

Higher Order Reasoning for Collaborative Communicationless Mobile Robot Operations

Jonathan Reasoner and Nicola Bezzo

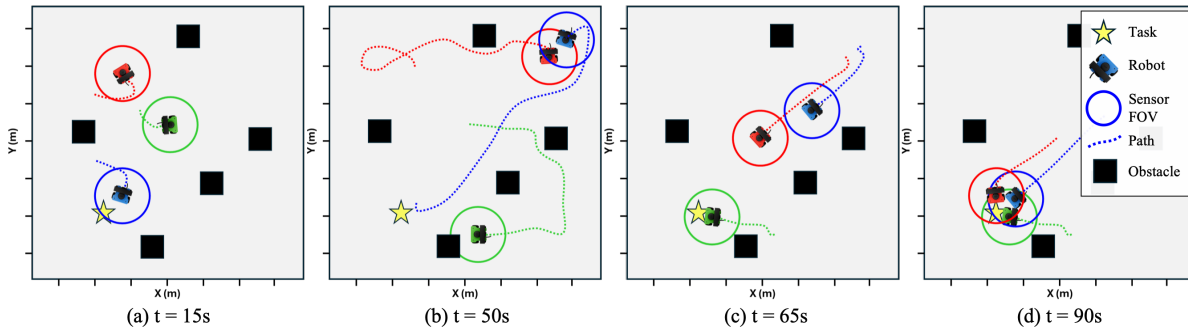


Fig. 1. Illustration of the proposed higher order reasoning for a 3 robot operation in a communicationless environment. *Blue* finds a task (a) and decides to go fetch *red* (b) because it reasons that *red* will take longer than *green* to find the task with its current plan. Meanwhile, *green* finds the task (c) and reasons that *blue* likely found it earlier and went to fetch *red*, thus triggering it to remain at the task to wait for the rest of the team (d).

Abstract—In communicationless environments, multi-robot systems must operate without the constant information exchange that many coordination strategies typically assume. This paper presents a novel dynamic epistemic planning framework that enables implicit coordination and long horizon planning through higher-order reasoning among robots. With our approach, robots form and propagate higher-order belief particles, update world beliefs using Bayesian inference, and select actions via a behavior tree that anticipates teammates’ likely decisions. A temporally aware Model Predictive Path Integral (MPPI) controller integrates this reasoning into low-level execution, allowing robots to plan intercepts and adapt trajectories under partial observability. The proposed framework is evaluated in both simulations and physical experiments, where it consistently reduces task completion time compared to a first-order baseline, demonstrating that epistemic logic can serve as a robust foundation for resilient coordination in communication-restricted domains.

I. INTRODUCTION

Multi-robot systems (MRS) provide advantages over single-robot systems for many task-oriented applications such as search and rescue, disaster relief, surveillance, and reconnaissance [1]. In such operations, task objectives are often unknown *a priori* and may require multiple robots to complete; therefore, deploying an MRS to localize these tasks is beneficial. However, MRS effectiveness often relies on unrealistic assumptions of constant communication [2] or carefully planned rendezvous strategies [3]. With these considerations in mind, this paper focuses on answering the following question: *How can we effectively coordinate an MRS to perform collaborative operations without explicit communication?* While this is a very strict constraint, we note that communication is often the bottleneck in MRS coordination, presenting delays and data loss. Additionally, communication may be not possible due to range constraints,

undesired as in military scenarios where detection is unwanted, or denied due to jamming or bandwidth congestion.

To this end, we propose a Theory of Mind (ToM)-based epistemic planning method that empowers each robot in the MRS with the ability to perform higher-order reasoning and predict and plan actions using the “I know that you know that I know” empathetic paradigm found in social/cognitive science. This approach is inspired by how humans form beliefs, reason about higher-order knowledge, and empathize with others’ perspectives to reach a consensus, enabling collaboration. ToM hypothesizes that through the combination of such belief and empathy modeling, humans can predict the behavior of others, adapt their own plans accordingly, and coordinate effectively, even in the absence of communication.

In this work, we extend the concept of ToM to autonomous systems, enabling robots to reason over both what they know and what they believe other robots may know. We will show that coordination strategies based on higher-order reasoning provide a robust basis for persistent MRS operations.

Fig. 1 illustrates the problem and outcome from our proposed ToM-based epistemic planning method. In this example, taken from the simulation, three robots are deployed to find and perform a task that requires the entire team to complete. The task location is unknown *a priori*. Using the proposed approach, the three robots begin by exploring the environment without communication using a frontier method. The *blue* robot finds the task, observes that it requires the rest of the team and evaluates what behavior to perform: 1) waiting to perform the task, 2) continuing exploring, or 3) seeking help. By predicting the behavior and future states of the other teammates, *blue* decides to proceed with 3) and fetch *red* who will take longer than *green* to find the task. *Blue* reasons that it can shorten the time for *red* to reach the task and complete the entire mission by selecting this behavior. While *blue* is fetching *red*, *green* discovers the

Jonathan Reasoner and Nicola Bezzo are with the Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA 22904, USA. Email: {vqh7rx, nb6be}@virginia.edu

task and reasons about what may have happened in the past and what the likely behaviors of the other robots are. *Green* “rewinds” time and predicts that *blue* may have found the task and it may have gone to fetch *red*, thus it chooses to wait at the task where it expects to see *blue* and *red* approaching in a near time horizon. *Blue* can similarly predict the expected time that it takes for *green* to reach the task and *green’s* behavior after task discovery. The same applies to *red* when it reaches the task. The example in Fig. 1 highlights the core higher-order reasoning principles behind our approach that allow robots to: ‘think’ about what has happened and what will happen over a horizon, evaluate different actions and the possible outcomes, and select the best behavior while considering the entire team’s capabilities.

We achieve this implicit higher-order reasoning and coordination through the use of: 1) Dynamic Epistemic Logic (DEL) which creates and propagates belief and empathy particles, allowing a robot to model the evolution of states and behaviors over a long horizon, 2) Bayesian belief updates, which allows the fusing of local observations with previously held epistemic state beliefs, and 3) a behavior tree search process which uses previously formed belief and empathy particles to determine the optimal behavior from a set of possible behavior primitives. Additionally, we propose: 4) a modified MPPI controller that produces stochastically optimal control inputs under a non-linear, time-varying objective function to fetch other robots, and 5) a belief-informed frontier exploration process which is used for efficient exploration under communication constraints.

Our contribution is two-fold: 1) we introduce a long-horizon estimation–planning–coordination framework for non-communicative multi-robot teams that couples DEL-based higher-order reasoning with belief-informed frontier exploration and an MPPI to achieve path planning in dynamic environments with non-linear cost functions; 2) we show that leveraging higher-order reasoning enables behavior-tree–driven long-horizon plans that adapt to unexpected situations and accelerate task completion across both simulations and experiments.

The rest of the paper is organized as follows. Sec. II reviews the state of the art in MRS coordination; Sec. III formally defines the problem; Sec. IV presents the MRS estimation–planning–coordination framework for communicationless systems; Secs. V and VI present simulation and experimental results; and Sec. VII provides conclusions and outlines future work.

II. RELATED WORK

Interest in multi-robot exploration and task allocation is rising, driven by affordable platforms and goals of reducing human involvement in hazardous or menial tasks [4]. Previous efforts have attempted to solve the coordination challenge associated with MRS through market-based approaches that leverage an auctioneer to determine the optimum agent to explore an area [5] or through decision-theoretic approaches [6], but these methods require assured communication between robots and the central controller. To

overcome the shortcomings of centralized systems, decentralized approaches have been increasingly proposed due to their flexibility and robustness. In [3], [6], global communication assumptions are relaxed and a decentralized approach is implemented that accounts for a limited communication range. While these approaches allow for more realistic environments, systems can still fail due to malfunction or intentional interference such as jamming or signal degradation [2]. Prior non-communicative work have leveraged biomimicry for coordination [7], using visual cues, predefined sensorimotor behaviors, and a priori search patterns, with simple evaluations in obstacle-free environments.

In order to address the challenges of coordinating communicationless MRS, we implement a ToM inspired approach. ToM focuses on how humans infer the beliefs, intentions, and desires of other humans and use that knowledge to improve our understanding of the world around us [8], [9]. To integrate Theory of Mind (ToM) into autonomous systems, epistemic planning has been proposed as a way to generate plans that consider what a robot knows and believes, as well as what it does not know [10], [11]. Dynamic epistemic logic (DEL) extends epistemic logic with a framework for information change, specifying how observations transform an agent’s belief state [12]. This approach enables higher-order reasoning: robots form beliefs about other robots, about what other robots believe about them, and about what those robots believe about third robots. Following [9], we refer to these as first-, second-, and third-order reasoning.

Through the use of DEL-based higher-order reasoning, this work advances the state of the art in decentralized MRS coordination for unknown, partially observable environments where explicit communication is not possible.

III. PROBLEM FORMULATION

Let us consider an MRS team, \mathcal{R} , of n_r robots operating in a partially known, communicationless environment $\mathcal{W} \subseteq \mathbb{R}^2$ in which the goal is to discover and complete a task g in the shortest amount of time.

Let $\mathcal{R}_g \subseteq \mathcal{R}$ denote the subset of robots required to complete g . For any robot $i \in \mathcal{R}_g$ we define \tilde{T}_i as the estimated task completion time, measured as the time lapsed from the start of the mission to the completion of g . With this, the total mission time \mathcal{T} can be written as:

$$\mathcal{T} = \mathbb{E} \left[\max_{i \in \mathcal{R}_g} \tilde{T}_i \right]. \quad (1)$$

We can then formally define our problem as:

Problem 1: Communicationless Collaborative Task Discovery and Completion. Design a strategy to dynamically select the best motion planning policy $\pi_i \in \Pi$ for each robot $i \in \mathcal{R}$ such that g is discovered and completed in the shortest amount of time possible:

$$\pi_i = \arg \min_{\pi \in \Pi} \mathcal{T} \quad \forall i \in \mathcal{R}. \quad (2)$$

As we will show next, by leveraging higher-order reasoning and planning - which for brevity we call epistemic planning - paired with future and backward horizon predictions,

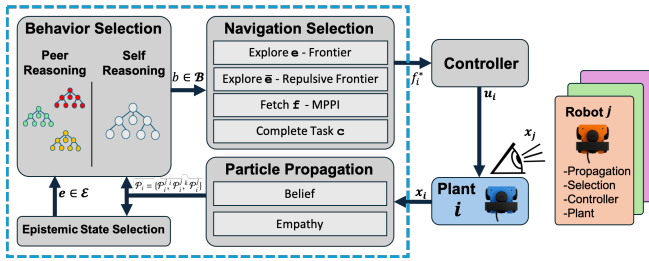


Fig. 2. Diagram of our proposed approach. Our contribution is shown in the blue dashed box.

each robot is able to minimize this global cost by selecting behaviors that others can predict in the absence of explicit direct communication.

IV. APPROACH

Our approach is summarized in Fig. 2. Each robot begins by exploring the environment in search of tasks using a Frontier method [13]. Once a robot discovers a task, it performs an epistemic belief assessment, predicting the likely behaviors of other teammates by propagating forward and backward belief particles. These particles inform the behavior tree, which applies high-level reasoning to evaluate all actions and predict associated outcomes to best accomplish the task. The optimal behavior is then selected, and the process continues until all tasks are completed or the environment has been fully explored. We now present how each component of the approach is implemented, starting with a preliminaries section. We then discuss how belief and empathy particles are created and propagated to predict teammates' behaviors and select the optimal behavior to execute.

A. Preliminaries


1) *Notation*: Throughout the paper, predicted quantities are denoted by a tilde (\tilde{x}), unit vectors by a hat (\hat{v}), and vectors are boldface (\mathbf{x}). The state vector, $\mathbf{x}_i(t) = [x_i(t), y_i(t), \theta_i(t)]^T \in \mathbb{R}^2 \times \mathbb{S}$, consists of the (x, y) position and orientation, θ for robot i . Lastly, the unit bearing between two vectors is denoted as:

$$\hat{d}(\mathbf{x}_a, \mathbf{x}_b) = \frac{\mathbf{x}_b - \mathbf{x}_a}{\|\mathbf{x}_b - \mathbf{x}_a\|}. \quad (3)$$

2) *Dynamic Epistemic Logic (DEL)*: DEL is a reasoning framework that allows robots to translate observations into beliefs, considering uncertainty and incompleteness of information, and reason over those beliefs for planning in systems with distributed knowledge and capabilities [10].

In order to use DEL in our approach, we construct 3 vectors which hold information that robot i has attained through observations and reasoning, specifically: the robot's pose \mathbf{x}_i , the epistemic state e_i , and the behavior state \mathbf{b}_i .

The epistemic state, $e_i(t) \in \mathbb{R}^{n_r}$ encodes team knowledge with entries $e_i^j(t) \in \{0, 1\}$, where $e_i^j(t) = 1$ if robot i believes j has discovered the task. For our work, we only consider environments with a single task, but our approach can be expanded to n_g tasks by expanding the epistemic state to $e_i(t) \in \mathbb{R}^{n_r \times n_g}$. The behavior state, $\mathbf{b}_i(t) = [\tilde{b}_i^0(t), \dots, \tilde{b}_i^{n_r}(t)]$, describes the predicted behavior of all robots within \mathcal{R} .



(a) Frontier selection process at t

(b) Frontiers at $t + 10$

Fig. 3. Example of the frontier selection process for exploration.

For collaborative multi-robot operations, most required behaviors can be decomposed into a set of behavior primitives. For this work, we define these primitives as i) *explore* (e): exploring an area in search of the task, ii) *complete task* (c): going to an identified task or waiting for other robots to arrive at the task, iii) *fetch* (\mathcal{F}_j): deliberately creating an interaction with another robot, j with the intention of driving it towards the task, or iv) *modified explore* (\bar{e}): a temporarily modified search pattern resulting from an unexpected interaction with another robot. This modified search pattern drives a robot away from a specific robot before resuming normal frontier exploration, e, behavior. With this minimal set of behavior primitives, each robot can select a policy according to (2) for the class of collaborative tasks addressed herein.

B. Belief and Empathy Formulation and Propagation

Each robot i starts the operation in exploration behavior $b_i = e$ and with an epistemic state $e_i = 0$ until a task is discovered in which case $e_i = 1$. Within our epistemic framework, each robot propagates belief states \mathbf{x} for all robots over a predictive time horizon H . In this way, each robot i can plan according to its belief about other robots and empathize with where other robots expect i to be located and what behavior they expect i to be doing, even without communicating. We define a set of particles $\mathcal{P}_i = \{\mathcal{P}_i^j, {}^i\mathcal{P}_i^j, {}^k\mathcal{P}_i^j\}$ to represent these belief and empathy states from the perspective of i , where:

- $\mathcal{P}_i^j(t) = \{\tilde{\mathbf{x}}_i^j(t) | \tilde{b}_i^j\}_t^H$ is the first order belief indicating i 's estimated state of j given predicted behavior of j over a time horizon H ;
- ${}^i\mathcal{P}_i^j(t) = \{\tilde{\mathbf{x}}_i^j(t) | {}^i\tilde{b}_i^j\}_t^H$ is the second-order empathy that indicates what i believes that j believes about i ; and
- ${}^k\mathcal{P}_i^j(t) = \{\tilde{\mathbf{x}}_i^j(t) | {}^k\tilde{b}_i^j\}_t^H$ represents what i believes that j believes about k .

These particles are used to inform robot i 's frontier exploration strategy and higher-order reasoning.

Exploration is achieved via a frontier-based approach leveraging first-order particle propagation. Traditional techniques rely on communication to decide what actions to perform. In contrast, here each robot is predicting the expected state of the other teammates to decide where to go next.

To start, an i^{th} robot will update its and the expected map of every other robot j at the state described by the j^{th} belief particle. Using this belief map, the i^{th} robot can evaluate its frontier set \mathcal{F}_i defined as the set of all cells that are adjacent

to unknown cells (i.e., outside of i^{th} field of view and not visited by i yet) and pick a frontier f_i^* to aim for, chosen among all $f_i \in \mathcal{F}_i$. This is illustrated in Fig. 3. To enforce a uniform exploration, in this work we decided to leverage a homogeneous frontier selection to track the f_i^* that is in the direction most opposite to the other robots positions:

$$f_{i,e}^* = \arg \min_{f_i \in \mathcal{F}_i} \frac{(\mathbf{x}(f_i) - \mathbf{x}_i)^\top}{\|\mathbf{x}(f_i) - \mathbf{x}_i\|} \hat{\mathbf{d}}_i, \quad (4)$$

$$\mathbf{d}_i = \frac{1}{n_r - 1} \sum_{\substack{j=1 \\ j \neq i}}^{n_r} (\mathbf{x}_j - \mathbf{x}_i) \implies \hat{\mathbf{d}}_i = \frac{\mathbf{d}_i}{\|\mathbf{d}_i\|}, \quad (5)$$

where \mathbf{d}_i is the mean displacement vector from robot i to the estimated positions of all other robots, $\|\mathbf{d}_i\|$ is the length of \mathbf{d}_i , and $\hat{\mathbf{d}}_i$ is the normalized vector. The position of the frontier cell is $\mathbf{x}(f_i)$, $f_i \in \mathcal{F}_i$.

i propagates these particles over time by predicting the control input and the resulting list of frontiers that robot j would possess, \mathcal{F}_j if j continues to behave according to b_j^i . This evolves the state over time according to:

$$\tilde{\mathbf{x}}_i^j(t + \ell + 1) = h(\tilde{\mathbf{x}}_i^j(t + \ell), \tilde{\mathbf{u}}_j(t + \ell)), \quad \ell = 1, \dots, H, \quad (6)$$

where $h(\cdot)$ is assumed to be a known dynamical model function of robot j .

Empathy particles are formed by propagating the nested state beliefs over H using the previously described process and are used by i to understand how predictable it is acting. For instance, if i must avoid an obstacle that is unknown to j , i will be forced to deviate from the location held within ${}^i\mathcal{P}_i^j(t)$. By creating ${}^i\mathcal{P}_i^j(t)$, robot i can control back to the location that i believes j believes i should be.

1) *Repulsive Frontier*: In addition to forming belief and empathy particles, a robot must recognize when its beliefs are incorrect and update its predictions. To do so, it compares its sensor readings with their predicted values to determine when an observation is unexpected. Such unexpected readings may occur naturally during exploration or be deliberately induced by peers. We formalize unexpectedness via a residual measure:

$$\lambda_{i,j}(t) = \|\tilde{\mathbf{x}}_i^j(t) - \mathbf{x}_i^j(t)\|, \quad (7)$$

where $\mathbf{x}_i^j(t)$ is the observed state of j . If $\lambda_{i,j}(t)$ exceeds a predefined threshold, i recalculates its frontier selection by selecting f_i^* as the frontier most opposite of \mathbf{x}_j using:

$$f_{i,\bar{e}}^* = \arg \min_{f_j \in \tilde{\mathcal{F}}_i^j} \hat{\mathbf{d}}(\mathbf{x}_i, \mathbf{x}(f_j))^\top \hat{\mathbf{d}}(\mathbf{x}_i, \mathbf{x}_j). \quad (8)$$

i then propagates belief and empathy particles assuming that $b_j = \bar{e}$ over H as in (6).

C. Bayesian Epistemic State Updating

During frontier exploration, a robot can experience a variety of events ζ that provide evidence about whether a robot has discovered a task. In collaborative robot operations, this list of events can be narrowed to the following: (i) *unexpected observation of another robot*, (ii) *prediction that another*

robot has discovered a task, (iii) *expected observation of another robot*, (iv) *failing to observe another robot when expected to*, (v) *successfully intercepting another robot*, (vi) *failing to intercept another robot*, and (vii) *task discovery*. For example, if the j^{th} robot unexpectedly detects another robot, i , that has deviated from its predicted path, $\mathbf{x}_i(t) \neq \tilde{\mathbf{x}}_i(t)$, j will reason that i probably possesses knowledge not held by j , increasing the likelihood that i knows the location of the task since j has not yet discovered it. In this section, we explain how a robot i updates its epistemic state hypothesis, e_i about every other robot, based on evidence collected during these events.

Each robot i maintains a distribution of probabilities P for every possible combination of epistemic states $e_i^j \in E$ (e.g., for a 3 robot case $E_i = \{e_i^1 = [0, 0, 0], e_i^2 = [0, 0, 1], \dots, e_i^3 = [1, 1, 1]\}$). Each probability is initialized to a user-defined value, setting a prior that gets updated as events occur to obtain a posterior via a Bayesian update process. In this work, an event occurrence signifies the discovery of a task by a robot k which will increase the probability of all the epistemic states that contain $e_k = 1$. For ease, probabilities are increased (or decreased) using a fixed increment (or decrement) as follows. Given $e_i = [e_1, e_2, \dots, e_k, \dots, e_{n_r}]$:

$$P(e_i^j) = \begin{cases} P(e_i^j) \cdot P^+ & \text{if } e_k = 1 \wedge \zeta_k = 1 \quad \forall e_i^j \in E \\ P(e_i^j) \cdot P^- & \text{if } e_k = 1 \wedge \zeta_k = 0 \quad \forall e_i^j \in E \end{cases} \quad (9)$$

where $1 > P^+ > P^- > 0$ are the fixed increment and decrement values respectively and ζ_k is an indicator binary variable set to 1 if an event is detected associated with k . Similar to occupancy grid mapping, this approach will reinforce probabilities related to epistemic states in which k is believed to have found the task.

Once an event has occurred and the probabilities have been set, we search among all epistemic state combinations for the most likely one. Here we compare the predicted map \mathcal{M}_j – associated with the states from the first order particle set \mathcal{P}_i^j for each robot j – with the current map \mathcal{M}_i discovered by i and task location (if known) to assess if any robot j has discovered a task. Formally, if i has discovered a task g :

$$e_j = \begin{cases} 1 & \text{if } e_i = 1 \wedge \mathcal{M}_j \cap \{\mathbf{x}_g\} \neq \emptyset \quad \forall j \in \mathcal{R} \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

We then use (9) to update the likelihood associated with each $e_i^j \in E$.

Finally, robot i updates its hypothesis of e_i by selecting the maximum-a-posteriori hypothesis from the set of all possible hypotheses, E :

$$e_i^* = \arg \max_{e_i \in E} P(e_i^j). \quad (11)$$

D. Behavior Selection via Depth-Limited Tree Search

Once an event is detected, the Bayesian update process outlined in the previous section leads to an updated e_i^* . In turn, this epistemic state is used to decide what behavior $b \in \mathcal{B}$ to perform to complete the task in the shortest amount of time, i.e., to satisfy (2). This multi-step process involves: i) updating the predicted teammate's behavior state vector b_j

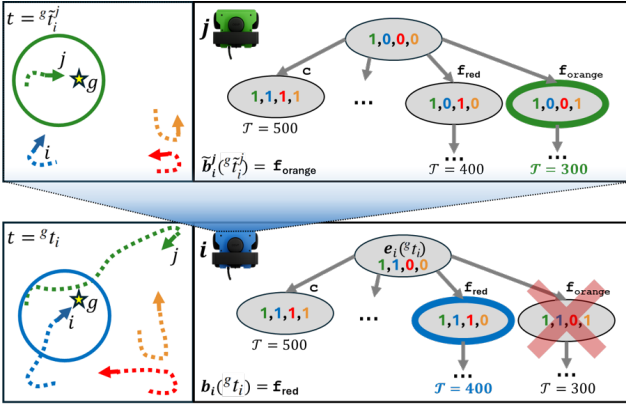


Fig. 4. Illustration of the tree search-based behavior selection method.

via higher-order reasoning informed by robot i 's belief and empathy particles, and ii) using b_i to inform selection of the $b_i \in \mathcal{B}$ that minimizes (1).

1) *Behavior State Formation via Higher-order Reasoning:* Behavior selection starts with estimating peers' behaviors. e_i^* indicates to robot i whether any $j \in \mathcal{R}$ is aware of the task, but does not encode information about the time in which j discovered it or j 's behavior afterward. Robot i discerns the time that j found the task by leveraging the historical pose estimates $\tilde{x}_j(t)$ encoded in $\{\mathcal{P}_i^j(\tau)\}_{\tau=0}^t$.

Specifically, the estimated discovery time, s_i^j , is the first time that i predicts x_g was within j 's sensing range, S_j :

$$s_i^j = \min \left\{ \tau \in [0, t] \mid \|\tilde{x}_j^j(\tau) - x_g\| \leq S_j \right\}. \quad (12)$$

If the i^{th} robot believes that j has discovered the task, i infers the behavior j would select after the discovery, \tilde{b}_i^j , by performing a depth-limited behavior tree search initialized at s_i^j . This behavior tree is constructed using state information contained within i 's second- and third-order beliefs to allow i to reason about how j believes the mission will progress based on any behavior j selected. Each branch of the behavior tree represents the resulting epistemic state and estimated total mission time, $\tilde{\mathcal{T}}_i^j(b)$ for all $b \in \mathcal{B}$.

Consider the example illustrated in Fig. 4 in which the blue robot discovers the task at s_i . Blue reasons that green has also seen the task. Using its second- and third-order empathy particles, blue determines that green would have likely selected the behavior, f_{orange} after task discovery. Blue then prunes this branch from its behavior tree and selects the remaining option that will minimize \mathcal{T} (i.e. f_{red}).

To estimate $\tilde{\mathcal{T}}_i^j(b)$, i must calculate when j believes that each $k \in \mathcal{R}$ will complete the task $\forall b \in \mathcal{B}$. We denote this higher-order belief as ${}^k\tilde{\mathcal{T}}_i^j$, read as j 's belief of the time required for k to complete the task as computed by i . In other words here i is rewinding time after discovering that j was at the same task to recreate j 's reasoning process given its knowledge at the time about all $k \in \mathcal{R}$.

For $b \in \{c, e, e^*\}$, ${}^k\tilde{\mathcal{T}}_i^j$ is computed by casting the problem as a forward reachability analysis with

$${}^k\tilde{\mathcal{T}}_i^j = \min\{t \mid (\mathcal{S}_k, t \in [s_i^j, s_i^j + H]) \cap \{x_g\} \neq \emptyset\} \quad (13)$$

where \mathcal{S} is the reachable set computed by integrating forward

the dynamics of k considering the sensing field of view area described by k while moving [14].

If $b = f$, ${}^k\tilde{\mathcal{T}}_i^j$ is estimated using a Dubins reachability analysis [15] to estimate the time required for j to reach the point that will lead to a successful fetch of k and return to the task (more details about fetching is provided in Sec. IV-E).

With ${}^k\tilde{\mathcal{T}}_i^j$, the maximum mission time is expressed as:

$$\tilde{\mathcal{T}}_i^j(b) = \left[\max_{k \in \mathcal{R}_g} {}^k\tilde{\mathcal{T}}_i^j \right], \quad (14)$$

and the behavior of j is obtained by:

$$\tilde{b}_i^j = \arg \min_{b \in \mathcal{B}} \tilde{\mathcal{T}}_i^j(b). \quad (15)$$

2) *Behavior Selection:* Robot i , that has discovered the task, now has a belief of who else discovered the task (e) and what behavior they are now performing (b). With this information, i can select its optimal behavior using the same behavior tree search process above. i can prune the behavior tree to remove behaviors that it believes have already been performed (i.e., not fetching robot k because i has predicted that j likely decided to fetch k when j previously found the task). Given this pruned tree, i can then select the best behavior which minimizes cost according to:

$$b_i = \arg \min_{b \in \mathcal{B}} \mathcal{T}_i. \quad (16)$$

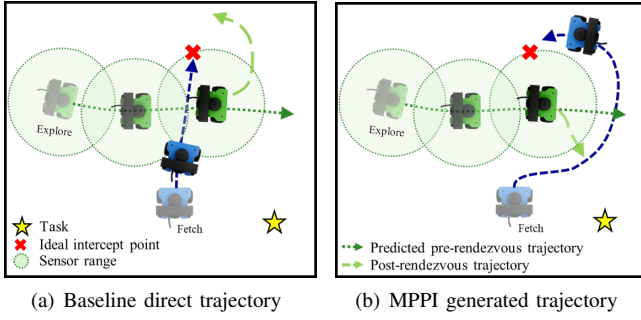
E. Interception via Temporally Aware MPPI

Upon discovering a task, if robot i selects *fetch* f , it uses $\mathcal{P}_i^j(t)$ to plan an intercept that redirects j 's frontier choice toward faster discovery of the task. This is possible because i knows that an unexpected appearance in j 's field of view raises the residual, prompting j to switch to the modified frontier approach \bar{e} in (8). Executing a fetch behavior f involves choosing where and when to meet j , generating a trajectory to reach the intercept point and assessing success by updating the predicted trajectory of j . We implement this with a modified MPPI controller, [16], which efficiently samples diverse candidates and returns a stochastically optimal intercept path. As shown in Fig. 5(a), a non-optimal intercept can be counter-productive to blue's goal of pushing green's search toward the task; with our MPPI (Fig. 5(b)), blue avoids unintended rendezvous and arrives at the right place and time. The next section details intercept-point selection, trajectory generation, and success evaluation.

1) *Finding the ideal intercept point:* To execute a fetch behavior, i first identifies x_{int}^* , the earliest reachable position that causes j to detect i and select repulsive frontiers that drive j toward the task. The set of candidate intercept points, $\mathcal{I} = \{x_{\text{int}}(\tau)\}_{\tau=\tau}^{\tau+H}$, is generated over a finite horizon H , and the earliest candidate that is deemed reachable by i via Dubins reachability analysis is selected. Each $x_{\text{int}}(\tau)$ is constructed as the point just inside j 's sensing range, S_j , on the side opposite from x_g at $t = \tau$:

$$x_{\text{int}}(\tau) = \tilde{x}_j(\tau) + (S_j - \sigma) \hat{d}(x_g, \tilde{x}_j(\tau)), \quad (17)$$

where σ is a small offset, placing i within j 's sensor range. We compute the time required for robot i to reach each



(a) Baseline direct trajectory (b) MPPI generated trajectory
 Fig. 5. Illustration of the fetching behavior: (a) a naive approach that leads to early rendezvous causes a counter-productive reaction by *green*; (b) with MPPI, rendezvous is avoided until the optimal time.

candidate intercept point $\mathbf{x}_{\text{int}}(\tau) \in \mathcal{I}$ using Dubins time [15], which combines the straight line travel time and heading alignment time to \mathbf{x}_{int} :

$$D(\mathbf{x}_{\text{int}}(\tau), \mathbf{x}_i(t)) = \frac{\|\mathbf{x}_{\text{int}}(\tau) - \mathbf{x}_i(t)\|}{v_{\max}} + \frac{1}{\omega_{\max}} \arccos(\hat{\boldsymbol{\theta}}_i(t) \cdot \hat{\mathbf{d}}(\mathbf{x}_i(t), \mathbf{x}_{\text{int}}(\tau))). \quad (18)$$

where $\hat{\boldsymbol{\theta}}_i(t) = [\cos(\theta_i(t)), \sin(\theta_i(t))]$ is the unit orientation vector, v_{\max} is robot i 's maximum linear velocity, and ω_{\max} is the maximum angular velocity.

The earliest feasible candidate is selected by minimizing the index τ subject to reachability:

$$\tau^* = \arg \min_{\tau \in \{t, \dots, t+H\}} (\tau - t) \quad \text{s.t.} \quad D \leq \alpha(\tau - t). \quad (19)$$

Because D assumes a direct route at maximum speed, we introduce a slack factor $\alpha \leq 1$ to account for longer, non-direct trajectories that may be generated by the MPPI.

The optimal intercept point is then set as $\mathbf{x}_{\text{int}}^* = \mathbf{x}_{\text{int}}(\tau^*)$, which we send to the MPPI for trajectory generation.

2) *MPPI Trajectory Selection*: Given $\mathbf{x}_{\text{int}}^*$, we use a modified MPPI to iteratively generate a stochastically optimal trajectory $\mathbf{u}^* = \{\mathbf{u}_1, \dots, \mathbf{u}_{H_u}\}$ over horizon H_u . Following [16], at each iteration we create K candidate control sequences $\{\mathbf{u}_k\}_{k=1}^K$, and use the robot's dynamic model to obtain candidate trajectories $\{\mathbf{r}_k(h)\}_{k=1}^K$ with $h = \{0, \dots, H_u\}$ being the discrete time index. We evaluate each candidate trajectory with a three-part cost function that penalizes misalignment between i and j , distance from $\mathbf{x}_{\text{int}}^*$, and safety:

$$J(\mathbf{r}_k) = \sum_{h=0}^{H_u} m(\mathbf{r}_k, h) \Phi(\mathbf{r}_k, h) + \|\mathbf{r}_k(h) - \mathbf{x}_{\text{int}}^*\|^2 + c_s(\mathbf{r}_k(h), \mathcal{M}). \quad (20)$$

The first term encourages an intercept geometry in which robot i positions itself on the side of j opposite the task \mathbf{x}_g (Fig. 5(b)). It activates only when i is within j 's sensor range S_j and then penalizes same-side alignment with respect to the task. Concretely, $m = \eta$ if $\|\tilde{\mathbf{x}}_j(h) - \mathbf{r}_k(h)\| \leq S_j$, otherwise 0, where η is a constant gain applied to the alignment cost. Φ is defined as:

$$\Phi(\mathbf{x}_k, h) = \begin{cases} (\Psi(\mathbf{r}_k, h) + 2)^2, & \text{if } \Psi(\mathbf{r}_k, h) > \gamma, \\ 0, & \text{else,} \end{cases} \quad (21)$$

where $\Psi(\cdot) = \hat{\mathbf{d}}(\tilde{\mathbf{x}}_j(h), \mathbf{x}_g)^\top \hat{\mathbf{d}}(\tilde{\mathbf{x}}_j(h), \mathbf{r}_k(h))$, measures the dot product between the vector from $\tilde{\mathbf{x}}_j(h)$ to \mathbf{x}_g and the vector from $\tilde{\mathbf{x}}_j(h)$ to $\mathbf{r}_k(h)$. γ is a user-defined threshold.

The second term in $J(\mathbf{r}_k)$ promotes liveness by penalizing deviation from the intercept point, while the third term penalizes unsafe proximity to known obstacles via $c_s(\mathbf{r}_k(h), \mathcal{M})$ when the separation between $\mathbf{r}_k(h)$ and any $\mathbf{x}_{\text{obs}} \in \mathcal{M}$ falls below a safety margin.

After evaluating J_k for all K candidate trajectories, the optimal sample index is selected as:

$$k^* = \arg \min_k J^{(k)}. \quad (22)$$

The corresponding control sequence is \mathbf{u}_{k^*} .

After executing the MPPI trajectory \mathbf{x}_i^* , robot i enters j 's field of view, prompting j to evaluate the residual (7); if $\lambda_{i,j}(t) > \delta$, j replans with the modified exploration policy $\bar{\mathbf{e}}$ and selects a new frontier. Anticipating this, i propagates an updated belief particle to predict j 's new trajectory with $b_j = \bar{\mathbf{e}}$. The intercept is declared successful if:

$${}^s t_i^j(t) \leq \kappa D(\tilde{\mathbf{x}}_j(t), \mathbf{x}_g), \quad (23)$$

where D is the Dubins time to task for j , as calculated in (18) and κ is a safety factor that allows indirect routes. If unsuccessful, i selects a new $\mathbf{x}_{\text{int}}^*$ and retries. If i determines the intercept was successful, i updates its epistemic state and selects a new behavior. If the intercept was unsuccessful, robot i repeats the MPPI process by selecting a new $\mathbf{x}_{\text{int}}^*$ and attempts another intercept.

V. SIMULATIONS

Our approach was evaluated in a simulated 150m² environment. Sensing was performed using a simple lidar sensor model with a range of 20m. Robots were modeled as differential-drive with nominal speed of 1.0m/s, and the ability to increase speed up to 1.25m/s when required for fetching or returning to a robot's second-order empathy particle. This formulation reflects a common practice in long-duration robotic operations, where a nominal endurance speed is maintained for efficiency, but higher speeds can be used for short bursts when needed [17]. Simulations were performed with teams of 2, 3, and 4 robots over 100 randomized trials per team size, sampling task locations and robots' initial poses uniformly within the environment. For each task, all robots were needed to complete the task, i.e., $\mathcal{R}_g = \mathcal{R}$ and task completion was defined as all team members reaching within 1.0m of the task. We considered both: (i) empty environments and (ii) cluttered environments containing five square obstacles at random locations.

As a comparison baseline, we implemented a first-order reasoning policy where each robot predicts peers' positions using only first-order beliefs biasing its exploration toward regions opposite to those peers, as explained in Sec. IV-B.

Across 600 trials (100 per setting), our approach matched or outperformed the baseline in 93.3% of cases, as shown in Fig. 6. A "match" occurs when a robot that finds the task predicts that all teammates will arrive at the task sooner than




Fig. 6. Outcome distribution across the 6 simulated scenarios. Our approach is deemed successful if the time to completion is the same (yellow) or improved (green) compared to the baseline.

a fetch maneuver could be accomplished. The time-optimal decision is therefore to abstain from fetching, as fetching would not reduce completion time. When our policy included a fetch behavior, it reduced completion time by an average of 169s (20.8%), shown in Fig. 7.

These improvements arise from the higher-order reasoning capability of our framework. An illustrative example is shown in Fig. 8. In this scenario, the *green* robot detects the task in Fig. 8(a), recognizes that it is the only robot aware of the task, and identifies that fetching *red* yields the greatest time reduction. As shown in 8(b), *green* uses the MPPI, to position itself to intercept *red* and, upon confirming a successful interception, *green* repeats this process. The next most beneficial action is to fetch *orange*. While *green* fetches *orange*, *red* discovers the task, as shown in Fig. 8(c). Leveraging its second- and third-order beliefs about *green*, *red* predicts that *green* would have fetched both *red* and *orange* and reasons that fetching *blue* is the most beneficial behavior remaining. *Red* completes the intercept in Fig. 8(d) while *green* completes its intercept in Fig. 8(e). Through this sequence of higher-order-based reasoning, all robots converge on the task in 529s: 333s faster than the baseline, shown in the supplementary material.

Figure 8(b) illustrates additional behaviors arising from higher-order reasoning. When the *orange* robot encounters an obstacle, it adjusts its trajectory while maintaining a second-order belief of its perceived location by teammates, allowing it to return gradually to the position consistent with its second-order empathy particle. This adjustment enables the *green* robot to successfully intercept *orange* later in the simulation at $t = 431$ s. Such behaviors demonstrate how higher-order epistemic reasoning facilitates coordination, even in unknown and cluttered environments. These results show that our long-horizon estimation–planning–coordination framework based on higher-order reasoning significantly reduces task completion time relative to first-order baselines. We next present experiments with multiple ground robots to validate the feasibility of our approach in physical systems.

VI. EXPERIMENTS

Our approach was validated through laboratory experiments with teams of 2 and 3 robots placed at random locations within a 4m^2 area. Because of the smaller operating region, the sensing range was decreased to 0.75 m and the

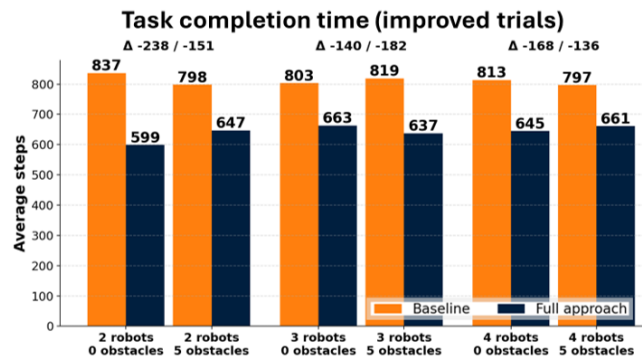


Fig. 7. A comparison of the task completion time for our approach versus the baseline when a decision to fetch was made by a robot within the MRS (improved trials from Fig. 6).

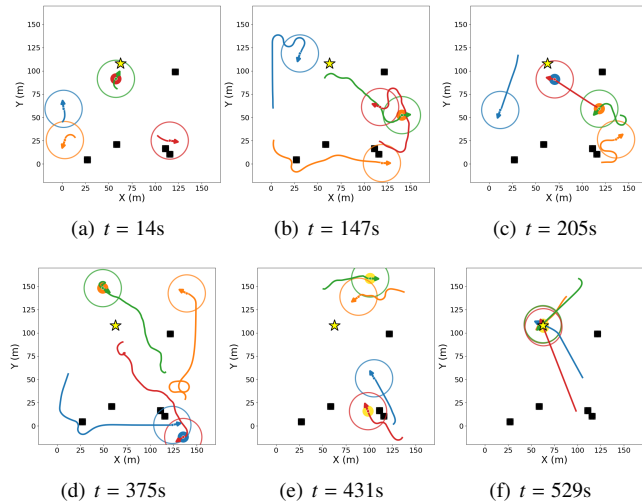


Fig. 8. Simulation snapshots of a 4-robot case in a 5-obstacle environment. Colored rings denote sensing range; circles encode behavior as: yellow = complete task; robot's color = fetching that robot; none = explore.

nominal speed was set to 0.1 m/s. As in the simulations, robots could increase their velocity up to 0.125 m/s during fetching or while returning to their second-order empathy particle after avoiding an obstacle. The task requires $\mathcal{R}_g = \mathcal{R}$.

One example is shown in Fig. 9, with the task placed in the center at (2.0 m, 2.0 m), and two ($0.3\text{m} \times 0.3\text{m}$) obstacles placed within the environment. As illustrated in Fig. 9(a), the team of robots begins exploring the area, at which point the *green* robot locates the task and reasons that no other robots have detected it. As shown in Fig. 9(b), *green* chooses to fetch *blue* using an MPPI. After evaluating the effectiveness of the intercept, the *green* robot determines that it succeeded fetching and selects a new behavior using the behavior tree process. This time, it finds that it can further minimize (1) by choosing to fetch *red*.

While *green* is fetching *red*, *blue* discovers the task. Upon doing so, it applies higher-order reasoning to infer that *green* had already found the task and decided to fetch *blue*. *Blue* then reasons that this behavior was completed and that *green* would have chosen to fetch *red* as its next action. Based on this reasoning, *blue* ultimately selects complete task as its behavior, shown in Fig. 9(c).

The *green* robot successfully intercepts the *red* robot, shifts its search pattern toward the task, and enables mission

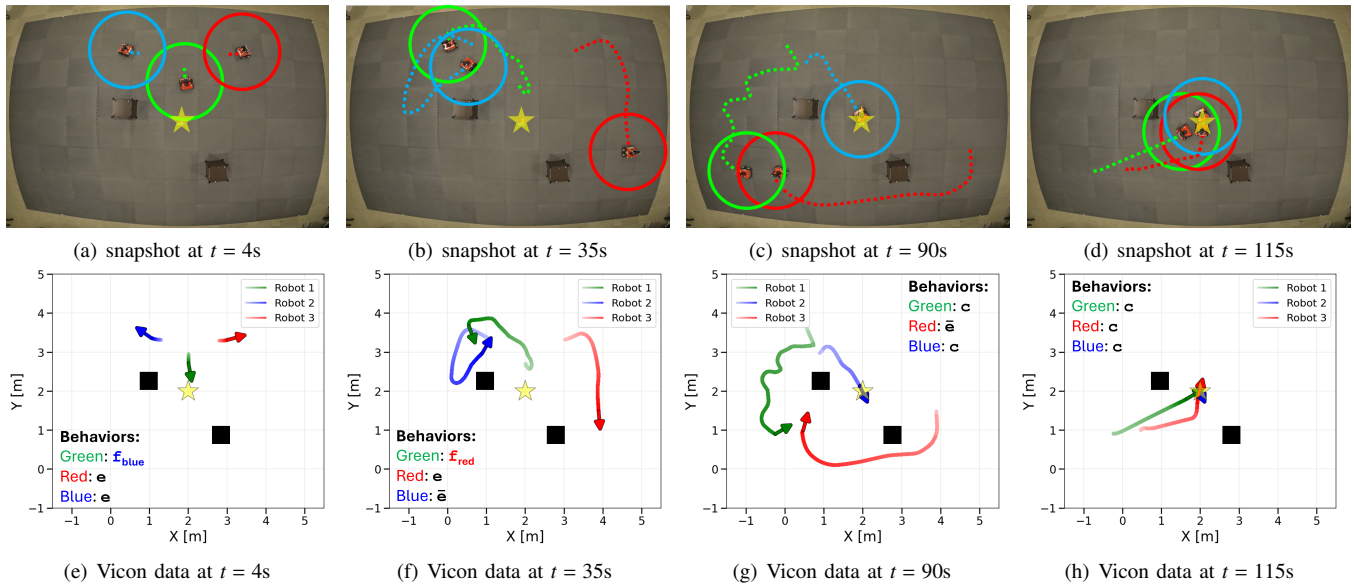


Fig. 9. Experiment results for a 3 robots case.

completion, shown in Fig. 9(d). This mission was completed in $t = 115s$, a reduction of 84s (57.8%) compared to the baseline time of $t = 199s$. Full experimental videos are provided in the supplemental material.

VII. CONCLUSIONS AND FUTURE WORK

This paper presented a decentralized framework for implicit multi-robot coordination in communicationless settings. Using dynamic epistemic logic, robots maintain belief and empathy particles to interpret unexpected events and reason about teammates; a higher-order reasoning-based behavior tree selects actions that minimize task completion time and an MPPI enables long-horizon planning under partial observability. Simulations and experiments show consistent reductions in task completion time versus first-order baselines, validating epistemic reasoning for resilient coordination. Although we focus on communicationless cases, this method naturally extends to limited-range communication systems, in which case the MPPI process can be simplified.

Future work aims to scale the system to larger and more complex environments with multiple tasks and to extend it to heterogeneous teams where robots must reason about the capabilities of other robots and task requirements. We also plan to investigate adversarial scenarios, where higher-order reasoning and epistemic logic may be leveraged not only for cooperation but also for countering deceptive or competitive behaviors and in human-robot scenarios where communication between human and robot is often not possible.

VIII. ACKNOWLEDGMENT

We gratefully acknowledge DARPA YFA award #D25AC00371 and the Commonwealth Cyber Initiative for partial support of this work.

REFERENCES

- [1] Y. Cao, A. Fukunaga, A. Kahng, and F. Meng, "Cooperative mobile robotics: antecedents and directions," in *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*, vol. 1, 1995, pp. 226–234.
- [2] Y. Wu, X. Ren, H. Zhou, Y. Wang, and X. Yi, "A survey on multi-robot coordination in electromagnetic adversarial environment: Challenges and techniques," *IEEE Access*, vol. 8, pp. 53 484–53 497, 2020.
- [3] L. Bramblett, R. Peddi, and N. Bezzo, "Coordinated multi-agent exploration, rendezvous, & task allocation in unknown environments with limited connectivity," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 12 706–12 712.
- [4] Z. Zhou, J. Liu, and J. Yu, "A survey of underwater multi-robot systems," *IEEE/CAA Journal of Automatica Sinica*, vol. 9, no. 1, pp. 1–18, 2022.
- [5] M. Berhault, H. Huang, P. Keskinocak, S. Koenig, W. Elmaghaby, P. Griffin, and A. Kleywegt, "Robot exploration with combinatorial auctions," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, 2003, pp. 1957–1962.
- [6] W. Burgard, M. Moors, C. Stachniss, and F. Schneider, "Coordinated multi-robot exploration," *IEEE Transactions on Robotics*, vol. 21, no. 3, pp. 376–386, 2005.
- [7] Y. Li, Y. Gao, S. Yang, and Q. Quan, "Swarm robotics search and rescue: A bee-inspired swarm cooperation approach without information exchange," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1127–1133.
- [8] D. Premack and G. Woodruff, "Does the chimpanzee have a theory of mind?" *The Behavioral and Brain Sciences*, vol. 4, pp. 515–526, 1978.
- [9] A. Valle, D. Massaro, I. Castelli, and A. Marchetti, "Theory of mind development in adolescence and early adulthood: The growing complexity of recursive thinking ability," *Europe's journal of psychology*, vol. 11, no. 1, p. 112, 2015.
- [10] T. Bolander, "A gentle introduction to epistemic planning: The del approach," *Electronic Proceedings in Theoretical Computer Science*, vol. 243, p. 1–22, Mar. 2017.
- [11] T. Engesser, T. Bolander, R. Mattmüller, and B. Nebel, "Cooperative epistemic multi-agent planning for implicit coordination," *Electronic Proceedings in Theoretical Computer Science*, vol. 243, p. 75–90, 2017.
- [12] T. Bolander, T. Charrier, S. Pinchinat, and F. Schwarzentruher, "Del-based epistemic planning: Decidability and complexity," *Artificial Intelligence*, vol. 287, p. 103304, 2020.
- [13] W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun, "Collaborative multi-robot exploration," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings*, vol. 1, 2000, pp. 476–481.
- [14] E. Yel, T. Lin, and N. Bezzo, "Computation-aware adaptive planning and scheduling for safe unmanned airborne operations," *Journal of Intelligent & Robotic Systems*, vol. 100, no. 2, pp. 575–596, 2020.
- [15] M. E. Buzikov and A. A. Galyaev, "Time-optimal interception of a moving target by a dubins car," *Automation and Remote Control*, vol. 82, pp. 745–758, 2021.
- [16] G. Williams, A. Aldrich, and E. A. Theodorou, "Model predictive path integral control: From theory to parallel computation," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [17] L. Bauersfeld and D. Scaramuzza, "Range, endurance, and optimal speed estimates for multicopters," 2024. [Online]. Available: <https://arxiv.org/abs/2109.04741>