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

Nicolas Hoischen*[1], Marius Wiggert*[2], and Claire J. Tomlin[2]

***Abstract*— Autonomous Multi-Agent Systems are increasingly being deployed in environments where winds and ocean currents can exert a significant influence on their dynamics. Recent work has developed powerful control policies for single agents that can leverage flows to achieve their objectives in dynamic environments. However, in the context of 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 solely dedicated to avoiding collisions and maintaining inter-agent communication. Subsequently, we design a low-interference safe interaction (LISIC) policy that trades-off the performance policy and the safety controller to ensure safe and optimal 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 underactuated Autonomous Surface Vehicles (ASV) operating in dynamical ocean currents where the assumptions do not always hold.**

I. Introduction

Autonomous Multi-Agent systems from drones, to balloons, to ocean surface vessels and underwater robots are increasingly explored for various applications from inspection and surveying, to providing internet access, collecting data and open ocean aquaculture [1]–[3]. In many applications the agents communicate amongst each other for various purposes such as coordination to achieve a joint objective, to ensure internet coverage [2], or to reduce the amount of communication needed to an external centralized controller. 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 an agent operates in the oceans and air it is exposed to winds and currents. Most 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 to nudge itself into favorable flows, the agent is able to achieve its objective with very little energy, be it station-keeping of balloons or navigating in the oceans [4]–[7].

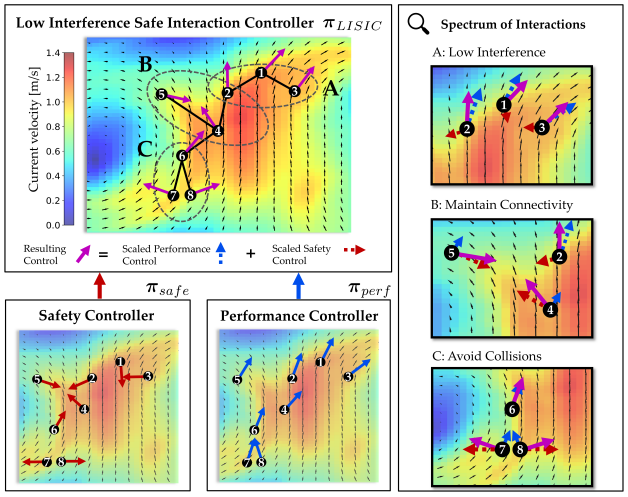

Fig. 1: Our Low Interference Safe Interaction Controller (LISIC) policy blends a single agent performance control input with a flocking-based safety control input to avoid connectivity losses and collisions in a multi-agent network, while minimally interfering with the performance objective of each agent.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

Inspired by these advances, we want to bring these powerful controllers to multi-agent networks operating in complex flows. Dynamic Programming (DP) approaches provide a value function, that yields optimal individual agent controls for an arbitrarily high number of agents without any additional cost beyond a cheap gradient computation. This is especially powerful for multi-agent problems where the objective can be decomposed in the sum of independent single agent objectives, e.g. in floating seaweed farms, the seaweed growth of each agent does not depend on the growth of other agents [5]. Given such individual agent performance controllers, our goal is to develop a method that ensures safe interaction of the multiple agents. Specifically, we focus on network connectivity and avoiding collisions.

Our insight is that this structure allows us to tackle the multi-agent problem with three different controllers in a Hierarchical Control of Multi-Agent-Systems (H-MAS) approach (Fig. 1). From the control perspective, this is challenging because of two key reasons. First, in the underactuated setting disconnections are sometimes unavoidable as the non-linear, time-varying flows can push agents in

* Both authors contributed equally to this research. [1]N.H. is with ETH Zurich, Switzerland, [2] 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 support of the C3.ai Digital Transformation Institute, the NASA ULI Grant on Safe Aviation Autonomy, the DARPA Assured Autonomy Program, the SRC CONIX Center, the ONR BRC program and the Zeno Karl Schindler Foundation.

opposing directions. The safe interaction controller needs to be resilient and recover connectivity after it was lost. Second, constraint satisfaction needs to be traded-off intelligently with the performance objective of each agent as they can often be in conflict. For example, a time-optimal controller for an agent would prefer staying in strong flows whereas the safe interaction controller needs to trade this off with maintaining connected to other agents (Fig. 1).

H-MAS can be viewed from the perspective of hierarchical reinforcement learning [8], which shares substantial similarities with our approach of breaking down complex tasks into two more manageable sub-tasks. 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 vehicles coordination separately with a leader-follower hierarchical structure. Graph methods provide a solid theoretical foundation for analyzing connectivity properties in distance-based control applications [11]. Our interest is drawn to flocking techniques incorporating a single agent navigation task via a dynamic virtual leader γ-agent [12], which maintain connectivity by influencing the agents' behavior to follow the movement of their neighbors, while avoiding collisions. Extensions to multiple virtual leaders are proposed in [13]–[15] holding promise for the generalization of virtual leaders as performance controllers. While most of the flocking schemes generally assume simple double integrator dynamics, adaptive flocking also emerged to handle to non-linear uncertain dynamics [16]–[19]. More recently, Model Predictive Control (MPC) has been successful in ensuring connectivity and collision-free operation in Multi-Agent-Systems (MAS) [20].

MPC approaches to connectivity or safety [20], [21], in addition to being computationally intensive, often assume position-invariant dynamics, which do not hold in dynamic ocean environments. This motivates a cheaper reactive strategy approach in this context. Determining the follower position with respect to the leader in [10] requires solving an additional optimization problem to achieve formation control, which adds complexity and feasibility issues for underactuated agents. On the other hand, adaptive and robust flocking schemes rely on overactuated agents to overcome the disturbances. Adaptive flocking in underactuated systems has been considered using a Radial Basis Function Neural Network (RBF-NN) to approximate uncertain non-linear dynamics [22], but can easily overfit the training data, which is a significant concern given the sparse and stochastic nature of ocean data. Existing literature on flocking has primarily focused on tracking a reference trajectory of a virtual leader or multiple virtual leaders. However, in complex flows, an optimal feedback control policy leads to significantly better results than tracking a reference trajectory [4].

To address the above shortcomings, we propose a safe interaction control policy Low Interference Safe Interaction Controller (LISIC), blending a performance single agent controller with a flocking-inspired safe controller. In addition to ensuring the network safe operation in terms of collisions and communication maintenance, our approach also enables recoveries in case of connectivity failures. The dynamics of communication and information sharing could then later be handled on another level with a Plug-and-Play control scheme [23], but is not in the scope of this work. In Sec. II, we introduce introduce important background and relevant metrics to measure connectivity in terms of time and network topology, as well as a single agent performance trade-off. We show our results theoretically in Sec. III and Sec. IV. Finally we assess the performance of our approach by conducting a large-scale empirical evaluation with agents that are underactuated in the sense that their propulsion is smaller than the surrounding flows.

II. Problem Formulation

In this section we first describe the systems dynamics and give a brief summary of 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. The dynamics for each agent $i \in \mathcal{V}$ are given by:

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

$\boldsymbol{q}_i \in \mathbb{R}^n$ 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 the agent i depends on the drift of the surrounding flow $v(\boldsymbol{q}_i, t)$ and the bounded control $\boldsymbol{u}_i \in \mathbb{U} \in \mathbb{R}^{n_u}$ where n_u is the dimensionality of the control. Let the agent trajectory resulting from Eq. 1 be described by $\boldsymbol{\xi}_i$ with $\boldsymbol{\xi}_i(t)$ the state $\boldsymbol{q}_i$ at t. For the global system of all N agents we use $\boldsymbol{q} = [\boldsymbol{q}_1^\top, \boldsymbol{q}_2^\top, \ldots, \boldsymbol{q}_N^\top]^\top$, $\boldsymbol{u} = [\boldsymbol{u}_1^\top, \boldsymbol{u}_2^\top, \ldots, \boldsymbol{u}_N^\top]^\top$, and $\boldsymbol{\xi}$ respectively to describe the state, control, and trajectory.

B. Communication Graph Preliminaries

The network topology of our Multi-Agent-Systems with state $\boldsymbol{q}$ can be represented with a graph abstraction in order to model interactions among agents. The communication graph $G(t)$ can be built from a set of finite vertices $\mathcal{V} = \{1, 2 \ldots N\}$ representing individual agents and a time-varying set of edges $\mathcal{E}(t) \subseteq \{(i, j) \in (\mathcal{V} \times \mathcal{V}), j \neq i\}$ representing direct communication between agents. We focus on *undirected* graphs implying that information can flow between agents in both directions. We further assume that only neighbors that are spatially close with respect to a distance measure $d(\boldsymbol{q}_i, \boldsymbol{q}_j)$ can communicate directly with each other. Given an upper communication threshold R_{com}, the pair of vertices i, j is connected by an edge $d(\boldsymbol{q}_i, \boldsymbol{q}_j) < R_{com} \iff (i, j) \in \mathcal{E}(t)$. The graph $G(t)$ is said to be *connected* if there exists an undirected path between every pair of distinct vertices. We can analyze the connectivity of the graph with its Laplacian matrix L which is a symmetric and positive semi-definite matrix based on the adjacency and the degree matrix which we define next. The adjacency matrix $A(t)$ is an $n \times n$ binary matrix that encodes which

vertices are connected to each other $A(G(t)) = [a_{ij}(t)] \in \{0,1\}$ with $a_{ij}(t) = 1 \iff (i,j) \in \mathcal{E}(t)$.

The valency or degree of a vertex i is denoted by $deg(i,t)$ and represents the number of its incident edges which is the row-sum of the adjacency matrix $deg(i,t) = \sum_{j=1}^{N} A_{ij}(t)$. The degree matrix D is then defined as the diagonal matrix $D(G(t)) = \text{diag}\,(deg(i,t))$. The Laplacian matrix L can then be inferred as $L(G(t)) = D(G(t)) - A(G(t)$, which is a symmetric and positive semi-definite matrix. The eigenvalues of L let us measure the graphs connectivity. In particular, the second smallest eigenvalue $\lambda_2(L(G(t))$ is commonly called the algebraic connectivity or Fiedler value of the network, and captures the robustness of the network to link failures. The graph $G(t)$ is connected only if and only if it is strictly positive i.e. $\lambda_2(L(G(t))) > 0$ [20], [24], [25].

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:

$$\min_{\pi} \quad \sum_{i=1}^{N} \mathbb{P}_i(\boldsymbol{\xi}_i, \boldsymbol{u}_i(\cdot)) \tag{2a}$$

$$\text{s.t.} \quad \forall\, t \in [t_0, T]$$

$$\dot{\boldsymbol{\xi}}(t) = v(\boldsymbol{\xi}(t), t) + \pi(\boldsymbol{\xi}(t)) \quad \text{global dynamics Eq. 1}$$

$$d(\boldsymbol{\xi}_i(t), \boldsymbol{\xi}_j(t)) > R_{coll} \quad (i,j) \in V \times V, i \neq j \tag{2b}$$

$$\lambda_2(L(G(\boldsymbol{\xi}(t), R_{com}))) > 0 \tag{2c}$$

The goal is to find the control policy π. The agents are coupled in only two constraints: the collision constraint (Eq. 2b) where $d(\boldsymbol{q}_i, \boldsymbol{q}_j)$ represents the distance between agent i and j and R_{coll} the minimum safe distance, and second in maintaining a graph where all agents are connected to each other based on the communication range R_{com} (Eq. 2c). Our insight is that in this settings 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 $(\pi_{perf})_i = \arg\min_{\pi_i} \mathbb{P}_i(\boldsymbol{\xi}_i, \boldsymbol{u}_i(\cdot))$ only considering its own dynamics $\boldsymbol{q}_i$ (Eq. 1). π_{perf} can be an arbitrary control policy from a fixed control signal to a feedback controller based on learning or dynamic programming (Sec. 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 are satisfied (Eq. 2b, 2c). Lastly, based on the control inputs $\boldsymbol{u}_{perf}$ and $\boldsymbol{u}_{safe}$ from the respective policies, the safe interaction controller decides the agents final control inputs $\boldsymbol{u} = \pi_{LISIC}(\boldsymbol{u}_{perf}, \boldsymbol{u}_{safe})$. To achieve good performance the safe interaction controller should not interfere too much with $\boldsymbol{u}_{perf}$ while still ensuring connectivity and avoiding collisions.

In this work we focus on designing π_{safe} and π_{LISIC} for an arbitrary π_{perf}. In Sec. IV we prove that our method guarantees constraint satisfaction under certain conditions on the maximum magnitude of the control $\mathbb{U}$ and the flow field velocities v across the agents. Additionally, we test how our method performs in realistic ocean currents where these conditions are not always met and quantify the constraint violations with the metrics we defined below.

D. Constraint Violation Metrics

As in some settings it is impossible to guarantee constraint satisfaction, we now define the metrics we use to evaluate how much the constraints are violated in our experiments in Sec. V. 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.

As we are interested in the network robustness against connectivity losses or link failures we additionally measure the minimum Fiedler value over time, the higher the more robust the communication network (Sec. II-B):

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

Lastly, it often matters for how long an agent is isolated from all other agents. Therefore, we introducing a new measure that we call Isolated Platform Metric (IPM)

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

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

In the realistic setting, we compare the constraint violation of controllers empirically over a large, representative set of missions $\mathbb{M}$ by evaluating the collision rate $\mathbb{E}_{\boldsymbol{q}(t_0), t_0 \sim \mathbb{M}}[\mathbb{I}_{coll}]$, the disconnection rate $\mathbb{E}_{\boldsymbol{q}(t_0), t_0 \sim \mathbb{M}}[\mathbb{I}_{disconn}]$, 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). We first introduce our flocking-inspired safety controller based on potential functions and then detail our design for π_{LISIC}. For ease of

understanding we assuming holonomic actuation $g(x,u,t) = u$ but note that the method can be generalized.

A. Flocking-Inspired Safety Controller

The sole objective of the safety controller is to ensure proper distances between the agents such that their communication graph is connected and they do not collide with each other. We are inspired by the reactive flocking approaches for achieving ideal inter-agent distances 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 just two agents i and j that are connected and are at an inter-agent distance $\left\|\boldsymbol{q}_{ij}\right\|_2 = \left\|\boldsymbol{q}_i - \boldsymbol{q}_j\right\|_2$. Consider the following bowl shaped potential function

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

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 $\nabla_{\boldsymbol{q}_i}\psi_{\text{connected}}(\left\|\boldsymbol{q}_{ij}\right\|_2)$ and for j $\nabla_{\boldsymbol{q}_j}\psi_{\text{connected}}(\left\|\boldsymbol{q}_{ji}\right\|_2 = -\nabla_{\boldsymbol{q}_i}\psi_{\text{connected}}(\left\|\boldsymbol{q}_{ij}\right\|_2)$. When those two agents are getting too close $\left\|\boldsymbol{q}_{ij}\right\|_2 \to 0$ the potential $\psi(\left\|\boldsymbol{q}_{ij}\right\|_2)$ go to infinite, 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 $\left\|\boldsymbol{q}_{ij}\right\|_2 \to R_{com}$ then $\psi(\left\|\boldsymbol{q}_{ij}\right\|_2) \to \inf$ 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 $\boldsymbol{q}_{ij}$.

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. Therefore, our final potential function is augmented with a second term accounting for agents outside their communication range R_{com} to encourage achieving connectivity between disconnected agents [26]. This results in our final potential function $\psi(z) : \mathbb{R}_{\geq 0} \to \mathbb{R}_{>0}$ that is also visualized in Fig. 2:

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

where σ_{ij} is an edge indicator similar to a_{ij} in Sec. II-B but with an hysteresis parameter ϵ defined below in 7. Hence, $\psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right)$ switches between two terms whether the pair of agents (i,j) are within communication range ($\sigma_{ij} = 1$) or disconnected ($\sigma_{ij} = 0$) [26]. The 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 [27].

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

where $\epsilon > 0$ is the switching threshold inducing a hysteresis in the process of adding new links to the flock.

Our approach yields a relatively low attraction force for agents far outside of their communication range. This is a design choice in the context of *underactuated* agents in dynamic oceanic environment, where remote flock members can experience strong divergent flows and direct connectivity may be infeasible or undesirable to achieve.

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

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

B. Low Interference Safe Interaction Controller

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

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

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

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

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

The function α can be thought as a monotonically increasing safety activation function taking values between $[0,1]$ depending on the (unbounded) magnitude of it's argument. From the definition of $\psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right)$ in 6, $\lim_{\left\|\boldsymbol{q}_{ij}\right\|_2 \to 0} \psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right) = \infty$ and $\lim_{\left\|\boldsymbol{q}_{ij}\right\|_2 \to R_{com}} \psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right) = \infty$. Hence, in critical situations $\left\|\sum_{j=1}^{N} \nabla_{\boldsymbol{q}_i}\psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right)\right\|_2$ gets very large so that $c_i^{(1)}$ saturates to 1 and $c_i^{(2)}$ to 0, thus prioritizing the network safety for the concerned agents $i \in \mathcal{V}$ i.e. $\boldsymbol{u}_i \to \boldsymbol{u}_i^{safe}$, over each agent individual objective $\boldsymbol{u}_i^{perf}$.

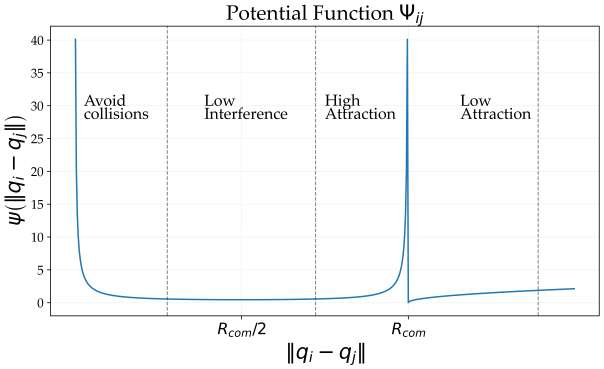

Fig. 2: Augmented potential function, with two terms to account for agents within the communication range and agents outside the communication range R_{com}. A high κ parameter is shown to increase the steepness of the slope around $\frac{R_{com}}{2}$, depending how achieving the exact ideal distance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

In other words, $\psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right)$ has a contractivity property for agent inter-distances at the boundary of the safe set, defined by 0 and R_{com}, similarly to Control Barrier Functions (CBFs) [28]. With this design, we ensure that agents coming from a disconnected status $\sigma_{ij} = 0$ to a connected status $\sigma_{ij} = 1$, experience a strong attracting gradient $\boldsymbol{u}_i^{safe}$ 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 $\left\|\sum_{j=1}^{N} \nabla_{\boldsymbol{q}_i} \psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right)\right\|_2$ is low, so that agent's i control input is dominated by the performance controller since $c_i^{(1)} \to 0$ and $c_i^{(2)} \to 1$.

IV. Theoretical Analysis

In this section, we propose to analyze under which conditions our safe interaction controller is able to maintain connectivity and avoid collisions [29]. Energy-based analysis and LaSalle's invariance principle are commonly utilized to obtain analytical guarantees for flocking. These methods establish the stability of the system and demonstrate that the flock converges to a lattice structure while preventing inter-agent collisions, as demonstrated in [12]. The structural collective dynamics can be derived by using a moving referential [12] with respect to the flock centroid $\boldsymbol{q}_c$. The relative coordinates are given by $\tilde{\boldsymbol{q}}_i = \boldsymbol{q}_i - \boldsymbol{q}_c$ and $\tilde{\boldsymbol{q}}_{ij} = \tilde{\boldsymbol{q}}_i - \tilde{\boldsymbol{q}}_j = \boldsymbol{q}_{ij}$. Therefore, $\psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right) = \psi\left(\left\|\tilde{\boldsymbol{q}}_{ij}\right\|_2\right)$, and the *total tension energy* or potential energy for the structural dynamics in the relative coordinates yields

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

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 $l = 0, 1, 2 \ldots$ and $\dot{H} \leq 0$ on each $[t_l, t_{l+1})$. Then, at switching time k, $H(t_k) = H(t_k^-) + \psi\left(\|R_{com} - \epsilon\|\right)$ [26]. As the graph topology becomes fixed after a certain time and only a finite number of maximum edges can be added, the energy can be shown to be bounded for any subsequent time. The time-derivative of H_i along the trajectory of agent i yields

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

where we exploited the relation $\nabla_{\boldsymbol{q}_i} \psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right) = -\nabla_{\boldsymbol{q}_j} \psi\left(\left\|\boldsymbol{q}_{ij}\right\|_2\right)$. Substituting $\dot{\tilde{\boldsymbol{q}}}_i = \dot{\boldsymbol{q}}_i - \dot{\boldsymbol{q}}_c$ in 11 and by applying the Cauchy–Schwarz inequality after some iterations, we obtain that $\dot{H}_i \leq 0$ holds if:

$$\left\|c_i^{(2)} \boldsymbol{u}_i^{perf}(\boldsymbol{q}_i) + v\left(\boldsymbol{q}_i\right) - \mathrm{Ave}\left(\dot{\boldsymbol{q}}_{Ni}\right)\right\|_2 \leq c_i^{(1)} U_{max,i} \tag{12}$$

where $\mathrm{Ave}\left(\cdot\right)$ denotes the average and the set $N_i = \mathcal{V} \setminus \{i\}$ the neighboring agents of i. The dynamics of the other agents are also defined by their surrounding flow and their individual control inputs such that $\mathrm{Ave}\left(\dot{\boldsymbol{q}}_{Ni}\right) = \mathrm{Ave}\left(v(\boldsymbol{q}_{Ni})\right) + \mathrm{Ave}\left(\boldsymbol{u}_{Ni}\right)$.

The agents do not necessarily need to be overactuated despite strong flows to achieve a local energy decrease ($\dot{H}_i \leq 0$). If the currents experienced by agent i are of similar magnitude and direction than the average current experienced by the neighboring agents, then $v\left(\boldsymbol{q}_i\right)$ compensates $\mathrm{Ave}\left(v(\boldsymbol{q}_{Ni})\right)$ and 12 can be fulfilled even if $\|v\left(\boldsymbol{q}_i\right)\|_2 > U_{max,i}$. The neighboring flocking control inputs $\mathrm{Ave}\left(\boldsymbol{u}_{Ni}\right)$ also helps accounting for the current difference term $v\left(\boldsymbol{q}_i\right) - \mathrm{Ave}\left(v(\boldsymbol{q}_{Ni})\right)$. If $\|v\left(\boldsymbol{q}_i\right) - \mathrm{Ave}\, v(\boldsymbol{q}_{Ni})\|_2 \gg U_{max,i}$, satisfying 12 becomes challenging, which can happen if agents experience strong divergent currents. Under these assumptions, we can show that $\dot{H} \leq 0$, which allows to bound the maximum energy and apply LaSalle's Invariance Principle [26], [17], thus ensuring that no collisions or disconnections happen, since $\psi(\left\|\boldsymbol{q}_{ij}\right\|_2) \to \infty$ when $\left\|\boldsymbol{q}_{ij}\right\|_2 \to 0$ or $\left\|\boldsymbol{q}_{ij}\right\|_2 \to R_{com}$. Condition 12 is sufficient but not necessary to guarantee $\dot{H} < 0$, as negative local energies can compensate positive ones.

V. Simulation Study

In the following section, the proposed flocking control scheme is evaluated on realistic ocean currents. We use Multi-Time Hamilton Jacobi Reachability (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 holonomic actuation of fixed thrust magnitude $\|u\|_2 = 0.1$ m/s. The control input in this context is the thrust angle θ. We consider a group of identical $n = 30$ ASVs with omnidirectional communication capabilities, navigating in strong ocean currents $v(\boldsymbol{q}, t) \in [0.3m/s, 2m/s]$, where each agent aims to reach a common pre-defined target, so that the group objective can be considered as the flock center reaching the target. In the following, we describe the creation

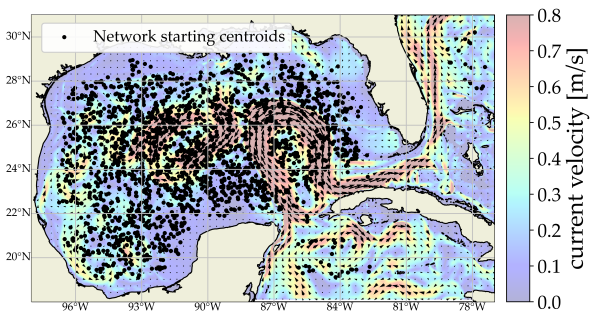

Fig. 3: We sample a large set of missions $|\mathbb{M}| = 1000$ in the Gulf of Mexico that are spatially and temporally representative of realistic scenarios.

of simulation experiments in a realistic, uncertain ocean environment and how we obtain a large set of missions to best illustrate trade-offs between single agent performance and flock connectivity maintenance.

a) Realistic Simulation of Ocean Conditions: Inspired by [4], we focus on the Gulf of Mexico region (Fig. 3), as it presents interesting and challenging currents. Moreover, we employ two ocean current data sources, that we refer to as HYCOM *hindcast* [30] and Copernicus hindcast [31] that we use as *forecast* 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. The ocean current data and the forecast error introduced are particularly relevant to the multi-time Hamilton-Jacobi (HJ) reachability controller, as it uses the currents to plan on, which impacts the performance of the time-optimal nominal controller. We propose two settings to investigate our approach, namely (a) performance HJR planning on *hindcast* and multi-agent simulation on *hindcasts* (HC-HC) and (b) performance HJR planning on *forecast* and multi-agent simulation on *hindcasts* (FC-HC). The first allows us to assess performance in an idealized setting where true flows are known whilst the second reflects a realistic application in dynamic ocean environments.

b) Large Representative Set of Missions: We assume that all agents start a navigation mission to a target region $\mathcal{T}$ at the same time t_0. The navigation objective is to drive the ASVs from their start states $(\boldsymbol{q}_1(t_0) \dots \boldsymbol{q}_n(t_0))$ to $\mathcal{T}$ within a maximum allowed time $T_{timeout}$. The target $\mathcal{T}$ is defined as a circular region with center coordinates $\boldsymbol{q}_{\mathcal{T}}$ 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 they can reach the target in $[72, 144]h$ 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 a total of $|\mathbb{M}| = 1000$ missions of initially connected and collision-free networks, see Fig. 3.

B. Baseline controllers

We build on recent work that proposed a reliable multi-time HJR controller for underactuated agents utilizing complex flows [4]. This approach directly extends to multiple agents with little extra compute, and the feedback controller for agent i can be obtained from an optimal value function $\mathcal{J}^*$ at time t as $\boldsymbol{u}_i(t)^* = \arg\min_{u_i \in \mathbb{U}} g(\boldsymbol{q}_i, \boldsymbol{u}_i, t) \cdot \nabla_{\boldsymbol{q}_i} \mathcal{J}^*(\boldsymbol{q}_i, t)$. All evaluated controllers use the multi-time HJR formulation as a single agent performance control. Our baseline scheme, called HJR-Baseline, involves each agent only utilizing its time-optimal performance control HJR without considering multi-agent interactions. Thus our baseline also provides a good likelihood estimation of collisions and communication losses if each agent were to rely solely on it's performance control. In addition, we define a second baseline controller from [32], denoted as Hamilton Jacobi Reachability Reactive Control (HJR-Reactive). 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 HJR performance controller. This approach is easily integrated with the time-optimal control HJR, and the reactive control term can be implemented in a decentralized manner.

Finally, we define our LISIC as Hamilton Jacobi Reachability Flocking Control (HJR-Flocking) with the safe interaction controller 8, where we also implement HJR as the single agent performance controller $\boldsymbol{u}_i^{perf}$. 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 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 9 as a SOFTMAX-like function

$$c_1^{(i)} = \frac{e^{\left\|\sum_{j=1}^{n} \nabla_{\boldsymbol{q}_i} \psi(\|\boldsymbol{q}_{ij}\|_2)\right\|_2}}{e^{\left\|\sum_{j=1}^{n} \nabla_{\boldsymbol{q}_i} \psi(\|\boldsymbol{q}_{ij}\|_2)\right\|_2} + e^{\rho}}, \ \forall i \in \mathcal{V}. \tag{13}$$

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

C. Additional Evaluation Metrics

The upper connectivity bound R_{com} in 2c and 7 is set to 9km, which corresponds approximately to radio communication capabilities for ASV and we set the collision lower threshold from 2b to $R_{coll} = 100$m, such that the ASVs would still have some margin in real conditions. Moreover, we also set the edge hysteresis parameter from 7 to $\epsilon =$ 300m. We use the euclidean norm to measure inter-agent distances $d(\boldsymbol{q}_i, \boldsymbol{q}_j)$ and the minimum flock center distance to target $d_{min}(\mathcal{T})$.

D. Numerical results

The results over a-priori known true currents (HC-HC) and realistic scenario (FC-HC) are presented in Table I.

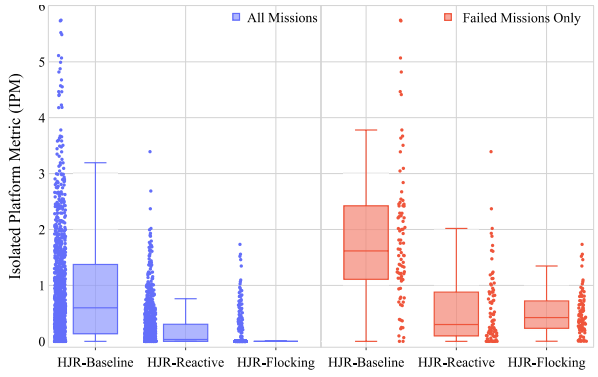

Fig. 4: IPM. Left evaluated on the set of all missions $\mathbb{M}$, right on the missions where HJR-Flocking failed only. Due to its low IPM, HJR-Flocking typically has both a low disconnection time and a low number of disconnected agents.

Both HJR-Flocking and HJR-Reactive exhibit superior performance in terms of connectivity and collision metrics compared to the baseline HJR. Thus we conduct statistical testing to compare HJR-Reactive and HJR-Flocking. Regarding the disconnection and collision rate, we perform a one-sided two-sample z proportion test for HJR-Flocking against HJR-Reactive.

Let Γ be the rate collision or disconnection over $\mathbb{M}$ with the null hypothesis be $H_0 : \Gamma_{\text{HJR-Flocking}} = \Gamma_{\text{HJR-Reactive}}$ to reject in favor of the alternative hypothesis $H_A : \Gamma_{\text{HJR-Flocking}} < \Gamma_{\text{HJR-Reactive}}$. HJR-Flocking is statistically significantly better than HJR-Reactive 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 HJR-Reactive at avoiding collisions. To compare the means over $|\mathbb{M}|$ of $\mu(\text{IPM})$ and $\mu(\lambda_2^{min})$ for connectivity and $d_{min}(\mathcal{T})$ for the performance trade-off, we perform a Welch's t-test due to the unequal variances of HJR-Reactive and HJR-Flocking. HJR-Flocking leads to statistically significantly better results for the network connectivity with $p < 1e^{-30}$ for $\mu(\text{IPM})$ and $\mu(\lambda_2^{min})$ for both (HC-HC) and (FC-HC) scenarios. Because of its higher value of $\mu(\lambda_2^{min})$, HJR-Flocking is more robust against disconnections, see Fig. 5 and should be the preferred control choice for communication maintenance

Moreover, we plot the IPM for the three controllers in Fig 4 for two cases, (1) the IPM evaluated on the full set of missions $|\mathbb{M}|$ (2) on a subset of missions where flocking failed to maintain connectivity. Among the three evaluated controllers, HJR-Flocking has the lowest IPM. Considering the missions where HJR-Flocking failed to maintain connectivity, it still achieves a smaller time of disconnection or fewer disconnected agents than the HJR-Baseline, but it is not as distinguishable from the HJR-Reactive controller. Interestingly, HJR-Reactive yields a statistically significantly better outcome for the objective trade-off $\mu(d_{min}(\mathcal{T}))$ with p-values $p < 1e^{-40}$ in both (HC-HC) and (FC-HC). Finally, Fig. 6 illustrates a navigation mission, comparing a naive multi-agent approach (HJR-Baseline) to our safe interaction controller, HJR-Flocking.

	Coll.	Disconn.	μ(IPM) ↓	$\mu(\lambda_2^{min})$ ↑	$\mu(d_{min}(\mathcal{T}))$ ↓
π_{perf} plans on true flows					
HJR-Baseline	68.5%	50.1%	0.37	0.39	**0** km
HJR-Reactive	**0**%	44.8%	0.19	0.42	0.14 km*
HJR-Flocking	0.7%	**9.9**%*	**0.05***	**1.15***	5.90 km
π_{perf} plans on forecast					
HJR-Baseline	39.1 %	70.5%	0.92	0.23	**10.55** km
HJR-Reactive	**0**%	58%	0.23	0.30	10.84 km*
HJR-Flocking	0.7%	**9.9**%*	**0.043***	**1.15***	13.96 km

TABLE I: We compare the performance of multiple controllers in two forecast settings. The * marks a statistically significant better performance of either HJR-Reactive or HJR-Flocking for the connectivity metrics $\mu(\text{IPM})$, $\mu(\lambda_2^{min})$ and the performance metric $d_{min}(\mathcal{T})$. Regarding the collision rate (Coll.) and disconnection rate (Disconn.), * denotes a statistically significant conclusion that HJR-Flocking leads to a lower collision or disconnection rate respectively.

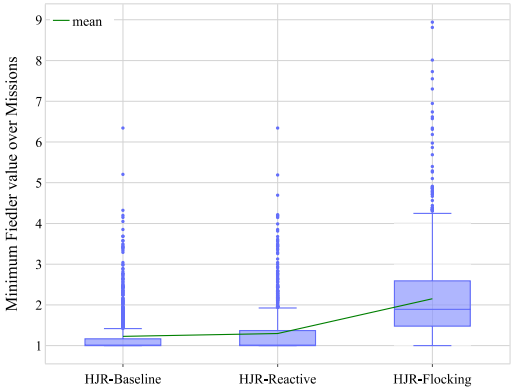

Fig. 5: The minimum Fiedler value λ_2^{min} can be used as a graph connectivity measure. HJR-Flocking has the highest minimum Fiedler value, which ensures better robustness against connectivity failures.

E. Discussion

It is clear that HJR-Flocking outperforms HJR-Reactive and the HJR-Baseline in terms of connectivity metrics. Interestingly, HJR-Flocking leads to a slightly higher collision rate in Table I than HJR-Reactive. We believe that it is mainly due to two reasons: (1) In HJR-Reactive the expected risk of collisions is inherently lower as each agent can achieve connectivity with a maximum amount of two other agents while HJR-Flocking 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 terms of collision rate for application to autonomous ASVs, where each agent maximizes an objective along it's trajectory [5]. The discrepancy between the performance trade-off with each agent target reaching objective $d_{min}(\mathcal{T})$ in Table I is less noticeable in the (FC-HC) setting, since the 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

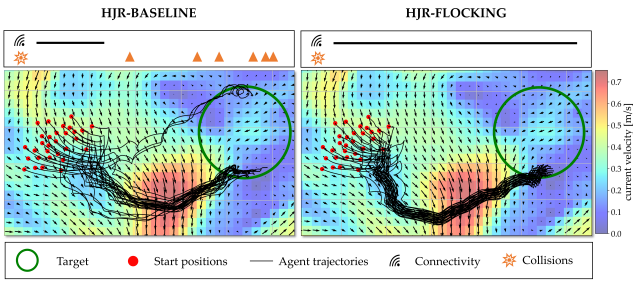

Fig. 6: Comparison of the HJR-Baseline with the low interference safe interaction controller HJR-Flocking. HJR-Flocking (right) guarantees communication through the full length of the mission, avoids collisions and ensures that all agent reach the target.

method blends a network safety controller for collisions and connectivity maintenance with a performance control policy, which allows us to decompose a complex multi-agent problem effectively. Our Low Interference Safe Interaction Controller prioritizes a safe control input from a flocking-inspired potential function in critical scenarios. We showed that connectivity can be maintained and collision avoided in underactuated agents, as long as the flow dynamics divergence between neighboring agent can be compensated. Our empirical results in realistic ocean dynamics showed that our method efficiently maintains connectivity and avoids collisions in most of the scenarios, while reasonably trading off with each agent's performance objective. Future work includes leveraging the agent's dynamics with forecast flows to predict future disconnections or collisions using predictive methods [20]. We anticipate that these methods will perform well on the true flow scenario (HC-HC) but may exhibit a performance drop when stochastic error is present as in the (FC-HC) scenario. It will be interesting to evaluate whether the additional computational cost of predictive methods pays off.

References

[1] P. Sujit *et al.*, "Uav and auvs coordination for ocean exploration," in *Oceans 2009-Europe*. IEEE, 2009, pp. 1–7.

[2] "Project loon," X, the moonshot factory. [Online]. Available: https://x.company/projects/loon/

[3] "Floating seaweed farms," Phykos PBC. [Online]. Available: https://www.phykos.co

[4] M. Wiggert *et al.*, "Navigating underactuated agents by hitchhiking forecast flows," in *2022 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 preprint 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 *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2443–2450.

[8] S. Pateria *et al.*, "Hierarchical Reinforcement Learning: A Comprehensive Survey," vol. 54, no. 5, pp. 109:1–109:35. [Online]. Available: https://dl.acm.org/doi/10.1145/3453160

[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] F. Chen and W. Ren, "On the Control of Multi-Agent Systems: A Survey," vol. 6, no. 4, pp. 339–499. [Online]. Available: http://www.nowpublishers.com/article/Details/SYS-019

[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] 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

[14] 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

[15] 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.

[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] 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

[18] 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

[19] 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

[20] A. Carron *et al.*, "Multi-agent distributed model predictive control with connectivity constraint," *arXiv preprint arXiv:2303.06957*, 2023.

[21] K. P. Wabersich and M. N. Zeilinger, "A predictive safety filter for learning-based control of constrained nonlinear dynamical systems," vol. 129, p. 109597. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0005109821001175

[22] 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.

[23] 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.

[24] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*, ser. Princeton Series in Applied Mathematics. Princeton University Press.

[25] 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/

[26] 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

[27] M. Ji and M. Egerstedt, "Distributed Coordination Control of Multiagent Systems While Preserving Connectedness," vol. 23, no. 4, pp. 693–703.

[28] A. D. Ames *et al.*, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference (ECC)*, 2019, pp. 3420–3431.

[29] 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

[30] 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.

[31] 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

[32] G. A. Pereira *et al.*, "Decentralized motion planning for multiple robots subject to sensing and communication constraints," *Departmental Papers (MEAM)*, p. 45, 2003.