

Safe Connectivity Maintenance of Underactuated Multi-Agent Networks in Dynamic Oceanic Environments

Nicolas Hoischen^{1,2,*}, Marius Wiggert^{2,*}, and Claire J. Tomlin²

Abstract—Autonomous multi-agent systems are increasingly being deployed in environments where winds and ocean currents have a significant influence. Recent work has developed control policies for single agents that leverage flows to achieve their objectives in dynamic environments. However, in multi-agent systems, these flows can cause agents to collide or drift apart and lose direct inter-agent communications, especially when agents have low propulsion capabilities. To address these challenges, we propose a hierarchical multi-agent control approach that allows arbitrary single-agent performance policies that are unaware of other agents to be used in multi-agent systems while ensuring safe operation. We first develop a safety controller using potential functions, solely dedicated to avoiding collisions and maintaining inter-agent communication. Next, we design a low-interference safe interaction (LISIC) policy that trades off the performance policy and the safety control to ensure safe and performant operation. Specifically, when the agents are at an appropriate distance, LISIC prioritizes the performance policy while smoothly increasing the safety controller when necessary. We prove that under mild assumptions on the flows experienced by the agents, our approach can guarantee safety. Additionally, we demonstrate the effectiveness of our method in realistic settings through an extensive empirical analysis with simulations of fleets of underactuated autonomous surface vehicles operating in dynamic ocean currents where these assumptions do not always hold.

I. INTRODUCTION

Autonomous multi-agent systems, from drones to balloons and ocean surface vessels, are increasingly being explored for various applications, including inspection, collecting data, or scaling ocean aquaculture [1], [2]. In many applications, the agents communicate with each other for various purposes: to achieve a joint objective, to ensure internet coverage [3], or to share information amongst each other to improve operations. Local communication often relies on limited-range systems, e.g., sonar or radar, requiring agents to stay close to each other for network connectivity (see Fig. 1).

When a robotic system operates in the oceans and air, it is exposed to winds and currents. Most control approaches consider these as disturbances for which an overactuated control needs to compensate. What if, instead, the agent takes advantage of these flows? Recent work demonstrated that by *going with the flow* and using small actuation strategically

* Both authors contributed equally to this research. ¹ N.H. is with the School of Computation, Information and Technology of the Technical University of Munich, Germany, ² N.H., M.W., and C.J.T. are with EECS at the University of California, Berkeley, USA. For inquiries contact: mariuswiggert@berkeley.edu

The authors gratefully acknowledge the C3.ai Digital Transformation Institute, NASA ULI on Safe Aviation Autonomy, DARPA Assured Autonomy Program, and the Zeno Karl Schindler Foundation.

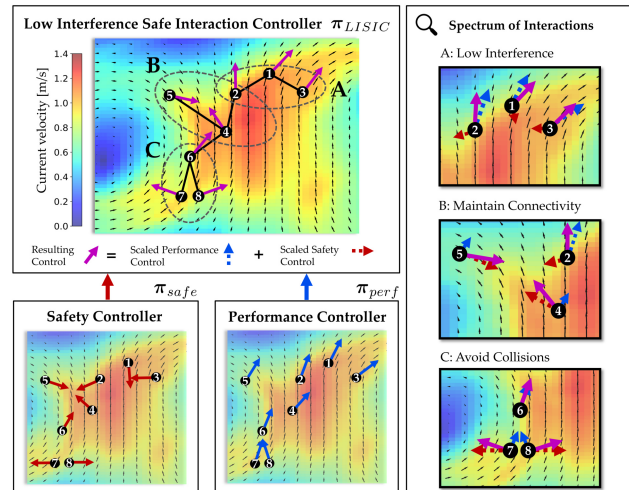


Fig. 1: Our Low Interference Safe Interaction Controller (LISIC) policy blends a single-agent performance control input with a flocking-based safety control input to avoid connectivity losses and collisions in a multi-agent network while minimally interfering with the performance objective of each agent. This ensures safe performance in ocean environments with strong ocean currents affecting the low-powered agents.

to nudge itself into favorable flows, an agent can achieve its objective with very little energy [4]–[8].

Given such individual agent performance controllers [4], we aim to develop a method that extends to multi-agent systems operating in complex flows while ensuring network connectivity and avoiding collision among agents. From the control perspective, this is challenging because of two key reasons. First, disconnections are sometimes unavoidable in the underactuated setting, where the agent’s individual propulsion is smaller than the surrounding flows, as the nonlinear, time-varying flows can push agents in opposing directions. The safe interaction controller needs to be resilient and recover connectivity after losing it. Second, constraint satisfaction needs to be traded off intelligently with the performance objective of each agent. For example, a time-optimal controller for an agent would prefer staying in strong flows, which can conflict with the network connectivity objective. Our insight is that we can simplify this multi-agent problem using three different controllers in a Hierarchical Control of Multi-Agent-Systems (H-MAS) approach (Fig. 1).

Related Literature. In H-MAS, agents are organized into multiple levels of hierarchy, with higher-level agents having more authority and control over lower-level agents, designated as followers [9]. For instance, [10] solves path planning

and ocean vehicle coordination separately with a leader-follower structure. For distance-based control tasks, such as for the Safety Controller in Fig. 1, flocking techniques can maintain connectivity by influencing the agent’s behavior to follow the movement of their neighbors while avoiding collisions [11], [12]. Recent advancements in Model Predictive Control (MPC) have also achieved connectivity and collision-free operation within Multi-Agent-Systems (MAS) [13], [14] and successfully approached control of varying-topology networks [15]. Nevertheless, a notable limitation of these approaches is their reliance on the assumptions of position invariance [13], fixed neighbor sets [14] and time invariance [13]–[15] of the flows, which do not apply in dynamic ocean environments. Thus far, the mentioned literature offers limited applicability to time-varying, uncertain flows predicted by forecasts. Since the agents are often underactuated, they cannot reliably compensate for disturbances. In this context, optimization-based approaches such as MPC frequently lead to increased computational complexity, convergence to local minima, and infeasibility with respect to the constraints. Hence, we opt for computationally efficient flocking approaches to achieve distance-based safety control, which are always feasible and, when coupled with effective single-agent planners [4], can be run at high update rates. While many flocking schemes only assume simple double integrator dynamics, adaptive flocking has also been applied to nonlinear dynamics [16]–[19]. However, due to the time-varying dynamics, adaptive approaches may not generalize well, especially given the diverse nature of currents.

Contributions. To address the above shortcomings, we propose a Low Interference Safe Interaction Controller (Fig. 1). This framework is more general in three dimensions. First, it takes an arbitrary performance policy, in contrast to [12], [20], [21], where the flock can only track reference trajectories of single or multiple virtual leaders. In fact, a feedback control policy can optimize objectives besides navigation and, in complex flows, leads to significantly better results than tracking a reference trajectory [4]. This enables the use of Dynamic Programming (DP) approaches where the value function yields optimal individual agent controls for an arbitrarily high number of agents without additional cost beyond a cheap gradient computation. This is especially powerful for multi-agent problems where the objective can be decomposed into the sum of independent single-agent objectives. Second, we provide design choices to modulate the aggressiveness of pursuing safety versus performance. Third, our approach also enables recoveries in case of connectivity losses. Finally, we investigate our method in the context of a promising approach to Carbon Dioxide Removal (CDR): utilizing robotic seaweed farms [2].

Organization. Section II introduces relevant background and metrics to evaluate our LISIC in complex flows. In Section III, we present our LISIC approach, whereas in Section IV, we prove that our method guarantees safe network interactions under certain conditions on the maximum

magnitude of the control and flow field velocities across the agents. Finally, we assess the performance of our approach in realistic ocean currents where these conditions are not always met with the metrics defined in Section II.

II. PROBLEM FORMULATION

In this section, we first describe the system’s dynamics and briefly summarize connectedness in communication graphs. Then, we define our problem statement and the metrics we use to measure constraint violation.

A. System Dynamics

We consider a swarm of N agents and use \mathcal{V} to describe the set of all agents. Let the actuation signal of each agent i be denoted by \mathbf{u}_i from a bounded set $\mathbb{U} \in \mathbb{R}^{n_u}$ where n_u is the dimensionality of the control. Then, the dynamics for each agent $i \in \mathcal{V}$ are given by:

$$\dot{\mathbf{q}}_i = v(\mathbf{q}_i, t) + g(\mathbf{q}_i, \mathbf{u}_i, t), \quad t \in [0, T] \quad (1)$$

$\mathbf{q}_i \in \mathbb{R}^n$ denotes the position of agent i in the n dimensional state space, where $n = 2$ for a surface vessel on the ocean. The movement of agent i depends on the time-varying non-linear flow field $v(\mathbf{q}_i, t) \rightarrow \mathbb{R}^n$ and its control $g(\mathbf{q}_i, \mathbf{u}_i, t)$. Although our method works for arbitrary g , we will focus on situations where the agent can directly actuate in each dimension, i.e., $g(\mathbf{q}_i, \mathbf{u}_i, t) = \mathbf{u}_i$, in line with our experiments discussed in Section V. Let the agent trajectory resulting from Eq. (1) be described by ξ_i with $\xi_i(t)$ the state \mathbf{q}_i at t . For the global system of all N agents, we use $\mathbf{q} = [\mathbf{q}_1^\top, \mathbf{q}_2^\top, \dots, \mathbf{q}_N^\top]^\top$, $\mathbf{u} = [\mathbf{u}_1^\top, \mathbf{u}_2^\top, \dots, \mathbf{u}_N^\top]^\top$, and ξ respectively to describe the state, control, and trajectory.

Remark 1. While our method also works in known currents, we focus on realistic settings, where only coarse forecasts \hat{v} are available to the planner, which differ significantly from the true flows v .

B. Communication Graph Preliminaries

The network topology of our MAS with state \mathbf{q} can be represented by an undirected communication graph $G(t)$, allowing information to flow bidirectionally between agents. The set of finite vertices $\mathcal{V} = \{1, 2 \dots N\}$ denotes individual agents, while the time-varying set of edges $\mathcal{E}(t) \subseteq \{(i, j) \in (\mathcal{V} \times \mathcal{V}), j \neq i\}$ represents direct communication between agents. Given R_{com} an upper communication threshold, $(i, j) \in \mathcal{E}(t) \iff d(\mathbf{q}_i, \mathbf{q}_j) < R_{com}$. In other words, agents i and j can communicate directly with each other if they are spatially close with respect to a distance measure $d(\mathbf{q}_i, \mathbf{q}_j)$. The graph $G(t)$ is said to be *connected* if an undirected path exists between every pair of distinct vertices. Next, we define the the $N \times N$ adjacency matrix $A(G(t)) = [a_{ij}(t)] \in \{0, 1\}$ encoding the connectivity between vertices, i.e. $a_{ij}(t) = 1 \iff (i, j) \in \mathcal{E}(t)$.

The degree of a vertex i at time t , $deg(i, t) = \sum_{j=1}^N A_{ij}(t)$ represents the number of incident edges to vertex i . The degree matrix D is then defined as the diagonal matrix $D(G(t)) = \text{diag}(deg(i, t))$. To measure the graph’s connectivity, we can compute the eigenvalues of the Laplacian

positive semi-definite matrix $L = D(G(t)) - A(G(t))$. The second smallest eigenvalue $\lambda_2(L(G(t)))$, commonly referred to as the algebraic connectivity or Fiedler value, captures the robustness of the network to link failures. In particular, $G(t)$ is connected if and only $\lambda_2(L(G(t))) > 0$ [11].

C. Problem Statement

We focus on multi-agent problems where the joint objective is the sum of independent objectives \mathbb{P}_i , which can be sketched out as:

$$\begin{aligned} \min_{\pi} \quad & \sum_{i=1}^N \mathbb{P}_i(\boldsymbol{\xi}_i, \mathbf{u}_i(\cdot)) \quad (2a) \\ \text{s.t.} \quad & \forall t \in [t_0, T] \\ & \dot{\boldsymbol{\xi}}(t) = v(\boldsymbol{\xi}(t), t) + \pi(\boldsymbol{\xi}(t)) \quad \text{global dynamics (1)} \\ & d(\boldsymbol{\xi}_i(t), \boldsymbol{\xi}_j(t)) > R_{coll} \quad (i, j) \in \mathcal{V} \times \mathcal{V}, i \neq j \quad (2b) \\ & \lambda_2(L(G(\boldsymbol{\xi}(t), R_{com}))) > 0 \quad (2c) \end{aligned}$$

We aim to find a control policy π that approximately solves Eq. (2) while being computationally cheap and always feasible. The agents are coupled in only two constraints: the collision constraint (2b) where $d(\mathbf{q}_i, \mathbf{q}_j)$ represents the distance between agent i and j and R_{coll} the minimum safe distance, and second, Eq. (2c), in maintaining a graph where all agents are connected based on the communication range R_{com} .

Assumption 1. In real-world scenarios, the initial network may begin in a disconnected state or transition to a disconnected state when underactuated agents are involved. Hence, it is assumed that each agent possesses an *emergency communication backup* to a central unit, e.g., via satellite, in case its closest distance to any other agent exceeds R_{com} . The objective is to minimize such instances, as these forms of communication are typically expensive.

Our insight is that in this setting, we can decompose the problem and handle the objectives and constraints on different levels with (1) a performance controller π_{perf} for each agent, (2) a safety controller π_{safe} , and (3) a low-interference safe interaction controller π_{LISIC} trading-off the two (Fig. 1).

The performance controller of an agent i minimizes its $(\pi_{perf})_i = \arg \min_{\pi_i} \mathbb{P}_i(\boldsymbol{\xi}_i, \mathbf{u}_i(\cdot))$ only considering its own dynamics (1). π_{perf} can be an arbitrary control policy from a fixed control signal to a feedback controller based on learning or dynamic programming (Section V). In challenging settings like ours with non-linear, time-varying dynamics, it is easier to design single-agent feedback controllers than solving the coupled multi-agent problem above, e.g., for time-optimal navigation, reference tracking, or optimizing seaweed growth [5]. The safety controller, π_{safe} , determines the control for all agents to ensure the interaction constraints (2b), (2c) are satisfied. Lastly, based on the control inputs \mathbf{u}_{perf} and \mathbf{u}_{safe} from the respective policies, the safe interaction controller decides the agents final control inputs $\mathbf{u} = \pi_{LISIC}(\mathbf{u}_{perf}, \mathbf{u}_{safe})$. To achieve good performance, the safe interaction controller should not interfere too much

with \mathbf{u}_{perf} while still ensuring connectivity and avoiding collisions. This work focuses on designing π_{safe} and π_{LISIC} for an arbitrary π_{perf} .

D. Evaluation Metrics

Due to the underactuated nature of the agents, it is impossible to guarantee network connectivity or collision avoidance in some scenarios. Hence, we define evaluation metrics to assess the safety of our control schema with respect to Eq. (2c) and (2b). A collision happens between any of the agents in the swarm if $\exists t \in [0, T]$ at which Eq. (2b) is violated. We denote this with the collision indicator $\mathbb{I}_{coll} \in \{0, 1\}$. To measure various aspects of losing connectivity, we use three metrics. First, for a binary measure, if disconnections occur, we define the disconnection indicator $\mathbb{I}_{disconn} \in \{0, 1\}$ which is 1 if $\exists t \in [0, T]$ at which Eq. (2c) is violated and zero otherwise. Additionally, we measure the minimum Fiedler value over time; the higher, the more robust the communication network (Section II-B):

$$\lambda_2^{min} = \min_{t \in [0, T]} \lambda_2(L(G(\boldsymbol{\xi}(t), R_{com}))) \quad (3)$$

Lastly, as single-agent backup communication is costly, it matters how long an agent is isolated from all other agents. Therefore, we are introducing a new measure called Isolated Platform Metric (IPM).

$$\text{IPM} = \frac{1}{T} \int_{t=0}^T M(\text{deg}(i, t) = 0) dt \quad (4)$$

where $M(\text{deg}(i, t) = 0)$ counts the number of disconnected vertices, which corresponds to the number of zeros in the diagonal of the graph degree matrix $D(G(\boldsymbol{\xi}(t), R_{com}))$ (Section II-B).

In Section V, we compare different controllers empirically over a large, representative set of missions \mathbb{M} by evaluating the collision rate $\mathbb{E}_{\mathbf{q}(t_0), t_0 \sim \mathbb{M}} [\mathbb{I}_{coll}]$, the disconnection rate $\mathbb{E}_{\mathbf{q}(t_0), t_0 \sim \mathbb{M}} [\mathbb{I}_{disconn}]$, as well as the distributions of IPM and λ_2^{min} . In our setting where the performance objectives \mathbb{P}_i are minimum time-to-target for each agent i , the connectivity constraint often leads to a trade-off with the performance objective. Hence, we also quantify the degradation of the performance controller by quantifying the minimum distance the swarm center got to the target area \mathcal{T} over the mission time $t \in [0, T]$ as $d_{min}(\mathcal{T})$.

III. METHOD

Our method tackles the multi-agent problem with a hierarchical control approach. The low interference safe interaction controller π_{LISIC} ensures performance and safe control based on an arbitrary performance controller π_{perf} and a safety controller π_{safe} (see Fig. 1). As explained in Section I, our approach is more general than [20]–[22], as an arbitrary performance policy is chosen and our framework allows for designing the aggressiveness of the safety flocking-based controller following the magnitude of the performance control input. We first introduce our flocking-inspired safety controller based on potential functions and then detail our design for π_{LISIC} . Note that this control scheme is applicable in fully actuated agents, even though our primary focus

lies in applications involving underactuated agents, which initially motivated a feasible reactive approach to collision avoidance and connectivity.

A. Flocking-Inspired Safety Controller

The sole objective of the safety controller is to ensure adequate distances between the agents without prescribing a formation. Hence, we design our safety controller π_{safe} based on the gradients of a potential function ψ .

To explain the principle, let us first focus on two connected agents i and j at an inter-agent distance $\|\mathbf{q}_{ij}\|_2 = \|\mathbf{q}_i - \mathbf{q}_j\|_2$. Consider the following bowl-shaped potential function

$$\psi_{\text{connected}}(\|\mathbf{q}_{ij}\|_2) = \frac{\kappa R_{com}}{\|\mathbf{q}_{ij}\|_2 (R_{com} - \|\mathbf{q}_{ij}\|_2)}, \quad (5)$$

where $\kappa > 0$ is a tuning factor to adjust the bell shape (see left of R_{com} in Fig. 2). Let the safety controllers for i be $\nabla_{\mathbf{q}_i} \psi_{\text{connected}}(\|\mathbf{q}_{ij}\|_2)$ and for j $\nabla_{\mathbf{q}_j} \psi_{\text{connected}}(\|\mathbf{q}_{ij}\|_2) = -\nabla_{\mathbf{q}_i} \psi_{\text{connected}}(\|\mathbf{q}_{ij}\|_2)$. When those two agents are getting too close $\|\mathbf{q}_{ij}\|_2 \rightarrow 0$, the potential $\psi(\|\mathbf{q}_{ij}\|_2)$ goes to infinity, so the gradient controllers are a strong repulsive force that pushes them away from each other. Conversely, when the two connected agents are at risk of losing their communication link $\|\mathbf{q}_{ij}\|_2 \rightarrow R_{com}$, then $\psi(\|\mathbf{q}_{ij}\|_2) \rightarrow \infty$, which means the gradient-controllers result in a strong, attractive force that brings them closer again. For multiple agents, the control becomes the sum of gradient potential terms of the other agents, and the magnitude of the gradients helps prioritize the critical inter-agent distances \mathbf{q}_{ij} .

When the agents are disconnected, which is sometimes unavoidable in underactuated settings where strong flows push them apart, we want them to be able to reconnect. Given the assumption of an emergency communication backup outlined in Section II-C, we incorporate a second term to restore connectivity among disconnected agents. To the best of our knowledge, this concept was introduced in [23]. While the augmented potential function in [23] uses a square function of distance to R_{com} for disconnected agents, we implement the second term in Eq. (6) as a square-root function. This is a design choice in the context of *underactuated* agents in a dynamic oceanic environment, where remote flock members can experience strong divergent flows and direct connectivity may be infeasible or undesirable to achieve. Thus, our approach yields a relatively low attraction force for agents beyond their communication range.

This results in our final potential function $\psi(z) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{> 0}$ that is also visualized in Fig. 2:

$$\begin{aligned} \psi(\|\mathbf{q}_{ij}\|_2) &= \sigma_{ij} \underbrace{\frac{\kappa R_{com}}{\|\mathbf{q}_{ij}\|_2 (R_{com} - \|\mathbf{q}_{ij}\|_2)}}_{\text{for connected agents}} \\ &+ (1 - \sigma_{ij}) \underbrace{\sqrt{(\|\mathbf{q}_{ij}\|_2 - R_{com} + \epsilon)}}_{\text{for disconnected agents}}. \end{aligned} \quad (6)$$

where σ_{ij} is an edge indicator similar to a_{ij} in Section II-B, but with a switching threshold $\epsilon > 0$ inducing a hysteresis when adding new edges, see Eq. (7). Hence, $\psi(\|\mathbf{q}_{ij}\|_2)$ switches between two terms whether the pair of agents (i, j) are within communication range ($\sigma_{ij} = 1$) or disconnected ($\sigma_{ij} = 0$). Following the notation of [23], we define:

$$\sigma_{ij}[t] = \begin{cases} 0, & \text{if } ((\sigma_{ij}[t^-] = 0) \cap (\|\mathbf{q}_{ij}\| \geq R_{com} - \epsilon)) \\ & \cup ((\sigma_{ij}[t^-] = 1) \cap (\|\mathbf{q}_{ij}\| \geq R_{com})), \\ 1, & \text{if } ((\sigma_{ij}[t^-] = 1) \cap (\|\mathbf{q}_{ij}\| < R_{com})) \\ & \cup ((\sigma_{ij}[t^-] = 0) \cap (\|\mathbf{q}_{ij}\| < R_{com} - \epsilon)), \end{cases} \quad (7)$$

This hysteresis mechanism avoids constant switching of the dynamical network with multiple agents for edges close to R_{com} and helps preserve connectivity in reactive control schemes [24].

The final safe interaction controller for each agent i with maximum propulsion $U_{max,i}$ is then defined as

$$(\pi_{safe})_i(\mathbf{q}) = - \frac{\sum_{j=1}^N \nabla_{\mathbf{q}_i} \psi(\|\mathbf{q}_{ij}\|_2)}{\left\| \sum_{j=1}^N \nabla_{\mathbf{q}_i} \psi(\|\mathbf{q}_{ij}\|_2) \right\|_2} U_{max,i} \quad (8)$$

B. Low Interference Safe Interaction Controller

For our π_{LISIC} that trades off the performance inputs \mathbf{u}_{perf} with the safety input \mathbf{u}_{safe} , we propose an approach that weights these control vector inputs for each agent i depending on the risk of losing connectivity or colliding.

$\mathbf{u}_i = (\pi_{LISIC})_i(\mathbf{u}_{perf}, \mathbf{u}_{safe}) = c_i^{(1)} \mathbf{u}_i^{safe} + c_i^{(2)} \mathbf{u}_i^{perf}$, $\forall i \in \mathcal{N}$ where $c_i^{(1)}$ and $c_i^{(2)}$ are weighting factors. Note that $\mathbf{u}_i^{safe} = (\pi_{safe})_i(\mathbf{q})$ depends on the other agents' positions to guarantee safe interactions.

When collisions or connectivity losses are imminent, \mathbf{u}_i should be able to rapidly tend to \mathbf{u}_i^{safe} to prioritize the safe interaction safety over performance, i.e. $c_i^{(1)} \rightarrow 1$ and $c_i^{(2)} \rightarrow 0$ (Fig. 1 B, C). Conversely, when the network is well connected and there is low danger of collisions, \mathbf{u}_i should align with \mathbf{u}_i^{perf} to have low interference with the agent's performance control, i.e. $c_i^{(1)} \rightarrow 0$ and $c_i^{(2)} \rightarrow 1$ (Fig. 1 A).

Hence we defined a weighting function $\alpha(\mathbf{q}) : \mathbb{R}^N \rightarrow [0, 1]$ such that $c_i^{(1)} = \alpha(\mathbf{q})$ and $c_i^{(2)} = 1 - \alpha(\mathbf{q})$, see an example in Section V. This function $\alpha(\mathbf{q})$ measures the urgency of \mathbf{u}_i to converge to \mathbf{u}_i^{safe} and we define it

$$c_i^{(1)} = \alpha \left(\left\| \sum_{j=1}^N \nabla_{\mathbf{q}_i} \psi(\|\mathbf{q}_{ij}\|_2) \right\|_2 \right) \quad (9)$$

The function α can be thought of as a monotonically increasing safety activation function taking values between $[0, 1]$ depending on its argument's (unbounded) magnitude. From the definition of $\psi(\|\mathbf{q}_{ij}\|_2)$ in Eq. (6), $\lim_{\|\mathbf{q}_{ij}\|_2 \rightarrow 0} \psi(\|\mathbf{q}_{ij}\|_2) = \infty$ and $\lim_{\|\mathbf{q}_{ij}\|_2 \rightarrow R_{com}} \psi(\|\mathbf{q}_{ij}\|_2) = \infty$. Hence, in critical situations $\left\| \sum_{j=1}^N \nabla_{\mathbf{q}_i} \psi(\|\mathbf{q}_{ij}\|_2) \right\|_2$ gets very large so that $c_i^{(1)}$ saturates to 1 and $c_i^{(2)}$ to 0, thus prioritizing the network

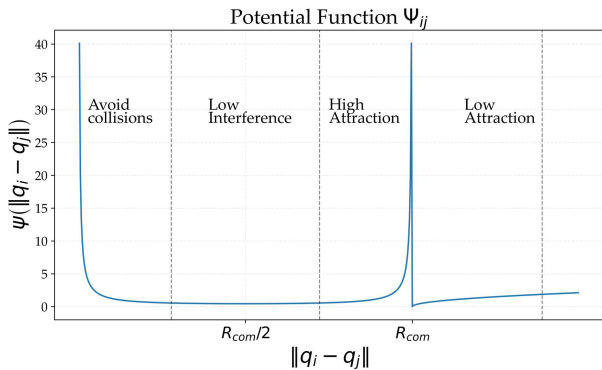


Fig. 2: Augmented potential function, with two terms to account for agents within and outside the communication range R_{com} . A high κ parameter is shown to increase the steepness of the slope around $\frac{R_{com}}{2}$, depending on how achieving the exact ideal distance is valued.

safety for the concerned agents $i \in \mathcal{V}$ i.e. $\mathbf{u}_i \rightarrow \mathbf{u}_i^{safe}$, over each agent individual objective \mathbf{u}_i^{perf} .

In other words, $\psi(\|\mathbf{q}_{ij}\|_2)$ has a contractivity property for agent inter-distances at the boundaries of the safe set, defined by 0 and R_{com} , similarly to Control Barrier Functions (CBFs) [25]. With this design, we ensure that agents coming from a disconnected status $\sigma_{ij}[t^-] = 0$ to a connected status $\sigma_{ij}[t] = 1$ experience a strong attracting gradient \mathbf{u}_i^{safe} to avoid escaping the communication range again. From Fig. 2, it is also clear that when the network is close to being ideally connected, the gradient norm of the potential function $\left\| \sum_{j=1}^N \nabla_{\mathbf{q}_i} \psi(\|\mathbf{q}_{ij}\|_2) \right\|_2$ is low so that agent's i control input is dominated by the performance controller since $c_i^{(1)} \rightarrow 0$ and $c_i^{(2)} \rightarrow 1$.

IV. THEORETICAL ANALYSIS

This section analyzes under which conditions our safe interaction controller can maintain connectivity and avoid collisions [26]. Our analysis follows the common approach to demonstrate that a flock converges to a lattice structure while preventing inter-agent collisions using energy-based analysis and LaSalle's invariance principle [12]. We start by introducing a moving referential frame for the structural collective dynamics [12] with respect to the flock centroid \mathbf{q}_c . The relative coordinates are given by $\tilde{\mathbf{q}}_i = \mathbf{q}_i - \mathbf{q}_c$ and $\tilde{\mathbf{q}}_{ij} = \tilde{\mathbf{q}}_i - \tilde{\mathbf{q}}_j = \mathbf{q}_{ij}$. Therefore, $\psi(\|\mathbf{q}_{ij}\|_2) = \psi(\|\tilde{\mathbf{q}}_{ij}\|_2)$, and the total tension energy or potential energy for the structural dynamics in the relative coordinates yields

$$H(\tilde{\mathbf{q}}) = \frac{1}{2} \sum_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N \psi(\|\tilde{\mathbf{q}}_{ij}\|_2) \quad (10)$$

A possible approach, although conservative, is to show that a global tension energy decrease of the system $\dot{H} = \sum_{i=1}^n \dot{H}_i \leq 0$ can be achieved by guaranteeing local tension energy decrease $\forall i \in \mathcal{V}$. Assume that $G(t)$ switches at time t_l for $l = 0, 1, 2, \dots$ and $\dot{H} \leq 0$ on each $[t_l, t_{l+1})$. Then, $H(t_k) = H(t_k^-) + m_k \psi(\|R_{com} - \epsilon\|)$ [23] with m_k the number of edges added at switching time k . The energy can

be bound for any subsequent time as the graph topology becomes fixed after a certain time, and only a finite number of maximum edges can be added. The time-derivative of H_i along the trajectory of agent i yields

$$\dot{H}_i = \dot{\tilde{\mathbf{q}}}_i^\top \sum_{\substack{j=1 \\ j \neq i}}^N \nabla_{\tilde{\mathbf{q}}_i} \psi(\|\tilde{\mathbf{q}}_{ij}\|_2) \quad (11)$$

where we exploited the relation $\nabla_{\mathbf{q}_i} \psi(\|\mathbf{q}_{ij}\|_2) = -\nabla_{\mathbf{q}_j} \psi(\|\mathbf{q}_{ij}\|_2)$. We seek a condition linking the maximum actuation power of each agent $U_{max,i}$ to the dynamics of the flock, subject to the nonlinear flow v . For ease of understanding, assume holonomic actuation i.e. $g(\mathbf{q}_i, \mathbf{u}_i, t) = \mathbf{u}_i$, then $\mathbf{u}_i^{safe} = (\pi_{safe})_i(\mathbf{q})$ can be directly substituted with Eq. (8). Using $\tilde{\dot{\mathbf{q}}}_i = \dot{\mathbf{q}}_i - \dot{\mathbf{q}}_c$ and the Cauchy-Schwarz inequality in Eq. (11) yields:

$$\left\| c_i^{(2)} \mathbf{u}_i^{perf}(\mathbf{q}_i) + v(\mathbf{q}_i) - \text{Ave}(\dot{\mathbf{q}}_{N_i}) \right\|_2 \leq c_i^{(1)} U_{max,i} \quad (12)$$

where $\text{Ave}(\cdot)$ denotes the average and the set $N_i = \mathcal{V} \setminus \{i\}$ the neighboring agents of i . More details about this proof can be found in [26]. A similar inequality to Eq. (12) can be derived for the general dynamics defined in Eq. (1) if $g(\cdot)$ is a linear map.

Let us interpret Eq. (12). The dynamics of the neighboring agents of i depend on their surrounding flows and respective individual control inputs, i.e. $\text{Ave}(\dot{\mathbf{q}}_{N_i}) = \text{Ave}(v(\mathbf{q}_{N_i})) + \text{Ave}(\mathbf{u}_{N_i})$. Despite strong flows, the agents do not necessarily need to be overactuated to meet a local energy decrease $\dot{H}_i \leq 0$. Eq. (12) can be fulfilled even if $\|v(\mathbf{q}_i)\|_2 > U_{max,i}$, since $v(\mathbf{q}_i)$ can be compensated by $\text{Ave}(v(\mathbf{q}_{N_i}))$. In other words, agents in strong flows could still maintain connectivity and avoid collisions as long as the currents experienced by each agent and its neighbors are of similar direction and magnitude. The neighboring flocking control inputs average $\text{Ave}(\mathbf{u}_{N_i})$ also helps accounting for the current difference term $v(\mathbf{q}_i) - \text{Ave}(v(\mathbf{q}_{N_i}))$. Under these assumptions, we can show that $\dot{H} \leq 0$, which allows to bound the maximum energy and apply LaSalle's Invariance Principle [23], [27], thus ensuring that no collisions or disconnections happen, since $\psi(\|\mathbf{q}_{ij}\|_2) \rightarrow \infty$ when $\|\mathbf{q}_{ij}\|_2 \rightarrow 0$ or $\|\mathbf{q}_{ij}\|_2 \rightarrow R_{com}$. However, for strong divergent flows between agents, it can happen that $\|v(\mathbf{q}_i) - \text{Ave}(v(\mathbf{q}_{N_i}))\|_2 \gg U_{max,i}$ due to the underactuated nature of the agents, which makes satisfaction of Eq. (12) challenging. Note that Eq. (12) is sufficient but not necessary to guarantee $\dot{H} < 0$, as negative local energies can compensate for positive ones.

V. SIMULATION STUDY

The proposed scheme is evaluated on realistic ocean currents, seeking a close reproduction of an innovative CDR approach [2], [5]. We use multi-time Hamilton Jacobi Reachability (mt-HJR) as a performance single agent controller since it generates a value function yielding the time-optimal control everywhere [4].

A. Experimental Set-Up

We study the effectiveness of different controllers in maneuvering a two-dimensional Autonomous Surface Vehicle (ASV), with a thrust angle θ as a control input. The actuation can be assumed holonomic of fixed thrust magnitude $\|u\|_2 = 0.1$ m/s, as the vessel can turn to a desired θ within seconds while our sampling times are 10 minutes. We consider a group of identical $N = 30$ ASVs with omnidirectional communication capabilities, navigating in strong ocean currents $v(q, t) \in [0.3, 2]$ m/s. Each single agent’s objective is to reach a target region common to all ASVs, which could be identified as an ideal seaweed growth region in the context of floating robotic farms [2]. Next, we detail the creation of an extensive set of missions to illustrate trade-offs between single-agent objectives and flock connectivity maintenance in a realistic ocean environment.

a) *Realistic Simulation of Ocean Conditions:* Inspired by [4], we focus on the Gulf of Mexico region (Fig. 3), as it presents challenging currents. Moreover, we employ two ocean current data sources, which we refer to as HYCOM *hindcasts* [28] and Copernicus *hindcasts* [29] that we use as *forecasts* for realistic scenarios. In our context, the ocean forecast data represents predicted currents \hat{v}_{FC} while the hindcast ocean data are true flows v . While the forecast error affects the optimality of the performance mt-HJR controllers, the advantage of our reactive safety controller design over predictive schemes [13], [14] is that it is not affected by the error on forecasted currents. We propose two settings to investigate our approach, namely (a) performance mt-HJR planning on *hindcasts* and simulation on *hindcasts* (HC-HC) and (b) performance mt-HJR planning on *forecasts* and simulation on *hindcasts* (FC-HC). The first allows us to assess performance in an idealized setting where true flows are known, while the second reflects a realistic application in dynamic ocean environments.

b) *Large Representative Set of Missions:* We assume all agents start simultaneously at time t_0 a navigation mission to a target region \mathcal{T} . The navigation objective of each ASV is to reach \mathcal{T} from their start states $(q_1(t_0) \dots q_n(t_0))$ within a maximum allowed time $T_{timeout}$. The target \mathcal{T} is defined as a circular region with center coordinates $q_{\mathcal{T}}$ and fixed radius $r_{\mathcal{T}} = 0.1^\circ$ around it. To obtain a diverse set of missions \mathbb{M} , the starting times t_0 are uniformly sampled between April 2022 and December 2022. $T_{timeout}$ is set to 144h, and the start points are sampled such that the ASVs can reach the target within $[72, 144]h$ to ensure that missions are by definition feasible on true flows and temporally representative enough of realistic scenarios. To prevent stranding side effects, we impose a minimum distance of 111km between the target area and the land and a minimum distance of 40km between each ASV’s initial position and the land. We generate $|\mathbb{M}| = 1000$ missions of initially connected and collision-free networks, see Fig. 3.

B. Baseline Controllers

We build on recent work that proposed a reliable mt-HJR controller for underactuated agents utilizing complex flows

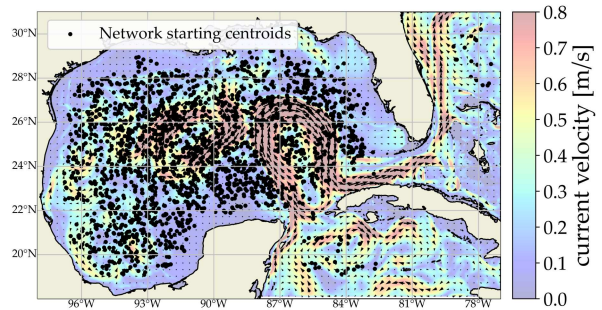


Fig. 3: We sample a large set of missions $|\mathbb{M}| = 1000$ in the Gulf of Mexico that are spatially and temporally representative of realistic scenarios.

[4]. This approach directly extends to multiple agents with little extra computation. The feedback controller for agent i can be obtained from an optimal value function \mathcal{J}^* at time t as $u_i(t)^* = \arg \min_{u_i \in \mathcal{U}} g(q_i, u_i, t) \cdot \nabla_{q_i} \mathcal{J}^*(q_i, t)$.

All evaluated controllers use the mt-HJR formulation as a single agent performance control. Our baseline scheme, called multi-time Hamilton Jacobi Reachability Baseline (mt-HJR-B), involves each agent only utilizing its time-optimal performance control mt-HJR without considering multi-agent interactions. This baseline provides a reasonable estimation of the likelihood of collisions and communication losses if each agent were to rely solely on its performance control. In addition, we define a second baseline controller, multi-time Hamilton Jacobi Reachability with multi-agent Reactive Control (mt-HJR-RC) adapted from [30]. This controller operates in three modes: ACHIEVECONNECTIVITY, MAINTAINCONNECTIVITY, and GOTOGOAL, which are selected based on the ASVs’ relative positions. The MAINTAINCONNECTIVITY and GOTOGOAL modes employ a general navigation function for each agent, which we instantiate to our mt-HJR performance controller. This approach is easily integrated with the time-optimal control mt-HJR, and the reactive control term can be implemented decentralized.

Finally, we denote our Low Interference Safe Interaction Controller (LISIC) approach from Section III as multi-time Hamilton Jacobi Reachability with multi-agent LISIC (mt-HJR-LISIC). The single agent performance controller u_i^{perf} is again mt-HJR. The trade-off between each agent’s navigational objective and the safe network interaction can be tuned with two parameters. First, the potential function shape (Fig. 2) can be more or less flat around the ideal distance $R_{com}/2$. In this application, we set $\kappa = 2$. Furthermore, we now detail our weighting scheme for $c_i^{(1)}$ and $c_i^{(2)}$ via the definition of α in Eq. (9) as a SOFTMAX-like function

$$c_i^{(1)} = \frac{e^{\|\sum_{j=1}^n \nabla_{q_i} \psi(\|q_{ij}\|_2)\|_2}}{e^{\|\sum_{j=1}^n \nabla_{q_i} \psi(\|q_{ij}\|_2)\|_2} + e^\rho}, \quad \forall i \in \mathcal{V}. \quad (13)$$

where the parameter $\rho \geq 0$ can be adjusted to achieve faster saturation of the potential function gradient term $u_i^{safe}(q)$.

C. Additional Parameters and Metrics

The upper connectivity bound R_{com} in Eq. (2c) and (7) is set to 9km, which corresponds approximately to radio communication capabilities for ASVs. The collision lower threshold from Eq. (2b) is set to $R_{coll} = 100\text{m}$, providing a practical margin, as one would typically do in a real-world implementation. Moreover, we define $\epsilon = 300\text{m}$ for the edge hysteresis parameter from Eq. (7). We use the Euclidean norm to measure inter-agent distances $d(\mathbf{q}_i, \mathbf{q}_j)$ and the minimum flock center distance to target $d_{min}(\mathcal{T})$. The parameters used in the experiments are summarized in Table I.

Symbol	Description	Value
U_{max}	ASVs maximum actuation	0.1m/s
v	Time-varying ocean currents	[0.3, 2]m/s
R_{com}	Upper connectivity bound	9000m
R_{coll}	Lower collision threshold	100m
$r_{\mathcal{T}}$	Radius of circular target region	0.1°
ρ	Saturation of potential function	1 (no units)
κ	Shape of potential function	2 (no units)
ϵ	Hysteresis for adding or removing edges	300m
$T_{timeout}$	Duration of a mission	144h

TABLE I: Relevant Simulation and Controller Parameters.

D. Numerical Results

The results over apriori known true currents (HC-HC) and realistic scenario (FC-HC) are presented in Table II. Both mt-HJR-LISIC and mt-HJR-RC exhibit superior performance in terms of connectivity and collision metrics compared to the baseline mt-HJR-B. Thus, we conduct statistical testing to compare mt-HJR-RC and mt-HJR-LISIC. Regarding the disconnection and collision rate, we perform a one-sided two-sample z proportion test for mt-HJR-LISIC against mt-HJR-RC.

Let Γ be the rate collision or disconnection over \mathbb{M} with the null hypothesis $H_0 : \Gamma_{\text{mt-HJR-LISIC}} = \Gamma_{\text{mt-HJR-RC}}$ to reject in favor of the alternative hypothesis $H_A : \Gamma_{\text{mt-HJR-LISIC}} < \Gamma_{\text{mt-HJR-RC}}$. mt-HJR-LISIC is statistically significantly better than mt-HJR-RC at avoiding disconnections in both (HC-HC) and (FC-HC) scenarios, with p-values of $p = 6.3e^{-69}$ and $p = 1.7e^{-114}$, respectively. However, it is not significantly better than mt-HJR-RC at avoiding collisions. We also perform a Welch's t-test due to the unequal variances of mt-HJR-RC and mt-HJR-LISIC to test (1) connectivity using the means over $|\mathbb{M}|$ of the Isolated Platform Metric, i.e., $\mu(\text{IPM})$ and the minimum Fiedler value recorded over time, i.e., $\mu(\lambda_2^{min})$, (2) performance trade-off with $\mu(d_{min}(\mathcal{T}))$. For both (HC-HC) and (FC-HC) scenarios, mt-HJR-LISIC leads to statistically significantly better results for the network connectivity with $p < 1e^{-30}$ for $\mu(\text{IPM})$ and $\mu(\lambda_2^{min})$ while mt-HJR-RC displays a better objective trade-off $\mu(d_{min}(\mathcal{T}))$ with p-values $p < 1e^{-40}$. Moreover, we plot the IPM and $\mu(\lambda_2^{min})$, evaluated on the full set of missions $|\mathbb{M}|$ for the three controllers in Fig. 4. Among the three evaluated controllers, mt-HJR-LISIC has the lowest IPM. Because of its higher value of $\mu(\lambda_2^{min})$

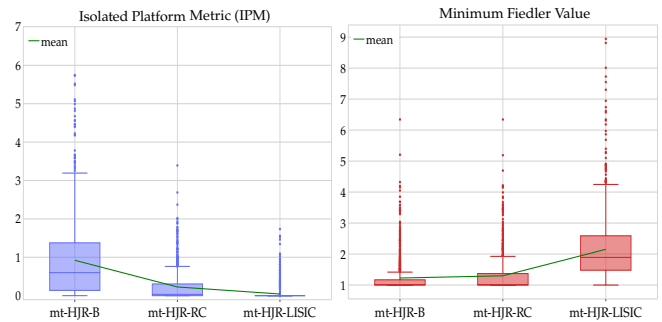


Fig. 4: Left: IPM evaluated on \mathbb{M} . Due to its low IPM, mt-HJR-LISIC typically has both a low disconnection time and a low number of disconnected agents. Right: The minimum Fiedler value λ_2^{min} can be used as a graph connectivity measure. A high λ_2^{min} ensures better robustness against connectivity failures.

(see Fig. 4, right), mt-HJR-LISIC is more robust against disconnections, and should be the preferred control choice when communication maintenance is prioritized. Finally, Fig. 5 illustrates a navigation mission, comparing a naive multi-agent approach (mt-HJR-B) to our safe interaction controller, mt-HJR-LISIC. Note that despite the initial strong currents pushing the ASVs away from the desired goal in Fig. 5, the currents eventually shift favorably, allowing the under-actuated ASVs to reach the target. The mt-HJR framework leverages this information through current forecasts to plan intelligently.

	Coll.	Disconn.	$\mu(\text{IPM}) \downarrow$	$\mu(\lambda_2^{min}) \uparrow$	$\mu(d_{min}(\mathcal{T})) \downarrow$
π_{perf} plans on true flows: HC-HC					
mt-HJR-B	68.5%	50.1%	0.37	0.39	0 km
mt-HJR-RC	0%	44.8%	0.19	0.42	0.14 km*
mt-HJR-LISIC	0.7%	9.9%*	0.05*	1.15*	5.90 km
π_{perf} plans on forecasts: FC-HC					
mt-HJR-B	39.1 %	70.5%	0.92	0.23	10.55 km
mt-HJR-RC	0%	58%	0.23	0.30	10.84 km*
mt-HJR-LISIC	0.7%	9.9%*	0.043*	1.15*	13.96 km

TABLE II: We compare the performance of three controllers in two forecast settings. The * symbol indicates a statistically significant better performance in terms of connectivity, collisions, and distance to the target.

E. Discussion

It is clear that mt-HJR-LISIC outperforms mt-HJR-RC and the mt-HJR-B in terms of connectivity metrics. Interestingly, mt-HJR-LISIC leads to a slightly higher collision rate in Table II than mt-HJR-RC. We believe that it is mainly due to two reasons: (1) In mt-HJR-RC, the expected risk of collisions is inherently lower as each agent can achieve connectivity with a maximum amount of two other agents [30] while mt-HJR-LISIC achieves a similar structure to a lattice configuration [12] (2) In our example, all agents navigate to the same target, which also increases the risk of collisions, as it is a common implicit regularizer. We expect improvement in collision rate for application to autonomous ASVs, where each agent maximizes an objective along its trajectory [5]. The discrepancy between the performance trade-off with each agent target reaching objective $d_{min}(\mathcal{T})$

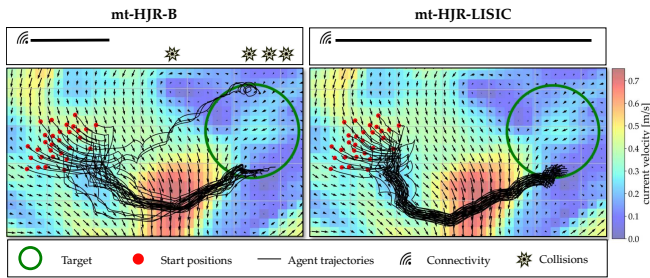


Fig. 5: Mission example with mt-HJR-B (left) versus mt-HJR-LISIC (right). Note that while the agents' trajectories are depicted for the interval $[t_0, T_{\text{timeout}}]$, the currents in the background represent a snapshot at time t_0 and evolve in both direction and magnitude over time. mt-HJR-LISIC guarantees communication through the full length of the mission, avoids collisions, and ensures that all agents reach the target.

in Table II is less noticeable in the (FC-HC) setting, since the mt-HJR performance is also degraded because of the stochastic error when planning on forecasts [4].

VI. CONCLUSION AND FUTURE WORK

In this work, we proposed a H-MAS approach to maintain network connectivity in complex dynamical flows while satisfying single agent level objectives when feasible. Our method blends a safety controller for collisions and connectivity maintenance with a performance control policy, which allows us to decompose a complex multi-agent problem effectively. Our empirical results in realistic ocean dynamics showed that our method efficiently maintains connectivity and avoids collisions in most scenarios while reasonably trading off with each agent's performance objective. Future work involves real-world testing of our experiments, as well as adapting predictive methods [13], [14] to time-varying flows to anticipate disconnections and collisions utilizing the only available coarse ocean forecasts.

REFERENCES

- [1] P. Sujit *et al.*, "Uav and auvs coordination for ocean exploration," in *Oceans 2009-Europe*. IEEE, 2009, pp. 1–7.
- [2] "Floating seaweed farms," Phykos PBC. [Online]. Available: <https://www.phykos.co>
- [3] "Project loon," X, the moonshot factory. [Online]. Available: <https://x.company/projects/loon/>
- [4] M. Wiggert *et al.*, "Navigating underactuated agents by hitchhiking forecast flows," in *IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2417–2424.
- [5] M. Killer *et al.*, "Maximizing seaweed growth on autonomous farms: A dynamic programming approach for underactuated systems navigating on uncertain ocean currents," *arXiv:2307.01916*, 2023.
- [6] M. G. Bellemare *et al.*, "Autonomous navigation of stratospheric balloons using reinforcement learning," *Nature*, vol. 588, no. 7836, pp. 77–82, 2020.
- [7] M. Doshi *et al.*, "Hamilton-jacobi multi-time reachability," in *IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2443–2450.
- [8] A. Doering *et al.*, "Stranding risk for underactuated vessels in complex ocean currents: Analysis and controllers," in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 7055–7060.
- [9] C.-X. Dou and B. Liu, "Multi-Agent Based Hierarchical Hybrid Control for Smart Microgrid," vol. 4, no. 2, pp. 771–778.
- [10] T. Lolla *et al.*, "Path planning in multi-scale ocean flows: Coordination and dynamic obstacles," vol. 94, pp. 46–66. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1463500315001225>

- [11] M. M. Zavlanos *et al.*, "Graph-theoretic connectivity control of mobile robot networks," vol. 99, no. 9, pp. 1525–1540. [Online]. Available: <http://ieeexplore.ieee.org/document/5948318/>
- [12] R. Olfati-Saber, "Flocking for Multi-Agent Dynamic Systems: Algorithms and Theory," vol. 51, no. 3, pp. 401–420. [Online]. Available: <http://ieeexplore.ieee.org/document/1605401/>
- [13] A. Carron *et al.*, "Multi-agent distributed model predictive control with connectivity constraint," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 3806–3811, 2023.
- [14] A. Filotheou *et al.*, "Decentralized Control of Uncertain Multi-Agent Systems with Connectivity Maintenance and Collision Avoidance," in *2018 European Control Conference (ECC)*. IEEE, pp. 8–13. [Online]. Available: <https://ieeexplore.ieee.org/document/8550343/>
- [15] A. Aboudonia *et al.*, "Reconfigurable plug-and-play distributed model predictive control for reference tracking," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 1130–1135.
- [16] W. Yu and G. Chen, "Robust adaptive flocking control of nonlinear multi-agent systems," in *2010 IEEE International Symposium on Computer-Aided Control System Design*, pp. 363–367.
- [17] Q. Zhang *et al.*, "Adaptive flocking of non-linear multi-agents systems with uncertain parameters," vol. 9, no. 3, pp. 351–357. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1049/iet-cta.2014.0471>
- [18] Y. Zou *et al.*, "Flocking of uncertain nonlinear multi-agent systems via distributed adaptive event-triggered control," vol. 465, pp. 503–513. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S092523122101359X>
- [19] Z. Peng *et al.*, "Neural adaptive flocking control of networked underactuated autonomous surface vehicles in the presence of uncertain dynamics," in *Proceedings of the 31st Chinese Control Conference*, pp. 2865–2870.
- [20] H.-S. Su, "Flocking in Multi-Agent Systems with Multiple Virtual Leaders Based Only on Position Measurements," vol. 57, no. 5, pp. 801–807. [Online]. Available: <https://iopscience.iop.org/article/10.1088/0253-6102/57/5/10>
- [21] Z. Li *et al.*, "Flocking of Multi-Agent Systems with Multiple Virtual Leaders Based on Connectivity Preservation Approach," in *Proceedings of the 2015 Chinese Intelligent Systems Conference*, ser. Lecture Notes in Electrical Engineering, Y. Jia *et al.*, Eds. Springer, pp. 465–473.
- [22] H. Su *et al.*, "Flocking in multi-agent systems with multiple virtual leaders," vol. 10, no. 2, pp. 238–245. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/asjc.22>
- [23] T. Yan *et al.*, "Flocking of Multi-agent Systems with Unknown Nonlinear Dynamics and Heterogeneous Virtual Leader," vol. 19, no. 9, pp. 2931–2939. [Online]. Available: <https://link.springer.com/10.1007/s12555-020-0578-3>
- [24] M. Ji and M. Egerstedt, "Distributed Coordination Control of Multi-agent Systems While Preserving Connectedness," vol. 23, no. 4, pp. 693–703.
- [25] A. D. Ames *et al.*, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference (ECC)*, 2019, pp. 3420–3431.
- [26] N. Hoischen, "Safe connectivity maintenance for a fleet of underactuated seaweed farms in dynamic oceanic environments," Master's thesis, 2023. [Online]. Available: <https://doi.org/10.3929/ethz-b-000626363>
- [27] M. Wang *et al.*, "Flocking of multiple autonomous agents with preserved network connectivity and heterogeneous nonlinear dynamics," vol. 115, pp. 169–177. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S092523121300101X>
- [28] 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.
- [29] 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>
- [30] G. A. Pereira *et al.*, "Decentralized motion planning for multiple robots subject to sensing and communication constraints," *Departmental Papers (MEAM)*, p. 45, 2003.