

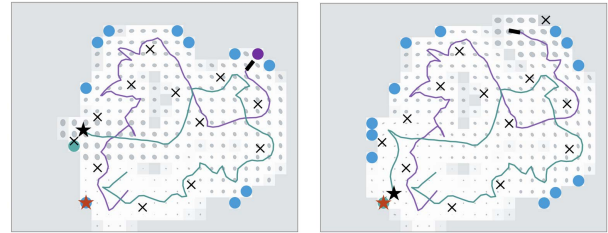
Multi-Robot Autonomous Exploration and Mapping Under Localization Uncertainty with Expectation-Maximization

Yewei Huang¹, Xi Lin¹ and Brendan Englot¹

Abstract—We propose an autonomous exploration algorithm designed for decentralized multi-robot teams, which takes into account map and localization uncertainties of range-sensing mobile robots. Virtual landmarks are used to quantify the combined impact of process noise and sensor noise on map uncertainty. Additionally, we employ an iterative expectation-maximization inspired algorithm to assess the potential outcomes of both a local robot’s and its neighbors’ next-step actions. To evaluate the effectiveness of our framework, we conduct a comparative analysis with state-of-the-art algorithms. The results of our experiments show the proposed algorithm’s capacity to strike a balance between curbing map uncertainty and achieving efficient task allocation among robots.

I. INTRODUCTION

Autonomous exploration and mapping describes a single robot or a group of robots navigating themselves in an unknown or partially known environment without human intervention. An accurate environment map constructed by the robots serves as the fundamental basis for all subsequent specific robot tasks [1]. While the maturity of autonomous exploration for ground and aerial robots has been notably demonstrated, its application within marine environments remains a subject of ongoing inquiry [2]. In spite of the widespread coverage of Global Navigation Satellite System (GNSS) signals in most maritime regions, there exist many marine locations where GNSS signals are vulnerable due to obstruction or attenuation. In these areas, robots face a higher risk of collisions, emphasizing the critical need for an accurate environment map. However, deploying exploration in such contexts is a particular challenge for robot teams. This is primarily due to the high uncertainty introduced by unique environmental factors pertaining to the marine environment, especially when operating below the ocean’s surface. Autonomous exploration and mapping have been a vigorously discussed subject for several decades. Early research has primarily centered around optimizing task distribution among team members [3], [4]. As Simultaneous Localization and Mapping (SLAM) techniques have advanced, various approaches [5], [6], [7], [8], [9] have incorporated the concept of map uncertainty into autonomous exploration. These strategies consider Gaussian noise sensor models and Gaussian noise kinematic models, aiming to strike a balance between exploration efficiency and managing uncertainties in the resulting maps. Although these techniques may differ in terms of their map representation and merging techniques, they share a similar utility function that takes into account both the information gain of the latest generated map and the



(a) Choosing a revisiting frontier (blue, with red star) driven by the uncertainty in the map. (b) Choosing an exploring frontier (green, with red star) due to relatively low map uncertainty.

Fig. 1: Problem setup. In this $100m \times 80m$ virtual map created by a two-robot team, gray ellipses depict the uncertainty of visited cells. The green robot’s position is denoted by a black star, and its newly selected target state is represented by a red star. The current position of the other robot on the team is marked as a black rectangle. Landmarks are expressed by black x’s. Potential frontiers emerge along the boundary between the explored and unexplored areas. The three types of frontiers—exploring, revisiting, and rendezvous—are denoted by their respective colors: green, blue, and purple.

required travel distance to the waypoints under consideration. However, in the process of computing information gain, these techniques either ignore uncertainty propagation or have only propagated uncertainty considering future steps, neglecting the history of past steps, which is inadequate for situations with high localization uncertainty.

Building upon our previous work on single robot expectation-maximization (EM) inspired exploration [10], we introduce an asynchronous EM exploration algorithm for both centralized and decentralized multi-robot teams. The algorithm is tightly coupled with a factor graph SLAM system. Upon a robot’s arrival at the target position while the exploration remains ongoing, a *virtual map* is constructed based on the SLAM result to gauge the prevailing map uncertainty. Subsequently, for each potential new target, we execute an expectation-maximization procedure to assess the potential information gain associated with the forthcoming actions of the robot and its interactions with neighboring robots. Selection of a new target goal position is then determined by considering both the collective information gain of the entire team within the virtual map, and the efficiency of task allocation among robots.

Our research introduces several innovative contributions:

- An asynchronous multi-robot exploration framework catering to both centralized and decentralized SLAM systems, taking into account efficient task allocation for exploration and addressing map uncertainty.
- Incorporating an expectation-maximization inspired technique to assess the future impact and interactions of a robot with its neighboring entities.

¹Y. Huang, X. Lin and B. Englot are with Stevens Institute of Technology, Hoboken, NJ, USA. {yhuang85, xlin26, benglot}@stevens.edu
This work was supported by ONR Grant N00014-21-1-2161.

- Introducing an efficient inter-robot and local map uncertainty propagation approach, tailored to scenarios involving multiple robots and localization uncertainty.

Our code for the proposed framework is publicly released¹. The subsequent sections of our paper are structured as follows: a review of background literature is presented in Sec. II, followed by the introduction of factor graph SLAM and the proposed expectation-maximization exploration method in Sec. III. The comprehensive process of our multi-robot exploration with expectation-maximization is outlined in Sec. IV. Our experimental outcomes are detailed in Sec. V, and the paper concludes with a summary in Sec. VI.

II. RELATED WORKS

In multi-robot exploration, the communication arrangement of the robot system holds significance. Charrow et al. [11] examines the exploration of a small area using a team of robots equipped with range-only sensors using a coordinated approach. Although a centralized system maximizes the utilization of data collected by robot teams, it experiences increased computation challenges as the size of the robot team expands. Numerous algorithms, such as [3], [7], [12], [13], [14] and [15] adopt a decentralized approach. In this approach, robots exchange a significant portion of their historical data while independently making decisions based on their most recent knowledge of the environment. The asynchronous nature of the decentralized system reduces the communication and computational load in comparison to a scenario involving a centralized computer. However, it can still face limitations imposed by the communication range of the robot team. Additionally, there exist distributed approaches, discussed in [4], [16], [17], [18] and [19], in which robots communicate only when they encounter each other, adhering to constraints on communication bandwidth. Their decisions depend on the limited information available from both the robot team and the environment. Such distributed approaches are designed for scenarios involving larger robot teams with constrained communication bandwidth, although they may involve some performance trade-offs to accommodate these conditions.

Another critical factor to consider is the criteria utilized for action selection. In earlier approaches such as [3], [4], the primary emphasis was on efficiently distributing tasks among team members. Later, the “Next Frontier” [12] introduced the concept of information potential to improve exploration efficiency. Jang et al. [5] utilize a Gaussian processes to characterize the environment, enabling obstacle avoidance and task allocation for a team of robots. In NeuralCoMapping [20], a multiplex graph neural network (GNN) is introduced to strike a balance between long-term and short-term performance optimization. Tzes et al. [21] present a novel approach that integrates the exploration under uncertainty problem into a learning framework using graph neural networks.

Given the uncertainty inherent in both robot control and sensing, some authors incorporate uncertainty into the exploration problem. Kontitsis et al. [8] incorporate Relative

Entropy into the utility function when dealing with a partially known map to reduce localization uncertainty. Similarly, in [22], the authors perform joint entropy minimization to actively explore over a long-term horizon. Other researchers utilize filter-based techniques to address the uncertainty associated with future steps. In the work by Atanasov et al. [7], a square-root information filter is employed to predict the impact of future actions. Schlotfeldt et al. [14] adopt a similar filter-based approach within a decentralized system. Meanwhile, Kantaros et al. [18] present a sampling-based method aimed at reducing cumulative uncertainty stemming from dynamic hidden states. Indelman [23] presents an approach that employs belief propagation to anticipate the performance of robot teams in future actions. Chen [24] considers graph topology when making action selections.

In contrast to other multi-robot exploration approaches, which may overlook localization uncertainty or prioritize the influence of future robot interactions to reduce map uncertainty, our expectation-maximization based approach emphasizes both the interactions within a robot team and the iterative re-visitation of the existing map. Furthermore, we use a *virtual map*, which serves as a robust tool for assessing not only the localization uncertainty of individual robots but also the collective uncertainty introduced by robots into the map.

III. PROBLEM FORMULATION AND APPROACH

We address an autonomous exploration problem that is tightly coupled with a SLAM factor graph for a team of n robots. We make the assumption that the initial states of all robots are sufficiently close to each other, enabling mutual observation among group members and facilitating an efficient map initialization process. Additionally, we impose a boundary on the exploration task, where the exploration process terminates upon fully exploring the enclosed environment.

A. Simultaneous Localization and Mapping

Let $N = \{1, 2, \dots, n\}$ be the set of n robots. For each robot $\alpha \in N$, we denote its state at timestamp i as $\mathbf{x}_{\alpha,i}$. The robot odometry observation between present state $\mathbf{x}_{\alpha,i}$ and previous state $\mathbf{x}_{\alpha,i-1}$ is described by the equation:

$$\mathbf{z}_{\alpha,i}^{\alpha,i-1} = f(\mathbf{x}_{\alpha,i-1}, \mathbf{x}_{\alpha,i}) + \epsilon_{\alpha,i}^{\alpha,i-1}. \quad (1)$$

Assume robot α observes a landmark state \mathbf{l}_j at timestamp i , we can describe the landmark observation:

$$\mathbf{z}_j^{\alpha,i} = g(\mathbf{x}_{\alpha,i}, \mathbf{l}_j) + \epsilon_j^{\alpha,i}. \quad (2)$$

If another robot β is observed at timestamp i by robot α , we refer to this as a *robot rendezvous observation*:

$$\mathbf{z}_{\beta,i}^{\alpha,i} = f(\mathbf{x}_{\alpha,i}, \mathbf{x}_{\beta,i}) + \epsilon_{\beta,i}^{\alpha,i}, \quad (3)$$

where $f(\cdot)$ denotes the state transformation between robot states, $g(\cdot)$ is the state transformation from robot state to landmark state, and $\epsilon_{\alpha,i}^{\alpha,i-1}$, $\epsilon_j^{\alpha,i}$ and $\epsilon_{\beta,i}^{\alpha,i}$ are zero-mean Gaussian noise variables.

At present timestamp t , $\mathcal{X} = \{\mathbf{x}_i^\alpha | \alpha \in N, i \in [0, t]\}$ represents the set containing the states of all n robots from the initial timestamp 0 to the present timestamp t .

¹<https://github.com/RobustFieldAutonomyLab/Multi-Robot-EM-Exploration>

$\mathcal{L} = \{\mathbf{l}_0, \mathbf{l}_1, \dots, \mathbf{l}_m\}$ is the set of landmarks observed until timestamp t . In this paper we define a landmark as a point with a unique identifier that is observable by the robots' on-board sensors. Additionally, let \mathcal{Z} be the set containing odometry observations, landmark observations and robot rendezvous observations from all timestamps. The SLAM problem can be framed as a maximum a posteriori estimation problem [25]:

$$\mathcal{X}^*, \mathcal{L}^* = \arg \max_{\mathcal{X}, \mathcal{L}} P(\mathcal{X}, \mathcal{L} | \mathcal{Z}). \quad (4)$$

B. Expectation-Maximization Exploration

During the exploration process, when a robot α reaches its current target state $\mathbf{a}_\alpha^{\text{this}}$, it becomes necessary to select a new target state $\mathbf{a}_\alpha^{\text{next}}$ from a set of potential new states $\mathcal{A}_\alpha^{\text{next}}$. Building upon our previous research on single-robot exploration [10], we consider a frontier-based strategy that incorporates two key factors: efficient task allocation among robots [3] and the maintenance of a low-uncertainty map.

When α reaches target state $\mathbf{a}_\alpha^{\text{this}}$, $\forall \mathbf{a}_\alpha^{\text{next}'} \in \mathcal{A}_\alpha^{\text{next}}$, we define:

$$\mathcal{X}^{\text{new}} = \mathcal{X}^{\text{old}} \cup \mathcal{X}^{\text{predict}} \cup \mathcal{X}^{\text{next}}. \quad (5)$$

\mathcal{X}^{old} contains the historical states of all n robots, $\mathcal{X}^{\text{predict}}$ denotes the set of predicted robot states for each current target state $\{\mathbf{a}_i^{\text{this}} | i \in N, i \neq \alpha\}$ and $\mathcal{X}^{\text{next}}$ represents the state sequence of robot α resulting from $\mathbf{a}_\alpha^{\text{next}'}$. A classification EM algorithm is used to predict the change of map uncertainty due to $\mathbf{a}_\alpha^{\text{next}'}$. To avoid the exponential expansion in potential virtual landmark states, we substitute the **E-step** with a classification step (**C-step**), in which we construct a virtual map using the historical data of our robot team:

$$\mathcal{V}^* = \arg \max_{\mathcal{V}} P(\mathcal{V} | \mathcal{X}^{\text{old}}, \mathcal{Z}^{\text{old}}), \quad (6)$$

$$= M(\mathcal{X}^{\text{old}*}, \mathcal{Z}^{\text{old}}). \quad (7)$$

Here, \mathcal{Z}^{old} represents the observations associated with \mathcal{X}^{old} . $\mathcal{X}^{\text{old}*}$ is the optimized value from the previous SLAM optimization. $M(\cdot)$ is the inverse observation model used to estimate the mean and covariance of the virtual map with the given SLAM estimate $\mathcal{X}^{\text{old}*}$. For a full explanation of the virtual map construction process, please refer to our prior work [26]. Subsequently, in the **M-step**, we again perform maximum a posteriori estimation to find \mathcal{X}^{new} (and choose $\mathcal{X}^{\text{next}}$):

$$\mathcal{X}^{\text{new}*} = \arg \max_{\mathcal{X}^{\text{new}}} P(\mathcal{X}^{\text{new}} | \mathcal{V}^*, \mathcal{Z}^{\text{new}}), \quad (8)$$

$$\mathcal{Z}^{\text{new}} = \mathcal{Z}^{\text{old}} \cup \mathcal{Z}^{\text{predict}}. \quad (9)$$

The updated observation set, \mathcal{Z}^{new} , is a combination of both previously gathered observations, \mathcal{Z}^{old} , and anticipated future observations, $\mathcal{Z}^{\text{predict}}$. These predicted observations are determined via the measurement models outlined in Eq. (1), Eq. (2), and Eq. (3), based on the virtual map \mathcal{V}^* and the target state assigned to each robot. Thus, we assess the overall local map uncertainty of robot α by computing the sum of the individual uncertainties for each map element in the virtual map $\mathcal{V}(\cdot)$ derived from Eq. (7), considering only the optimized

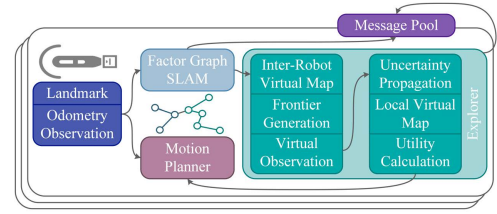


Fig. 2: **System Architecture.** The pipeline of the proposed approach.

states of robot α , $\mathcal{X}_\alpha^{\text{new}*}$:

$$U_M = \phi(\Sigma_M(\mathcal{X}_\alpha^{\text{new}*}, \mathcal{Z}^{\text{new}})), \quad (10)$$

$$= \sum_{\mathbf{v}_i \in \mathcal{M}(\mathcal{X}_\alpha^{\text{new}*}, \mathcal{Z}^{\text{new}})} \phi(\Sigma_{\mathbf{v}_i}). \quad (11)$$

Each element \mathbf{v}_i contributes to the sum based on its covariance $\Sigma_{\mathbf{v}_i}$, and in this paper, we employ the A-Optimality metric as our uncertainty criterion ϕ .

To optimize the distribution of exploration tasks among the robots, we utilize a distance evaluation metric inspired by the methodology presented in [3]. For each potential new target state \mathbf{a}_α^t assigned to robot α at time t , we quantify the effectiveness of task allocation using U_T . A higher U_T results in the robots being closer to their teammates, increasing the likelihood that they repeatedly explore the same area.

$$U_T = \sum_{\mathbf{a}_\beta^i \in \mathcal{A}, \beta \neq \alpha} h(\|\mathbf{a}_\alpha^t - \mathbf{a}_\beta^i\|_2), \quad (12)$$

$$h(d) = \begin{cases} 1 - \frac{d}{d_{\max}} & d < d_{\max} \\ 0 & d \geq d_{\max} \end{cases}. \quad (13)$$

Here, $\|\cdot\|_2$ represents the L2 norm, and \mathcal{A} denotes the collection of all historical target states for all robots. Based on this, we can define the new target state for robot α as one of the potential target states that optimally balances the factors U_M , U_T and the Euclidean distance to target state factor, U_D , with scale factors λ_0 , λ_1 , λ_2 :

$$\mathbf{a}_\alpha^{\text{next}} = \arg \min_{\mathbf{a}_\alpha^t} (\lambda_0 U_M + \lambda_1 U_T + \lambda_2 U_D). \quad (14)$$

IV. PROPOSED MULTI-ROBOT EXPLORATION ALGORITHM

Here we discuss the details of the proposed algorithm. Fig. 2 shows the pipeline of the proposed algorithm. We consider a situation where a group of robots collaborates, and each individual robot within the group is furnished with both motion and perception sensors. A centralized or decentralized factor-graph based SLAM algorithm is assumed to operate at a specific frequency. The robot team shares their optimized SLAM trajectories, landmark positions and history of chosen target states among its members. This simulation does not take into account any limitations on communication bandwidth. However, individual robots maintain virtual maps locally using the latest SLAM estimates.

After a robot α successfully reaches its current target state, its virtual map is updated and a group of potential new target states $\mathcal{A}_\alpha^{\text{next}}$ is chosen from the virtual map. Subsequently, the virtual observation is synthesized utilizing the prevailing environmental knowledge. Next, the expectation-maximization based uncertainty propagation computes the potential impact

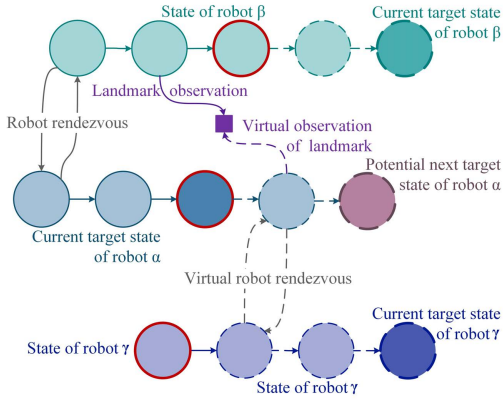


Fig. 3: **EM-based uncertainty propagation with virtual observations.** Nodes representing the current robot states are distinguished by red edges. For every potential next target state (shown in pink), a trajectory simulation is executed, leading to the generation of a sequence of virtual observations indicated by dashed arrows. Simultaneously, we model the future states and observations of other robots as they approach their individual current target states.

of the target state under consideration, leading to an update in the virtual map’s covariances. The new target goal is then selected according to the utility function (Eq. (14)). Finally, the motion planner is initiated to formulate a series of actions leading robot α to the new target state.

A. Virtual Map

As depicted in Eq. (7), the virtual map \mathcal{V} of a finite environment is generated from the robot states \mathcal{X} and their associated observations \mathcal{Z} . Assuming that \mathcal{V} comprises b map cells \mathbf{v}_i referred to as *virtual landmarks*, the posterior can be redefined as follows:

$$P(\mathcal{V}|\mathcal{X}, \mathcal{Z}) = \prod_{\mathbf{v}_i \in \mathcal{V}} P(\mathbf{v}_i|\mathcal{X}, \mathcal{Z}). \quad (15)$$

The likelihood of virtual landmark \mathbf{v}_i being observed is $q(\mathbf{v}_i) = \mathbb{E}[P(\mathbf{v}_i|\mathcal{X}, \mathcal{Z})]$. When this virtual landmark is observed by multiple robot states, the procedure for updating $q(\mathbf{v}_i)$ is similar to updating map cell values in an occupancy grid map [27]. We assume the virtual landmark \mathbf{v}_i is observed by a robot state $\mathbf{x}_{\alpha,j}$, with its estimated value denoted $\hat{\mathbf{x}}_{\alpha,j}$, and a marginal covariance $\Sigma_{\mathbf{x}_{\alpha,j}}$. We can compute the covariance: $\Sigma_{\mathbf{v}_i} = \mathbf{H} \cdot \Sigma_{\mathbf{x}_{\alpha,j}} \cdot \mathbf{H}^\top$. Here, $\mathbf{H} = \frac{\partial g(\mathbf{x}_{\alpha,j}, \mathbf{v}_i)}{\partial \mathbf{x}_{\alpha,j}} |_{\hat{\mathbf{x}}_{\alpha,j}}$ represents the Jacobian matrix obtained by differentiating the landmark observation model $g(\mathbf{x}_{\alpha,j}, \mathbf{v}_i)$ with respect to estimated robot state $\hat{\mathbf{x}}_{\alpha,j}$. We utilize Covariance Intersection to compute the covariance of a virtual landmark that is observed by multiple robot states, as detailed in [26]. Fig. 1 shows inter-robot virtual maps built from both local robot states and neighbors’ robot states received by local robot α . The observed regions are highlighted in white; gray ellipses show covariances describing the uncertainty of the map’s cells.

B. Uncertainty Propagation

As depicted in Fig. 1, most potential target states are chosen from the perimeters of the observed regions. Three types of frontiers are identified for selection: *exploration frontiers* close to the robot’s latest position, *revisiting frontiers* near previously visited landmarks, and *rendezvous frontiers*,

Algorithm 1: Uncertainty Propagation

Global: Latest SLAM Graph \mathcal{G} , and Inter-robot Virtual Map \mathcal{V}^*
Input: Potential frontier state $\mathbf{a}_{\alpha'}^t$, Current robot states \mathcal{X}^t , Current Target States of Neighbors \mathcal{A}^t
Output: Optimized robot states $\mathcal{X}^{\text{new}*}$
 $\mathcal{X}^{\text{predict}} \leftarrow \emptyset$
foreach $\mathbf{a}_\gamma^t \in \mathcal{A}^t, \mathbf{x}_\gamma^t \in \mathcal{X}^t$ **do**
 $\mathcal{X}_\gamma^{t:*} \leftarrow \text{GenerateVirtualWaypoints}(\mathbf{a}_\gamma^t, \mathbf{x}_\gamma^t)$
 $\mathcal{X}^{\text{predict}} \leftarrow \mathcal{X}_\gamma^t \cup \mathcal{X}^{\text{predict}}$
 $\mathcal{X}^{\text{next}} \leftarrow \text{GenerateVirtualWaypoints}(\mathbf{a}_{\alpha'}^t, \mathbf{x}_\gamma^t)$
 # Calculate virtual observations
 $\mathcal{Z}^{\text{predict}} \leftarrow \text{VirtualObserve}(\mathcal{V}^*, \mathcal{X}^{\text{predict}} \cup \mathcal{X}^{\text{next}})$
 # M-step: graph optimization
 $\text{UpdateGraph}(\mathcal{G}, \mathcal{X}^{\text{predict}} \cup \mathcal{X}^{\text{next}})$
 $\mathcal{X}^{\text{new}*} \leftarrow \text{OptimizeGraph}(\mathcal{G})$
return $\mathcal{X}^{\text{new}*}$

which are the current target positions of neighboring robots. Subsequently, for each potential target, a set of waypoints \mathcal{X}^{new} is uniformly sampled along the shortest path connecting the robot’s current state and the next potential target state. Additionally, waypoints $\mathcal{X}^{\text{predict}}$ are generated to connect the present states of all robots with their respective target states. The process of generating virtual observations $\mathcal{Z}^{\text{predict}}$ along the paths to these target states is depicted in Fig. 3. Virtual odometry measurements are created between adjacent waypoints using Eq. (1). For virtual landmark observations, Eq. (2) is employed, generating observations between previously observed landmarks and nearby waypoints. When two robots are within each other’s sensing range at the same timestep, a virtual robot observation is produced using Eq. (3). Following this, virtual observations are added into the SLAM graph to propagate uncertainty. The procedure is outlined in Alg. 1.

C. Map Uncertainty Utility Computation

We compute the uncertainty of map cells by considering both the likelihood of the existing virtual map \mathcal{V}^* and the optimized robot states of robot α , $\mathcal{X}_\alpha^{\text{new}*} \subset \mathcal{X}^{\text{new}*}$. Referencing Fig. 1(b), it is evident that despite the purple robot’s trajectory being impacted by localization uncertainty due to accumulated odometry error, this uncertainty is partially mitigated by the historical trajectory of the green robot, which exhibits relatively lower uncertainty. The construction of an inter-robot virtual map for map uncertainty estimation could potentially lead to conflicts in decision-making for a local robot. As outlined in Alg. 2, we generate a local virtual map denoted as $\mathcal{V}_\alpha^{\text{new}*}$ for robot α . This map is constructed using $\mathcal{X}_\alpha^{\text{new}*}$ and corresponding virtual observations $\mathcal{Z}_\alpha^{\text{new}*}$. Then, we compute the uncertainty utility of the map by considering the overlapping observed regions common to both the current inter-robot virtual map \mathcal{V}^* and the predicted local virtual map $\mathcal{V}_\alpha^{\text{new}*}$, via Eq. (10).

D. Complexity Analysis

In this section, we analyze the time complexity of the Expectation-Maximization Explorer shown in Fig. 2. For a team of n robots, each with a maximum of N_x historical robot states, updating a robot state in the virtual map takes T_v time,

Algorithm 2: Compute U_M of potential frontier \mathbf{a}'_α

Input: Inter-robot Virtual Map \mathcal{V}^* , Local Virtual Map $\mathcal{V}_\alpha^{\text{new}^*}$

Output: Map uncertainty utility factor U_M

$U_M \leftarrow 0$

for $i \leftarrow 0$ **to** b **do**

$\mathbf{v}_{i,g} \in \mathcal{V}^*$, $\mathbf{v}_{i,l} \in \mathcal{V}_\alpha^{\text{new}^*}$

 # q_{\min} : minimum accepted probability of observed.

if $q(\mathbf{v}_{i,l}) > q_{\min}$ **and** $q(\mathbf{v}_{i,g}) > q_{\min}$ **then**

$U_M \leftarrow U_M + \text{trace}(\Sigma_{\mathbf{v}_{i,l}})$ # A-Optimality

return U_M

resulting in a potential inter-robot virtual map construction time of $n \cdot N_x \cdot T_v$. During frontier generation, detecting exploration bounds (time T_b) containing N_c virtual landmarks and observing each landmark (time T_l) contributes to a step duration of $T_b + N_c \cdot T_l$. Frontier selection generates N_f frontiers. Crafting virtual waypoints (up to N_w) and conducting virtual robot rendezvous checks (time T_r) per frontier results in a total frontier generation time of $(N_f + n - 1) \cdot (T_l + T_r)$. Utilizing iSAM2 [25] for uncertainty propagation (time T_u per update) leads to a combined propagation time of $N_f \cdot T_u$. Constructing a local virtual map requires $N_x \cdot T_v$ time. The most time-intensive steps involve iSAM2's uncertainty propagation ($O(n^{2.36})$ complexity for n states) and covariance intersection during virtual map creation ($O(s^3)$ complexity with s nonzero matrix block size) [28].

V. EXPERIMENTS AND RESULTS

We perform simulated experiments within environments of varied size, employing randomly generated landmarks. The errors described below define 95% confidence intervals. Similar to our previous work [10], we assume each robot is furnished with a sonar with range error 0.002m, bearing error 0.5° , and max. sensing range 7.5m. Each robot also performs inertial dead reckoning with translational and rotational components in a 2D space. The gyro and accelerometer measurements introduce errors of 0.5 deg. and 0.05 meters, respectively. The robot has the capability to rotate by 15° during each action, maintaining a constant speed of $1m/s$. It is restricted to moving solely in the direction of its current heading. At the beginning of each trial, we ensure that all robots are located within the sensing range of their teammates to guarantee a proper initialization of the SLAM framework. To navigate towards uncharted territory, we utilize the Artificial Potential Field (APF) method, detailed in [29], to avoid collision. The boundaries of the environment are restricted from selection as frontiers, preventing robots from exiting the mission area.

We compare our proposed approach, denoted EM, with two advanced multi-robot exploration algorithms: the coordinated multi-robot exploration method by Burgard et al. [3], referred to as CE in subsequent sections, and the cooperative multi-robot belief space planning technique introduced by Indelman et al. [23], denoted as BSP subsequently. Identical frontier selection criteria and virtual observation generation techniques are employed for both methods.

The CE planner [3] places emphasis on optimizing task distribution among robots, employing a utility function that

takes task allocation into account:

$$U_{CE} = \lambda_0 U_D + \lambda_1 U_T. \quad (16)$$

In this context, both U_D and U_T are computed the same way as outlined in Eq. 14. In our experiment, we select $\lambda_0 = 1$ and $\lambda_1 = 10$. **The BSP planner** [23] takes into account both exploration efficiency and localization uncertainty, employing a similar virtual observation strategy as ours. This leads to the formulation of the following utility function:

$$U_{BST} = \lambda_0 U_D + \lambda_1 U_M, \quad (17)$$

$$U_M = \sum_{\mathbf{x}_i \in \mathcal{X}^{\text{predict}^*} \cup \mathcal{X}^{\text{next}^*}} \phi_A(\sqrt{\Sigma_{\mathbf{x}_i}}). \quad (18)$$

We choose λ_0 as 5 and λ_1 as 1. Here, $\Phi_A(\cdot)$ denotes the trace of the matrix, and $\mathcal{X}^{\text{predict}^*}$ and $\mathcal{X}^{\text{next}^*}$ are computed through the maximum a posteriori process:

$$\mathcal{X}^{\text{predict}^*}, \mathcal{X}^{\text{next}^*} = \arg \max_{\mathcal{X}^{\text{predict}}, \mathcal{X}^{\text{next}}} P(\mathcal{X}^{\text{predict}}, \mathcal{X}^{\text{next}} | \mathcal{Z}^{\text{predict}}, \mathcal{L}).$$

The BSP planner, focusing only on the localization uncertainty of future steps, ignores potential benefits from revisiting previously explored landmarks by the team. Since only the uncertainties at individual robot states are taken into account, U_M might become imbalanced and not effectively depict the overall localization uncertainty across the entire environment, especially in the presence of varying vehicle speeds.

A quantitative analysis using two distinct sizes of random generalized environments is conducted: one measuring $100m \times 100m$, and the other measuring $200m \times 200m$. The $200m \times 200m$ environment presents a unique challenge due to its larger size and sparser distribution of landmarks. Across 50 trials, each trial involves 20 landmarks with a radius of $1m$ within the environment. We introduce three robots into the environment to prevent rendezvous and landmark revisitation from becoming trivially achievable in scenarios involving a large number of robots. The landmark positions are selected randomly, ensuring a minimum distance of $10m$ between each pair of landmarks. For each trial, the robots initiate their positions from the middle left region of the environment, with their initial positions also being chosen randomly.

We test two distinct configurations for the proposed EM explorer: one with $\lambda_0 = 1, \lambda_1 = 0, \lambda_2 = 10$ referred to as EM.2, and the other with $\lambda_0 = 1, \lambda_1 = 20 \cdot (1 - r), \lambda_2 = 10$ referred to as EM.3, where r denotes the ratio of the explored area in the environment. For the $100m \times 100m$ environment, we employ a virtual map cell size of $c_v = 2m$, while for the $200m \times 200m$ environment, we choose a cell size of $c_v = 4m$ for the virtual map. We use three statistics to evaluate the performance of the proposed algorithm:

- Robot localization error: root-mean-square error (RSME) of the optimized trajectories of all robots from the SLAM framework.
- Landmark position error: RSME of the positions of all landmarks observed by the robot teams.
- Explored ratio: The proportion of the area considered as observed in the inter-robot virtual map.

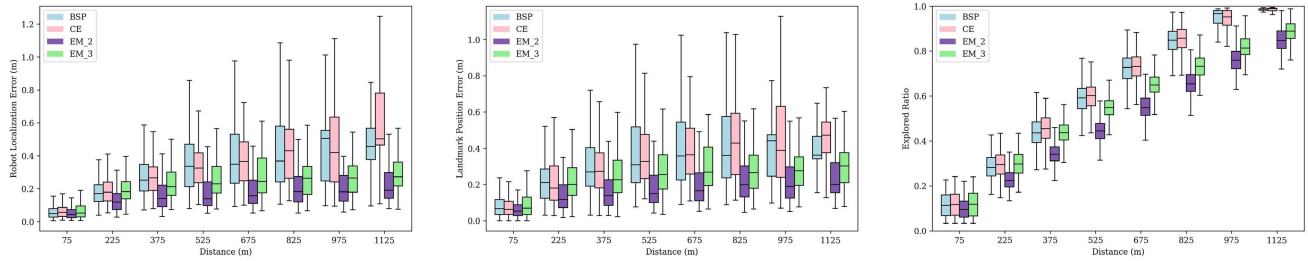


Fig. 4: 50 trials of three robots navigating in 100m x 100m environments with 20 landmarks, each with a radius of 1m. They are exploring the environment by constructing a virtual map with cell size of 2m. At left, the average robot localization error for each robot state, at center, the average landmark position error for each landmark, and at right, the explored ratio, all plotted against distance.

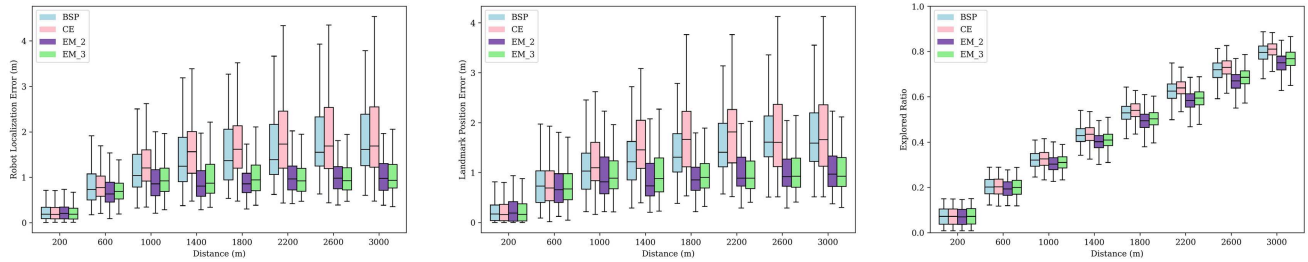


Fig. 5: 50 trials of three robots navigating in 200m x 200m environments with 20 landmarks, each with a radius of 1m. They are exploring the environment by constructing a virtual map with cell size of 4m. At left, the average robot localization error for each robot state, at center, the average landmark position error for each landmark, and at right, the explored ratio, all plotted against distance.

Fig. 4 and Fig. 5 depict the outcomes across different scenarios within smaller and larger environments, respectively. EM_2 and EM_3 exhibit superior accuracy in localizing both robot and landmark states when compared to CE and BSP. This improvement stems from the fact that the proposed approach takes into account not only the forthcoming interactions among robots but also the impact of the team’s historical actions. Nonetheless, it’s worth noting that the proposed method demonstrates lower exploration efficiency when contrasted with both CE and BSP, especially in smaller environments. This trade-off is made in favor of achieving higher accuracy. When considering the task allocation element, the proposed approach in the EM_3 configuration demonstrates better exploration efficiency compared to the EM_2 configuration.

Fig. 6 and Fig. 7 present qualitative results of the proposed algorithm’s EM_3 configuration. In Fig. 6, five robots navigate themselves in a 150m x 150m environment with 40 landmarks. The cyan robot aggregates uncertainties (left) and decides to choose a new target goal state close to a landmark that has been visited by its teammate previously. Three robots navigate a 200m x 200m environment in Fig. 7, with only one landmark observed thus far. After some time, the accumulated localization uncertainty is relatively high (left), prompting the green robot to decide to rendezvous with its pink neighbor. This strategic move leads to a reduction in uncertainty, thanks to the inter-robot loop closure (middle). Concurrently, the light blue robot also accumulates errors and opts to rendezvous with its teammates. Once all three robots meet each other (right), the loop is closed, and uncertainty is propagated and reduced.

VI. CONCLUSIONS

This paper introduces multi-robot exploration using expectation-maximization. We present an asynchronous exploration framework suitable for both centralized and decentralized robot teams, accounting for map uncertainty. The

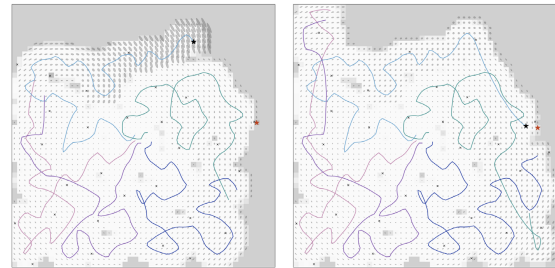


Fig. 6: Five robots navigating a 150m x 150m environment. The active robot seeking its next target state is highlighted by a black star, while its chosen goal is marked with a red star. In the left map, the cyan robot revisits a previously explored landmark due to high map uncertainty (gray ellipses in each grid cell). After a loop closure in the right map, overall localization uncertainty decreases, prompting the robot to select the nearest unexplored frontier for its next move.

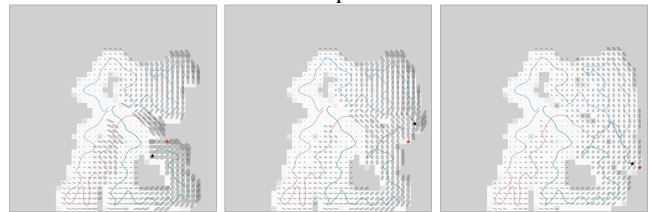


Fig. 7: Three robots explore a 200m x 200m environment. Only the cyan robot has discovered a landmark, prompting the team of robots to strategically rendezvous with their teammates one by one.

inclusion of rendezvous frontiers enhances the system’s adaptability to environments with sparse features. The incorporation of a virtual map enables the estimation of future influence and interactions within the robot teams. The utilization of inter-robot and local maps heightens the robot’s sensitivity to the accumulation of localization uncertainty in its trajectory. As a future direction, we aim to integrate this system with our previous decentralized multi-robot SLAM framework designed for underwater robot teams [30] and subsequently validate the algorithm through experimentation on real robots.

REFERENCES

- [1] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1686–1705, 2023.
- [2] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [3] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider, "Coordinated multi-robot exploration," *IEEE Transactions on robotics*, vol. 21, no. 3, pp. 376–386, 2005.
- [4] D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schulz, and B. Stewart, "Distributed multirobot exploration and mapping," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1325–1339, 2006.
- [5] D. Jang, J. Yoo, C. Y. Son, D. Kim, and H. J. Kim, "Multi-robot active sensing and environmental model learning with distributed gaussian process," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5905–5912, 2020.
- [6] J. Yu, J. Tong, Y. Xu, Z. Xu, H. Dong, T. Yang, and Y. Wang, "Smmr-explore: Submap-based multi-robot exploration system with multi-robot multi-target potential field exploration method," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 8779–8785.
- [7] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas, "Decentralized active information acquisition: Theory and application to multi-robot slam," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4775–4782.
- [8] M. Kontitsis, E. A. Theodorou, and E. Todorov, "Multi-robot active slam with relative entropy optimization," in *American Control Conference (ACC)*, 2013, pp. 2757–2764.
- [9] M. Corah, C. O'Meadhra, K. Goel, and N. Michael, "Communication-efficient planning and mapping for multi-robot exploration in large environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1715–1721, 2019.
- [10] J. Wang and B. Englot, "Autonomous exploration with expectation-maximization," in *Robotics Research: The 18th International Symposium (ISRR)*. Springer, 2017, pp. 759–774.
- [11] B. Charrow, V. Kumar, and N. Michael, "Approximate representations for multi-robot control policies that maximize mutual information," *Autonomous Robots*, vol. 37, pp. 383–400, 2014.
- [12] R. G. Colares and L. Chaimowicz, "The next frontier: Combining information gain and distance cost for decentralized multi-robot exploration," in *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, 2016, pp. 268–274.
- [13] T. Regev and V. Indelman, "Multi-robot decentralized belief space planning in unknown environments via efficient re-evaluation of impacted paths," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 5591–5598.
- [14] B. Schlotfeldt, D. Thakur, N. Atanasov, V. Kumar, and G. J. Pappas, "Anytime planning for decentralized multirobot active information gathering," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 1025–1032, 2018.
- [15] S. Chen, L. Zhao, S. Huang, and Q. Hao, "Multi-robot active slam based on submap-joining for feature-based representation environments," in *Australasian Conference on Robotics and Automation*, 2022.
- [16] M. Lauri, E. Heinänen, and S. Frintrop, "Multi-robot active information gathering with periodic communication," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 851–856.
- [17] L. Freda, M. Gianni, F. Pirri, A. Gawel, R. Dubé, R. Siegwart, and C. Cadena, "3d multi-robot patrolling with a two-level coordination strategy," *Autonomous Robots*, vol. 43, pp. 1747–1779, 2019.
- [18] Y. Kantaros, B. Schlotfeldt, N. Atanasov, and G. J. Pappas, "Asymptotically optimal planning for non-myopic multi-robot information gathering," in *Robotics: Science and Systems*, 2019, pp. 22–26.
- [19] M. F. Ginting, K. Otsu, J. A. Edlund, J. Gao, and A.-A. Agha-Mohammadi, "Chord: Distributed data-sharing via hybrid ros 1 and 2 for multi-robot exploration of large-scale complex environments," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5064–5071, 2021.
- [20] K. Ye, S. Dong, Q. Fan, H. Wang, L. Yi, F. Xia, J. Wang, and B. Chen, "Multi-robot active mapping via neural bipartite graph matching," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 14 839–14 848.
- [21] M. Tzes, N. Bousias, E. Chatzipantazis, and G. J. Pappas, "Graph neural networks for multi-robot active information acquisition," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 3497–3503.
- [22] M. Ossenkopf, G. Castro, F. Pessacq, K. Geihs, and P. De Cristóforis, "Long-horizon active slam system for multi-agent coordinated exploration," in *European Conference on Mobile Robots (ECMR)*, 2019, pp. 1–6.
- [23] V. Indelman, "Cooperative multi-robot belief space planning for autonomous navigation in unknown environments," *Autonomous Robots*, vol. 42, pp. 353–373, 2018.
- [24] Y. Chen, L. Zhao, K. M. B. Lee, C. Yoo, S. Huang, and R. Fitch, "Broadcast your weaknesses: Cooperative active pose-graph slam for multiple robots," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2200–2207, 2020.
- [25] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, "isam2: Incremental smoothing and mapping using the bayes tree," *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, 2012.
- [26] J. Wang, F. Chen, Y. Huang, J. McConnell, T. Shan, and B. Englot, "Virtual maps for autonomous exploration of cluttered underwater environments," *IEEE Journal of Oceanic Engineering*, vol. 47, no. 4, pp. 916–935, 2022.
- [27] S. Thrun, "Probabilistic robotics," *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002.
- [28] M. Kaess and F. Dellaert, "Covariance recovery from a square root information matrix for data association," *Robotics and autonomous systems*, vol. 57, no. 12, pp. 1198–1210, 2009.
- [29] X. Fan, Y. Guo, H. Liu, B. Wei, and W. Lyu, "Improved artificial potential field method applied for AUV path planning," *Mathematical Problems in Engineering*, vol. 2020, pp. 1–21, 2020.
- [30] J. McConnell, Y. Huang, P. Szenher, I. Collado-Gonzalez, and B. Englot, "DRACo-SLAM: Distributed robust acoustic communication-efficient SLAM for imaging sonar equipped underwater robot teams," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 8457–8464.