Stochastic Path Planning for Autonomous Underwater Gliders with Safety Constraints

Chanyeol Yoo[1], Stuart Anstee[2] and Robert Fitch[1]

***Abstract*— Autonomous underwater gliders frequently execute extensive missions with high levels of uncertainty due to limitations of sensing, control and oceanic forecasting. Glider path planning seeks an optimal path with respect to conflicting objectives, such as travel cost and safety, that must be explicitly balanced subject to these uncertainties. In this paper, we derive a set of recursive equations for state probability and expected travel cost conditional on safety, and use them to implement a new stochastic variant of FMT* in the context of two types of objective functions that allow a glider to reach a destination region with minimum cost or maximum probability of arrival given a safety threshold. We demonstrate the framework using three simulated examples that illustrate how user-prescribed safety constraints affect the results.**

I. Introduction

An autonomous underwater glider is an energy-efficient mobile platform that is capable of monitoring the ocean for long periods of time. Underwater gliders are used extensively for applications such as environmental surveying [1], exploration [2] and defence [3].

Gliders have minimal sensing capability and operate in the open ocean. They only have precise localisation when they are on the surface and can access global navigation satellite systems such as GPS. While submerged, they are subject to growing uncertainties that arise from control noise, uncertain forecasts of ocean currents and imprecise underwater localisation, which inhibit closed-loop navigation. These uncertainties can be resolved by surfacing, at the cost of reduced speed of advance and increased risk of collision; hence surfacing must be minimised. It is therefore important to account for noisy state estimation to utilise the full capability of a glider.

The uncertainty of a submerged glider's state (most importantly, position) grows over time, placing the glider at increasing risk of catastrophic failure. While collisions with reefs, facilities and vessels are important, political and regulatory boundaries must also be considered. In fact, virtual obstacles may arise more frequently than physical obstacles in the ocean and are critical for commercial operation of gliders. Hence a path planning algorithm should consider the likelihood of those failures while maximising its objective [4]. These considerations require evaluation of the state probability and the expected cost (i.e., energy consumption) conditional on the path remaining collision-free.

Our approach to finding an optimal path with safety constraints models failure as a singular event from which the glider does not recover. We optimise the expected cost by taking the expectation over paths in which failure does not occur and constraining the expected probability of failure. An exhaustive approach would be computationally expensive and not scalable in the length of the control sequence. Further, such an approach could not benefit from existing computationally-efficient sampling-based path planners since they require *recursivity*.

In this paper, we propose a stochastic variant of a sampling-based planning algorithm for autonomous underwater gliders. The key idea is a set of recursive equations that calculate safety and cost distributions after executing $k+1$ controls using only the distributions after k controls. The full trajectory history is not required. We exploit this property by integrating the recursive state expansion within a sampling-based planner called *fast marching tree** (FMT*) [5] which guarantees asymptotic optimality. We present the stochastic variant of FMT* in the context of two objective functions. First, we consider a control policy that minimises the conditional expected cost given a user-defined safety threshold. Then, we consider a control policy that maximises the probability of arriving inside a designated region, conditional on collision-free transitions.

Our paper contributes a principled method to balance between safety and energy cost using a recursive form of state expansion that can be embedded in a variety of sampling-based algorithms. We demonstrate three simulated examples to show how the proposed framework can be used for different settings. The first example shows how a control policy is evaluated to measure the probability of safety and the conditional expected cost. We then use the framework to solve for the two objective functions and discuss how safety constraints affect the results.

II. Related work

Path planning for autonomous underwater gliders often aims to find an optimal path with respect to a single objective, such as safety or energy cost. Since the glider is designed to operate over an extended duration of time, minimising energy cost has been of primary interest [6, 7]. In this work, both physical and virtual obstacles are implicitly avoided while finding the cost-optimal path. Treatments that neglect safety underestimate the risk that an otherwise optimal may travel dangerously close to obstacles [8]. On the other hand, safety

This research is supported in part by Australia's Defence Science and Technology Group and the University of Technology Sydney

[1]Authors are with the University of Technology Sydney, Ultimo, NSW 2007, Australia {Chanyeol.Yoo, Robert.Fitch}@uts.edu.au

[2]Stuart Anstee is with the Defence Science and Technology Group, Department of Defence, Australia Stuart.Anstee@dst.defence.gov.au

can also be optimised [9, 10]. In existing work, the planner seeks the safest possible path considering risks from shipping traffic and bathymetry [11]. Safety and cost are, however, conflicting criteria; finding an optimal path for one aspect may critically reduce its counterpart, as is shown in [4]. In this paper, we derive a framework that constrains risk, allowing for an explicit balance between safety and cost.

A variety of planning algorithms have been used to find optimal paths for underwater gliders [12–14]. In particular, sampling-based algorithms have been widely used for efficient computation [15–17]. Our recent work introduces FMT* [5] as a promising method for glider path planning due to its efficiency and asymptotic optimality [8, 18]. However, extending these algorithms to the stochastic and safety constrained case is not straightforward, because safety and cost at a state depend not only on the previous state but also on the entire state history. Here, we derive a set of recursive equations that satisfies the requirements for use in sampling-based algorithms and thus can be widely applied.

III. Autonomous Underwater Glider: Models and Notation

A. Trim-based glider model

A trim-based model for representing glider dynamics as a kinematic model is presented in our previous work [8, 18] and summarised here for convenience. The dynamics of an underwater glider G in the presence of ocean currents that vary slowly in space and time relative to its size and rate of motion can be represented as

$$\dot{\mathbf{x}} = f\left(\mathbf{x}(t), \mathbf{u}(t)\right) + \mathbf{V}_c(t), \tag{1}$$

where $\mathbf{x}(t)$ is the 12-dimensional glider state at time t, $\mathbf{u}(t)$ is the control vector and $\mathbf{V}_c$ is the 3-dimensional ocean flow vector. Further details can be found in [6].

The attitude of the glider is controlled by moving a component of its mass relative to its centre of buoyancy and it is propelled by changing its net buoyancy (e.g., by pumping water in and out of a ballast tank). Wings convert a component of the net vertical force acting on the glider to forward horizontal thrust as the glider moves up and down in the z-direction. Depending on the glider's design, the control vector $\mathbf{u}(t)$ may consist, for example, of forces $\mathbf{u}_{r_p}$ acting on the moving mass and u_{m_b}, the scalar rate at which the ballast mass is changing, thus $\mathbf{u}(t) = \left[\mathbf{u}_{r_p}(t)\ u_{m_b}(t)\right]^T$.

Controlling the glider in this way is energy expensive and therefore limits the endurance of the glider. We seek to minimise control expenditure. Instead of continuously varying the control over time, we use the concept of *trim states*. A trim state is a state of dynamic equilibrium in which the glider will continue indefinitely in the absence of further control inputs [6]. Essentially, the glider cycles between upward and downward motion and at each net buoyancy inversion point or waypoint, which is triggered by reaching a target depth, we set the controls and let the glider go until it reaches the next waypoint (target depth).

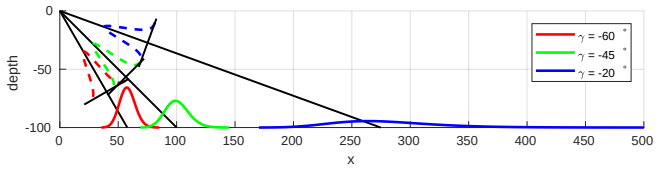

Fig. 1. Change in probability variances in with respect to different glide angles and Gaussian control noise. A 2D ($x - z$) slice is shown for illustration, with three different glide angles values γ. Dashed lines represent control noise, solid black lines show deterministic paths in the absence of control noise, and coloured lines indicate resulting distribution over position at target depth.

Using trim states, the dynamic problem is reduced to a kinematic problem, where the kinematic state $\mathbf{x}_k$ of the glider at the kth waypoint has of 6 degrees of freedom,

$$\mathbf{x}_k = \left[x_k, y_k, z_k, \gamma_k, \theta_k, m_{bk}\right]^T, \tag{2}$$

where (x_k, y_k, z_k) is the position of the glider at control index k, $\gamma_k \in [\pm\gamma_{\min}, \pm\gamma_{\max}]$ is the glide angle, θ_k is the heading angle and m_b is the ballast mass. We assume the ballast tank is either empty or full (i.e., $m_b \in \{0, m_{b\,\max}\}$). Changing trim state from $\mathbf{x}_k$ to $\mathbf{x}_{k+1}$ at each waypoint incurs a cost r_k. We assume that trim transitions are instantaneous.

In this paper, we denote the sequences of glider states and controls as

$$\begin{aligned} \mathbf{X}^n &= \mathbf{x}_0\mathbf{x}_1\cdots\mathbf{x}_n = \mathbf{X}^{n-1}\mathbf{x}_n \\ \mathbf{U}^n &= \mathbf{u}_0\mathbf{u}_1\cdots\mathbf{u}_n = \mathbf{U}^{n-1}\mathbf{u}_n. \end{aligned} \tag{3}$$

Given two consecutive glider states $\mathbf{x}_k$ and $\mathbf{x}_{k+1}$ with control $\mathbf{u}_k$, a safety variable $\mathbf{v}_k$ is true if the transition is collision free. We denote $\mathbf{V}^n$ as the sequence of safety variables over $\mathbf{X}^n$. The transition cost from $\mathbf{x}_k$ to $\mathbf{x}_{k+1}$ is denoted as $r(\mathbf{x}_k, \mathbf{x}_{k+1})$. In this paper, we use Euclidean distance as a simple cost metric to demonstrate the proposed framework; a more realistic approach would normally apply.

We denote by $\mathbf{P}^n = \mathbf{p}_0\cdots\mathbf{p}_n$, where $\mathbf{p}_k = [x_k, y_k, z_k]^T$, the sequence of glider positions at states $\mathbf{X}^n$. The trim state τ_k that manoeuvres the glider from $\mathbf{p}_k$ to $\mathbf{p}_{k+1}$ in the presence of ocean flow $\mathbf{V}_c$ is

$$\tau_k = \left[v_k, \gamma_k, \theta_k, m_{bk}\right]^T, \tag{4}$$

where v_k is the velocity, which is determined by the glide angle, ballast mass and oceanic current. We refer to [6, 8] for a description of how this trim state is calculated.

Let the control policy be π, which is a sequence of trim states that is fixed during execution. This is because the glider controller is open-looped and can only access its true final position and forecast oceanic currents while surfaced. The control instance at control index k is of the form

$$\pi_k = [\gamma_k^d, \theta_k^d, m_{bk}^d, z_k]^T, \tag{5}$$

where superscript d indicates 'demanded' and z_k is the target depth for the next transition. We slightly abuse previous notation here to denote by $\mathbf{X}^\pi$, $\mathbf{U}^\pi$ and $\mathbf{V}^\pi$ the sequences of state, control and safety induced by policy π.

B. Transition probability in trim state

The glider G is subject to noise in glide and heading angles that arises from control and sensing uncertainties. In this paper, we neglect noise in the depth z, since in practice the associated uncertainty is small compared to other sources.

Given the desired control $\pi_k = [\gamma_k^d, \theta_k^d, m_{bk}^d, z_k]^T$ at the k-th control instant, the glide and heading angles are subject to Gaussian noise, such that

$$\begin{aligned} \hat{\gamma}_k &\sim \mathcal{N}(\gamma_k^d, \sigma_{\gamma_k}^2) \\ \hat{\theta}_k &\sim \mathcal{N}(\theta_k^d, \sigma_{\theta_k}^2). \end{aligned} \tag{6}$$

We denote the transition probability from state $\mathbf{x}_k$ to $\mathbf{x}_{k+1}$ given policy π as $\mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{x}_k)$ and the probability of safe transition as $\mathcal{P}(\mathbf{v}_{k+1} \mid \mathbf{x}_k\mathbf{x}_{k+1})$, where $\mathbf{v}_{k+1}$ is true when all $\mathbf{x} \in (\mathbf{x}_k, \mathbf{x}_{k+1}]$ subject to control instance and ocean currents are safe. For demonstration purposes, we omit the influence of oceanic currents and the associated uncertainties.

Figure 1 illustrates how the control noise in (6) affects the transition probability. The important properties are: 1) the smaller the glide angle γ, the more distributed the transition probability, and 2) the transition probability is not Gaussian over the horizontal plane. We discuss the former property later, in detail.

IV. Problem formulation

We consider a stochastic path planning problem for an autonomous underwater glider subject to control noise as described in Sec. III. Given a glider G with Gaussian control noise in glide and heading angles γ and θ, initial surface position x_0 and final surface position x_g, we seek an optimal control policy π^* with respect to objective functions that are defined over probabilities and expected costs. We denote by $\mathcal{P}_\pi(\cdots)$ and $\mathcal{R}_\pi(\cdots)$ the probability and expected cost functions given policy π.

In [4], we showed that solving for a single objective often induces a solution that is overly biased towards that objective, while other parameters are ignored. For instance, maximising the probability of success $\mathcal{P}_\pi(\mathbf{V})$ may have unacceptably high cost, and minimising the expected cost $\mathcal{R}_\pi(\mathbf{V})$ may induce a low probability of success. In this paper, we formulate objective functions in such a way that one aspect is optimised while the other acts as a constraint. We present two types of such objective functions. We consider two problems in this paper as follows.

Problem 1 (Minimise manoeuvre cost given satisfaction threshold)**.** *Find the optimal control policy π^* that minimises the expected cost given satisfaction threshold α,*

$$\begin{aligned} \pi^* &= \arg\min_\pi \mathcal{R}_\pi(\mathbf{V}^\pi) \\ &\textit{s.t. } \mathcal{P}_\pi(\mathbf{V}^\pi) > \alpha, \end{aligned} \tag{7}$$

where $\mathcal{R}_\pi(\mathbf{V}^\pi)$ is the expected cost given safe completion.

Problem 2 (Maximise the conditional in-region probability)**.** *Find the optimal control policy π^* that maximises the probability of surfacing within the desired region $R_g \subset \mathbb{R}^2$,*

$$\pi^* = \arg\max_\pi \mathcal{P}_\pi(\mathbf{x}_n \in R_g \mid \mathbf{V}^n), \tag{8}$$

where n is the final trim control in policy π.

V. Probabilistic evaluation of control policy

We present algorithms for recursively evaluating a control policy given a noisy controller. The functions critical for solving the problems in Sec. IV are the *conditional state probability* $\mathcal{P}_\pi(\mathbf{x}_{k+1} = \mathrm{x} \mid \mathbf{V}^{k+1})$, the probability of being positioned at x after $k+1$ control instants given collision-free paths and the *conditional state cost* $\mathcal{R}_\pi(\mathbf{x}_{k+1} = \mathrm{x} \mid \mathbf{V}^{k+1})$, the corresponding expected cost.

The important aspect of these functions is that they are defined recursively, in that the $k+1$-th control can be computed using only the result of the k-th control. Also, the functions are over two known probability functions including the transition probability $\mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{x}_k)$ and the probability of safe transition $\mathcal{P}_\pi(\mathbf{v}_{k+1} \mid \mathbf{x}_k\mathbf{x}_{k+1})$.

In this paper, we assume that the Markov property holds true over the environment. Notably, the transition and transition safety probabilities at control index $k+1$ only depend on those at k and $k+1$, such that

$$\begin{aligned} \mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{x}_k) &\equiv \mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{x}_k, \cdots) \\ \mathcal{P}_\pi(\mathbf{v}_{k+1} \mid \mathbf{x}_k, \mathbf{x}_{k+1}) &\equiv \mathcal{P}_\pi(\mathbf{v}_{k+1} \mid \mathbf{x}_k, \mathbf{x}_{k+1}, \cdots). \end{aligned} \tag{11}$$

Then, we have the conditional safety probability $\mathcal{P}_\pi(\mathbf{V}^{k+1} \mid \mathbf{V}^k)$ that is derived as

$$\begin{aligned} \mathcal{P}_\pi(\mathbf{V}^{k+1} \mid \mathbf{V}^k) &\equiv \frac{\mathcal{P}_\pi(\mathbf{V}^k \mid \mathbf{V}^k, \mathbf{v}_{k+1})\mathcal{P}_\pi(\mathbf{V}^{k+1})}{\mathcal{P}_\pi(\mathbf{V}^k)} \\ &\equiv \frac{\mathcal{P}_\pi(\mathbf{V}^{k+1})}{\mathcal{P}_\pi(\mathbf{V}^k)} \end{aligned} \tag{12}$$

Intuitively, the conditional safety probability is the probability of safety during $\mathbf{X}^{k+1}$ given safety during $\mathbf{X}^k$.

A. Probability over glider state given success

We are interested in determining the probability distribution of the state of the glider after the k-th control instant, namely $\mathcal{P}_\pi(\mathbf{x}_k)$. However, this form is not the right representation because it also considers failed paths (i.e., obstructed paths). With safety in mind, we consider the probability conditional on success, $\mathcal{P}_\pi(\mathbf{x}_k \mid \mathbf{V}^k)$.

The conditional probability at control index $k+1$ is computed recursively by integrating over a conditional product probability function for $\mathbf{x}_k$ and $\mathbf{x}_{k+1}$. Using the Markov property, the conditional product probability function is shown in (9), The conditional safety probability $\mathcal{P}_\pi(\mathbf{V}^{k+1} \mid \mathbf{V}^k)$ serves as a normaliser to the integral in (9), which can be easily computed. It is important to note that the conditional state probability at $k+1$ is recursively computed only with transition and safety probability.

Using the Markov property, we can express the conditional satisfaction property at control index k as

$$\begin{aligned} \mathcal{P}_\pi(\mathbf{V}^k) &= \mathcal{P}_\pi(\mathbf{v}_0) \cdot \prod_{i=1}^{k} \mathcal{P}_\pi(\mathbf{V}^i \mid \mathbf{V}^{i-1}) \\ &= \mathcal{P}_\pi(\mathbf{V}^{k-1}) \cdot \mathcal{P}_\pi(\mathbf{V}^k \mid \mathbf{V}^{k-1}). \end{aligned} \tag{13}$$

$$
\begin{aligned}
\mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{V}^{k+1}) &= \int_{\mathbf{x}_k \in \mathbb{R}^2} \mathcal{P}_\pi(\mathbf{x}_k, \mathbf{x}_{k+1} \mid \mathbf{V}^{k+1}) \cdot d\mathbf{x}_k \\
&= \frac{1}{\mathcal{P}_\pi(\mathbf{V}^{k+1} \mid \mathbf{V}^k)} \int_{\mathbf{x}_k \in \mathbb{R}^2} \mathcal{P}_\pi(\mathbf{v}_{k+1} \mid \mathbf{x}_k, \mathbf{x}_{k+1}) \cdot \mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{x}_k) \cdot \mathcal{P}_\pi(\mathbf{x}_k \mid \mathbf{V}^k) \cdot d\mathbf{x}_k
\end{aligned} \tag{9}
$$

$$
\begin{aligned}
\mathcal{R}_\pi(\mathbf{x}_{k+1} \mid \mathbf{V}^{k+1}) &= \mathbb{E}_\pi\left[r_0 + \cdots + r_{k+1} \mid \mathbf{V}^{k+1}, \mathbf{x}_{k+1}\right] \\
&= \mathbb{E}_\pi\left[\sum_{i=0}^{k+1} r_i \mid \mathbf{V}^{k+1}, \mathbf{x}_{k+1}\right] \\
&= \int_{\mathbf{x}_0 \in \mathbb{R}^2} \cdots \int_{\mathbf{x}_k \in \mathbb{R}^2} \mathcal{P}_\pi(\mathbf{x}_0 \cdots \mathbf{x}_k \mid \mathbf{V}^{k+1}, \mathbf{x}_{k+1}) \cdot \sum_{i=0}^{k} r(\mathbf{x}_i, \mathbf{x}_{i+1}) \cdot d\mathbf{x}_0 \cdots d\mathbf{x}_k \\
&= \int_{\mathbf{x}_0 \in \mathbb{R}^2} \cdots \int_{\mathbf{x}_k \in \mathbb{R}^2} \mathcal{P}_\pi(\mathbf{x}_0 \cdots \mathbf{x}_k \mid \mathbf{V}^{k+1}, \mathbf{x}_{k+1}) \cdot \left(r(\mathbf{x}_k, \mathbf{x}_{k+1}) + \sum_{i=0}^{k-1} r(\mathbf{x}_i, \mathbf{x}_{i+1}) \cdot d\mathbf{x}_0\right) \cdots d\mathbf{x}_k \\
&= \int_{\mathbf{x}_k \in \mathbb{R}^2} \frac{\mathcal{P}_\pi(\mathbf{v}_{k+1} \mid \mathbf{x}_k, \mathbf{x}_{k+1}) \cdot \mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{x}_k) \cdot \mathcal{P}_\pi(\mathbf{x}_k \mid \mathbf{V}^k)}{\mathcal{P}_\pi(\mathbf{V}^{k+1} \mid \mathbf{V}^k) \cdot \mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{V}^{k+1})} \\
&\cdot \left(\int_{\mathbf{x}_0 \in \mathbb{R}^2} \cdots \int_{\mathbf{x}_{k-1} \in \mathbb{R}^2} \mathcal{P}(\mathbf{X}^{k-1} \mid \mathbf{V}^k, \mathbf{x}_k) \cdot \left(r(\mathbf{x}_k, \mathbf{x}_{k+1}) + \sum_{i=0}^{k-1} r(\mathbf{x}_i, \mathbf{x}_{i+1}) \cdot d\mathbf{x}_0\right) \cdot d\mathbf{x}_0 \cdots d\mathbf{x}_{k-1}\right) \cdot d\mathbf{x}_k \\
&= \int_{\mathbf{x}_k \in \mathbb{R}^2} \frac{\mathcal{P}_\pi(\mathbf{v}_{k+1} \mid \mathbf{x}_k, \mathbf{x}_{k+1}) \cdot \mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{x}_k) \cdot \mathcal{P}_\pi(\mathbf{x}_k \mid \mathbf{V}^k)}{\mathcal{P}_\pi(\mathbf{V}^{k+1} \mid \mathbf{V}^k) \cdot \mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{V}^{k+1})} \cdot \left(r(\mathbf{x}_k, \mathbf{x}_{k+1}) + \mathcal{R}(\mathbf{x}_k \mid \mathbf{V}^k)\right) \cdot d\mathbf{x}_k \\
&= \frac{\displaystyle\int_{\mathbf{x}_k \in \mathbb{R}^2} \mathcal{P}_\pi(\mathbf{v}_{k+1} \mid \mathbf{x}_k, \mathbf{x}_{k+1}) \cdot \mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{x}_k) \cdot \mathcal{P}_\pi(\mathbf{x}_k \mid \mathbf{V}^k) \cdot \left(r(\mathbf{x}_k, \mathbf{x}_{k+1}) + \mathcal{R}(\mathbf{x}_k \mid \mathbf{V}^k)\right) \cdot d\mathbf{x}_k}{\mathcal{P}_\pi(\mathbf{V}^{k+1} \mid \mathbf{V}^k) \cdot \mathcal{P}_\pi(\mathbf{x}_{k+1} \mid \mathbf{V}^{k+1})}
\end{aligned} \tag{10}
$$

Note that the probability is recursively computed from the sequence of conditional safety probabilities described in (12). The algorithm for finding the probability is shown in Alg. 1.

B. Conditional state cost

The conditional state cost $\mathcal{R}(\mathbf{x}_{k+1} \mid \mathbf{V}^{k+1})$ is computed recursively at control index $k+1$ given that at control index k, using the conditional probability $\mathcal{P}_\pi(\mathbf{x}_k \mid \mathbf{V}^k)$. The proof and method of computation are shown in (10), where $r(\mathbf{x}_k, \mathbf{x}_{k+1})$ is the state transition cost.

Similarly, computing the expected cost $\mathcal{R}_\pi(\mathbf{x}_{k+1})$ given a policy may not make sense, because the result includes failed paths. Hence, we use the conditional state cost $\mathcal{R}(\mathbf{x}_{k+1} \mid \mathbf{V}^{k+1})$. The expected conditional cost given success is

$$
\mathcal{R}(\mathbf{V}^k) = \int_{\mathbf{x}_k \in \mathbb{R}^2} \mathcal{R}(\mathbf{x}_k \mid \mathbf{V}^k) \cdot \mathcal{P}(\mathbf{x}_k \mid \mathbf{V}^k) \cdot d\mathbf{x}_k, \tag{14}
$$

similar to $\mathcal{P}_\pi(\mathbf{V}^{k+1})$. The computational algorithm is given in Alg. 1.

Algorithm 1 Expected cost of action given success

Inputs: $\mathbf{X}$, $P_{\mathbf{v}}$, $R_{\mathbf{v}}$, $\mathbf{P}_{\mathbf{xv}}$, $\mathbf{R}_{\mathbf{xv}}$, $u = \{\lambda, \theta, \Delta z\}$
Outputs: $\mathbf{X}'$, $P_{\mathbf{v}'}$, $R_{\mathbf{v}'}$, $\mathbf{P}_{\mathbf{x}'\mathbf{v}'}$, $\mathbf{R}_{\mathbf{x}'\mathbf{v}'}$

enumerate X1 and get 5-sd rectangle to get pre-X1
calculate non-collision prob over pre-X1
shrink and compute X1 and dX1 (or dA1 for area)
$\mathbf{X}' \leftarrow$ Evenly spaced such that $\mathcal{P}(\mathbf{x}' \mid \mathbf{x}, u \pm \sigma) > 0$ for all $\mathbf{x} \in \mathbf{X}$
$n' = |\mathbf{X}'|$
$\mathbf{P}_{\mathbf{x}'\mathbf{v}'}[0 \cdots n'] \leftarrow 0$
$\mathbf{R}_{\mathbf{x}'\mathbf{v}'}[0 \cdots n'] \leftarrow 0$
for all $\mathbf{x} \in \mathbf{X}$ s.t. $\mathbf{P}_{\mathrm{xv}}[\mathbf{x}] > 0$ **do**
  for all $\mathbf{x}' \in \mathbf{X}'$ s.t. $\mathbf{P}_{\mathrm{x'v'}}[\mathbf{x}'] > 0$ **do**
    $p_{\mathbf{x}'\mathbf{x}} \leftarrow \mathcal{P}(\mathbf{x}' \mid \mathbf{x}, u)$
    $c_{\mathbf{x}'\mathbf{x}} \leftarrow \mathcal{C}(\mathbf{x}' \mid \mathbf{x})$
    $r_{\mathbf{x}'\mathbf{x}} \leftarrow \mathcal{R}(\mathbf{x}' \mid \mathbf{x})$
    $\mathbf{P}_{\mathbf{x}'\mathbf{v}'}[\mathbf{x}'] \leftarrow c_{\mathbf{x}'\mathbf{x}} \cdot p_{\mathbf{x}'\mathbf{x}} \cdot \mathbf{P}_{\mathbf{xv}}[\mathbf{x}]$
    $\mathbf{R}_{\mathbf{x}'\mathbf{v}'}[\mathbf{x}'] \leftarrow c_{\mathbf{x}'\mathbf{x}} \cdot p_{\mathbf{x}'\mathbf{x}} \cdot \mathbf{P}_{\mathbf{xv}}[\mathbf{x}] \cdot (r_{\mathbf{x}'\mathbf{x}} + \mathbf{R}_{\mathbf{xv}}[\mathbf{x}])$
$\mathbf{P}_{\mathbf{x}'\mathbf{v}'} \leftarrow \mathbf{P}_{\mathbf{x}'\mathbf{v}'} / P_{\mathbf{v}'\mathbf{v}}$
$\mathbf{R}_{\mathbf{x}'\mathbf{v}'} \leftarrow \mathbf{R}_{\mathbf{x}'\mathbf{v}'} / P_{\mathbf{v}'\mathbf{v}} / \mathbf{P}_{\mathbf{x}'\mathbf{v}'}$
$P_{\mathbf{v}'\mathbf{v}} \leftarrow \sum_{\mathbf{x}' \in \mathbf{X}'} \mathbf{P}_{\mathbf{x}'\mathbf{v}'}[\mathbf{x}']$
$P_{\mathbf{v}'} \leftarrow P_{\mathbf{v}} \cdot P_{\mathbf{v}'\mathbf{v}}$
$R_{v'} \leftarrow \sum_{\mathbf{x}' \in \mathbf{X}'} (\mathbf{R}_{\mathbf{x}'\mathbf{v}'}[\mathbf{x}'] \cdot \mathbf{P}_{\mathbf{x}'\mathbf{v}'}[\mathbf{x}'])$
reduce X'

VI. Stochastic sampling-based planning with safety constraints

In this section, we present a stochastic planning framework based on the recursive equations introduced in Sec. V. We present a stochastic FMT*implementation, which is a stochastic variant of the sampling-based algorithm FMT*. The deterministic components of the original algorithm are replaced with recursive functions to introduce uncertainty.

Algorithm 2 Stochastic FMT*

Inputs: sample set V, starting sample $v_0 \in V$, destination sample $v_f \in V$, transition probability $\mathcal{P}(\mathbf{x}' \mid \mathbf{x})$, collision probability $\mathcal{P}(\mathbf{v}_i \mid \mathbf{x}_i\mathbf{x}_{i+1})$, constrained objective Ω

Initialise $\mathcal{P}_0(\mathbf{x}_0 \mid \mathbf{V}) = 1$, 0 otherwise
$V_{open} \leftarrow v_0$
$V_{unvisited} \leftarrow V \backslash v_0$
for all $z \in V_{open}$ such that $z \leftarrow \arg\Omega^*(V_{open})$ **do**
  for all $x \in V_{unvisited}$ **do**
    Find $Y \subseteq V_{open}$ such that $\mathcal{P}(\mathbf{v} \mid (y \to x)) > 0$ and y satisfies constraint in Ω
    Find control instants $\pi(y, x)$ for all $y \in Y$
    Find $\mathcal{P}_{y\to x}(\cdot)$ and $\mathcal{R}_{y\to x}(\cdot)$ for all $y \in Y$
    Find $y^* \in Y$ such that $y^* \leftarrow \arg\Omega^*(V_{open})$
  Add x to V_{open} and remove from $V_{unvisited}$
  Add z to V_{closed} and remove from V_{open}
  if $v_f \in V_{open}$ **then**
    Find control policy π^* from v_0 to v_f
    return

We then discuss how we modified the original algorithm to achieve energy-oriented manoeuvres and how the change affects the algorithm.

A. FMT* overview

The FMT*algorithm starts with a set of sample nodes V in the state space (as in PRM*). The samples are initially labelled as *unvisited* aside from the initial position $\mathbf{v}_0$, which is labelled as *opened*. Detailed FMT*implementations for deterministic gliders are shown in [8, 18].

The algorithm incrementally grows a branching tree until no opened nodes remain. The procedure eventually finds an optimal solution. At every branching instant, we first find an open node $z \in Z$ with the lowest overall cost. From z, we find a set of unvisited nodes $xNear$ with the least transition cost. For each unvisited node $x \in xNear$, we find the open node y for which the sum of the cumulative cost at y and the transition cost from y to x is the minimum. After we find such an open node y, we then check if the transition is collision free. If it is, we set z to *closed* and any newly connected states in $xNear$ become *opened*.

B. Stochastic variant of FMT*

In stochastic FMT*, which we present as Alg. 2, we find a set of sample positions $p_i = [x_i, y_i, z_i]^T \in \mathbb{R}^3$. Each sample position serves as a reference node from which we derive the corresponding control instant. A node $v_i \in V$ is given as

$$v_i = \left[x_i \; y_i \; z_i \; p_{vi} \; r_{vi} \; p_{xvi} \; r_{xvi}\right], \tag{15}$$

where p_{vi}, r_{vi}, p_{xvi} and r_{xvi} correspond to $\mathcal{P}(\mathbf{V})$, $\mathcal{R}(\mathbf{V})$, $\mathcal{P}(\mathbf{x} \mid \mathbf{V})$ and $\mathcal{R}(\mathbf{x} \mid \mathbf{V})$ that achieve $\Omega^*(v_i)$, where Ω is the objective function of interest.

For Problem 1, we impose an additional constraint on Alg. 2 so that an edge is built when the safety condition is met. For Problem 2, we add an additional value p_{ri} to all $v_i \in V$, which denotes the conditional in-region probability $\mathcal{P}(\mathbf{x} \in R_g \mid \mathbf{V})$ that we maximise (i.e., $\Omega^*(v_{goal}) = p_{r,goal}$).

TABLE I

EVALUATION OF A GIVEN CONTROL POLICY IN FIG. 2 FOR CONDITIONAL SAFETY PROBABILITY AND EXPECTED COST

k	0	1	2	3	4
$\mathcal{P}_\pi(\mathbf{V}^k)$	1	0.7971	0.7616	0.7587	0.7001
$\mathcal{R}_\pi(\mathbf{V}^k)$	0	494.917	667.709	960.576	1303.407

Algorithm 2 starts with one open node $V_{open} = v_0$ at the starting position given prior distributions $\mathcal{P}_0(\cdot)$ and $\mathcal{R}_0(\cdot)$. At each open node $z \in V_{open}$, the algorithm finds the set of unvisited nodes $x \in V_{unvisited}$ with non-zero transition probability (i.e., $\mathcal{P}(\mathbf{v} \mid (y \to x)) > 0$). We enumerate the set and compute $\mathcal{P}_{y\to x}(\cdot)$ and $\mathcal{R}_{y\to x}(\cdot)$ and find $y^* \in Y$ that is optimal with respect to the given constrained objective function Ω^*. The edge from y^* to x is added to the tree and the distributions are updated (i.e., $\mathcal{P}_x(\cdot) = \mathcal{P}_{y^*\to x}(\cdot)$ and $\mathcal{R}_x(\cdot) = \mathcal{R}_{y^*\to x}(\cdot)$). The algorithm finishes when the best $z \in V_{open}$ is the goal.

C. Analysis

Stochastic FMT*ignores the branching bound in FMT*, where branching from a node is bounded by distance or number of candidate nodes (i.e., k-nearest neighbours). Although the computational complexity increases from $\mathcal{O}(n \log n)$ to $\mathcal{O}(n^2 \log n)$, this modification is beneficial in the context of energy consumption. Recalling that energy is consumed when changing trim state, manoeuvring in a straight line is more efficient than an industry-standard 'sawtooth'-like path generated by fixed glide angles.

We fully exploit the benefits of recursive functions by integrating them naturally into the existing FMT*framework. With this property, the computational complexity for each iteration is constant and invariant with respect to time horizon or number of control instants, because we use results from the previous iteration to update the next. Unlike deterministic FMT*, where collision checking is done after finding an optimal solution, we check for collision beforehand. This is due to two reasons: 1) computing $\mathcal{P}_{y\to x}(\cdot)$ and $\mathcal{R}_{y\to x}(\cdot)$ is much more expensive than collision checking, and 2) the uncertainty of the environment means that probability of success is not solely determined by collision.

Although we forego lazy collision checking, this FMT*variant is still computationally efficient relative to other sampling-based methods such as PRM. A stochastic variant of PRM using the same recursive framework would require enumeration of every valid transition in the environment, whereas stochastic FMT*only enumerates a portion of the transitions, similar to the deterministic case.

VII. CASE STUDIES

We present simulated examples of the underwater glider using our stochastic planner. We demonstrate how a control policy is evaluated and present an example where we compare our framework to a industry-standard approach with

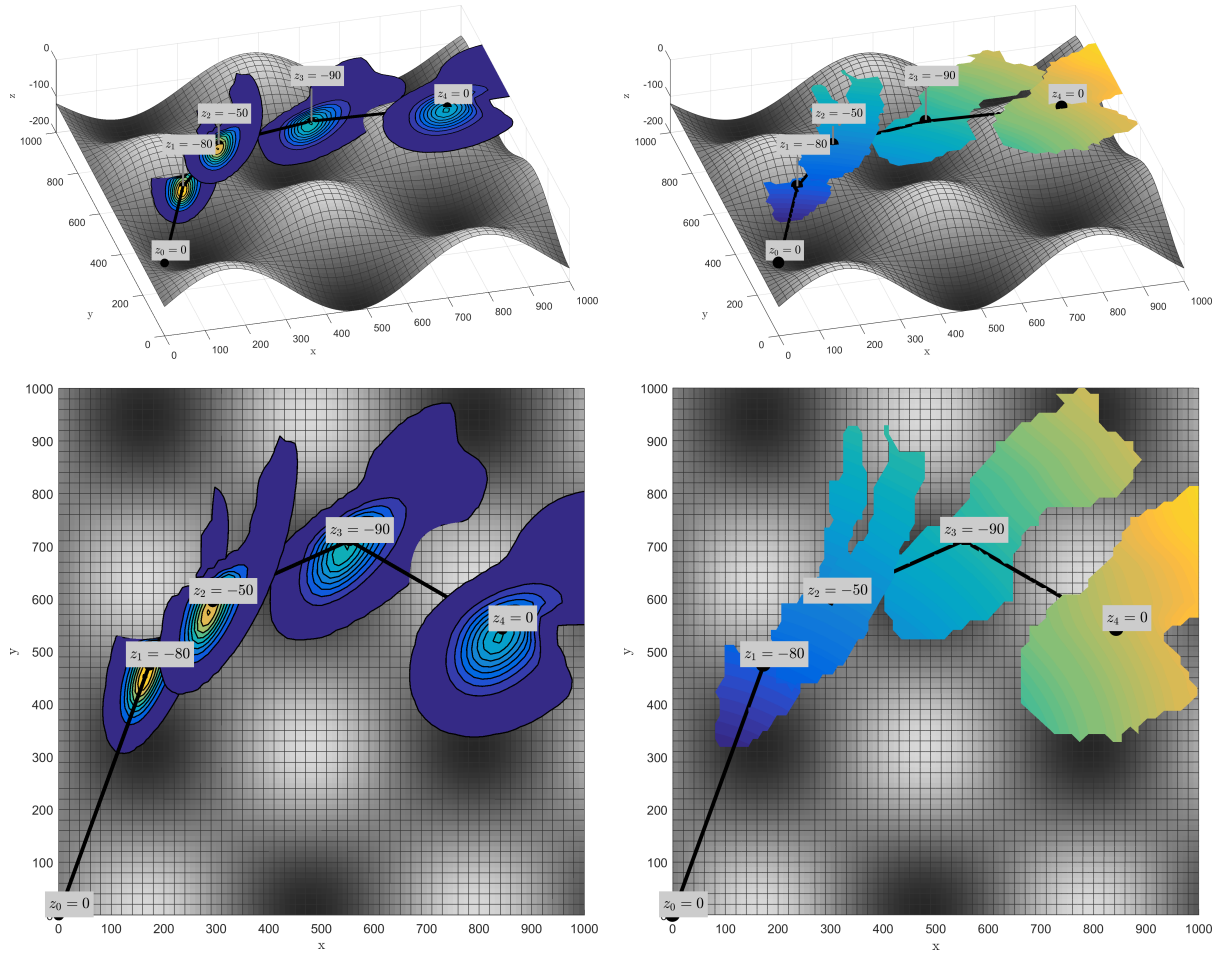

(a) Conditional state probability over control policy

(b) Conditional expected cost over control policy

Fig. 2. Evaluation of control policy (16) conditional to collision-free path. The glider starts at $[0, 0, 0]^T$ and the ocean floor is shown in shaded black. Black lines indicate deterministic paths in the absence of control noise. Coloured regions indicate conditional state probability (a) and expected cost (b) distributions. Blue represents low probability/cost; yellow represents high probability/expected cost. Target depths z_k are labelled. Oblique and top-down views are shown.

respect to probabilistic satisfaction and expected cost. For demonstration purposes, we set the actuation noise for glide and heading angles to 0.5 rad and 0.8 rad, respectively. As mentioned in Sec. III-B, we omit the ocean currents $\mathbf{V}_c$ for clearer demonstration of the recursive framework, which is otherwise capable of accounting for stochastic ocean currents.

A. Evaluating a control policy

In this example, we evaluate the conditional state probability and cost after completing each control instance in a policy π in the presence of underwater obstacles. We consider a policy

$$\pi(k) = \begin{array}{c|ccc} k & \lambda_k & \theta_k & z_k\ (m) \\ \hline 0 & \text{-}0.1571 & 1.2217 & \text{-}80 \\ 1 & 0.1745 & 0.7854 & \text{-}50 \\ 2 & \text{-}0.1396 & 0.4189 & \text{-}90 \\ 3 & 0.2618 & \text{-}0.5236 & 0 \end{array}, \qquad (16)$$

where λ_k, θ_k and z_k are the desired glide angle, heading and depth at the k-th control instant.

Starting from $\mathbf{p}_0 = [0, 0, 0]^T$, the conditional state probabilities and costs after each control instance are shown in Fig. 2. Table I shows the probability of success $\mathcal{P}_\pi(\mathbf{V}^k)$ and the expected cost $\mathcal{R}_\pi(\mathbf{V}^k)$ at the k-th time step. The result shows that the success probability is 0.7001 and the expected cost is 1303.407 at the end of the manoeuvre.

The figures illustrate clearly how conditional probability propagates over the control instants. The position estimates show how obstacles affect the estimates of dive-out position. Unlike existing work that approximates position estimates using a Gaussian assumption with mean and variance, our framework is non-Gaussian and therefore provides more accurate results for underwater gliders where safety of the platform is critical.

B. Minimise the expected cost given probability threshold

In this example, we synthesise control policies that minimise conditional state cost given various probability thresholds. The objective function is described in Problem 1. In Fig. 3, we demonstrate the conditional state probabilities $\mathcal{P}_{\pi^*}(\mathbf{x} \mid \mathbf{V})$ for the case with no threshold (i.e.,

TABLE II

MINIMISING THE CONDITIONAL STATE COST GIVEN SAFETY CONSTRAINTS IN FIG. 3. THE SATISFACTION PROBABILITIES $\mathcal{P}_{\pi^*}(\mathbf{V}^k)$ AND THE EXPECTED COSTS $\mathcal{R}_{\pi^*}(\mathbf{V}^k)$ ARE SHOWN FOR THREE DIFFERENT SAFETY THRESHOLD α.

	Safety Constraint	$k=0$	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$
$\mathcal{P}_{\pi^*}(\mathbf{V}^k)$	No safety constraint	1	0.6624	0.5584	0.5373	0.5288	–	–
	$\mathcal{P}_{\pi^*}(\mathbf{V}^k) > 0.6$	1	0.9627	0.8009	0.7209	0.6315	0.6296	–
	$\mathcal{P}_{\pi^*}(\mathbf{V}^k) > 0.9$	1	0.9987	0.9548	0.9395	0.9225	–	–
	$\mathcal{P}_{\pi^*}(\mathbf{V}^k) > 0.99$	1	1	1	1	1	1	1
$\mathcal{R}_{\pi^*}(\mathbf{V}^k)$	No safety constraint	0	288.3	590.0	802.4	1017.0	–	–
	$\mathcal{P}_{\pi^*}(\mathbf{V}^k) > 0.6$	0	321.0	630.0	876.2	986.5	1120.9	–
	$\mathcal{P}_{\pi^*}(\mathbf{V}^k) > 0.9$	0	222.2	537.4	762.6	1081.4	–	–
	$\mathcal{P}_{\pi^*}(\mathbf{V}^k) > 0.99$	0	230.7	556.2	773.9	890.4	1003.2	1115.7

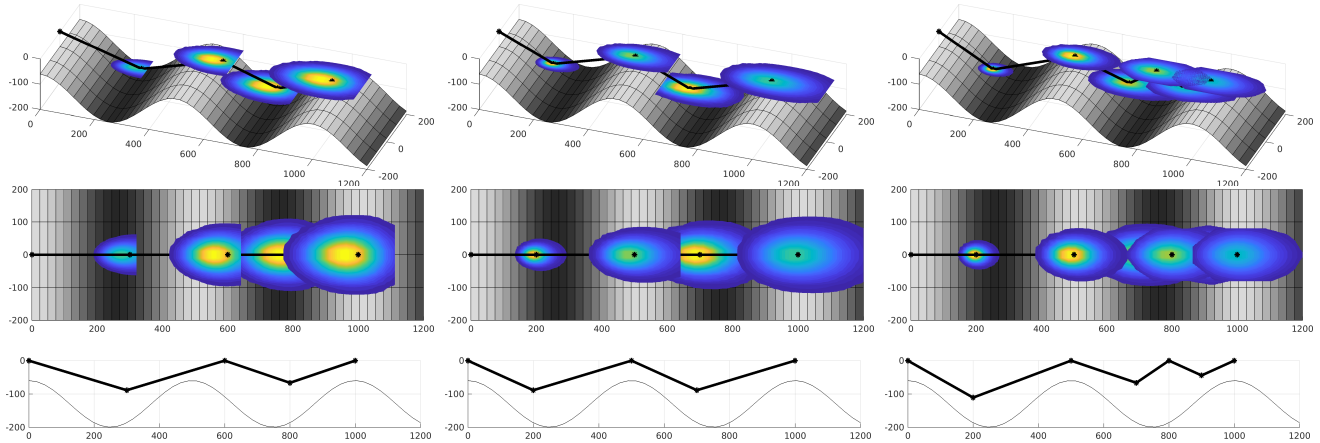

(a) No safety constraint (i.e., $\alpha = 0$) (b) Given $\mathcal{P}_{\pi^*}(\mathbf{V}) > 0.9$ (c) Given $\mathcal{P}_{\pi^*}(\mathbf{V}) > 0.99$

Fig. 3. Minimum conditional state costs given different probability threshold α. Glider paths and state distributions are shown as in Fig. 2. Changes in depth are also shown; dark lines indicate deterministic paths and shaded lines indicate ocean floor geometry.

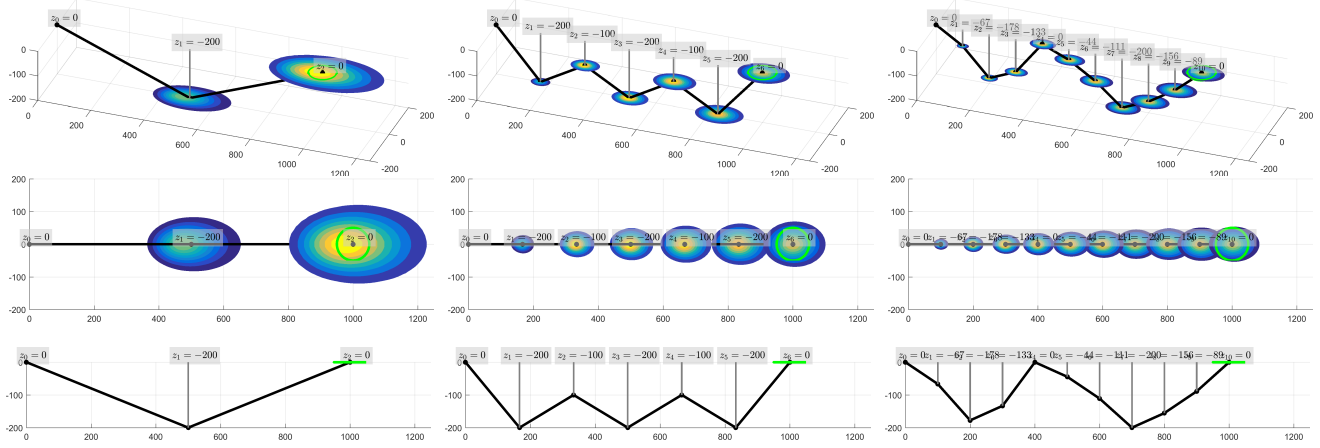

(a) 1-cycle, $\mathcal{P}_{\pi_1}(\mathbf{x}_2 \in R_g \mid \mathbf{V}^2) = 0.1865$ (b) 3-cycle, $\mathcal{P}_{\pi_3}(\mathbf{x}_6 \in R_g \mid \mathbf{V}^6) = 0.5540$ (c) FMT*, $\mathcal{P}_{\pi^*}(\mathbf{x}_{10} \in R_g \mid \mathbf{V}^{10}) = 0.7098$

Fig. 4. Conditional in-region probability distributions for the industry-standard (a-b) and the FMT*-based method (c). The target region is shown as a green circle with radius $r = 50m$. Glider paths and state distributions are shown as in Fig. 2. Changes in depth are also shown; dark lines indicate deterministic paths and target depths z_k are labelled.

minimising the cost without considering safety) and the cases with thresholds 0.6, 0.9 and 0.99. The satisfaction probabilities $\mathcal{P}_{\pi^*}(\mathbf{V}^k)$ and the expected costs $\mathcal{R}_{\pi^*}(\mathbf{V}^k)$ are presented in Table II. Note that the 'sawtooth'-like motion is inevitable since the glide angle cannot go below its minimum as described in Sec. III (i.e., $|\gamma| \geq \gamma_{\min}$).

The result shows that the satisfaction probability for each scenario is always higher than the minimum constraints. With no safety constraint, the safety probability is $\mathcal{P}_{\pi^*}(\mathbf{V}) = 0.5288$. This safety probability is the minimum guaranteed by the framework. Note that high conditional state probability at a position given a policy does not necessarily mean that the glider is more likely to reach the the same position with a different policy. This is because the probability is conditional on overall safety; all the failed path are neglected in the probability distributions. For example, the conditional probability for $\alpha = 0$ seems to be higher than that for $\alpha = 0.99$ whilst the safety probabilities are 0.5288 and 1, respectively.

An important observation is that the expected cost generally increases as the safety threshold increases. This is due to two reasons. First, the state probability distribution becomes more spread out while reducing the transition cost when the glider angle is small, as shown in Fig. 1. As a result, the path at $\alpha = 0.99$ makes a number of sharp transitions to maintain the distribution density. Second, the glider tries to move further away from the obstacle, which could incur higher cost. This observation validates our claim that the proposed framework is able to balance between two conflicting objective functions.

C. Maximise the conditional in-region probability

In this example, we synthesise a control policy that maximises the probability of surfacing within a pre-specified region. Formally, we maximise $\mathcal{P}_{\pi^*}(\mathbf{x}_n \in R_g \mid \mathbf{V}^n)$, where n is the number of control instants and R_g is the goal region, as presented in Problem 2. We compare our result with the industry-standard approach where the glider follows a predefined sawtooth-like path with a fixed glide angle.

Starting from $[0, 0, 0]^T$, the glider is required to reach a circle located at $[1000, 0, 0]^T$ with a radius of $50\ m$ as drawn in green in Fig. 4. We demonstrate the conditional state probabilities for three different policies: 1) sawtooth-based 1-cycle policy (Fig. 4a-c), 2) sawtooth-based 3-cycle policy (Fig. 4d-f), and 3) FMT*-based optimal policy (Fig. 4g-i).

The result shows that the proposed framework outperforms the sawtooth-based approach. As discussed in Sec. VII-B, this is because the conditional state probability becomes more dense with higher glide angle. Since the 1-cycle policy inherently has smaller glide angles, the resulting distribution is much less concentrated than that with greater glide angles. The in-region probabilities are 0.1865 and 0.5540, respectively, while the FMT*-based approach achieves a success probability of 0.7098. It is important to note that cost, however, is generally minimised when glide angles are small, resulting in a shorter path.

VIII. CONCLUSION

This paper has presented a stochastic variant of a navigation planner for autonomous underwater gliders. We have addressed the inherent computational bottleneck in balancing safety and cost, and derived a recursive framework that accurately computes the probability and expected cost conditional on safety constraints. The recursive functions allow for utilisation of existing sampling-based planners such as FMT*. We have presented two types of objective function, where we 1) minimise conditional expected cost for a given safety threshold, and 2) maximise conditional in-region goal probability. We argue that the proposed framework is general enough to formulate more complex objectives.

Although our framework is fully capable of admitting stochastic ocean currents, there are important extensions to be addressed. Finding the transition and transition safety probabilities subject to stochastic ocean currents is computationally difficult. This bottleneck can be addressed with an efficient streamline-based algorithm [7]. Also, replanning can be necessary when the forecast ocean current varies over time [19]. Beyond path planning, we also would like to combine our framework with task planning algorithms [20, 21].

REFERENCES

[1] D. L. Rudnick, R. E. Davis, C. C. Eriksen, D. M. Fratantoni, and M. J. Perry, "Underwater gliders for ocean research," *Mar. Technol. Soc. J.*, vol. 38, no. 2, pp. 73–84, 2004.

[2] L. M. Russell-Cargill, B. S. Craddock, R. B. Dinsdale, J. G. Doran, B. N. Hunt, and B. Hollings, "Using Autonomous Underwater Gliders for Geochemical Exploration Surveys," *APPEA J.*, vol. 58, pp. 367–380, 2018.

[3] H. Johannsson, M. Kaess, B. Englot, F. Hover, and J. Leonard, "Imaging sonar-aided navigation for autonomous underwater harbor surveillance," in *Proc. of IEEE/RSJ IROS*, 2010, pp. 4396–4403.

[4] C. Yoo, R. Fitch, and S. Sukkarieh, "Provably-correct stochastic motion planning with safety constraints," in *Proc. of IEEE ICRA*, 2013, pp. 981–986.

[5] L. Janson, E. Schmerling, A. Clark, and M. Pavone, "Fast Marching Tree: A fast marching sampling-based method for optimal motion planning in many dimensions," *Int. J. Rob. Res.*, vol. 34, no. 7, pp. 883–921, 2015.

[6] N. Leonard and J. Graver, "Model-based feedback control of autonomous underwater gliders," *IEEE J. Ocean. Eng.*, vol. 24, no. 4, pp. 633–645, 2001.

[7] K. Y. C. To, K. M. B. L. Lee, C. Yoo, S. Anstee, and R. Fitch, "Streamlines for motion planning in underwater currents," in *Proc. of IEEE ICRA*, 2019.

[8] J. J. H. Lee, C. Yoo, R. Hall, S. Anstee, and R. Fitch, "Energy-optimal kinodynamic planning for underwater gliders in flow fields," in *Proc. of ARAA ACRA*, 2017.

[9] J. Binney, A. Krause, and G. S. Sukhatme, "Informative path planning for an autonomous underwater vehicle," in *Proc. of IEEE ICRA*, 2010, pp. 4791–4796.

[10] A. A. Pereira, J. Binney, G. A. Hollinger, and G. S. Sukhatme, "Risk-aware path planning for autonomous underwater vehicles using predictive ocean models," *J. Field Robot.*, vol. 30, no. 5, pp. 741–762, 2013.

[11] A. A. Pereira, J. Binney, B. H. Jones, M. Ragan, and G. S. Sukhatme, "Toward risk aware mission planning for autonomous underwater vehicles," in *Proc. of IEEE/RSJ IROS*, 2011, pp. 3147–3153.

[12] D. Kularatne, S. Bhattacharya, and M. A. Hsieh, "Time and Energy Optimal Path Planning in General Flows," in *Proc. of RSS*, 2016, pp. 47–57.

[13] E. Fernndez-Perdomo, J. Cabrera-Gmez, D. Hernndez-Sosa, J. Isern-Gonzlez, A. C. Domnguez-Brito, A. Redondo, J. Coca, A. G. Ramos, E. . Fanjul, and M. Garca, "Path planning for gliders using Regional Ocean Models: Application of Pinzn path planner with the ESEOAT model and the RU27 trans-Atlantic flight data," in *Proc. of IEEE OCEANS*, 2010, pp. 1–10.

[14] J. Isern-Gonzalez, D. Hernandez-Sosa, E. Fernandez-Perdomo, J. Cabrera-Gamez, A. C. Domnguez-Brito, and V. Prieto-Maranon, "Path planning for underwater gliders using iterative optimization," in *Proc. of IEEE ICRA*, 2011, pp. 1538–1543.

[15] I. Ko, B. Kim, and F. C. Park, "Randomized path planning on vector fields," *Int. J. Rob. Res.*, vol. 33, no. 13, pp. 1664–1682, 2014.

[16] J. Witt and M. Dunbabin, "Go with the flow: Optimal AUV path planning in coastal environments," in *Proc. of ARAA ACRA*, 2014.

[17] D. Rao and S. B. Williams, "Large-scale path planning for Underwater Gliders in ocean current," in *Proc. of ARAA ACRA*, 2009.

[18] K. M. B. Lee, J. J. H. Lee, C. Yoo, B. Hollings, and R. Fitch, "Active perception for plume source localisation with underwater gliders," in *Proc. of ARAA ACRA*, 2018.

[19] K. M. B. Lee, C. Yoo, B. Hollings, S. Anstee, S. Huang, and R. Fitch, "Online estimation of ocean current from sparse GPS data for underwater vehicles," in *Proc. of IEEE ICRA*, 2019.

[20] C. Yoo, R. Fitch, and S. Sukkarieh, "Probabilistic temporal logic for motion planning with resource threshold constraints," in *Proc. of RSS*, 2012.

[21] ——, "Online task planning and control for fuel-constrained aerial robots in wind fields," *Int. J. Robot. Res.*, vol. 35, no. 5, pp. 438–453, 2016.