2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Stranding Risk for Underactuated Vessels in Complex Ocean Currents: Analysis and Controllers

Andreas Doering[1,2,*], Marius Wiggert[1,*], Hanna Krasowski[2], Manan Doshi[3]
Pierre F.J. Lermusiaux[3] and Claire J. Tomlin[1]

2023 62nd IEEE Conference on Decision and Control (CDC) | 979-8-3503-0124-3/23/$31.00 ©2023 IEEE | DOI: 10.1109/CDC49753.2023.10383383

***Abstract*— Low-propulsion vessels can take advantage of powerful ocean currents to navigate towards a destination. Recent results demonstrated that vessels can reach their destination with high probability despite forecast errors. However, these results do not consider the critical aspect of safety of such vessels: because their propulsion is much smaller than the magnitude of surrounding currents, they might end up in currents that inevitably push them into unsafe areas such as shallow waters, garbage patches, and shipping lanes. In this work, we first investigate the risk of stranding for passively floating vessels in the Northeast Pacific. We find that at least 5.04% would strand within 90 days. Next, we encode the unsafe sets as hard constraints into Hamilton-Jacobi Multi-Time Reachability to synthesize a feedback policy that is equivalent to re-planning at each time step at low computational cost. While applying this policy guarantees safe operation when the currents are known, in realistic situations only imperfect forecasts are available. Hence, we demonstrate the safety of our approach empirically with large-scale realistic simulations of a vessel navigating in high-risk regions in the Northeast Pacific. We find that applying our policy closed-loop with daily re-planning as new forecasts become available reduces stranding below 1% despite forecast errors often exceeding the maximal propulsion. Our method significantly improves safety over the baselines and still achieves a timely arrival of the vessel at the destination.**

I. Introduction

Autonomous systems are increasingly deployed for long-term tasks where energy-efficient operation is critical. For systems operating in the oceans or in the air, this leads to a growing interest in utilizing the dynamics of the surrounding flows as a means of propulsion. Stratospheric balloons and airships utilize wind fields [1], [2], while ocean gliders and active drifters exploit ocean currents [3]–[7].

Our recent work [5], [8] has demonstrated that a vessel with just $0.1\,\mathrm{m\,s^{-1}}$ propulsion can reliably navigate to a target region by *hitchhiking* on ocean currents of up to $2.0\,\mathrm{m\,s^{-1}}$. This work has further been extended to the application of floating farms, which maximize the growth of seaweed over long time horizons [9], [10] and to multiple agents that want to stay in proximity to each other to stay connected in local communication networks [11].

* A.D. and M.W. have contributed equally to this work.
[1] A.D., M.W., and C.J.T. are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, USA. For inquiries contact: mariuswiggert@berkeley.edu
[2] A.D. and H.K. are with the School of Computation, Information and Technology of the Technical University of Munich, Germany
[3] M.D. and P.F.J.L. are with the Department of Mechanical Engineering at the Massachusetts Institute of Technology, USA.
The authors gratefully acknowledge the support of the C3.ai Digital Transformation Institute, the IFI fellowship of the German Academic Exchange Service (DAAD) , the research group ConVeY funded by the German Research Foundation under grant GRK 2428, the DARPA Assured Autonomy Program, and the ONR BRC program.

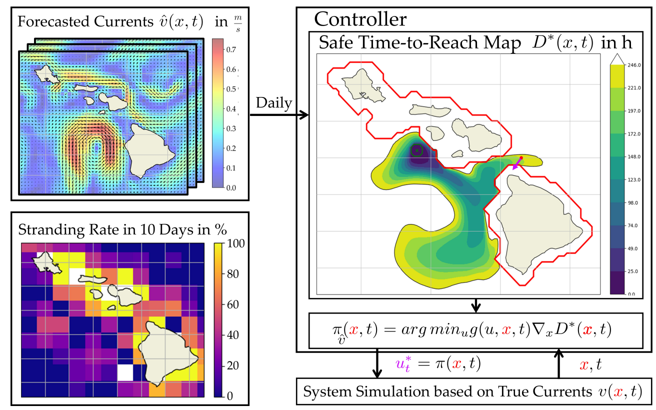

Fig. 1: Our method for safe navigation with imperfect dynamics is based on re-planning on two timescales: 1) compute the safe Time-to-Reach value function D^* daily as new forecasts become available; 2) For every timestep, e.g. 10 min, we obtain $\boldsymbol{u}^*$ from $\pi(\boldsymbol{x},t)$, which is a computationally cheap and equivalent to re-planning. This is necessary, as the true currents $v(\boldsymbol{x},t)$ differ from the forecasted currents $v \neq \hat{v}$, thus, the vessel will be in a different spatial state $\boldsymbol{x}$ than anticipated. Further motivation is provided in the lower left image which displays the stranding rate of passively floating vessels over 10 days.

However, these approaches do not factor in safety aspects, although the use of Autonomous Surface Vehicles (ASV) in unmanned and long-term operations may pose crucial safety risks. In the event of significant damage, the ASV may become inoperable and may be abandoned or sunk, resulting in financial losses and environmental impacts. One important safety hazard are shallow waters, especially near strong currents, as the ASV can easily strand. Another significant safety hazard is entering a garbage patch that has a high concentration of marine debris, which can cover an area of up to 1.6 million km^2 [12] as in the case of the Great Pacific Garbage Patch (GPGP). The garbage can get entangled in the ASV rotors or damage other components, resulting in loss of control. Furthermore, collisions with other vessels may cause damage to the ASV and endanger the crew of the other vessel. Thus, shipping lanes are another area of increased risk to the ASV as they are used by large, fast-moving vessels. Next, we present the related work on safe motion planning for autonomous vessels in maritime environments.

a) Related Work: Prior research focuses on collision avoidance [13], [14] and compliance with the Convention on the International Regulations for Preventing Collisions at

Sea (COLREGS) [15]–[20] for safe planning of autonomous motions. For example, Zhao et al. [17] use Reinforcement Learning (RL) to achieve COLREGS-compliant motion planning for encounters with multiple vessels. Other works addressed dynamic obstacles and forbidden regions [21]–[23]. Most research on safe motion planning for autonomous vessels considers fully actuated vessels. Since our vessel has restricted maneuverability due to its low propulsion, complying with the rules for power-driven vessels from the COLREGS is not critical. Therefore, we focus on safety and collision avoidance with largely static obstacles such as shallow areas, shipping lanes, or garbage patches.

Although research on maritime safety considers underactuated vessels, most consider underactuation due to non-holonomic actuation of vessels such as [24]–[27]. In this paper, we focus on systems that have a maximum propulsion that is less than the magnitude of forecasted flow and often also the forecast error, posing severe challenges for the safety of ASVs. We use the term *underactuated* in this paper to describe such systems, and note that we are using the term in a way that is closer to its use in robotics, rather than its typical control-theoretic meaning.

Agents operating in three-dimensional flows can avoid obstacles or strong currents by utilizing the third dimension [28]. This was demonstrated for low-propulsion stratospheric balloons in [29]. Safety was ensured with high probability by formulating the problem as Markov Decision Process using a heuristic cost function and training a RL agent. This approach relies on a realistic uncertainty distribution of stratospheric winds.

Robust Model Predictive Control (MPC) approaches can guarantee safety under disturbances by ensuring that the system is always in a state from which it can reach a robust control invariant set within a finite time horizon [30], [31]. In this robust control-invariant set, there always exists a control input that ensures that the system can stay in this safe set indefinitely. However, in our problem setting with underactuated vessels and imperfect, deterministic ocean current forecasts, no such control invariant set exists; hence robust control with realistic bounds is infeasible, see Sec. II.

Our paper makes two main contributions: First, we perform an empirical evaluation of stranding risk for passively floating vessels in the Northeast Pacific. Second, we present our methods of Hamilton-Jacobi Multi-Time Reachability (HJ-MTR) with re-planning on two timescales for safe motion planning of underactuated ASVs in a setting with realistic ocean currents and daily forecasts. Furthermore, we evaluate our controller with several baseline controllers over a large set of simulated missions. Since we focus on realistic conditions, we chose to empirically evaluate our controller in ocean currents with realistic forecast errors.

In what follows, we define our problem statement in Sec. II and motivate the need for a safety controller with a stranding study in Sec. III. In Sec. IV we introduce our method and summarize HJ-MTR. We present our extensive simuation experiments in Sec. V and discuss them in Sec. VI. We conclude and present future work in VII.

II. Problem Statement

We now define the problem of collision avoidance for underactuated vessels by introducing the flow model, vessel model, and representation of obstacles. We then introduce the notion of stranding as our key performance measure.

A. System Dynamics, Obstacles and Target

We consider moving in a general time-varying non-linear flow field $v(\boldsymbol{x}, t) \rightarrow \mathbb{R}^n$, with $\boldsymbol{x} \in \mathbb{R}^n$ representing the spatial state, $t \in [0, T]$ the time and n the dimension of the spatial domain. In our case, $n = 2$ as we regard an ASV operating on the ocean surface. We denote the actuation signal by $\boldsymbol{u}$ from a bounded set $\mathbb{U} \in \mathbb{R}^{n_u}$ with n_u the dimension of the control. Let $\boldsymbol{x}(t) \in \mathbb{R}^n$ denote the position of our ASV at time t. Our model for the dynamics of the ASV is given by

$$\dot{\boldsymbol{x}}(t) = \boldsymbol{f}(\boldsymbol{x}(t), \boldsymbol{u}(t), t) \tag{1}$$

$$= v(\boldsymbol{x}, t) + g(\boldsymbol{u}, \boldsymbol{x}, t) \quad \forall t \in [0, T], \tag{2}$$

The actuation $\boldsymbol{u}$ influences the relative velocity $g(\boldsymbol{u}, \boldsymbol{x}, t)$ of the ASV with respect to the ocean. Hence the absolute velocity of the vessel is given by the vector sum of the ocean currents at the location of the vessel and the relative velocity of the vessel with respect to the currents. The maximum actuation of the vessel is constrained by $||g(\boldsymbol{u}, \boldsymbol{x}, t)||_2 \leq \boldsymbol{u}_{\text{max}}$. We define the target and obstacle as sets $\mathcal{T} \subset \mathbb{R}^n$ and $\mathcal{O} \subset \mathbb{R}^n$ respectively. We assume that these sets are not time-dependent. However, note that our results can be extended to time-dependent cases using schemes given in [32].

B. Problem Setting

The agent's goal is to navigate *safely* and *reliably* from a start state $\boldsymbol{x}_0$ at start time t_0 to a target region $\mathcal{T} \subset \mathbb{R}^n$. We employ the same empirical definition of *reliability* as [8] defining it as the success rate of a controller navigating from $\boldsymbol{x}_0$ at t_0 to $\mathcal{T}$ within a maximum allowed time T_{max}, over a representative set of start-target missions $(\boldsymbol{x}_0, t_0, \mathcal{T}, \mathcal{O}) \in \mathbb{M}$. We define stranding as an agent entering the obstacle set $\mathcal{O}$ before T_{max}. We then quantify *safety* as the stranding rate of a controller over the set of missions $\mathbb{M}$. The obstacle set in our experiments are shallow waters, but this can be any static obstacles such as garbage patches or areas with high shipping traffic density. If the true currents are known *a* priori and there exists a trajectory that prevents stranding, our method guarantees safety. However, we are interested in realistic settings where only daily current forecasts $\hat{v}$ are available and these differ from the true currents v by the forecast error δ. The forecasts provided by HYCOM [33] have a global forecast error of Root Mean Squared Error (RMSE) $\text{RMSE}(\delta) = 0.2\,\text{m}\,\text{s}^{-1}$. In this setting with a complex empirical distribution of forecast errors $\delta(\boldsymbol{x}, t)$ and severe underactuation e.g. in our experiments $||g(\boldsymbol{u}, \boldsymbol{x}, t)||_2 = 0.1\,\text{m}\,\text{s}^{-1} \lll$ Root Mean Squared Error (RMSE)$(\delta) \approx 0.2\,\text{m}\,\text{s}^{-1}$ and currents $||max(v)||_2 \approx 1.4\,\text{m}\,\text{s}^{-1}$ safety despite worst-case forecast errors is impossible. Hence, in Sec. V we evaluate the performance of our method empirically over a large set of missions $\mathbb{M}$ under realistic currents and forecasts $\mathbb{V}$.

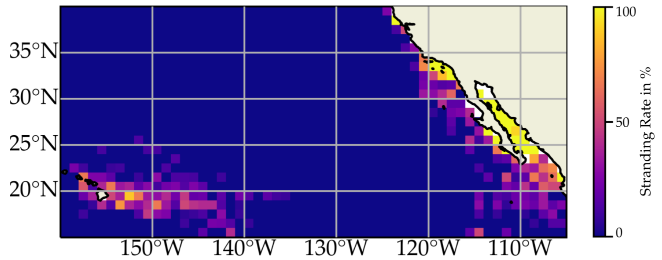

Fig. 2: Rate of passively floating vessels stranding, which we define as entering waters with depth less than $-150\,\mathrm{m}$. 5.04% of 10 000 simulated vessels starting in the region strand within 90 days. Some vessels float over $20.0°$ before stranding.

We evaluate the performance based on the binary indicator functions $\mathbb{I}_{Suc}, \mathbb{I}_{Obs}$ that evaluate to 1 if in the $\mathcal{T}$ and $\mathcal{O}$ sets respectively and to 0 otherwise:

$$\mathbb{E}_{\underbrace{\boldsymbol{x}, t \sim \mathbb{M}}_{\text{initial condition}};\ \underbrace{v, \hat{v} \sim \mathbb{V}}_{\substack{\text{real and forecasted} \\ \text{ocean currents}}}} \underbrace{\{\mathbb{I}_{Suc}, \mathbb{I}_{Obs}\}}_{\text{Success and Stranding}}$$

III. Stranding Study

To illustrate the need for our safety controller, we analyze the rate of stranding for passively floating vessels off the coast of California and Mexico between N$15°$ and N$40°$ and W$160°$ and W$105°$. We define entering an area with a depth of less than $-150\,\mathrm{m}$ as stranding.

The stranding study can be conducted analytically or experimentally. Here, we perform it empirically by sampling 10 000 missions each consists of a uniformly sampled starting location, outside of the stranding area, and a uniformly sampled starting time. We simulate the trajectories over a time horizon of 10 and 90 days, using Copernicus data for the year 2022. We observe that 1.67% missions strand within 10 days and 5.04% strand within 90 days. In Figure 2 we show the spatial distribution of the stranding rate over 90 days. We stop the simulation for vessels leaving the investigated region and assume they do not strand, this occurred for 1.78% missions within 10 days and 18.75% within 90 days.

IV. Safe HJ Controller

The forecasts provided by many operational ocean forecasting systems are deterministic. This prohibits applying probabilistic methods that would require a realistic distribution of currents [34]. As alluded to in Sec. II, robust control is infeasible in our setting. Hence we ensure safety with imperfect forecasts by first solving an optimal control problem assuming the forecated currents are accurate (Sec. IV-A) and then developing a schema for re-planning on two timescales to ensure safe and reliable operation (Sec. IV-B).

A. Multi-Time Hamilton Jacobi Reachability

To obtain our controller we first solve the continuous-time optimal control problem formalized in HJ-MTR assuming the forecasted currents $\hat{v}$ are accurate. For the sake of completeness and to focus on the relevant aspects, we summarize the technique here. HJ-MTR [32] uses dynamic programming to derive a controller that avoids obstacles and (a) if possible, will get the system to the target in the minimum time, and (b) if not, will get as close to the target as possible. To achieve this behavior, we first define the modified dynamical system f_a such that the state $\boldsymbol{x}$ of the vessel remains in the target or obstacle once it enters either of them.

$$\dot{\boldsymbol{x}} = f_a(\boldsymbol{x}, \boldsymbol{u}, t) = \begin{cases} 0, & \boldsymbol{x} \in \mathcal{O} \cup \mathcal{T} \\ \hat{v}(\boldsymbol{x}, t) + \boldsymbol{u}(t), & \text{otherwise}. \end{cases} \quad (3)$$

Next, we define a cost function consisting of a running cost l and a terminal cost l_{term} over the time-horizon t to T:

$$J(\boldsymbol{x}, \boldsymbol{u}(\cdot), t) = l_{term}(\boldsymbol{x}(T)) + \int_t^T l(\boldsymbol{x}, s) ds \quad (4)$$

The terminal cost function $l_{term}(\boldsymbol{x})$ is infinitely high if the ASV terminates in an obstacle and is otherwise equal to the distance from the target set.

$$l_{term}(\boldsymbol{x}) = \begin{cases} \infty, & \boldsymbol{x} \in \mathcal{O} \\ d(\boldsymbol{x}, \mathcal{T}), & \text{otherwise}. \end{cases} \quad (5)$$

Our running cost $l(\boldsymbol{x}, s)$ provides a reward for every time step once the target is reached:

$$l(\boldsymbol{x}, s) = \begin{cases} -1, & \boldsymbol{x} \in \mathcal{T} \text{ and } \boldsymbol{x} \notin \mathcal{O} \\ 0, & \text{otherwise}. \end{cases} \quad (6)$$

Using dynamic programming we then obtain the Hamilton-Jacobi Partial Differential Equation (PDE) which lets us solve for the optimal value function J^*:

$$\frac{\partial J^*(\boldsymbol{x}, t)}{\partial t} = \begin{cases} 1, & \boldsymbol{x}(t) \in \mathcal{T} \cap (\mathcal{O})^c \\ 0, & \boldsymbol{x}(t) \in \mathcal{O} \\ -\min_{\boldsymbol{u}} \left[\nabla_{\boldsymbol{x}} J^* \cdot f(\boldsymbol{x}, \boldsymbol{u}, t)\right], & \text{otherwise}. \end{cases}$$
$$J^*(\boldsymbol{x}, T) = l_{term}(\boldsymbol{x}). \quad (7)$$

To better interpret the the meaning of the value $J^*(\boldsymbol{x}, t)$ we use a simple transformation to obtain the safe Time-to-Reach (TTR) map D^*. An example is plotted in Fig. 3.

$$D^*(\boldsymbol{x}, t) = T + J^*(\boldsymbol{x}, t) - t,\ \forall (\boldsymbol{x}, t) \text{ s.t., } J^*(\boldsymbol{x}, t) \leq 0$$

For example, if $D^*(\boldsymbol{x}, t)$ is equal to 3, it means that a vessel starting at $\boldsymbol{x}$ at time t can reach the target in 3 time units when following the time-optimal control which can be easily computed from D^* as detailed next. The value functions J^* and D^* allow us to compute the time-optimal policy for the system assuming the forecasted currents are accurate

$$\pi_{\hat{v}}(\boldsymbol{x}, t) = \arg\min_{\boldsymbol{u}} \left[\nabla_{\boldsymbol{x}} D^*(\boldsymbol{x}, t) \cdot f(\boldsymbol{x}, \boldsymbol{u}, t)\right] \quad (8)$$

Under holonomic dynamics $g(\boldsymbol{u}, \boldsymbol{x}, t) = \boldsymbol{u}$ as in our experiments, this simplifies to the negative gradient of D^*, as this is the direction that will lead the ASV to the lowest time-to-reach value:

$$\pi_{\hat{v}}(\boldsymbol{x}, t) = -\frac{\nabla_{\boldsymbol{x}} D^*(\boldsymbol{x}, t)}{\|\nabla_{\boldsymbol{x}} D^*(\boldsymbol{x}, t)\|_2} \boldsymbol{u}_{\max}.$$

B. Replanning to Operate With Imperfect Forecasts

The derived HJ-MTR policy $\pi_{\hat{v}}$ is time-optimal and guarantees safety when the forecasted currents are accurate. However, in realistic settings forecasts are imperfect and differ from the true currents. To achieve safe and reliable operation we replan on two timescales as depicted in Fig. 1.

First, on a daily basis we receive new forecasts $\hat{v}$ and solve the optimal control problem to obtain the value function

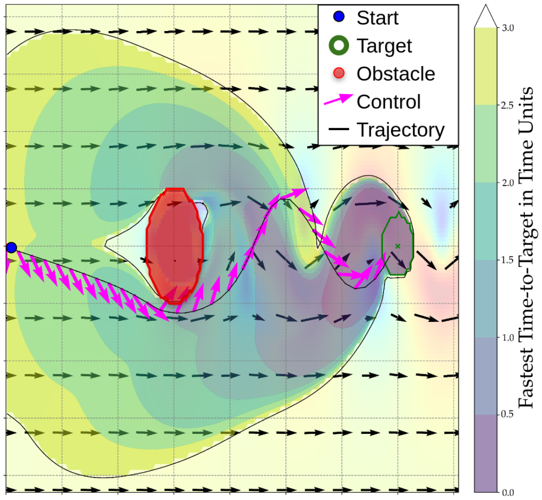

Fig. 3: Trajectory of our method in analytical currents evading obstacle. The safe Time-to-Reach map spares out obstacles and areas where the underactuated agent is inevitably pushed into the obstacle.

using HJ-MTR. Second, we use the policy $\pi_{\hat{v}}$ (Eq. 8) as feedback policy. This means that at each timestep when we are at a state $\boldsymbol{x}$, we take the control that is time-optimal assuming the forecast is true from that time onwards. This is equivalent to full-time-horizon replanning on the forecast at every step. This is necessary because the true currents v differ from the forecast currents $\hat{v}$, therefore, at each step, we will be in a spatial state $\boldsymbol{x}$ different from that anticipated. The full schema is sketched out in Algorithm 1.

In summary, there are three core advantages of our method compared to classical MPC with non-linear programming. The solution of the HJ-MTR guarantees time-optimality under known arbitrary non-linear dynamics over the full-time horizon and does not get stuck in local optima. The online computation for replanning at every step is very low, for holonomic systems it's a simple gradient computation. Lastly, this approach enables graceful degradation: in case the vessel cannot reach the destination our method will attempt to minimize the terminal distance to the target, while non-linear programming would not provide a solution.

Algorithm 1: Replanning on Two Timescales

Input: Forecast Currents $\hat{v}(\boldsymbol{x}, t),\ t = 0,\ \boldsymbol{x}(t) = \boldsymbol{x}_0$
while $t \leq T_{max}$ *and* $\boldsymbol{x}(t) \notin \mathcal{T}$ *and* $\boldsymbol{x}(t) \notin \mathcal{O}$ **do**
  if *new forecast available* **then**
    compute new safe time-to-reach map $\mathcal{D}^*$
  $\boldsymbol{u}_t = \pi_{\hat{v}}(\boldsymbol{x}(t), t) =$
  $\arg\min_{\boldsymbol{u}} g(\boldsymbol{u}, \boldsymbol{x}(t), t) \cdot \nabla_x \mathcal{D}^*(\boldsymbol{x}(t), t)$
  Apply $\boldsymbol{u}_t$ for Δt:
  $x(t + \Delta t) = x(t) + \int_t^{t+\Delta t} f(\boldsymbol{u}_t, x(s), s)\, ds$
  $t = t + \Delta t$

V. Experiments

To evaluate if our method can reduce the stranding rate while also reaching the target, we simulate a large number of missions on realistic ocean currents and compare the performance of our method to baseline methods. In this section, we first detail the system dynamics and simulation setup and then how we ascertain the realism of our simulations. Further, we explain the creation of our obstacle sets, the generation of a representative set of missions, and lastly the baselines and controllers we compare to and the evaluation metrics.

A. Experimental Setup

Our experiments investigate the stranding rate and reliability of several controllers for navigating a two-dimensional ASV with holonomic actuation $g(\boldsymbol{u}, \boldsymbol{x}, t) = \boldsymbol{u}$ of fixed magnitude $||\boldsymbol{u}||_2 = 0.1\,\mathrm{m\,s^{-1}}$. The control input is the angle θ for steering the ASV in ocean currents $v(\boldsymbol{x}, t) \in [0\,\mathrm{m\,s^{-1}}, 1.4\,\mathrm{m\,s^{-1}}]$, which the vessel utilizes to reach its target region. We simulate the ASV for $T_{max} = 240\,\mathrm{h}$. If the ASV collides with an obstacle, we terminate the mission and count it as stranded, if it reaches the target region within T_{max} we count it as success, if it does neither we count it as timeout.

a) Realistic Ocean Forecast Simulation: In a real-world setting, a vessel can receive the most recent forecast daily and provide it to the control method to perform replanning. To simulate the system dynamics based on the true currents $v(\boldsymbol{x}, t)$ we use Copernicus hindcasts [35]. For simulating the daily forecasts we use a series of 5 days of HYCOM [33] hindcasts that become available daily. It should be noted that, unlike HYCOM, Copernicus incorporates tidal currents into its forecasts. We want to ensure a realistic simulation of the forecast error δ. The forecast error can be measured with various metrics such as RMSE, vector correlation, and separation distance [33], [35]. In our simulations, these are on average $0.18\,\mathrm{m\,s^{-1}}$ RMSE, which is close to the RMSE of the HYCOM forecast error [36] at $0.19\,\mathrm{m\,s^{-1}}$. In terms of vector correlation, with a value of 2 representing perfect correlation and 0 no correlation, we measured 0.63 compared with 0.64 for the HYCOM forecasts [36] and 0.62 for Copernicus [35], each measured at $t = 71\,\mathrm{h}$. Thus, our simulation setup represents realistic situations well.

b) Obstacles derived from Bathymetry: The bathymetry data we employ is the GEBCO 2022 grid [37]. It is a global, continuous terrain grid with a resolution of $15''$. We coarsen it to the same resolution of the current data $5'$ by using the maximum in each grid cell to overestimate the elevation in each grid cell.

c) Representative Set of Missions: We generate a representative set of 1146 start-target missions $(\boldsymbol{x}_0, t_0, \mathcal{T}, \mathcal{O}) \in \mathbb{M}$ with high stranding risk based on the following procedure. First, we select suitable target regions $\mathcal{T}$ using rejection sampling: we uniformly sample the center points for target regions $x_{\mathcal{T},center}$ spatially in the region introduced in Sec. III. We then reject points based on three conditions: if they are too close to the simulation region boundary (distance below $0.5°$), if they are too close to the obstacles (distance below $0.025°$), and if they are too far from the obstacle (maximum distance to an obstacle above $3°$). The high-risk target regions $\mathcal{T}$ are then the circular regions with radius

$r_{\mathcal{T}} = 0.1°$ around $x_{\mathcal{T},center}$. To ensure diverse ocean current conditions we then uniformly sample times t_0 from 2022. To select the start-point $\boldsymbol{x}_0$ for each mission, we need to make sure that it is fundamentally feasible to reach the target from $\boldsymbol{x}_0$. For that, we calculate the set of all states from which there exists a control such that the ASV can reach the target within a specific timeframe knowing the true currents v and assuming no obstacles. Note that if we calculate the TTR map D^* based on v as illustrated in Fig. 3, we directly get the sets of states that can reach the target within a specific timeframe $[t_{lower}, t_{upper}]$ as using the values of D^*, $\{\boldsymbol{x} \mid D^* \in [t_{lower}, t_{upper}]\}$. We then sample the start position $\boldsymbol{x}_0$ from these sets so that the ASV could reach the target within 5-9 days assuming known currents and no obstacles.

d) Baselines and Controllers: We evaluate our method *Multi-Time HJ Reachability Closed-Loop Controller With Obstacles* (MTR) in comparison to three safety controllers and two baselines. The first baseline is *Passively Floating* (Floating) to lower bound the performance of controllers and to show the complexity of safely completing the missions. The second baseline is the *MTR With No Obstacles Controller* (MTR-no-Obs) [8]. The first safety controller we compare against is *Switch-MTR-no-Obs*, a simple solution that adds safety functionality to the baseline MTR-no-Obs. It does not reason about currents for safety, instead, it is a switching controller that uses MTR-no-Obs when it is far away from obstacles and switches to an away-from-obstacles controller if the ASV gets closer than $20\,\mathrm{km}$ to the obstacle. The away-from-obstacle controller actuates with full actuation into the direction of the largest distance to obstacles and once the distance to obstacles is larger than the threshold, the navigation controller MTR-no-Obs takes over again. The second and third safety controllers are ablations to our main method MTR. The *Switching Controller With Obstacles* (Switch-MTR) is the same switching as before but uses MTR for navigation. Lastly we compare with *MTR With Small Disturbance* (SmallDist-MTR), a robust control version of MTR with an unrealistically low disturbance of $d = 0.05\,\mathrm{m\,s^{-1}}$. We obtain the robust value function by formulating the optimal control problem as game between u and d and then solving the HJ PDE [38].

e) Evaluation Metrics: As detailed in Sec. II we have two metrics. The main metric *stranding rate* measures safety as the rate of entering the obstacle set $\mathcal{O}$ over the set of missions $\mathbb{M}$. The success rate measures the reliability of reaching the target.

B. Experimental Results

In the complex realistic flows with large forecast errors and close proximity to obstacles, our controller MTR has a stranding rate of only 0.96%, compared to 4.71% of the baseline MTR-no-Obs (Table I) and outperforms the other safety controllers.

We evaluate if the stranding rate of the controllers is lower than the baseline of MTR-no-Obs in a statistically significant manner with a one-sided two-sample z proportion test. Let Γ be the stranding rate of a controller and our null hypothesis be: $H_0 : \Gamma_{\text{MTR-no-Obs}} = \Gamma_{\text{controller}}$. With the alternate hypothesis: $H_A : \Gamma_{\text{MTR-no-Obs}} > \Gamma_{\text{controller}}$. The stranding rate of all safety controllers is higher than MTR-no-Obs in a statistically significant way with p-values Switch-MTR-no-Obs $p = 2.6e^{-3}$, MTR $p = 3.1e^{-8}$, Switch-MTR $p = 9.3e^{-7}$, and SmallDist-MTR $p = 9.5e^{-5}$. The success rate of MTR is not reduced by safety: in fact it shows the highest success rate.

TABLE I: We compare the performance of multiple controllers, the arrows indicate if high, or low is preferred. Our MTR outperforms other controllers in terms of safety and reliability. * marks a statistically significant lower stranding rate compared to MTR-no-Obs.

Controller	Stranding Rate ↓	Success Rate ↑
MTR(ours)	**0.96**%*	**37.26**%
Switch-MTR	1.31%*	37.17%
SmallDist-MTR	1.92%*	33.16%
Switch-MTR-no-Obs	2.53%*	36.82%
MTR-no-Obs	4.71%	36.91%
Floating	4.89%	2.53%

VI. Discussion

We note that all our safety controllers have a success rate that is lower than the 82.3% as reported for MTR-no-Obs in [8]. We believe this is due to three differences in the experimental setup. First, our missions are longer and have smaller time buffers to reach the target. The time-to-target for each mission in [8] is between 20-120h with $T_{max} = 150\,\mathrm{h}$, while our sampled time-to-target is 120-216h with $T_{max} = 240\,\mathrm{h}$. In extreme cases, their missions are expected to finish in $20\,\mathrm{h}$ with 130 additional hours to reach the target, while in the worst case, our missions can have a $216\,\mathrm{h}$ mission with a buffer of $24\,\mathrm{h}$. Second, some of our missions may be infeasible due to stranding on obstacles as we sampled the starting points $\boldsymbol{x}_0$ without considering obstacles. Third, we sample missions with a maximum distance to shore of $3°$, which exposes the vessels to more tidal currents near the shore which implies a higher forecast error.

VII. Conclusion and Future Work

In this work, we have demonstrated that HJ-MTR with obstacles can be used to reduce the rate of stranding even in complex flows using daily forecasts with large errors. We evaluated our method over a large set of 5-9 day start-to-target missions distributed spatially near the Coast of California, Hawaii, and the Baja California area and temporally across the year 2022 using realistic ocean currents. In our experiments, our method has achieved a stranding rate of 0.96%, which is significantly lower than that for the baseline controllers and also has a slightly higher success rate. While we have demonstrated the ability of our method with two-dimensional ocean currents, we emphasize that it is also applicable in a three-dimensional setting [28] underwater or in the air. Furthermore, HJ-MTR is able to handle time-varying constraints [32]. However, including dynamic obstacles such as ships that move fast and change their course would require a higher frequency of re-planning to account for course changes, resulting in higher computational costs.

In the future, we plan to model zones of a potential hazard, e.g. shipping lanes and garbage patches, as soft constraints, where instead of preventing entering altogether it would be beneficial to, e.g. minimize the time spent therein. By reducing time in shipping lanes an ASV could avoid many vessels. As of now, it's unknown how underactuated ASVs would be classified under the COLREGS and if evasion is necessary or if they should stop their propulsion to be floated along a vessel [39], [40]. Getting the rotors of an ASV entangled in the garbage can render it inoperable, hence it is beneficial to avoid areas with a larger density of garbage such as the center of the GPGP, while not making the whole 1.6 million km^2 [12] of the GPGP an obstacle to avoid. To solve this, we can investigate using a risk-based extension of a soft-edge and dynamic forbidden region [23], [41].

REFERENCES

[1] M. G. Bellemare *et al.*, "Autonomous navigation of stratospheric balloons using reinforcement learning," *Nature*, vol. 588, no. 7836, pp. 77–82, 2020.

[2] M. Manikandan and R. S. Pant, "Research and advancements in hybrid airships - a review," *Progress in Aerospace Sciences*, vol. 127, p. 100741, 2021.

[3] P. F. J. Lermusiaux *et al.*, "Science of autonomy: Time-optimal path planning and adaptive sampling for swarms of ocean vehicles," in *Springer Handbook of Ocean Engineering: Autonomous Ocean Vehicles, Subsystems and Control*, T. Curtin, Ed. Springer, 2016, ch. 21, pp. 481–498.

[4] D. Meyer, "Glider technology for ocean observations: A review," *Ocean Science Discussions*, pp. 1–26, 2016.

[5] D. N. Subramani *et al.*, "Time-optimal path planning: Real-time sea exercises," in *Oceans '17 MTS/IEEE Conference*, Aberdeen, Jun. 2017.

[6] A. Molchanov *et al.*, "Active drifters: Towards a practical multi-robot system for ocean monitoring," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 545–552.

[7] P. F. J. Lermusiaux *et al.*, "A future for intelligent autonomous ocean observing systems," *Journal of Marine Research*, vol. 75, no. 6, pp. 765–813, Nov. 2017, the Sea. Volume 17, The Science of Ocean Prediction, Part 2.

[8] M. Wiggert *et al.*, "Navigating underactuated agents by hitchhiking forecast flows," in *IEEE Conference on Decision and Control (CDC)*, 2022, pp. 2417–2424.

[9] M. S. Bhabra *et al.*, "Optimal harvesting with autonomous tow vessels for offshore macroalgae farming," in *OCEANS 2020 IEEE/MTS*. IEEE, Oct. 2020, pp. 1–10.

[10] M. Killer *et al.*, "Maximizing seaweed growth on autonomous farms: A dynamic programming approach for underactuated systems navigating on uncertain ocean currents," *arXiv preprint arXiv:2307.01916*, 2023.

[11] N. Hoischen *et al.*, "Safe connectivity maintenance in underactuated multi-agent networks for dynamic oceanic environments," *arXiv preprint arXiv:2307.01927*, 2023.

[12] L. Lebreton *et al.*, "Evidence that the great pacific garbage patch is rapidly accumulating plastic," *Scientific Reports*, vol. 8, no. 1, p. 4666, 2018.

[13] X. Geng *et al.*, "Motion plan of maritime autonomous surface ships by dynamic programming for collision avoidance and speed optimization," *Sensors*, vol. 19, no. 2, 2019.

[14] M. Junmin *et al.*, "Mechanism of dynamic automatic collision avoidance and the optimal route in multi-ship encounter situations," *Journal of Marine Science and Technology*, vol. 26, pp. 141–158, 2021.

[15] "COLREGs: Convention on the International Regulations for Preventing Collisions at Sea," International Maritime Organization, 1972.

[16] E. Meyer *et al.*, "COLREG-compliant collision avoidance for unmanned surface vehicle using deep reinforcement learning," *IEEE Access*, vol. 8, pp. 165 344–165 364, 2020.

[17] L. Zhao and M. I. Roh, "COLREGs-compliant multiship collision avoidance based on deep reinforcement learning," *Ocean Engineering*, vol. 191, pp. 106 436–106 450, 2019.

[18] L. Zhao *et al.*, "Control method for path following and collision avoidance of autonomous ship based on deep reinforcement learning," *Journal of Marine Science and Technology*, vol. 27, no. 4, pp. 293–310, 2019.

[19] T. A. Johansen *et al.*, "Ship collision avoidance and COLREGS compliance using simulation-based control behavior selection with predictive hazard assessment," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 12, pp. 3407–3422, 2016.

[20] A. Lazarowska, "A trajectory base method for ship's safe path planning," *Procedia Computer Science*, vol. 96, pp. 1022–1031, 2016.

[21] T. Lolla *et al.*, "Time-optimal path planning in dynamic flows using level set equations: Theory and schemes," *Ocean Dynamics*, vol. 64, no. 10, pp. 1373–1397, 2014.

[22] ——, "Time-optimal path planning in dynamic flows using level set equations: Realistic applications," *Ocean Dynamics*, vol. 64, no. 10, pp. 1399–1417, 2014.

[23] ——, "Path planning in multiscale ocean flows: Coordination and dynamic obstacles," *Ocean Modelling*, vol. 94, pp. 46–66, 2015.

[24] S. Moe and K. Y. Pettersen, "Set-based line-of-sight (LOS) path following with collision avoidance for underactuated unmanned surface vessels under the influence of ocean currents," in *IEEE Conf. on Control Tech. and Applications (CCTA)*, 2017, pp. 241–248.

[25] M. Wu *et al.*, "Augmented safety guarantee-based area keeping control for an underactuated USV with environmental disturbances," *ISA Transactions*, vol. 127, pp. 415–422.

[26] J. Liu *et al.*, "Collision avoidance for underactuated ocean-going vessels considering COLREGs constraints," *IEEE Access*, vol. 9, pp. 145 943–145 954.

[27] E. Borhaug *et al.*, "Integral LOS control for path following of underactuated marine surface vessels in the presence of constant ocean currents," in *IEEE Conference on Decision and Control (CDC)*, 2018, pp. 4984–4991.

[28] C. S. Kulkarni and P. F. J. Lermusiaux, "Three-dimensional time-optimal path planning in the ocean," *Ocean Modelling*, vol. 152, Aug. 2020.

[29] A. Alexander *et al.*, "Loon library: Lessons from building loon's stratospheric communications service," Loon LLC, Tech. Rep., 2021.

[30] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control. Lecture Notes in Control and Information Sciences*, vol. 245, 1999, pp. 207–226.

[31] M. Lazar *et al.*, "On input-to-state stability of min–max nonlinear model predictive control," *Systems & Control Letters*, vol. 57, no. 1, pp. 39–48, 2008.

[32] M. Doshi *et al.*, "Hamilton-Jacobi multi-time reachability," in *IEEE Conference on Decision and Control (CDC)*, 2022, pp. 2443–2450.

[33] E. P. Chassignet *et al.*, "US GODAE: global ocean prediction with the HYbrid Coordinate Ocean Model (HYCOM)," *Oceanography*, vol. 22, no. 2, pp. 64–75, 2009.

[34] P. F. J. Lermusiaux *et al.*, "Quantifying uncertainties in ocean predictions," *Oceanography*, vol. 19, no. 1, pp. 92–105, 2006.

[35] EU Copernicus Marine Service Information, "Global ocean 1/12 physics analysis and forecast updated daily product [data set]." [Online]. Available: https://doi.org/10.48670/moi-00016

[36] E. Metzger *et al.*, "Validation test report for the global ocean forecast system 3.5-1/25 degree hycom/cice with tides," Naval Research Lab Washington DC United States, Tech. Rep., 2020.

[37] GEBCO Bathymetric Compilation Group 2022, "The GEBCO_2022 grid - a continuous terrain model of the global oceans and land." 2022.

[38] S. Bansal *et al.*, "Hamilton-Jacobi reachability: A brief overview and recent advances," in *IEEE Conference on Decision and Control (CDC)*, 2017, pp. 2242–2253.

[39] W. Sun *et al.*, "Pursuit-evasion games in dynamic flow fields via reachability set analysis," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 4595–4600.

[40] W. Sun and P. Tsiotras, "Multiple-pursuer one-evader pursuit-evasion game in dynamic flowfields," *Journal of guidance, control, and dynamics*, vol. 40, no. 7, pp. 1627–1637, 2017.

[41] D. N. Subramani and P. F. J. Lermusiaux, "Risk-optimal path planning in stochastic dynamic environments," *Computer Methods in Applied Mechanics and Engineering*, vol. 353, pp. 391–415, Aug. 2019.