Optimal Policies for Autonomous Navigation in Strong Currents Using Fast Marching Trees

Bernardo Martinez Rocamora Jr.^{1*} and Guilherme A. S. Pereira¹

¹Department of Mechanical, Materials and Aerospace Engineering, Benjamin M. Statler School of Engineering and Mineral Sciences, West Virginia University, 1306 Evansdale Dr., Morgantown, 26506, WV, USA.

> *Corresponding author(s). E-mail(s): bernardo.rocamora@gmail.com; Contributing authors: guilherme.pereira@mail.wvu.edu;

Abstract

Several applications require that unmanned vehicles, such as UAVs and AUVs, navigate environmental flows. While the flow can improve the vehicle's efficiency when directed towards the goal, it may also cause feasibility problems when it is against the desired motion and is too strong to be counteracted by the vehicle. This paper proposes the flow-aware fast marching tree algorithm (FlowFMT*) to solve the optimal motion planning problem in generic three-dimensional flows. Our method creates either an optimal path from start to goal or, with a few modifications, a vector field-based policy that guides the vehicle from anywhere in its workspace to the goal. The basic idea of the proposed method is to replace the original neighborhood set used by FMT* with two sets that consider the reachability from/to each sampled position in the space. The new neighborhood sets are computed considering the flow and the maximum speed of the vehicle. Numerical results that compare our methods with the state-of-the-art optimal control solver illustrate the simplicity and correctness of the method.

Keywords: Path planning, motion policy, strong currents, strong winds, fast marching tree, holonomic robot.

1 Introduction

In diverse applications, autonomous aerial and aquatic vehicles are required to work in environments with strong natural flows. In these applications, the flow can both assist or constrain the vehicle's motion. For instance, NASA's IceNode buoyant robots exploit ocean currents to move along the basal ice-ocean interface of the Antarctic ice shelves, acquiring long-duration melt rate measurements (Rossi et al., 2021). In another example, balloons developed by Loon provide internet to remote places (Nagpal & Samdani, 2017). Rather than flying against the winds, these balloons take advantage of the atmospheric winds to navigate without propulsion.

In an extraterrestrial application, NASA has considered the use of aerial robots to explore the cloud layer of the Venusian atmosphere (VEXAG, 2019), where recent observations have found evidence of phosphine (Greaves et al., 2020), a compound associated with microbial presence. The strong winds of Venus, which may be as fast as 100 m s^{-1} (Rossi, Saboia, Krishnamoorthy, & Vander Hook, 2023), can help the vehicle to circulate the planet in five days but can also prevent some latitudes from being reached sooner than that (Martinez Rocamora Jr., Juan, & Pereira,



Fig. 1: Given an environment with a time-invariant flow (\rightarrow) and prohibited regions (\square) , the Policy-FlowFMT* algorithm proposed in this paper creates a tree (represented in the center by vertices colored by their costs) that encodes the optimal motion policy to a specified goal (\bullet) . Besides being directly used to provide optimal trajectories to the goal (-), the tree can also generate a vector field (\rightarrow) , which can be applied online for robot control (\cdots) .

2022). To help with these applications, the methods proposed in this paper and illustrated in Figure 1 compute optimal trajectories or motion policies for vehicles navigating natural flows.

The motion planning problem under the influence of environmental flows was approached from many different perspectives. Alvarez, Caiti, and Onken (2004) developed a method using a genetic algorithm for minimum energy path planning of autonomous underwater vehicles (AUVs) in twoand three-dimensional (2D and 3D) environments for time and space-varying ocean currents. While the method considers complex ocean currents, it does not constrain the vehicle's maximum speed, assuming that the vehicle can always counteract the ocean currents to reach a nominal speed. Garau, Alvarez, and Oliver (2005) applied a gridbased search (A^*) method to find minimum energy paths for AUVs in 2D ocean environments. The velocity of the flow was considered to be always smaller than the velocity of the vehicle relative to the current. When the vehicle's speed relative to the fluid is smaller than the speed of the fluid relative to the inertial frame we say that the vehicle is subjected to strong winds or currents. To handle incorrectness and incompleteness issues caused by the presence of strong currents, a sliding wavefront

expansion method was proposed by Soulignac (2011). More recently, other path planning methods based on A* for such 2D environments were introduced by Koay and Chitre (2013); Kularatne, Bhattacharya, and Hsieh (2018). Minimum-time and minimum-energy cost functions were developed by Kularatne et al. (2018) using floworiented coordinates, which considered the feasibility problem implicitly. Most of the existing research for AUVs has considered 2D or quasitwo-dimensional environments, but Kulkarni and Lermusiaux (2020) proposed a level-set solution based on exact differential equations for timevarying and 3D scenarios. Tri-dimensional path planning for AUVs was also considered by Zhai, Hou, Zhang, and Zhou (2022), who developed an efficient method that partitions the flow into a set of regions with constant fields, thus transforming the problem into a mixed integer optimization problem (MIP), and uses a combination of branchand-bound breath-first-search and the Method of evolving junctions (Li, Lu, Zhou, & Chow, 2017) to compute the final optimal path.

The problem was also tackled in the realm of Unmanned Aerial Vehicles (UAVs). Path planners for a gliding aircraft under complex wind fields were studied by Chakrabarty and Langelaan (2013); Langelaan (2008). The proposed planners grow a tree of feasible trajectories from a discrete set of allowable inputs and weigh the branches using a cost function that accounts for changes in total energy and distance to the goal. A real-time environment-aware planner has been proposed by Oettershagen, Achermann, Müller, Schneider, and Siegwart (2017) based on a variation of the Rapidly Exploring Random Trees star (RRT^{*}) (Karaman & Frazzoli, 2011) algorithm that approximates the dynamics of the aircraft using Dubins' Airplane (Chitsaz & LaValle, 2007) and incorporates non-uniform wind field data in its heuristics. The method calculates the Dubins Airplane paths considering zero wind and then deforms the paths using the wind field. Later, it uses an iterative method to generate a virtual goal that will lead to the actual goal position after the path is deformed by the wind. However, this iterative method fails for strong currents.

In a previous work of our group, we used a sampling-based approach, based on the RRT algorithm, to create trajectories considering the strong winds of the atmosphere of Venus (Martinez Rocamora Jr. et al., 2022). We first calculate connections without considering the wind and then integrate the wind along the calculated paths to calculate a wind drift. Although efficient, our method finds sub-optimal trajectories. An extension of the method to find optimal trajectories would require rewiring the search tree, as the RRT* does, which is difficult to perform under the influence of strong winds. Very recently, a few research groups focused their attention on applying the Fast Marching Tree star (FMT*) algorithm (Janson, Schmerling, Clark, & Pavone, 2015), which can find optimal paths faster and without tree rewiring. FMT* is used by Lee, Yoo, Hall, Anstee, and Fitch (2017) for energyoptimal planning of underwater gliders navigating time-invariant flows. The connection between nodes is performed by calculating the "trim states" required to achieve straight-line motion. Bonin, Guitart, Delahaye, Prats, and Feron (2023) adapted the FMT^{*} algorithm to time-optimal motion planning of ultralight gliders navigating in the time-invariant wind flows. The connections between nodes are calculated using a simplified optimal control problem that relies on a two-point boundary solver. Both previous papers consider 3D scenarios and raise reachability problems due

to strong opposing flows. However, different from the analytical procedure proposed in the present paper, these methods numerically calculate the reachability set by solving the vehicle's equation of motion.

This paper presents a different sampling-based algorithm based on FMT^{*} to solve the optimal motion planning problem for holonomic vehicles navigating strong 3D environmental flows. Unlike previous FMT*-based methods, our algorithm explicitly handles the feasibility problem created by situations where the flow speed is larger than the vehicle's maximum speed relative to the flow. To achieve that, we extended the use of reachability cones defined by Soulignac (2011) to 3D and used them to remove unreachable vertices in constructing the search tree. We also use reachability cones to determine accurate minimum-time and minimum-energy cost functions for the problem. Additionally, based on a few modifications to the proposed path planning algorithm, we present a method that finds a motion policy for the vehicle. The policy is represented as a vector field that serves as a feedback planner that optimally guides the vehicle from any valid configuration to its goal. An illustration of the method is shown in Figure 1. In summary, the main contributions of this paper are (1) a motion planner based on FMT^{*} that accounts for the feasibility of the trajectory in the presence of strong flows and prohibited regions (e.g., obstacles), (2) a method to obtain minimumtime and minimum-energy cost functions based on the reachability condition, (3) a method that creates a vector-field policy for optimal navigation in strong flows. We provide the source code of our methods and, for comparison, the setup environment for the equivalent optimal control solver (OCS).

The rest of this paper is organized as follows. The next section introduces the problem we are solving in the paper. The proposed methods and algorithms are explained in Section 3 and experiments are provided in Section 4. Conclusions and future work are presented in Section 5.

2 Problem Definition

2.1 Environment

In this paper, a single robot navigates within a fluid (most commonly air or water) situated in a

tri-dimensional (3D) environment $\mathcal{W} \subset \mathbb{R}^3$. The fluid is assumed to be flowing, and its motion is modeled as a time-independent vector field $F_c : \mathcal{W} \mapsto \mathbb{R}^3$ that maps each position $\vec{x} =$ $[x, y, z]^{\mathsf{T}} \in \mathcal{W}$ to a flow velocity vector $\vec{c}(\vec{x}) =$ $[c_x(\vec{x}), c_u(\vec{x}), c_z(\vec{x})]^{\mathsf{T}} \in \mathbb{R}^3$. The vector field F_c that represents the flow is continuous, i.e., given any point $\vec{x}_0 \in \mathcal{W}$, there is a ball $B_r(\vec{x}_0) = \{\vec{x} \in$ $\mathcal{W} \mid \|\vec{x} - \vec{x}_0\| < r\}$ of radius $r \in \mathbb{R}^+$ centered in \vec{x}_0 , where $||F_c(\vec{x}) - F_c(\vec{x}_0)|| < \epsilon$, for every arbitrary $\epsilon \in \mathbb{R}^+$. Intuitively, if we consider two points in space that are sufficiently close to each other, the velocities of the flow at these points are also similar. Finally, we assume that the environment may contain prohibited regions, \mathcal{O} , which may be seen as obstacles for the vehicle but not for the fluid (e.g., "no-fly zones" for an aircraft). Therefore, the region where the vehicle can safely navigate, or free space, is given by $\mathcal{W}_{free} = \mathcal{W} \setminus \mathcal{O}$.

2.2 Robot

We assume a vehicle whose state space is given by $\mathcal{X} = X \times X$, where $X \subseteq \mathcal{W}$ is the space of vehicle positions, and X is its respective tangent space. We consider that the vehicle is small (meter scale) compared to the environment and the distances traveled (kilometer scale). With this assumption in mind, we use a point with no orientation to represent the vehicle. In practice, this point can be seen as the vehicle itself or, for vehicles with more complicated dynamic models, as a setpoint that could be tracked with the aid of a non-linear controller. Therefore, the vehicle's configuration space is equal to \mathcal{W} , where the vehicle is represented by its position \vec{x} in the global reference frame. At position \vec{x} , the velocity of the vehicle in the global reference frame is given by

$$\vec{v}_g(\vec{x}) = \vec{v}_r(\vec{x}) + \vec{c}(\vec{x}),$$
 (1)

where $\vec{v}_r(\vec{x})$ is the velocity of the vehicle relative to the flow and $\vec{c}(\vec{x})$ is the velocity of the flow. The relative velocity is bounded by a maximum speed v_r^{max} , such that $\|\vec{v}_r(\vec{x})\| \leq v_r^{max}$ for all $\vec{x} \in \mathcal{W}$. We assume that the flow speed can be greater than the vehicle's speed relative to the flow (i.e., $\|\vec{c}(\vec{x})\| \geq v_r^{max}$), which means that the vehicle may not be able to counteract the flow all over the environment.

2.3 Problem statement

The motion planning problem solved in this paper is to find the best trajectory σ^* in the set Σ of all feasible trajectories $\sigma(t) \colon [0, t_f] \to \mathcal{W}_{free}$, parameterized by $t \in [0, t_f]$, that move the vehicle from an initial configuration \vec{x}_S at t = 0 to a set of admissible goal configurations $\vec{x}_G \in X_{goal}$ in a finite and arbitrary time $t = t_f$. Notice that t_f is not necessarily the same for all trajectories $\sigma \in \Sigma$. The cost function $\mathbf{C}(\sigma)$ to be minimized by our problem maps a trajectory to a positive real value $\mathbf{C} : \Sigma \mapsto \mathbb{R}^+$ and respects the constraints imposed on the vehicle by the environment. Time and energy cost functions will be minimized in this paper. Once these functions are defined, our problem is posed as:

$$\min_{\sigma} \quad \mathbf{C}(\sigma)
\text{s.t.} \quad \vec{v}_g(\sigma(t)) = \vec{v}_r(\sigma(t)) + \vec{c}(\sigma(t)) ,
\quad \sigma(0) = \vec{x}_S ,
\quad \sigma(t_f) = \vec{x}_G ,
\quad \sigma(t) \in \mathcal{W}_{free}.$$
(2)

3 Methodology

This section presents our motion planning algorithm. First, we introduce minimum-time and minimum-energy cost functions necessary to specify our problem completely. Then, we define neighborhood functions that are created using reachability cones. Finally, we describe the Flow-Aware Fast Marching Tree (FlowFMT^{*}) algorithm and an adaptation to compute motion policies.

3.1 Flow-Aware Cost Functions

The problem presented in the previous section is not completely specified without a cost function. The literature uses two main categories of cost functions, depending on the task (Kirk, 2004). The first one specifies the minimum time problem, which tries to guide the vehicle to its goal as fast as possible, regardless of the energy spent. The other type specifies the minimum energy problem, which is useful when the vehicle has a finite amount of energy available (e.g., battery) and is required to spend this energy efficiently. Functions in these two main categories are defined next.

3.1.1 Minimum-Time Cost Function

In early attempts to solve the minimum-time problem, the cost functions used were incorrect in the presence of strong flows (i.e., $\|\vec{c}(\vec{x})\| \ge \|\vec{v}_r(\vec{x})\|$ for any workspace position \vec{x}). That meant that trajectories found using such cost functions were not necessarily feasible. The limited reachability problem due to strong flows is illustrated in Figure 2. For simplicity, in this figure and along the rest of this paper the index \vec{x} is omitted from the vehicle's and flow's velocity components, i.e, $\vec{c}(\vec{x})$ is written as \vec{c} , $\vec{v}_r(\vec{x})$ as \vec{v}_r , and $\vec{v}_q(\vec{x})$ as \vec{v}_q .

In this paper, we provide an extension to 3D of the correct minimum-time cost function proposed by Soulignac (2011) to solve 2D problems (we recommend the reader to read the discussion about cost function incorrectness in (Soulignac, 2011)). Minimum-time cost functions are usually defined for continuous trajectories and n-waypoints discretizations, respectively, as:

$$C_t(\sigma) = \int_0^{t_f} dt = t_f \quad \text{and} \quad C_t(\sigma) \approx \sum_{k=0}^n \Delta t^{(k)} \,.$$
(3)

Although these equations are correct, previous work computed the travel time between two sufficiently close configurations (Δt) with simple heuristics that accounted for the traveled distance and an additional term that considered the relative heading of the flow and the vehicle (Petres et al., 2007) (thus assuming constant speed) or a ratio between the traveled distance and the sum of the vehicle speed and the speed of the flow opposing the movement (Blackmore et al., 2010; Garau et al., 2005). None of these heuristics were able to handle the strong flow condition.

When the flow speed gets larger than the maximum vehicle's speed relative to the flow, the drift caused by the flow cannot be compensated, thus creating regions that cannot be reached by the vehicle (Figure 2). To create discrete functions that consider this characteristic, we follow the same procedure adopted by Soulignac (2011) but extend it to 3D. We assume that the vehicle can change direction instantaneously and that the minimum time path between two waypoints under a constant flow vector is a straight line if the connection is feasible. We also assume that two consecutive waypoints are close enough so that the flow is constant between them. Notice that



Fig. 2: The reachable region for a vehicle moving on a constant flow is greatly affected by the ratio between the flow velocities and the vehicle's velocity relative to the flow. The circles represent iso-temporal lines given by the position reached by the vehicle after moving towards any direction for a fixed amount of time at a constant speed $\|\vec{v}_r\|$. When the flow velocity $\|\vec{c}\| = 0$, the circles are concentric and expand outwards with time, thus allowing the vehicle to reach the entire space (green). With $\|\vec{v}_r\| > \|\vec{c}\| > 0$, the circles are drifted accordingly, indicating that the vehicle travels less when moving against the flow but still reaches the entire space (green). When $\|\vec{c}\| \geq \|\vec{v}_r\| > 0$, the vehicle cannot counteract the flow, which creates a region that cannot be reached (red) even if the vehicle's velocity relative to the flow is directly opposed to the flow.

these assumptions are quite realistic in large environments sampled with a very high number of waypoints. Thus, considering two positions $\vec{x}(t_A)$ and $\vec{x}(t_B)$, we can define a directional edge as $\vec{d} = \vec{x}(t_B) - \vec{x}(t_A) = [d_x, d_y, d_z]^{\mathsf{T}}$. The time interval between these two positions is $\Delta t = t_B - t_A$. By integrating the model in (1) and considering the vehicle's relative velocity \vec{v}_r and flow velocity \vec{c} to be constant along this edge we have:

$$\vec{d} = \vec{v}_g \Delta t = (\vec{v}_r + \vec{c}) \Delta t \Leftrightarrow \begin{cases} (v_{r,x} + c_x) \Delta t = d_x \\ (v_{r,y} + c_y) \Delta t = d_y \\ (v_{r,z} + c_z) \Delta t = d_z. \end{cases}$$
(4)

By incorporating this equation on the inner product of the relative velocity $(\vec{v}_r \cdot \vec{v}_r = v_{r,x}^2 + v_{r,y}^2 + v_{r,z}^2)$, we obtain an equation for the 3D cone shown in Figure 3 as:

$$(\vec{c}\cdot\vec{c}-\vec{v}_r\cdot\vec{v}_r)\Delta t^2 - 2(\vec{d}\cdot\vec{c})\Delta t + \vec{d}\cdot\vec{d} = 0.$$
 (5)

An in-depth discussion about the possible outcomes of this equation is derived by Soulignac (2011). In the next few paragraphs, we try to summarize these outcomes and connect them to our method.

In the case that $\vec{c} \cdot \vec{c} = \vec{v}_r \cdot \vec{v}_r$, (5) degenerates to a first-order equation, and the solution is simply $\Delta t = (\vec{d} \cdot \vec{d})/(2(\vec{d} \cdot \vec{c}))$. When $\vec{c} \cdot \vec{c} \neq \vec{v}_r \cdot \vec{v}_r$, the time cost for the straight line path between the two vertices is obtained by solving the quadratic equation for Δt as:

$$\Delta t = \frac{\vec{d} \cdot \vec{c} \pm \sqrt{(\vec{d} \cdot \vec{c})^2 - (\vec{c} \cdot \vec{c} - \vec{v}_r \cdot \vec{v}_r)(\vec{d} \cdot \vec{d})}}{(\vec{c} \cdot \vec{c} - \vec{v}_r \cdot \vec{v}_r)}.$$
(6)

In the case where $\vec{c} \cdot \vec{c} < \vec{v_r} \cdot \vec{v_r}$, the determinant $((\vec{d} \cdot \vec{c})^2 - (\vec{c} \cdot \vec{c} - \vec{v_r} \cdot \vec{v_r})(\vec{d} \cdot \vec{d}))$ is positive, but the product of the two roots is negative, indicating that only one solution is valid (the negative sign solution, which results in positive Δt). Thus, in this case $(\vec{c} \cdot \vec{c} < \vec{v_r} \cdot \vec{v_r})$, the sign of the determinant defines whether or not the connection is feasible. When the computation of Δt in (6) returns complex numbers, the solution does not provide any physical meaning except indicating that the connection is unfeasible. In fact, for some given $\vec{v_r}, \vec{c}$, and \vec{d} , the cone angle β , as shown in Figure 3, is derived from the geometry of the problem as:

$$\cos\beta = \frac{\sqrt{\vec{c}\cdot\vec{c}-\vec{v}_r\cdot\vec{v}_r}}{\|\vec{c}\|} \,. \tag{7}$$

To check if a straight line connection between two points is feasible, we can check if $\cos(\beta) \leq \cos(\beta_{max})$, where $\cos(\beta_{max})$ is found when $\|\vec{v}_r\| = v_r^{max}$. Thus, the complex solutions can be ruled out *a priori* (without computing the determinant in (6)) by assessing this inequality.

Interestingly, when the flow velocity is stronger than the relative velocity $(\vec{c} \cdot \vec{c} > \vec{v_r} \cdot \vec{v_r})$, the two possible values for Δt are positive and valid. In this case, the lower-time solution should be chosen for the minimum-time problem. Therefore, considering that the path is feasible and the vehicle moves with its maximum speed (i.e., $\vec{v_r} \cdot \vec{v_r} = (v_r^{max})^2$), we obtain the minimum-time solution between the two points for both the 2D and 3D cases by taking the negative sign solution of (6). Finally, to compute the total cost of a trajectory, the time cost Δt calculated using (6) is used in (3) for each segment that composes the n piece-wise linear path traversed by the vehicle from start to goal. Notice that the resulting minimum-time cost function is equivalent to the one derived by Zhai et al. (2022), both higher-dimension extensions of the cost function proposed by Soulignac (2011).

3.1.2 Minimum-Energy Cost Functions

Common energy-based cost functions for continuous and discrete trajectories are, respectively, given by:

$$C_e(\sigma) = \int_0^{t_f} P(t)dt \quad \text{and} \\ C_e(\sigma) \approx \sum_{k=0}^n P^{(k)} \Delta t^{(k)}, \qquad (8)$$

where P is the power required to overcome the drag. This required power can be modeled as a polynomial of the vehicle's speed relative to the flow to represent diverse forms of drag (e.g., viscous, pressure, or lift-induced drag):

$$P(t) = \sum_{i=0}^{n} \alpha_i \|\vec{v}_r(t)\|^i.$$
 (9)

From (9) and (8), if the power model has n > 1, we observe that the influence of the vehicle speed is predominant in the energy cost. Thus, to minimize the cost function, it is necessary to minimize the vehicle's speed along the planned trajectory. However, this is not trivial as small vehicles' speeds relative to the flow may yield feasibility problems, since they reduce the reachable region as shown in Figure 2. Further, this type of energy-based cost function can create situations where the resultant energy cost could be as small as wanted as long as it is acceptable to have very long terminal times.

Some methods like Zhai et al. (2022) or Kularatne et al. (2018) address this issue by combining time- and energy-based cost functions. This is achieved by adding a component that is proportional to time (known as hotel cost), which penalizes larger final times, to the energy-cost function shown above. In this work, we propose an alternative to find the energy cost for each segment of the discrete trajectories without explicitly



Fig. 3: Computation of the reachable region. Given a time interval Δt , the displacement due to the flow velocity is given by $\vec{c}\Delta t$, while the displacement due to the vehicle's velocity relative to the flow is the surface of a sphere of radius $\vec{v}_r \Delta t$. Reachability cones are found by considering different values for these two parameters \vec{v}_r and Δt as shown in Figure 2. In (a), the largest reachable cone is defined by β_{max} , which is obtained from when $\|\vec{v}_r\| = v_r^{max}$ using (7). This case also corresponds to the minimum time Δt_{min} connecting the start position (•) to the goal position (•). In (b), considering $\|\vec{v}_r\| < v_r^{max}$, the cone angle β is smaller than β_{max} . In (c), a cone defined by β_{min} is the limit case containing the vector \vec{d} which connects the start position (•) to the goal (•) position. The geometry of this minimum reachable cone is used to calculate a v_r^{min} , which is used in our minimum-energy cost function. Notice that a lower speed increases the time that it takes to move by \vec{d} , i.e., $\Delta t < \Delta t'$.

including such a term. For a given pair of sampled positions (start and goal), our solution finds the minimum speed v_r^{min} that guarantees that the vehicle can still reach the goal position from the start position. To find the minimum speed, we use the limit situation in which the surface of the reachability cone contains the vector \vec{d} that connects start and goal samples. From the geometry of the problem, illustrated by Figure 3, we obtain:

$$\frac{\sqrt{\vec{c}\cdot\vec{c}-(v_r^{min})^2}}{\|\vec{c}\|} = \frac{\vec{d}\cdot\vec{c}}{\|\vec{d}\|\|\vec{c}\|} \Leftrightarrow$$
$$v_r^{min} = \sqrt{\vec{c}\cdot\vec{c} - \left(\frac{\vec{d}\cdot\vec{c}}{\|\vec{d}\|}\right)^2}, \qquad (10)$$

which can be used with (9) and (6) to calculate $P(t)^{(k)}$ and $\Delta t^{(k)}$, which are then used in (8) to calculate the energy cost for each edge of the path.

3.2 Flow-Aware Neighborhood Sets

Sampling-based motion planners, such as the one proposed in this paper, usually rely on the concept of neighborhood, which is usually used to specify a set of vertices (samples) that are candidates to be connected to the current vertices of a graph during its construction. One of the novelties of our method is the division of the neighborhood set into posterior and anterior sets. The posterior neighborhood set of a given vertex v, as shown in Figure 4(a), is the set of vertices that respect two conditions: 1) the vertices can be reached from v (i.e., they are contained in the reachability cone defined by v) and 2) the vertices are within a distance d_{max} from v. On the other hand, the anterior neighborhood set of a vertex v, as shown in Figure 4(b), is the set of vertices that respect two conditions: 1) the reachability cone defined by these vertices contains v and 2) the vertex v is within a distance d_{max} from these vertices. The definition of these two sets allows us to handle the reachability constraint caused by the presence of strong currents when creating a sample-based algorithm, as explained next.

3.3 Flow-Aware Fast Marching Tree

Using the previously defined cost functions and neighborhood sets, this section presents specialized Fast Marching Tree star algorithms (FMT*) (Janson et al., 2015) that consider the reachability constraints imposed by the strong currents. We first show the standard version of the algorithm, for which the tree grows from the



Fig. 4: Two types of neighborhood sets are defined for our method, substituting the standard Near() function used by the original FMT* algorithm. Given a vertex v, we define (a) the posterior neighborhood, with vertices that can be reached from v, and (b) the anterior neighborhood, with vertices that can reach the vertex v.

start position until it reaches the goal region. In the second part of the section, we slightly change the algorithm so it expands the tree from the goal position to every feasible point sampled in the environment and creates a policy.

3.3.1 Flow-Aware FMT*

The Flow-Aware FMT^{*} algorithm (FlowFMT^{*}) proposed in this paper is shown in Algorithm 1. Changes with respect to the original FMT^{*} algorithm (Janson et al., 2015) are displayed in red. The algorithm creates a tree graph with an empty edge set and the vertex set defined by m random

samples of the space, the start, and the goal positions (line 2). Three sets of vertices (lines 3–4) are created: the set of unvisited vertices, which is initialized with all vertices but the start position; the open set, which initially has the start position; and the closed set, which is empty at the beginning. Then, two structures (N^A and N^P) that save two sets of neighbors (see Section 3.2) for each vertex are initialized (lines 6–9).

Algorithm 1 Flow-Aware FMT* Algorithm				
1: function $\sigma = \text{FLOWFMT}^*(\mathbf{x}_S, \mathbf{x}_G, V_{nool})$				
2: $V \leftarrow \mathbf{x}_S \cup V_{pool} \cup \mathbf{x}_G, E \leftarrow \varnothing$				
3: $V_{unvisited} \leftarrow V \setminus \{\mathbf{x}_S\}, V_{open} \leftarrow \mathbf{x}_S$				
4: $V_{closed} \leftarrow \emptyset$				
5: $\mathbf{z} \leftarrow \mathbf{x}_S$				
6: $N_{\mathbf{z}}^{P} \leftarrow Post.Nbhd(\emptyset, V \setminus \{\mathbf{z}\}, \mathbf{z}, r_{n})$				
7: $N^P \leftarrow Save(\emptyset, N_z^P, z)$				
8: $N_{\mathbf{z}}^{A} \leftarrow Ant.Nbhd(\emptyset, V \setminus \{\mathbf{z}\}, \mathbf{z}, r_{n})$				
9: $N^A \leftarrow Save(\emptyset, N_z^A, \mathbf{z})$				
10: while $z \neq x_G$ do				
11: $V_{open,new} \leftarrow \emptyset$				
12: $X_{near} = N_z^P \cap V_{unvisited}$				
13: for $\mathbf{x} \in X_{near}$ do				
14: $N_{\mathbf{x}}^{A} \leftarrow Ant.Nbhd(N^{A}, V \setminus \{\mathbf{x}\}, \mathbf{x}, r_{n})$				
15: $N^A \leftarrow Save(N^A, N^A_{\mathbf{x}}, \mathbf{x})$				
16: $Y_{near} \leftarrow N_{\mathbf{x}}^A \cap V_{open}$				
17: $y_{min} \leftarrow \operatorname{argmin}_{y \in Y_n} \{c(y) + Cost(y, x)\}$				
18: if $CollisionFree(y_{min}, x)$ then				
19: $E \leftarrow E \cup \{(\mathbf{y}_{min}, \mathbf{x})\}$				
20: $V_{open,new} \leftarrow V_{open,new} \cup \{\mathbf{x}\}$				
21: $V_{unvisited} \leftarrow V_{unvisited} \cap \{\mathbf{x}\}$				
22: $c(\mathbf{x}) = c(\mathbf{y}_{min}) + Cost(\mathbf{y}_{min}, \mathbf{x})$				
23: end if				
24: end for				
25: $V_{open} \leftarrow (V_{open} \cup V_{open,new}) \setminus \{z\}$				
26: $V_{closed} \leftarrow V_{closed} \cup \{z\}$				
27: if $V_{open} = \emptyset$ then				
28: return \emptyset				
29: end if				
30: $z \leftarrow \operatorname{argmin}_{y \in V_{open}} \{c(y)\}$				
31: $N_{\mathbf{z}}^{P} \leftarrow Post.Nbhd(N^{P}, V \setminus \{\mathbf{z}\}, \mathbf{z}, r_{n})$				
32: $N^P \leftarrow Save(N^P, N_z^P, z)$				
33: end while				
34: return $GetPath(\mathbf{x}_G, T(V_{open} \cup V_{closed}, E))$				
35: end function				

The algorithm executes until no vertices are found on the open set or a path to the goal is found (lines 10 and 29). At each iteration, the minimum cost vertex of the open set is computed (line 30), and the posterior neighborhood of this vertex is calculated or retrieved (lines 31–32). This set is intersected with the set of unvisited vertices (line

Algorithm 2 Anterior/Posterior Neighborhood Set Computation Functions

1:	function $N_{q} = ANT./POST.NBHD(N, V, q, r_n)$
2:	if $\exists N(q)$ then
3:	$\mathbf{return} N(\mathbf{q})$
4:	end if
5:	$N_{\mathbf{q}} \leftarrow \varnothing$
6:	$\vec{q} \leftarrow GetPosition(\mathbf{q})$
7:	$ec{c} \leftarrow \mathbf{F}_c(ec{q})$
8:	$\cos\left(\beta_{max}\right) \leftarrow \sqrt{\vec{c} \cdot \vec{c} - (v_r^{max})^2} / \ \vec{c}\ $
9:	$\mathbf{for} \ \mathbf{v} \in V \ \mathbf{do}$
10:	$\vec{x} \leftarrow GetPosition(\mathbf{v})$
11:	$\vec{d} \leftarrow \vec{q} - \vec{x}$ (Ant.) or $\vec{d} \leftarrow \vec{x} - \vec{q}$ (Post.)
12:	$\cos\beta \leftarrow (\vec{d} \cdot \vec{c}) / (\ \vec{d}\ \ \vec{c}\)$
13:	$\mathbf{if} \ \vec{d}\ < r_n \mathbf{then}$
14:	if $\ \vec{c}\ \geq v_r^{max}$ then
15:	if $\cos \beta \geq \cos (\beta_{max})$ then
16:	$N_{\mathbf{q}} \leftarrow N_{\mathbf{q}} \cup \mathbf{v}$
17:	end if
18:	else
19:	$N_q \leftarrow N_q \cup \mathrm{v}$
20:	end if
21:	end if
22:	end for
23:	$\mathbf{return} \ N_{\mathbf{q}}$
24:	end function

12), and for each of the vertices x in the intersection, a search for a locally optimal one-step connection is performed by testing connections from all the vertices **y** in the intersection of the anterior neighborhood and the open set to x (lines 13–17). Up to this point, no collision checks are done. Once the locally-optimal one-step connection is found, it is tested for collision (line 18). If no collisions are detected, the edge is added to the tree (line 19). The newly added vertex is added to the open set and removed from the unvisited set (lines 20–21). After looping through all vertices in the posterior neighborhood, the query vertex is removed from the open set and added to the closed set (lines 25–26). The neighborhood sets are calculated using the AnteriorNeighborhood (shortened to Ant.Nbhd) and PosteriorNeighborhood (shortened to Post.Nbhd) functions as shown in Algorithm 2. The difference between the two functions is given by the definition of d (line 11).

3.3.2 Computing a Policy using Flow-Aware FMT*

More useful than simply obtaining a path from the start to the goal, a policy defines paths from all positions in the space to the goal. A policy can be represented by a vector field, which can be used as a feedback system to compensate for disturbances (LaValle, 2006). Since the FMT* method is a graph-based level-set method that computes increasing levels of the cost function as the tree marches to cover the environment, it can be easily used to compute a policy. We then propose a few modifications to the Flow-Aware FMT* algorithm (Algorithm 1) to make this possible: (i) the root of the tree (lines 2–5) is set to the goal position; (ii) the anterior (lines 14–15) and posterior (lines 31–32) neighborhood sets are swapped; (iii) the order of the vertices on the cost function and collision checking functions (lines 18–19, and 22) are inverted; (iv) the stopping criteria (lines 10 and 27) is changed to only interrupt the procedure if the open set becomes empty.

The outcome of these modifications is shown in Algorithm 3. Changes with respect to the standard Flow-Aware FMT* algorithm (Algorithm 1) are displayed in red. Notice that Algorithm 3 builds a tree from the goal position but still considers that the direction of motion of the vehicle is towards the goal. The dynamic programming procedure (line 17) considers only valid connections given the sets defined in Section 3.2.

A policy is encoded in the resulting tree. Once this tree is found, the vehicle's velocity in the global reference frame \vec{v}_g can be calculated using a multivariate interpolation around each position \vec{x} in the continuous space. We propose using an inverse distance weighting interpolation of the vertices of the tree that are close to \vec{x} to extract the policy vector field as:

$$\vec{v}_r(\vec{x}) = \overline{\vec{v}_g}(\vec{x}) - \vec{c}(\vec{x}) = \frac{\sum_{i=0}^k w_i(\vec{x}) \, \vec{v}_{g,i}}{\sum_{i=0}^k w_i(\vec{x})} - \vec{c}(\vec{x}) \,,$$
(11)

where $\vec{v}_{g,i}$ $(0 \leq i \leq k)$ are the vehicle's velocities in the global reference frame at the k vertices v_i of the tree that are within a distance r from \vec{x} (i.e., $v_i \in V \cap \mathcal{B}_r(\vec{x})$). The process is illustrated by Figure 5. The weight function is given by $w_i(\vec{x}) = 1/(\vec{x} - \vec{x}_i)^p$ where p is a parameter (commonly, p = 2). Notice that we assume that \vec{x} is located in a region from which the goal can be reached. We do not include any further development in this paper, but simple heuristics, like counting the number of vertices of the tree in the ball $\mathcal{B}_r(\vec{x})$ can help identify whether \vec{x} belongs to a feasible region or not. From (1) the vehicle's

Algorithm 3 Flow-Aware FMT* Algorithm (Policy Version)

1:	function $T = \text{POLICY-FLOWFMT}^*(\mathbf{x}_G, V_{pool})$
2:	$V \leftarrow \mathbf{x}_G \cup V_{pool}, E \leftarrow \emptyset$
3:	$V_{unvisited} \leftarrow V \setminus \{\mathbf{x}_G\}, V_{open} \leftarrow \mathbf{x}_G$
4:	$V_{closed} \leftarrow \emptyset$
5:	$\mathrm{z} \leftarrow \mathrm{x}_G$
6:	$N_{\mathbf{z}}^{P} \leftarrow Post.Nbhd.(\emptyset, V \setminus \{\mathbf{z}\}, \mathbf{z}, r_{n})$
7:	$N^P \leftarrow Save(N^P, N_z^P, z)$
8:	$N_{z}^{A} \leftarrow Ant.Nbhd.(\emptyset, V \setminus \{z\}, z, r_{n})$
g.	$N^A \leftarrow Save(N^A N^A z)$
10:	while $V_{open} \neq \emptyset$ do
11:	$V_{open,new} \leftarrow \emptyset$
12:	$X_{near} = N_z \cap V_{unvisited}$
13:	for $x \in X_{near}$ do
14:	$N_{\mathbf{x}}^{P} \leftarrow Post.Nbhd(N^{P}, V \setminus \{\mathbf{x}\}, \mathbf{x}, r_{n})$
15:	$N^P \leftarrow Save(N^P, N^P_{\mathbf{x}}, \mathbf{x})$
16:	$Y_{near} \leftarrow N_{\mathbf{x}}^P \cap V_{open}$
17:	$y_{min} \leftarrow \operatorname{argmin}_{\mathbf{y} \in Y_n} \{ c(\mathbf{y}) + Cost(\mathbf{x}, \mathbf{y}) \}$
18:	if $CollisionFree(v_{min}, x)$ then
19:	$E \leftarrow E \cup \{(\mathbf{y}_{min}, \mathbf{x})\}$
20:	$V_{open,new} \leftarrow V_{open,new} \cup \{\mathbf{x}\}$
21:	$V_{unvisited} \leftarrow V_{unvisited} \cap \{\mathbf{x}\}$
22:	$c(\mathbf{x}) = c(\mathbf{y}_{min}) + Cost(\mathbf{x}, \mathbf{y}_{min})$
23:	end if
24:	end for
25:	$V_{open} \leftarrow (V_{open} \cup V_{open,new}) \backslash \{z\}$
26:	$V_{closed} \leftarrow V_{closed} \cup \{\mathbf{z}\}$
27:	$\mathbf{z} \leftarrow \operatorname{argmin}_{\mathbf{y} \in V_{open}} \{ c(\mathbf{y}) \}$
28:	$N_{\mathbf{z}}^{A} \leftarrow Ant.Nbhd(N^{A}, V \setminus \{\mathbf{z}\}, \mathbf{z}, r_{n})$
29:	$N^A \leftarrow Save(N^A, N^A_z, z)$
30:	end while
31:	$\mathbf{return} \ T \left(V_{open} \cup V_{closed}, E \right)$
32:	end function

velocity relative to the flow can be computed by subtracting the interpolated value from the flow velocity at this point.

3.3.3 Analysis

In this section, we discuss the computational complexity, the optimality, and the completeness of the FlowFMT^{*}. Notice that FlowFMT^{*} is very similar to the original version of FMT^{*}, as can observed by the highlighted lines in algorithms 1 and 3. The main difference is that, instead of a single spherical neighborhood set of radius r_n for each sample, FlowFMT^{*} creates two spherical-conic neighborhood sets (intersection of a wind-driven elliptical cone and a sphere of radius r_n) for each sample. With this modification, we reduce the number of calls of functions Cost() and CollisionFree(), which highly improves the computation time of the algorithm, but does not



Fig. 5: Policy encoded by the tree T(V, E) resulting from Algorithm 3. The global velocity $\vec{v}_{g,i}$ at the vertices *i* of the tree around the position of the robot \vec{x} are combined to find the target global velocity $\overline{\vec{v}_g}$. By subtracting the local flow velocity \vec{c} , the robot's velocity \vec{v}_r , which is used to control the robot, can be retrieved as shown in (11).

change its computational complexity when compared to the original FMT^{*}. Therefore, for n samples, the computational complexity of FlowFMT^{*} is still $O(n \log(n))$ in expectation, as discussed by Janson et al. (2015). To achieve this computational complexity, Algorithm 2 which finds the neighborhood sets for each query node should be implemented using KDTrees, so that the filtering by r_n happens in $O(\log(n))$.

The total number of samples in the two neighborhood sets of FlowFMT^{*} can be larger or smaller than the number of samples of the spherical set of FMT^{*}, depending on the relative speed of the wind. This, however, does not change the space complexity of the method, which also remains $O(n \log(n))$. The multiplicative constant can make FlowFMT^{*} to use more memory for low-velocity flows and less memory for high-velocity flows.

For the original spherical neighborhood set with radius

$$r_n = \left(\frac{\lambda \log(n)}{n}\right)^{1/d}, \qquad (12)$$

where λ is a positive constant and d is the space dimensionality, FMT^{*} is proven to be asymptotically optimal (and therefore probabilistically complete) (Janson et al., 2015). The adoption of spherical-conic neighborhood sets by FlowFMT^{*} does not change this characteristic because, in essence, the use of these sets is equivalent to the use of the original spherical set and a cost function Cost(x, y) that would return infinite when sample y is unreachable from sample x. Thus, the idea behind FlowFMT^{*}, which does not change the asymptotic optimality proof of FMT^{*}, is to avoid several computations of the cost function (and also of CollisionFree()) by pre-selecting the neighboring samples that would not return an infinite cost (reachable neighbors). Hence, FlowFMT^{*} inherits the characteristics of FMT^{*} and remains asymptotically optimal.

The idea of reachable sets changes the overall behavior of FMT^{*} in a few situations. Depending on the flow configurations, placement of prohibited regions (obstacles), and also on the maximum speed of the vehicle, there may exist a set of samples in the environment from where the vehicle cannot reach the goal. If the initial configuration belongs to this set, the motion planning problem is infeasible. In most cases, this happens when the vehicle is not able to counteract the flow speeds as it tries to avoid the prohibited regions and environment limits. Situations like this will be exemplified in the experiments of the next section. Although the presence of a set of samples that leads to infeasible problems does not change the optimality and completeness of either FMT* or FlowFMT*, it makes FlowFMT* more interesting. FMT* would find an element x of this set when the cost to all neighbors in the spherical set is infinite. FMT* would then connect x to one of its neighbors and at the end would return a path of infinity cost. On the other hand, FlowFMT* would explicitly define empty neighborhood sets for each sample that cannot reach the goal. This is a major advantage of FlowFMT^{*}, which would immediately (without computing any cost or performing collision checks) return an empty path, instead of a path that can't be physically followed by the vehicle, as returned by FMT^{*}. Since FlowFMT^{*} returns a path when the initial position can reach the goal and returns an empty path when this is not possible, FlowFMT* is probabilistically complete (a condition necessary for asymptotic optimality), even when it faces an infeasible motion planning problem. Notice that, with small modifications, FMT^{*} could have the same behavior but at the cost of trying to connect samples in the infeasible region with all of their neighbors before detecting that the problem is actually infeasible.

The Policy-FlowFMT* algorithm has the same complexity as FlowFMT^{*} since the differences between the algorithms are only related to the in-line ordering of some of its computations. We can then analyze the post-processing step given by (11). After the search tree is computed offline, this equation is supposed to be used at runtime. It is assumed that the tree resulting from Algorithm 3 would be stored in memory and the robot would only need to query for the vertices that are close to it. A simple search using KDTrees, which is $O(\log(n))$, can be used to assess the distance between the current robot position and the vertices of the tree and calculate the policv at O(1). Therefore, the runtime complexity of Policy-FlowFMT* is $O(\log(n))$.

4 Experimental Results

In this section, the performance of the proposed planner is evaluated in 2D and 3D scenarios, with and without prohibited regions. The resulting paths are always compared to the paths obtained by a state-of-the-art optimal control solver (OCS), the Imperial College London Optimal Control Software (ICLOCS2) (Nie, Faqir, & Kerrigan, 2018). Additionally, by testing our method in environments previously published in the literature, we could also compare FlowFMT* with a Level Set Method (implemented by Zhai et al. (2022)), the Evolving Junction method (Zhai et al., 2022), true optimal solutions (Subramani, Wei, & Lermusiaux, 2018), and a competitive grid-based approach (Kularatne et al., 2018). All the simulations were performed using an Intel[®] $Core^{\mathbb{M}}$ i9-9900K CPU at 3.6 GHz, with 16 cores and 32GB of RAM.

4.1 Canonical Jet Flow Environment

The crossing of canonical jet flows are a set of illustrative problems described in prior works like Subramani et al. (2018) and Zhai et al. (2022), that bring possibilities to validate FlowFMT^{*} on different aspects. One specific problem of this validation scenario is that it has discontinuities in the flow speeds, which contradicts our assumption that, in the limit, as the sample density gets



Fig. 6: Adaptive Neighborhood Radius for the FlowFMT^{*}. The neighborhood size is decreased when the streamlines curve to improve on the assumption that the flow velocity is constant inside of the neighborhood.

higher, the flow speed between two neighboring samples tends to be constant. Although increasing the number of samples minimizes the effect of the discontinuity on the entire path, in our simulations we also included an adaptive neighborhood radius according to a vector-similarity metric based on the difference between the flow velocities in the start and goal vertices (\vec{c}_{parent} and \vec{c}_{child}). The adaptive neighborhood radius is given by

$$r_n^* = r_n \left(1 - \frac{\|\vec{c}_{parent} - \vec{c}_{child}\|}{2c_{max}} \right), \qquad (13)$$

where r_n is the original radius proposed for the FlowFMT^{*} as shown in Equation (12). Figure 6 illustrates how this process works. Notice that r_n^* tends to r_n as our assumption of constant flow within the neighborhood region is better satisfied. If discontinuities or large gradients are present, the smaller radius will reduce the chances of connecting samples far from each other.

4.1.1 Two-Dimensional Jet

In this subsection we consider a vehicle that can move at a maximum speed $v_r^{max} = 10.0 \,\mathrm{m\,s^{-1}}$ relative to the flow. We limit the environment to a box of side 100 m, i.e., $\mathcal{W} = [0, 100] \times [0, 100]$ (the dimensional unit 'meters' is omitted in the rest of this section for simplicity). To match the scenario proposed Subramani et al. (2018), we plan paths from the start position $\vec{x}_{start} = [20.0, 20.0]^{\mathsf{T}}$ to the goal position $\vec{x}_{goal} = [80.0, 80.0]^{\mathsf{T}}$. The flow environment is modeled as

$$\vec{c}(x,y) = \begin{cases} [20.0, 0.0]^{\mathsf{T}} \text{ if } y \ge 40 \text{ and } y \le 60, \\ [0.0, 0.0]^{\mathsf{T}} \text{ otherwise}. \end{cases}$$
(14)

FlowFMT* is compared to the optimal results found using the methodology explained in Subramani et al. (2018), where a simple gradient-based numerical method is used to find the optimal solution to the jet crossing problem. The optimization is solved using Matlab (fmincon function). In Figure 7, we show the optimal solution obtained numerically and solutions found using FlowFMT^{*} for different number of samples. The cost using 25,600, 102,400, and 409,600 samples, were, respectively, $C_t(\sigma) = 6.2671 \,\mathrm{s}, C_t(\sigma) =$ 6.2608 s, and $C_t(\sigma) = 6.2569$ s. The cost of the numerical solution was $C_t(\sigma) = 6.2523 \,\mathrm{s.}$ Notice that the discontinuity in the flow field causes errors in FlowFMT*'s solution (a small shift compared to the optimal solution), but as the number of samples increases, and with the inclusion of the adaptive neighborhood radius this impact is significantly decreased.

Figure 8 shows a solution computed using Policy-FlowFMT* with 102,400 samples. The tree starting at the goal position is shown in Fig 8(a) and the resultant vector field is shown in Fig 8(b). Similar to the previous result, the computed trajectory is very close to the optimal one. Also, notice that, differently from other methods, our vector field explicitly indicates a region of the space from where the vehicle cannot reach the goal (no solution).

4.1.2 Three-Dimensional Jet

In this subsection, we consider a vehicle that can move at a maximum speed $v_r^{max} = 3.0 \,\mathrm{m\,s^{-1}}$ relative to the flow. We set the environment limits to the box $\mathcal{W} = [-10, 10] \times [-10, 10] \times [0, 20]$ (the dimensional unit 'meters' is omitted in the rest of this section for simplicity). To match what is found in the literature, in this section, we plan paths from the start position $\vec{x}_{start} = [0.0, 0.0]^{\mathsf{T}}$ to the goal position $\vec{x}_{goal} = [0.0, 20.0]^{\mathsf{T}}$. The flow



Fig. 7: Motion planning for a vehicle crossing a 2D jet flow using FlowFMT^{*}. The robot moves from the start position (\bullet) to the goal (\bullet). We show the regions where the speed of the flow is greater (\square) and smaller (\square) than the maximum speed of the vehicle relative to the flow. The edges of the tree are colored based on the cost of the children's vertices. The numerical optimal solution is shown in black (-). The solution found with the adapted FlowFMT^{*} is shown in red (-) for (a) 25,600, (b) 102,400, and (c) 409,600 samples.



Fig. 8: Motion planning for a vehicle crossing a 2D jet flow using Policy-FlowFMT^{*}. (a) Policy-FlowFMT^{*} tree with the root at the goal position (\bullet) for a two-dimensional jet flow. (b) Vector field computed with Policy-FlowFMT^{*}. The figures show the following trajectories from start (\bullet) to the goal (\bullet): the optimal trajectory computed numerically (—), the trajectory from the tree computed by PolicyFlow^{*} (—), and the trajectory computed by integrating the vector field (\cdots). From the regions where the field is not computed, the robot cannot reach the goal (no solution).

environment is modeled as

$$\vec{c}(x, y, z) = \begin{cases} [0.5, 0.0, 0.0]^{\mathsf{T}} \text{ if } z \ge 0 \text{ and } z \le 10, \\ [2.0, 1.0, 0.0]^{\mathsf{T}} \text{ if } z \ge 10 \text{ and } z \le 15, \\ [0.0, 0.0, 0.0]^{\mathsf{T}} \text{ otherwise}. \end{cases}$$

FlowFMT^{*} is compared to the optimal control solver ICLOCS2 (OCS), the Evolving Junction method (EJ) (Zhai et al., 2022), and the Level

(15) 13

Set method (LS) when computing minimal-time paths. It is important to mention that LS is known to compute the shortest-time path over time-varying flows at the cost of longer computational times. On the other hand, EJ was proposed as an efficient solution that avoids errors induced by the discretization of the environment. Because it works on partitions of the flow, it explicitly handles discontinuous flows, as is the case of the environment used in this section.

Table 1 shows the comparison of the methods. The numerical results for EJ and LS came directly from Zhai et al. (2022). Besides the minimumtime cost, C_t , found by each method, Table 1 also shows the parameters of each path found, where angles θ_i and γ_i are elevation and azimuth angles of the path segments crossing the horizontal planes in z = 10 and z = 15 as defined by Zhai et al. (2022). Notice that all paths are very similar, showing that FlowFMT^{*} converges to the optimal path, even in an environment with discontinuous flow. Figure 9 compares the trajectory found by FlowFMT* using 204,800 samples with trajectories found by OCS. The trajectories found with EJ and LS were omitted from the plot due to their similarity to the other two trajectories. Since we did not implement the Evolving Junction (EJ) method, a time comparison between the methods cannot be made. Despite this, although EJ could be the best choice for environments with discontinuous and easy-to-segment flows, we believe that FlowFMT* would be a better choice for general flow environments, where the field cannot be easily partitioned into piece-wise constant subfields. The double-gyre flow environment, exploited in the next subsection, is an example of such an environment.

4.2 Double-Gyre Flow Environment

In this subsection, we consider a vehicle that can move at a maximum speed $v_r^{max} = 0.05 \,\mathrm{m \, s^{-1}}$ relative to the flow. The environment is modeled using a three-dimensional (3D) version of doublegyre flow, which is usually used to evaluate motion planning on flows (see (Kularatne et al., 2018; Lee

Table 1: Comparison of Evolving Junction (EJ), Level Set (LS), ICLOCS2 (OCS), and FlowFMT* methods for the 3D jet flow environment. The segment angles for the FlowFMT* were estimated by calculating elevation and azimuth angles to the interpolated path crossing the horizontal planes in z = 10 and z = 15.

	EJ	LS	OCS	$FlowFMT^*$
θ_1	82.7924	83.5659	83.0881	83.7492
θ_2	62.0255	63.3118	63.4041	64.8360
θ_3	73.7397	73.80277	74.7784	75.3569
γ_1	-136.0775	-135.6592	-134.8846	-133.9709
γ_2	30.2293	30.2407	30.9735	32.8730
γ_3	-161.6199	-161.2246	-161.6005	-158.1282
C_t	6.9096	6.9826	6.9092	6.9090

et al., 2017)), as:

$$\vec{c}(x,y,z) = \begin{cases} -\pi A \sin(\pi x/s) \cos(\pi y/s) \cos(\pi z/s) \\ \pi A \cos(\pi x/s) \sin(\pi y/s) \cos(\pi z/s) \\ \pi A \sin(\pi z/s) , \end{cases}$$
(16)

where A is the amplitude scaling factor that controls the maximum flow speed $(\max_{\vec{x}\in\mathcal{W}} \|\vec{c}(\vec{x})\|)$, and s determines the characteristic length of the gyres. As suggested by Kularatne et al. (2018), we choose A = 0.02 so that the maximum flow speed is $0.0625 \,\mathrm{m\,s^{-1}}$, which is greater than the vehicle's maximum speed relative to the flow. We limit the environment to a box of side 2 m, i.e., $\mathcal{W} = [0,2] \times [0,2] \times [0,2]$ (the dimensional unit 'meters' is omitted in the rest of this paper for simplicity), and choose s = 1. To obtain the standard two-dimensional (2D) double-gyre flow, the slice where z = 0 is used.

4.2.1 Two-Dimensional Gyre

In a 2D environment (z = 0 in (16)), FlowFMT^{*} is used to find paths from the start position $\vec{x}_{start} = [0.1, 0.1]^{\mathsf{T}}$ to a few different goal positions $X_{goal} = \{[0.1, 1.9]^{\mathsf{T}}, [1.5, 1.0]^{\mathsf{T}}, [1.9, 0.9]^{\mathsf{T}}, [1.9, 1.1]^{\mathsf{T}}, [1.9, 1.9]^{\mathsf{T}}\}$, as shown in Figure 10. This figure shows the optimal-time paths for both minimum-time and minimumenergy cost functions and the resultant trees used to obtain these paths. The edges of the tree were colored by the cost of the children's vertices. The solution from the equivalent optimal control problem is shown for comparison. The general



Fig. 9: Motion planning for a vehicle crossing a 3D jet flow. The robot moves from the start position (\bullet) to the goal (\bullet). The solution found with the adapted FlowFMT* using 204,800 samples is shown in red (—) and the optimal solution found using the OCS is shown in black (—).

topology of the path for both methods is the same. Obtaining the best path from the OCS requires some tuning and an appropriate initial guess. Without these two requirements, the OCS can either take too long to converge (several minutes) or converge to a local optimum. We noticed that as we increased the number of samples, the path generated by FlowFMT* approached the best path obtained by the OCS without the extra work.

The absolute values of the minimum-time cost obtained by FlowFMT* and the OCS solution agree with the results obtained by Kularatne et al. (2018), even though the piece-wise time cost between vertices was obtained using different formulations (ours is based on (Soulignac, 2011)). For comparison, considering the path from $[0.1, 0.1]^{\mathsf{T}}$ to $[1.9, 0.9]^{\mathsf{T}}$, the OCS solution obtained by Kularatne et al. (2018) presented a minimumtime cost of $C_t(\sigma) = 32.87 \,\mathrm{s}$ while their proposed solution had $C_t(\sigma) = 32.92 \,\mathrm{s.}$ Meanwhile, we observed a cost of $C_t(\sigma) = 32.86 \,\mathrm{s}$ using OCS (set up with ICLOCS2), and $C_t(\sigma) = 32.88 \,\mathrm{s}$ using FlowFMT^{*}. Our method ran with 40,000 samples to match the spatial resolution of Kularatne et al. (2018). This indicates that both approaches generate similar valid solutions for the minimum-time problem.

Finding the true absolute values for the optimal paths considering the minimum-energy cost function is a more difficult task in the scenario involving gyres. Since the double-gyre flow model has parallel streamlines that circulate the center of each gyre, a vehicle without actuation would just drift around the center of its current gyre, carried by the flow. To reach a position in a different streamline while minimizing the energy spent, the vehicle must move from streamline to streamline with the minimum speed possible. In an ideal scenario, this speed and its resultant energy cost could be as small as we wanted. Clearly, very small costs would also result in solutions with an excessive final time. A discussion related to this problem was made by Kirk (2004). We then conclude that the true minimum energy cost is dependent on the minimum speed that the vehicle produces relative to the flow. This issue is not elucidated by Kularatne et al. (2018) when their OCS solution was explained. To allow the OCS to find a solution with a finite and reasonable time, we set the inferior limit of the vehicle's speed relative to the flow to $0.0001 \,\mathrm{m\,s^{-1}}$ (0.02% of the maximum relative speed). For FlowFMT*, we do not enforce a minimum speed since it is calculated using (10), which will be zero only in the rare cases where the next waypoint is completely aligned with the flow.



Fig. 10: Paths planned by FlowFMT^{*} (—) in a two-dimensional double-gyre flow (\rightarrow) compared to the associated optimal control problem (—). On the left, we show the regions where the speed of the flow is greater (\blacksquare) and smaller (\blacksquare) than the maximum speed of the vehicle relative to the flow. On the right, we show FlowFMT^{*} trees, with vertices colored by their costs. The minimum-time cost function was used in (a) and (b) and the minimum-energy cost function in (c) and (d).

Figure 10(c) and (d) show that our method can find solutions that are similar to the ones found by the OCS. Numerical comparisons for all the paths shown are presented in Table 2. The average computational time spent by FlowFMT* to compute the paths in Figure 10 was 74 s, while the OCS was "generally" able to finish in a few seconds. However, we also observed that the OCS computational time can increase



Fig. 11: Comparison of FlowFMT^{*} and optimal control in a two-dimensional double-gyre flow with an increasing number of prohibited regions (\square): (a) 0 regions, (b) 5 regions, (c) 10 regions, (d) 20 regions, (e) 40 regions, and (f) 80 regions. The solution trajectory computed by FlowFMT^{*} (—) was found using 102,400 samples. For each scenario, two solutions were found with optimal control. The first was initialized with a simple straight trajectory (—) and the second using the trajectory found by FlowFMT^{*} (–).

Table 2: Performance metric of the tra-
jectories shown in Fig. 10 computed by
FlowFMT* and optimal control method.

	$FlowFMT^*$		ICLOCS2	
Goal \vec{x}_g	C_t [s]	$C_e \ [\mu J]$	C_t [s]	C_e [µJ]
$[1.9, 0.9]^{T}$	32.89	1.120	32.86	2.532
$[1.9, 1.1]^{T}$	35.12	1.779	35.06	1.747
$[1.5, 1.0]^{T}$	34.47	0.834	34.43	0.737
$[1.9, 1.9]^{T}$	30.15	1.681	30.11	1.204
$[0.1, 1.9]^{T}$	27.58	1.584	27.62	2.143

drastically (up to minutes) if obstacles (prohibited regions) are added and, consequently, more active constraints are in place, or if the problem is not set up correctly (poor initial guess, poor constraint limits, etc.). Furthermore, OCS solutions tend to converge to local minima, which does not happen with our graph-based solution. Depending on the initial guess, the OCS solutions either get stuck in suboptimal homotopies or show additional loops around gyres, introducing suboptimality.

This behavior is further explored in the experiment shown in Figure 11, where we incrementally add randomly generated prohibited regions to a scenario and test FlowFMT^{*} and OCS. The radii and center coordinates of the prohibited regions are sampled from uniform distributions, U(0.02, 0.08) and U(0.0, 2.0) respectively. In this experiment OCS is initialized in two forms: 1) with a straight trajectory from start to end, and 2) using the optimal trajectory found with FlowFMT^{*}. The optimal time

Table 3: Performance and computational time metrics (in seconds) of the trajectories shown in Fig. 11 computed by FlowFMT* and the optimal control method (1) initialized with a straight line trajectory between start and goal and (2) initialized with the trajectory computed by FlowFMT*.

	$\operatorname{FlowFMT}^*$		OCS(1)		OCS(2)	
#Obs.	C_t	Δt_c	C_t	Δt_c	C_t	Δt_c
0	30.13	204.9	30.11	1.735	30.11	3.738
5	30.13	206.0	30.11	2.360	30.11	5.893
10	30.53	210.9	38.89	2.734	30.45	7.212
20	31.30	221.0	40.95	3.031	31.89	21.15
40	31.61	227.7	46.14	47.45	33.41	140.3
80	31.80	290.1	-	>1500	32.39	1299

costs and the computational times for each test are shown in Table 3. We observe that in the absence or with sparse distribution of prohibited regions (Figures 11(a) and (b)), OCS can reach optimal trajectories in less time compared to trajectories of the same quality with FlowFMT*. However, as the number of prohibited regions grows, it is possible to see the trajectories found by OCS getting trapped in sub-optimal homotopies. Also, the computational time grows considerably. By comparing the computational time from situations with no prohibited regions (Figures 11(a)) to environments with 80 prohibited regions (Figures 11(f)) we noticed that the time grows $1.4 \times$ with FlowFMT*, mainly due to collision checking. On the other hand, the computational time grows $347.5 \times$ for the OCS initialized with FlowFMT^{*} trajectory. OCS with a straight line initialization did not converge before 25 min, when it was interrupted.

The Policy-FlowFMT^{*} in Algorithm 3 was also compared to the solutions obtained with ICLOCS2. First, we found trajectories from 20 random initial positions to a single final position $\vec{x}_{goal} = [1.9, 0.9]^{\mathsf{T}}$ in the double-gyre environment without prohibited regions. We provided a straight line between the start and goal as the initial guess for the OCS. In our algorithm, we tested different resolutions, doubling the number of samples from 100×2^0 to 100×2^{10} . The results are shown in Figure 12. It is possible to see that 2 of the 20 trajectories were dissonant, with FlowFMT^{*} finding lower-cost solutions (in fact, better solutions). As mentioned before, this could be solved by providing a better initial guess for the OCS. Figure 12(b) shows how close the cost found by FlowFMT* gets to the "true" cost found by ICLOCS2 on the other 18 trajectories in terms of cost ratio (i.e., $C_t(\sigma_{\text{FlowFMT}^*})/C_t(\sigma_{\text{OCS}})$). As expected, the cost ratio approaches one when the number of samples increases. Depending on the validity of the locally constant flow assumptions, FlowFMT* may not find trajectories with a cost that is necessarily higher than the OCS solutions. However, as the number of samples increases, the assumption becomes more valid, and the cost estimate becomes more realistic. Figure 12(b) also shows how the computational time to compute the policy increases as the number of samples increases. Despite that, remember that a policy is only computed once for each goal position, which can make this time diluted among the number of future trajectories that need to be computed using such a policy.

The Policy-FlowFMT* algorithm can also be used to compute a velocity vector field as shown in Figure 13(b) for an environment with prohibited regions. This figure also shows simulated trajectories that illustrate how the vector field can help recover the vehicle after an actuation fault, while the OCS solution may lead to a collision. We considered that the OCS solution and the vector field are both used in real-time to control the vehicle. We then simulated that vehicle lost power for 10s starting at t = 5 s. Figure 13(b) shows that at this moment, the vehicle drifts with the flow due to the lack of actuation. When the actuation is recovered, the OCS solution is not valid anymore, leading the vehicle to a prohibited region. In the best-case scenario, if a trajectory tracker was available, the vehicle maybe would be able to retake the path, at the cost of suboptimality of the resultant trajectory. In comparison, the vector field, no matter where the control is regained, would guide the vehicle through an optimal path from that position to the goal, as can be verified in Figure 13(b).

4.2.2 Three-Dimensional Gyres

Figure 14 shows the vehicle's trajectories from $\vec{x}_{start} = [0.1, 0.1, 0.1]^{\mathsf{T}}$ to $\vec{x}_{goal} = [1.9, 0.9, 0.9]^{\mathsf{T}}$ moving in a 3D double-gyre flow (as defined in (16)) with prohibited regions. We compared the



Fig. 12: Policy-FlowFMT^{*} computed over the 2D double-gyre flow environment. In (a) we compare the solution for 20 random initial positions using our method and ICLOCS2. In (b) we show that our solution approaches the optimal solution with more samples at the cost of more computational time.



Fig. 13: Comparison with optimal control: (a) Comparison of FlowFMT^{*} and optimal control in a twodimensional double-gyre flow with prohibited regions (\square). Both algorithms provide similar solutions when optimal control has a good initial guess (-). However, optimal control with a bad initial guess (-) can return paths with a different homotopy. (b) Comparison of Policy-FlowFMT^{*} and optimal control for real-time control in the presence of a disturbance. After an actuation fault, the vector field (\rightarrow) created using Policy-FlowFMT^{*} brings the vehicle to the goal following an optimal trajectory (\cdots) while optimal control (\cdots) leads the vehicle to cross a prohibited region.



Fig. 14: Navigating in a 3D double-gyre flow with prohibited regions. The Policy-FlowFMT* solution (—) using 102,400 samples is compared to the solution (…) obtained from a simulation using the policy vector field (\rightarrow), and the solution obtained from an optimal control problem solver (—).

performance of the optimal trajectories obtained by Policy-FlowFMT* using 102,400 samples and OCS (ICLOCS2). The OCS trajectory resulted in a time cost of $C_t(\sigma) = 28.18$ s. The cost of the trajectory obtained directly from the tree built using Policy-FlowFMT* is $C_t(\sigma) = 28.91$ s. By integrating the vector field resultant from the tree generated by Policy-FlowFMT* with a 0.1 s time step, we obtain a trajectory that reaches the goal in 29.4 s.

5 Conclusion

This paper presented a method for finding optimal paths and policies for holonomic vehicles moving in the presence of strong currents. We modified FMT* by defining new neighborhood sets that consider the directionality of the vehicle's movement due to the flow and reduce the number of calls to the cost function. These sets are defined using a reachability cone that restricts the space that can be reached when strong currents are found (i.e., the flow speed is greater than the vehicle's relative speed to the flow). The reachability cone is also used to define a minimum-energy cost function for this type of environment.

Additionally, we provide simple modifications to our original path-planning algorithm to calculate global policies to a goal. By using these methods, we can obtain paths with costs similar to the paths obtained using optimal control techniques without significantly increasing the computation time of the solution. Moreover, the proposed methods highly simplify the introduction of obstacles and other constraints, guaranteeing that the optimal homotopy is found when the number of samples is high enough. Finally, we show through a simulation that the policy generated by our method is inherently more resilient to actuation failures than optimal control. Therefore, the methods in this paper should be preferred over previous solutions, especially if the environment contains obstacles and when replanning is not desirable.

Our future work includes the application of the proposed methodology to the motion planning problem of aerial vehicles, also known as aerobots, for *in-situ* exploration of the atmosphere of Venus. These aerobots are up to three times slower than the superrotation winds present on the planet, and even if they offer 3D control, their reachability is constrained by the strong winds. To reach some specific locations, the motion planner needs to find trajectories that circumnavigate the planet. Consequently, for this application, the algorithms proposed in this paper need to be extended to allow for periods of time when the vehicle loses actuation (night side of the planet) due to the lack of solar energy. In these regions, the method needs to be able to let the aerobot drift passively with the flow.

Acknowledgments. The authors thank L5 Automation, Inc. for supporting B. M. Rocamora Jr. during the revision process of this manuscript.

Declarations

conflict of interest, availability of data, ethics approval, funding information and author contribution

Conflict of interest. The authors declare no competing interests.

Data availability. Not applicable.

Ethics approval and consent to participate. Not applicable.

Funding. This work was partially funded by NASA under grant #80NSSC23M0168. B. M. Rocamora Jr. was funded by the Statler Ph.D. Fellowship while he was a student at the Department of Mechanical and Aerospace Engineering at West Virginia University.

Code availability. The code for this research is available in the following GitHub repository: https://bitbucket.org/wvufarolab/flowfmtstar

Author contribution. Conceptualization, B.M.R.J. and G.A.S.P.; methodology, B.M.R.J. and G.A.S.P.; software, B.M.R.J.; validation, B.M.R.J. and G.A.S.P.; formal analysis, B.M.R.J. and G.A.S.P.; investigation, B.M.R.J. and G.A.S.P.; resources, G.A.S.P.; data curation, B.M.R.J.; writing—original draft preparation, B.M.R.J.; writing—review and editing, B.M.R.J. and G.A.S.P.; visualization, B.M.R.J.; supervision, G.A.S.P.; project administration, G.A.S.P.; funding acquisition, G.A.S.P. All authors have read and agreed to the published version of the manuscript.

References

- Alvarez, A., Caiti, A., Onken, R. (2004). Evolutionary path planning for autonomous underwater vehicles in a variable ocean. *IEEE Journal of Oceanic Engineering*, 29(2), 418–429, https://doi.org/10.1109/ JOE.2004.827837
- Blackmore, L., Kuwata, Y., Wolf, M.T., Assad, C., Fathpour, N., Newman, C., Elfes, A. (2010). Global reachability and path planning for planetary exploration with montgolfiere balloons. *Ieee intl. conference on robotics and automation* (pp. 3581–3588).
- Bonin, L., Guitart, A., Delahaye, D., Prats, X., Feron, E. (2023). Computing optimal trajectories for light soaring aircraft using Fast Marching Tree Star. Retrieved 2023-02-13, from https://hal-enac.archivesouvertes.fr/hal-03923994
- Chakrabarty, A., & Langelaan, J. (2013). UAV flight path planning in time varying complex wind-fields. 2013 American Control Conference (pp. 2568–2574).
- Chitsaz, H., & LaValle, S.M. (2007). Timeoptimal paths for a Dubins airplane. *IEEE Conference on Decision and Control* (pp. 2379–2384).
- Garau, B., Alvarez, A., Oliver, G. (2005). Path Planning of Autonomous Underwater Vehicles in Current Fields with Complex Spatial Variability: an A* Approach. *IEEE Intl. Conference on Robotics and Automation* (pp. 194–198).
- Greaves, J.S., Richards, A.M.S., Bains, W., Rimmer, P.B., Sagawa, H., Clements, D.L., ... Hoge, J. (2020). Phosphine gas in the cloud decks of Venus. *Nature Astronomy*, 5(7), 655–664, https://doi.org/10.1038/s41550 -020-1174-4 Retrieved 2023-04-04, from https://www.nature.com/articles/s41550-020-1174-4

- Janson, L., Schmerling, E., Clark, A., Pavone, M. (2015). Fast marching tree: a fast marching sampling-based method for optimal motion planning in many dimensions. *The International Journal of Robotics Research*, 34(7), 883–921,
- Karaman, S., & Frazzoli, E. (2011). Samplingbased algorithms for optimal motion planning. The International Journal of Robotics Research, 30(7), 846–894,
- Kirk, D.E. (2004). Optimal control theory: an introduction. Courier Corporation.
- Koay, T.-B., & Chitre, M. (2013). Energy-efficient path planning for fully propelled AUVs in congested coastal waters. 2013 MTS/IEEE OCEANS - Bergen (pp. 1–9).
- Kularatne, D., Bhattacharya, S., Hsieh, M.A. (2018). Going with the flow: a graph based approach to optimal path planning in general flows. Autonomous Robots, 42(7), 1369–1387, https://doi.org/10.1007/s10514 -018-9741-6 Retrieved 2022-12-07, from http://link.springer.com/10.1007/s10514-018-9741-6
- Kulkarni, C.S., & Lermusiaux, P.F. (2020). Three-dimensional time-optimal path planning in the ocean. Ocean Modelling, 152, 101644, https://doi.org/10.1016/j.ocemod .2020.101644
- Langelaan, J.W. (2008). Tree-based trajectory planning to exploit atmospheric energy. *American Control Conference* (pp. 2328– 2333).
- LaValle, S.M. (2006). *Planning algorithms*. Cambridge University Press.
- Lee, J.J.H., Yoo, C., Hall, R., Anstee, S., Fitch, R. (2017). Energy-optimal kinodynamic planning for underwater gliders in flow fields. Australasian conference on robotics and automation.

- Li, W., Lu, J., Zhou, H., Chow, S.-N. (2017). Method of evolving junctions: A new approach to optimal control with constraints. *Automatica*, 78, 72–78,
- Martinez Rocamora Jr., B., Juan, A.P.I., Pereira, G.A.S. (2022). Towards finding energy efficient paths for hybrid airships in the atmosphere of venus. *Intl. conference on* unmanned aircraft systems (p. 386-393).
- Nagpal, L., & Samdani, K. (2017). Project Loon: Innovating the connectivity worldwide. *IEEE Intl. Conference on Recent Trends in Electronics, Information & Communication Technology* (pp. 1778–1784).
- Nie, Y., Faqir, O., Kerrigan, E.C. (2018). ICLOCS2: Try this Optimal Control Problem Solver Before you Try the Rest. Intl. Conference on Control (pp. 336–336).
- Oettershagen, P., Achermann, F., Müller, B., Schneider, D., Siegwart, R. (2017). Towards Fully Environment-Aware UAVs: Real-Time Path Planning with Online 3D Wind Field Prediction in Complex Terrain. arXiv:1712.03608 [cs], Retrieved 2021-11-18, from http://arxiv.org/abs/1712.03608 (arXiv: 1712.03608)
- Petres, C., Pailhas, Y., Patron, P., Petillot, Y., Evans, J., Lane, D. (2007). Path planning for autonomous underwater vehicles. *IEEE Transactions on Robotics*, 23(2), 331–341,
- Rossi, F., Branch, A., Schodlok, M.P., Stanton, T., Fenty, I.G., Hook, J.V., Clark, E.B. (2021). Stochastic Guidance of Buoyancy Controlled Vehicles under Ice Shelves using Ocean Currents. *IEEE/RSJ Intl. Conference on Intelligent Robots and Systems* (pp. 8657–8664).
- Rossi, F., Saboia, M., Krishnamoorthy, S., Vander Hook, J. (2023). Proximal exploration of venus volcanism with teams of autonomous buoyancy-controlled balloons. *Acta Astronautica*, 208, 389–406,

- Soulignac, M. (2011). Feasible and Optimal Path Planning in Strong Current Fields. *IEEE Transactions on Robotics*, 27(1), 89– 98, https://doi.org/10.1109/TRO.2010 .2085790
- Subramani, D.N., Wei, Q.J., Lermusiaux, P.F. (2018, May). Stochastic time-optimal path-planning in uncertain, strong, and dynamic flows. Computer Methods in Applied Mechanics and Engineering, 333, 218–237, https://doi.org/10.1016/j.cma .2018.01.004 Retrieved 2021-09-07, from https://linkinghub.elsevier.com/retrieve/pii/ S0045782518300069
- VEXAG (2019). Roadmap for Venus Exploration (Tech. Rep.). Venus Exploration Analysis Group.
- Zhai, H., Hou, M., Zhang, F., Zhou, H. (2022). Method of evolving junction on optimal path planning in flows fields. *Autonomous Robots*, 46(8), 929–947,