is chosen by the user.

Next we define the transition probabilities $P_a(s_i, s'_i)$ that govern what state s'_i is entered after executing each action a from each state

s_i . Note that all uncertainty is limited to the horizontal direction; the vehicle will transition to a selected altitude with probability 1, provided that the action is allowable (i.e., that the altitude is reachable based on vertical actuation limits). Thus, for the discussion below, only the projections onto the horizontal plane are relevant.

Let us first address the situation where there is no horizontal vehicle actuation. In this case, the wind direction wholly determines the next state's horizontal position, and the action space consists only of selection of the next altitude from those possible under the constraints of v_{rise} and v_{sink} . So, for each horizontal slice within the allowable altitudes dictated by a , the probability of the wind forcing the Montgolfière from state s_i to a horizontally adjacent state s_j is

$$P_a(s_i, s_j) = \int_{\theta_{ij}}^{\theta_{ij} + \frac{\pi}{4}} f_{\text{VM}}(\theta | \phi_i, \kappa) d\theta, \quad (15)$$

where θ_{ij} is the smaller bordering angle of the one-eighth circular sector pointing from s_i to s_j . Note that this integration must be accomplished numerically, as no closed-form solution exists.

Next, let us address the use of horizontal vehicle actuation. As before, horizontal actuation is accounted for via an additive velocity on top of the wind velocity. In the stochastic case, we still treat the actuation as known, but the direction and magnitude of the wind are now uncertain. We are interested in the transition governed by the resultant velocity, $\tilde{\mathbf{v}}_{i,a} = \tilde{\mathbf{w}}(\mathbf{r}_i, t) + \mathbf{u}_a$.

A Monte Carlo method is used to determine the probabilities P of and expected rewards R from the resultant vector for each action, as follows. First, a set Θ_0 of N sample points are drawn from the von Mises distribution $f_{\text{VM}}(\theta | 0, \kappa)$ 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 ϕ_i , $\Theta_i = \Theta_0 + \phi_i$, and the wind magnitude samples are adjusted by $W_i = \sigma_i W_0 + \mu_i$ (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:

$$V_{i,a} = W_i \begin{bmatrix} \cos \Theta_i \\ \sin \Theta_i \end{bmatrix} + \mathbf{u}_a, \quad (16)$$

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 :

$$P_a(s_i, s_j) = \frac{1}{N} |V_{ij,a}|, \quad (17)$$

where $|\cdot|$ represents cardinality and:

$$V_{ij,a} = \left\{ \mathbf{v}_{i,a}^n \mid \angle \mathbf{v}_{i,a}^n \in [\theta_{ij}, \theta_{ij} + \frac{\pi}{4}] \right\}. \quad (18)$$

B. Action Space and Rewards

Minimum travel time is the goal of our MDP solution and thus travel time is still appropriate to use as a transition cost. 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:

$$\langle v_{i,a} \rangle = \frac{1}{N} \sum_{n=1}^N \|\mathbf{v}_{i,a}^n\|. \quad (19)$$

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

$$R_a(s_i, s_j) = -\frac{d_{ij}}{\langle v_{i,a} \rangle}, \quad (20)$$

where $d_{ij} = \sqrt{(x_j - x_i)^2 + (y_j - y_i)^2}$ is the horizontal distance from \mathbf{r}_i to \mathbf{r}_j . Note that the reward is negative to enforce elapsed time as a ‘‘cost’’. 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).

C. Solving for Minimum Time-to-Goal

To complete the problem setup, we define a goal location for the Montgolfière and create a sink state at this location. All transitions leaving the sink state s_g have zero probability—it is the only state where the the vehicle can remain stationary (or, more precisely, once in s_g , the vehicle returns to s_g with probability 1 and zero reward). 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^* . Note that this differs conceptually from the deterministic graph search method, which instead defines the initial state and determines the cost to reach each other cell from this location.

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 in our ‘‘reachability’’ plots, which, for the MDP case, now represent the expected time-to-goal (rather than time-from-start as in the deterministic case).

D. Results

We applied the above approach for planning paths in uncertain wind fields to simulations of a Montgolfière balloon in the atmosphere of Titan. Nominal wind field values were again taken from the TitanWRF model [5]. 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.

Figure 9 displays the expected time-to-goal (AKA time-to-go) of the vehicle from anywhere on the Titan globe. The expected time-to-goal is equivalent to the ‘‘value’’ $V(s_i)$ of the state. For each plot, only one horizontal slice (at $z = 1000m$) is shown. Note that these figures are different from ‘‘reachability’’ plots in previous section, which plot the time to reach the cell from a given start position; here, we are given the goal location and show how long we expect to take to reach it. Also, the time is an expected value, since the actual time is uncertain. The path is randomly generated from a manually selected starting location (20° N, 0° E), following an optimal policy, and moving according to the state transition probabilities P . The size of the circle along the path increases with time to show the progress of the vehicle, while the color inside the circle indicates vehicle altitude. Note that since the transitions are uncertain, visiting the same location multiple times may result in different transitions, as happens in the unactuated case.

It is clear that the use of actuation greatly aids the vehicle in reaching the goal. (Note that the color bars are plotted on different scales, so the colors cannot be compared directly.) Especially in the case of the goal nearest the southern pole, the goal cannot be reached from several locations without the aid of actuation. A more detailed description of algorithms is discussed in [25].

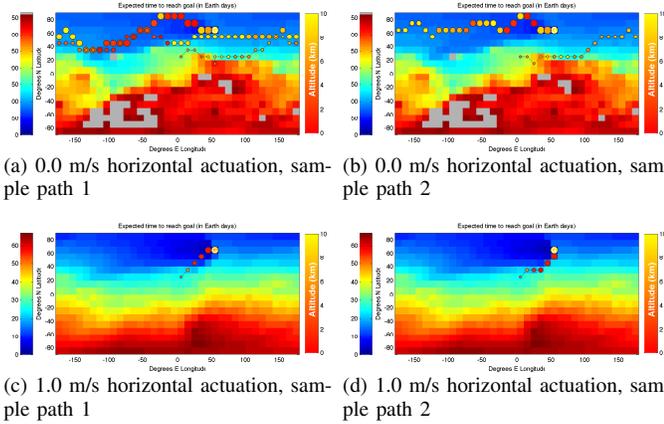


Fig. 9. Possible paths for goal location 68° N, 50° E and start location 20° N, 0° E, with and without horizontal actuation. The paths shown represent random samples of the transition probabilities generated along the optimal policy. Note that in the figure without actuation, the vehicle revisits some cells multiple times.

VII. LOITERING

The methods described in the preceding chapters enable path planning for a Montgolfière to reach the location of a desired science target, either through vertical actuation alone or combined with limited horizontal actuation. Once arrived at the target, however, the limitation or lack of horizontal actuation presents a new problem:

- Given the vehicle has reached its target location cell, how long can it remain in that cell to perform science observations?

One obvious solution – dropping an anchor – may not be practical for the Titan mission from altitude and mass considerations. Instead, we consider again using the wind field predictions. If the altitude is not constrained, then we can position the Montgolfière at the altitude of smallest horizontal wind magnitude to maximize the time spent within the target cell. Eventually, the Montgolfière may drift out of the science target cell. If the mission requires further observations of the target cell, it would be useful to determine if the Montgolfière could repeatedly return to the same cell given the predicted wind fields; i.e., find the shortest time or distance cycle path that includes the science target cell. Our reachability graphs can be used both to calculate hover times and to plan the shortest possible path to return to the target.

A. Minimum Wind Speeds and Maximum Hover Times

Wind speed extrema were determined by linearly interpolating the horizontal wind velocity components to 50 m intervals in altitude for each latitude and longitude grid point, computing the vector magnitude, and sorting by magnitude within each vertical column of cells. Figure 10 illustrates predicted minimum speeds for a stationary Titan wind field (at time $L_s = 0$), at the model’s highest spatial resolution ($5 \text{ deg} \times 5 \text{ deg} \times 500 \text{ m}$). The minimum horizontal speeds range from near zero to over 2 m/s. At this time of year, minimum wind speeds in the southern hemisphere are generally lower than in the northern hemisphere, and range up to around 1 m/s. For comparison purposes, the maximum wind speeds for the same point in time, shown in Figure 11, range to over 5 m/s.

To help illustrate the benefit of adding horizontal actuation, Figure 12 shows a coarser contoured version of the minimum speed data in Figure 10, binned to 0.25 m/s intervals. In Table I, the area of each colored contour was summed cumulatively and divided by the total area, to estimate the fraction of cells at each actuation

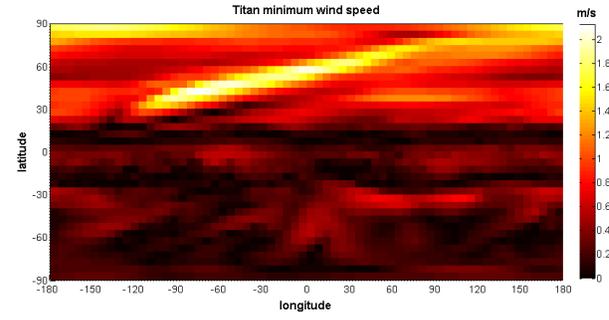


Fig. 10. Titan wind field minimum magnitudes at each latitude and longitude, over altitude up to 10 km. Grid resolution of $5 \text{ deg} \times 5 \text{ deg} \times 500 \text{ m}$, with static wind field at time $L_s = 0$ [deg].

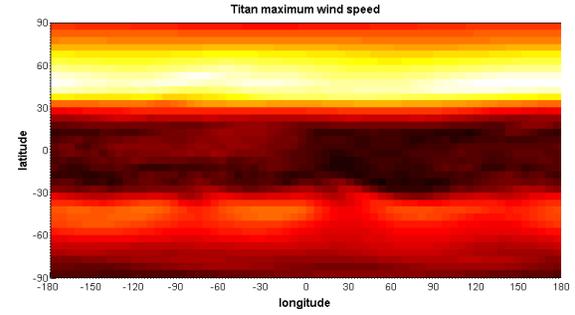


Fig. 11. Titan wind field maximum magnitudes at each latitude and longitude, over altitude up to 10 km. Grid resolution of $5 \text{ deg} \times 5 \text{ deg} \times 500 \text{ m}$, with static wind field at time $L_s = 0$ [deg].

level in which the vehicle can hover indefinitely. For example, with onboard actuation capability of 0.5 m/s, the vehicle would be able to hover indefinitely over greater than half of the total number of surface locations.

Without horizontal actuation, the smallest horizontal wind speed is usually found at the altitude corresponding to the longest duration adjacency arc in $A(i, j)$. However, $A(i, j)$ also depends on cell shape and size, so it is a better indicator than wind speed of the actual loitering time within a given cell. Using $\max A(i, j)$ at each latitude and longitude, the predicted maximum hover times for the stationary wind field are shown in Figure 13, for the case of no horizontal actuation. Calculating the maximum hover time is straightforward to extend to the nonstationary or stochastic wind field models: substitute the new search methods and weighted adjacency matrices as defined in the preceding chapters.

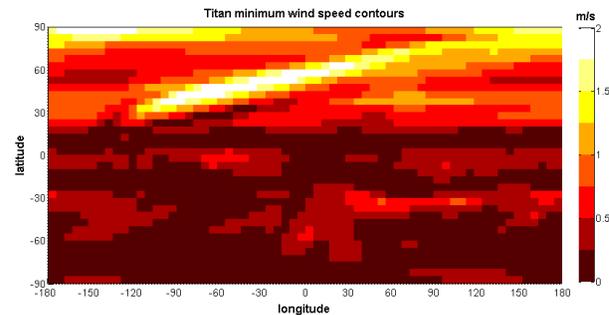


Fig. 12. Titan wind field minimum magnitudes at each latitude and longitude, over altitude up to 10 km, binned to 0.25 m/s intervals. Grid resolution of $5 \text{ deg} \times 5 \text{ deg} \times 500 \text{ m}$, with static wind field at time $L_s = 0$ [deg].

TABLE I
FRACTION OF GRID LOCATIONS ALLOWING INDEFINITE HOVER.

Actuation level (m/s)	0	0.50	0.75	1.00	1.50	2.00
Indefinite hover fraction	0	0.623	0.760	0.867	0.963	0.997

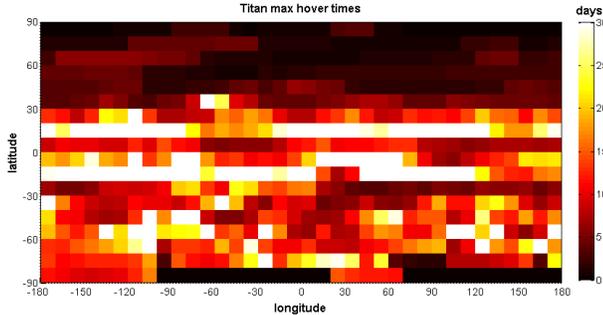


Fig. 13. Titan wind field maximum hover times within each latitude and longitude cell, for the case of no horizontal actuation. Grid resolution of $10 \text{ deg} \times 10 \text{ deg} \times 500 \text{ m}$, altitude up to 10 km, static wind field at time $L_s = 0$ [deg].

B. Minimum Closed Cycle in a Stationary Wind Field

After the Montgolfière drifts out of a science target cell, we can use our graph search tools to find the shortest time or distance cycle path to return to the target cell. For the general case of locating cycles in a directed graph, there are reported methods available. However, for our case we know the initial science target cell and can take advantage of our precomputed weighted adjacency matrix, which means that a more straightforward approach is feasible. Given a desired science target cell, s_i , and a weighted adjacency matrix $A(i, j)$ for a wind field grid, the minimum-time closed-cycle to return to cell s_i after drifting out can be determined as follows:

- 1) Determine $R(s_i)$, the set of all cells s_j adjacent to s_i , using the methods outlined in Section IV on Graph Construction.
- 2) For each reachable cell s_j in $R(s_i)$, run the shortest path algorithm (in Section IV) from s_j with s_i as the end goal.

If the science target requires us to be at a particular altitude, then the altitude of s_i is fixed. In this case, we must perform a path search from each of the m adjacent nodes of s_i , where m is the cardinality of $R(s_i)$, to s_i . For example, in the deterministic wind field without horizontal actuation, there is only one possible horizontal direction from s_i , but L possible adjacent altitudes, so $m = L$, and m path searches must be run. If, instead, the science target can be observed from any altitude, then we must also perform path searches for all of the L possible starting altitudes. Due to the overlap of the adjacent nodes, in the worst case we have to perform only $8L$ path searches. In practice, the computation is more tractable because not all 8 directions will be reachable from the science target location, and the set of adjacent cells for a given starting altitude will often overlap with adjacent cells of other starting altitudes.

We can increase the graph search efficiency by reducing the graph size with distance or time constraints. Assuming the Montgolfière must return to s_i within a desired time T_d , or travel no farther from s_i than distance D , then we can first determine the reachability map from s_i as the starting point, and remove all cells from the graph that are farther than T_d in time, or D in distance, before searching the remaining graph for the shortest possible cycles.

VIII. CONCLUSION

In this paper we presented a new path planning approach for local and global path planning of propelled and un-propelled hot-air

(Montgolfière-type) balloons operating in Titan’s atmosphere. This planning approach enables a Montgolfière to exploit variations in the wind field at different altitudes to achieve a desired horizontal motion. The approach was extended to the time-varying wind fields, which significantly increased the computational complexity. A new decomposition method has been developed that requires less memory and computation time compared to the original approach. To account for the uncertainties in the wind fields, a stochastic formulation was also presented. It was shown that by using a spatial and temporal discretization combined with existing graph search techniques we can determine the altitude profiles that reach a target in minimum time, can determine the set of all reachable targets from a given start location, and can provide a reachability analysis for the entire moon. This new approach was applied to different scenarios and the performance of a free floater, an unpropelled balloon with altitude control, and a propelled balloon was examined. It was shown that given a small amount of horizontal actuation, the balloon will be capable of reaching more targets on Titan given a limited mission duration, and of reaching all targets more quickly than a passive balloon. Moreover, in the cases where there is a need to repeatedly visit (or ‘loiter’ over) a science target, it was shown that the cycle times dramatically decrease when an actuated balloon is used. Therefore, use of actuation clearly helps the vehicle to reach its target faster and to remain there if desired.

Primarily, the focus of this study has been on trade studies and quantitative assessments of balloon guidance performance on Titan based on approximated dynamics of the balloon. These guidance trajectories are as valid as the wind field models, and using a finer grid for the wind fields one can generate finer guidance trajectories. It is important to note that the study did not focus on the control problem for the balloon to track these trajectories. A future work can be dedicated on studying the full dynamics of the system and control strategies that could drive the balloon to desired guidance trajectories in presence of all kinds of disturbances.

The proposed path planning method provides a quantitative guidance performance, science return assessment, and sensitivity of mission performance to key parameters, which will have significant impact on design choices and trade studies for cost effective future balloon missions. The reachability analysis is critical in performing trade studies to determine the right aerial system to use, and where and when such a system should be deployed for maximum science return. This path planning tool is a general purpose guidance algorithm that can be applied to exploration balloons on any moon/planet with atmosphere, including Titan, Mars, Venus and the gas giants, provided that wind field models are available.

ACKNOWLEDGMENTS

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [1] A. Elfes, J. L. Hall, J. F. Montgomery, C. F. Bergh, and B. A. Dudik, “Towards a substantially autonomous aerobot for exploration of Titan,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2004. 1
- [2] J. Jones, “Montgolfiere balloon missions from Mars and Titan,” in *Proceedings of the 3rd International Planetary Probe Workshop*, 2005. 1
- [3] J. O. Elliott, K. Reh, and T. Spilker, “Concept for Titan exploration using a radioisotopically heated montgolfiere,” in *Proceedings of the IEEE Aerospace Conference*, 2007. 1
- [4] D. Fairbrother, “Development of planetary balloons,” in *Proceedings of the NASA Space Technology Conference*, 2007. 1

- [5] 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. [1](#), [2](#), [3](#), [7](#)
- [6] R. M. Haberle, M. M. Joshi, J. R. Murphy, J. R. Barnes, J. T. Schofield, G. Wilson, M. Lopez-Valverde, J. L. Hollingsworth, A. F. C. Bridge, and J. Schaeffer, "General circulation model simulations of the mars pathfinder atmospheric structure investigation/meteorology data," *Journal of Geophysical Research*, vol. 104, pp. 8597–8974, 1999. [1](#)
- [7] S. R. Lewis, M. Collins, P. L. Read, F. Forget, F. Hourdin, R. Fournier, C. Hourdin, O. Talagrand, and J. P. Huot, "A climate database for mars," *Journal of Geophysical Research*, vol. 104, pp. 24 177–24 194, 1999. [1](#)
- [8] Y. O. Takahashi, H. Fujiwara, H. Fukunishi, M. Odaka, and Y. Hayashi, "Topographically induced north-south asymmetry of the meridional circulation in the martian atmosphere," *Journal of Geophysical Research*, vol. 108, 2003. [1](#)
- [9] T. Tokano, F. M. Neubauer, M. Laube, and C. P. McKay, "Seasonal variation of Titans atmospheric structure simulated by a general circulation model," *Planet. Space Sci.*, vol. 47, pp. 493–520, 1999. [1](#)
- [10] D. F. Luz, F. Hourdin, P. Rannou, and S. Lebonnois, "Latitudinal transport by barotropic waves in Titans stratosphere. ii. results from a coupled dynamics-microphysics-photochemistry gcm," *Icarus*, vol. 166, no. 2, pp. 343–358, 2003. [1](#)
- [11] M. Yamamoto and M. Takahashi, "The fully developed superrotation simulated by a general circulation model of a Venus-like atmosphere," *J. Atmos. Sci.*, vol. 60, no. 3, pp. 561–574, 2003. [1](#)
- [12] C. Lee, S. R. Lewis, and P. L. Read, "A numerical model of the atmosphere of Venus," *Adv. Space Res.*, vol. 11, no. 36, pp. 2142–2145, 2005. [1](#)
- [13] L. Blackmore, Y. Kuwata, M. Wolf, C. Assad, N. Fathpour, C. Newman, and A. Elfes, "Path Planning and Global Reachability for Planetary Exploration with Montgolfiere Balloons," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2010. [2](#)
- [14] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, pp. 269–271, 1959. [2](#)
- [15] 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. [2](#), [3](#)
- [16] L. A. Carlson and W. J. Horn, "New thermal and trajectory model for high-altitude balloons," *Journal of Aircraft*, vol. 20, no. 6, pp. 500–507, 1983. [2](#)
- [17] T. Kämpke and A. Elfes, "Optimal wind-assisted flight planning for planetary aerobots," in *Proceedings of the International Conference on Robotics and Automation*, 2004. [3](#)
- [18] A. Alvarez, A. Caiti, and R. Onken, "Evolutionary path planning for autonomous underwater vehicles in a variable ocean," *IEEE Journal of Ocean Engineering*, vol. 2, no. 29, pp. 418–429, 2004. [3](#)
- [19] B. Garau, A. Alvarez, and G. Oliver, "Path planning of autonomous underwater vehicles in current fields with complex spatial variability: an A* approach," in *Proceedings of the International Conference on Robotics and Automation*, 2005. [3](#)
- [20] C. Petres, Y. Pailhas, P. Patron, Y. Petillot, J. Evans, and D. Lane., "Path planning for autonomous underwater vehicles," *IEEE Transactions on Robotics*, vol. 2, no. 23, pp. 331–341, 2007. [3](#)
- [21] D. Kruger, R. Stolkin, A. Blum, and J. Briganti, "Optimal AUV path planning for extended missions in complex, fast-flowing estuarine environments," in *Proceedings of the International Conference on Robotics and Automation*, 2007. [3](#)
- [22] M. Soullignac, P. Taillibert, and M. Rueher, "Adapting the wavefront expansion in presence of strong currents," in *Proceedings of the International Conference on Robotics and Automation*, 2008. [3](#)
- [23] J. Witt and M. Dunbabin, "Go with the Flow: Optimal AUV Path Planning in Coastal Environments," in *Proceedings of the Australasian Conference on Robotics and Automation*, 2008. [3](#)
- [24] Y. Kuwata, L. Blackmore, M. Wolf, N. Fathpour, C. Newman, and A. Elfes, "Decomposition Algorithm for Global Reachability Analysis on a Time-Varying Graph with an Application to Planetary Exploration," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2009. [6](#)
- [25] M. Wolf, L. Blackmore, Y. Kuwata, N. Fathpour, C. Newman, and A. Elfes, "Probabilistic Motion Planning of Balloons in Strong, Uncertain Wind Fields," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2010. [7](#)