

# Communication-Constrained Multi-Robot Exploration with Intermittent Rendezvous

Alysson Ribeiro da Silva<sup>1</sup>, Luiz Chaimowicz<sup>1</sup>, Thales C. Silva<sup>2</sup>, and M. Ani Hsieh<sup>2</sup>

**Abstract**—Communication constraints can significantly impact robots’ ability to share information, coordinate their movements, and synchronize their actions, thus limiting coordination in Multi-Robot Exploration (MRE) applications. In this work, we address these challenges by modeling the MRE application as a DEC-POMDP and designing a joint policy that follows a rendezvous plan. This policy allows robots to explore unknown environments while intermittently sharing maps opportunistically or at rendezvous locations without being constrained by joint path optimizations. To generate the rendezvous plan, robots represent the MRE task as an instance of the Job Shop Scheduling Problem (JSSP) and minimize JSSP metrics. They aim to reduce waiting times and increase connectivity, which correlates to the DEC-POMDP rewards and time to complete the task. Our simulation results suggest that our method is more efficient than using relays or maintaining intermittent communication with a base station, being a suitable approach for Multi-Robot Exploration. We developed a proof-of-concept using the Robot Operating System (ROS) that is available at: <https://github.com/multirobotplayground/Noetic-Multi-Robot-Sandbox>.

## I. INTRODUCTION

The exploration of unknown environments by a team of robots is an important task for applications like search and rescue, environmental monitoring and patrolling, and mapping of unknown environments. Robots must define goals, navigate, and collect and share information in a distributed manner [1] to accomplish this efficiently. However, multi-robot exploration (MRE) requires coordination among team members and is often conducted in geometrically complex environments where robots are subject to constraints on their sensing, mobility, and communication capabilities.

There is a large body of work in the multi-robot coordination domain where the objective is to maintain a connected communication network at all times as the team executes various tasks in both known [2]–[4] and unknown environments [5]–[7]. However, keeping the robots connected can restrict exploration, especially when team sizes and communication ranges are small, leading to increased exploration times. Differently, intermittent communication strategies for teams of robots have been employed in various application settings [8]–[10]. These works suggest that MRE strategies that employ intermittent communication can be more performative since they allow the team to cover more space, possibly in less time [1]. To allow intermittent rendezvous or encounters, the

This work was supported by ARL DCIST CRA W911NF-17-2-0181. This work was also partially supported by CAPES, CNPq, and Fapemig.

Alysson Ribeiro da Silva and Luiz Chaimowicz are with the VeR-Lab, Universidade Federal de Minas Gerais, Brazil. Thales C. Silva and Ani Hsieh are with the GRASP Laboratory, University of Pennsylvania, USA. {alysson.silva, chaimo}@dcc.ufmg.br, {scthales, mya}@seas.upenn.edu.



Fig. 1: Robots meeting at rendezvous locations spread over a construction site. The R1, R2, and R3 are hypothetical rendezvous locations. Colored arrows indicate the rendezvous locations. Orange circles are places to explore and dashed lines represent the paths followed by robots.

authors in [8] and [10] propose formulating a navigation plan through joint-path planning using connectivity graphs [11]. Similarly, Kantaros et al. [9] propose an intermittent communication method that divides robots into sub-teams, and members in different sub-teams must rendezvous periodically following a predetermined schedule. In their method, robots update their paths given their next encounter to satisfy a global *Linear Temporal Logic* (LTL) formula every time they meet. Unfortunately, these approaches are infeasible for more general MRE applications in which the environment is unknown and paths must adapt as robots explore.

In this work, we propose a different approach where we model the MRE task as a Decentralized Partially Observable Markov Decision Process (DEC-POMDP) [12]. The objective is to design a policy considering intermittent communication that maximizes its reward related to exploring new places and sharing maps. Due to the large policy search space of the MRE task, we propose an induced joint policy composed of robot controllers. Each controller allows a robot to select promising places to explore and, at the same time, attend to agreed rendezvous locations scheduled in a rendezvous plan (e.g., Fig. 1). To ensure that robots meet the time constraints imposed by data collection and exploration sub-tasks, we propose representing the rendezvous plan as an instance of the Job Shop Scheduling Problem (JSSP) [13]–[15]. A significant advantage of employing instances of the JSSP is the ability for robots to coordinate tasks and estimate the resources needed to complete them while minimizing the time they wait at rendezvous locations due to navigation and exploration uncertainties.

Our contributions can be summarized as follows:

- We represent the MRE task as an instance of the JSSP to enable the generation of a rendezvous plan

for intermittent communication. The mapping of the MRE problem into a JSSP enables the minimization of wait times at rendezvous locations, maximizing the returns of the DEC-POMDP, and results in rendezvous encounters that can be updated during exploration.

- We propose a decentralized joint policy for the MRE application based on a DEC-POMDP, where each robot's controller aims to maximize the utility of visiting frontiers while considering the rendezvous plan to share maps. Unlike existing methods, our approach enables robots to explore freely and attend to rendezvous locations only when necessary, eliminating the necessity for joint-path optimization among team members.

We evaluated our method considering that robots have a short communication range and considered scenarios where the workspace's length scale is large to avoid incidental all-time connectivity when exploring. Due to the vast MRE literature, we design two baselines, based on [1] and [16], using a distributed frontier-based method that delivers information to a base station and maintains a relay network. We measured the time taken to complete the exploration and the rewards of our method's DEC-POMDP against the baselines in our simulator in different urban scenarios obtained from satellite images. Our results suggest that our method is able to explore faster and obtain higher rewards for the underlying DEC-POMDP without a base station or a relay network. In addition, we provide a proof of concept deployment using Gazebo Simulator and the Robot Operating System (ROS).

## II. BACKGROUND AND PROBLEM STATEMENT

We define a set of  $n$  robots as  $\mathcal{R}$ , where each one explores an unknown bounded environment  $\Omega$ . A robot  $i$  perceives the environment, builds a local map, estimates its pose, and constructs the frontier set at each time step  $t$ , where frontiers are the boundary between known and unknown regions. It then explores by maintaining its current state  $s_i(t)$ , represented by a tuple with the relative poses of nearby robots, its local map, and its pose. An action  $a_i(t) \in \mathcal{A}_i(t)$  consists in navigating towards a location in  $\Omega$ . For the MRE application, an action location refers to a frontier or a rendezvous location (*i.e.*, point of encounter between two or more robots). We define the utility of an action as  $u_i(a_i(t))$ , which reflects the benefit of navigating towards the location specified by  $a_i(t)$  considering its path cost. To coordinate exploration among multiple robots, a task allocation algorithm [1] helps selecting frontiers that maximize the sum of utilities.

We model the connectivity among robots at a time step  $t$  as a graph  $\mathcal{G}(t) = (\mathcal{R}, \mathcal{E}(t))$ , where an edge  $e(t) = (i, j)$  from  $\mathcal{E}(t)$  represents the connectivity between robots  $i$  and  $j$ . We use a communication model based on a communication range. Therefore, at each time step  $t$ , an edge  $e(t)$  between robots  $i$  and  $j$  exists only if they are close to each other. Importantly, we also define a sub-team as a set of two or more robots  $\mathcal{T}_s$  that periodically meet. Robots are intermittently connected if the union of the graphs  $\{\mathcal{G}(t_a), \dots, \mathcal{G}(t_b)\}$  forms a connected graph for every  $t_b - t_a$  steps. When robots are connected, they share and merge their maps to enhance efficiency. Robots engage in exploration by asynchronously selecting actions

and executing them until there are no more frontiers left on their local map

The underlying DEC-POMDP of the MRE application is defined as the tuple  $(\mathcal{R}, \mathbb{S}, \mathbb{A}, T, \mathbb{O}, O, R, h)$ , where  $\mathbb{S}$  is a set of all possible joint states,  $\mathbb{A}$  is a set of joint actions,  $T$  is the set of conditional transition probabilities,  $\mathbb{O}$  is the set of joint observations,  $O$  is the set of observations probabilities,  $R$  is the immediate reward, and  $h$  is the horizon of the problem (*i.e.*, time to finish the exploration). For simplicity, we omitted the initial system's conditions. We define a joint state  $s(t) \in \mathbb{S}$  as  $(s_1(t), \dots, s_n(t))$  and a joint action  $a(t) \in \mathbb{A}$  as  $(a_1(t), \dots, a_n(t))$ , where  $\mathbb{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ . A joint observation  $o(t) \in \mathbb{O}$  is defined as  $o(t) = s(t)$  at a time step  $t$ , and its probability is  $P(o(t)|a(t), s(t+1))$ . The transition and observation probabilities are usually unknown for the MRE application, and robots can not measure the true state of the system due to localization and mapping uncertainties.

The immediate reward  $R : \mathbb{A} \times \mathbb{S} \rightarrow \mathbb{R}$  is function that associates actions and states to real numbers, and we define the *return* of the DEC-POMDP as a bounded sum  $\sum_{t=0}^h R(s(t), a(t))$  because  $\Omega$  has a limited number of frontiers. For the MRE application, we approximate the reward  $R(s(t), a(t))$  as the sum of the utility of the actions  $\sum_{i=1}^n u_i(a_i(t))$  from the joint action  $a(t)$ . In this regard, we define a *joint policy* as  $\pi = \{\pi_1, \dots, \pi_n\}$ , where  $\pi$  is decentralized, and each  $\pi_i$  is a finite state controller [17] that aims at maximizing the return by taking appropriate actions. From the perspective of a decentralized policy, robots are unaware of the joint state  $s(t)$ , joint action  $a(t)$ , joint observation  $o(t)$ , and the return is available only when searching for a policy and before exploration takes place.

*Problem 1 (MRE Problem with Intermittent Rendezvous):* Design a set of decentralized robot controllers  $\pi = \{\pi_1, \dots, \pi_n\}$ , where each controller  $\pi_i$  directs robot  $i$  to explore frontiers through joint actions  $a(t)$ . Ensure that the union of the connectivity graphs  $\{\mathcal{G}(t_a), \dots, \mathcal{G}(t_b)\}$  is connected for every  $t_b - t_a$  steps until there are no more frontiers to explore. The objective is to maximize the return of the DEC-POMDP, defined as  $\sum_{t=0}^h \sum_{i=1}^n u_i(a_i(t))$ , where  $a_i(t)$  represents the action of robot  $i$  in the joint action  $a(t)$ . At each time step, robots operate based on their states  $s_i(t)$  and are unaware of the joint state  $s(t)$ , joint actions  $a(t)$ , and joint observations  $o(t)$ , and the return is available only during the planning phase.

## III. METHODOLOGY

Our solution to Problem 1 has two main parts: an exploration *plan* that defines sub-team structures and resource allocation for the exploration, and a decentralized joint policy that controls navigation to rendezvous locations and frontiers. Before exploring, one robot from  $\mathcal{R}$  calculates an initial rendezvous plan for the team, considering an estimate about the area of  $\Omega$  and the robots' sensing and traveling speed. Upon receiving their plans, the robots begin exploration by executing their local policy, which selects actions based on the current state.

We define the plan through agreements between robots, where agreements specify which sub-team  $\mathcal{T}_s \subseteq \mathcal{R}$  of robots

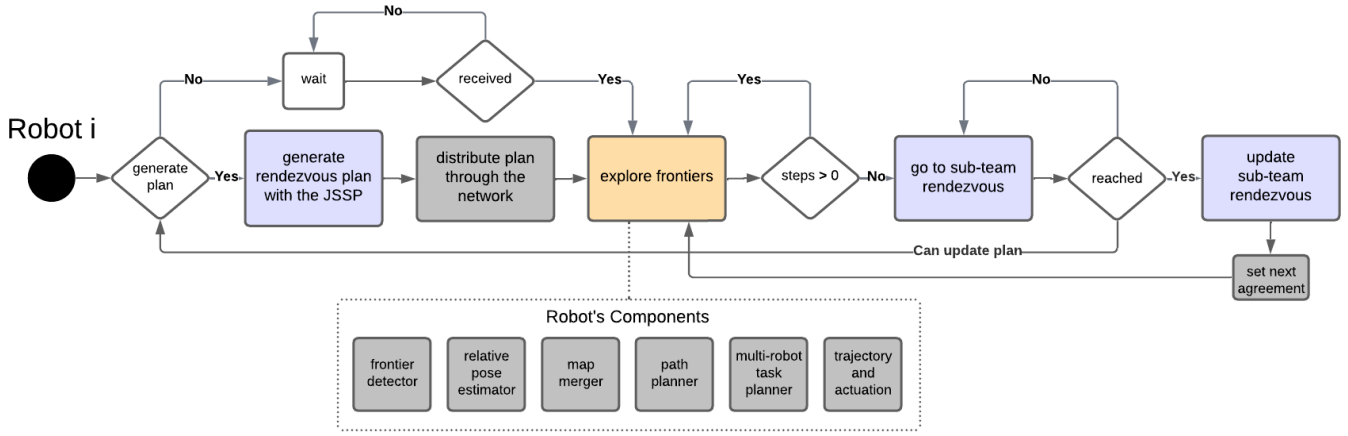



Fig. 2: Diagram for robot  $i$  controller integrated with the rendezvous plan generator and the robot's components. The purple components generate a rendezvous plan through a JSSP, decide when and with whom robots should meet, and update the rendezvous locations. Differently, the light orange component refers to the exploration task and it makes robots explore unknown places. The robot's components are used for deployment and help assemble the robot's state and execute actions.



$$K = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix} \quad W = \begin{bmatrix} 50 & 50 & 0 \\ 0 & 50 & 100 \\ 90 & 0 & 40 \\ 50 & 100 & 0 \end{bmatrix}$$

Fig. 3: Hypothetical agreements between 3 robots (red, green, blue). Robot 1 participates in the 1<sup>st</sup>, 3<sup>rd</sup>, and 4<sup>th</sup> agreements, while robot 2 participates in the 1<sup>st</sup>, 2<sup>nd</sup>, and 4<sup>th</sup> ones. In the 1<sup>st</sup> agreement, robots 1 and 2 should explore for 50 time steps before going to a meeting location to share maps.

should meet and how much each robot must explore before going to an assigned rendezvous location. We use time steps as a metric of how much resources a robot should spend before meeting with a sub-team. To fulfill an agreement, all robots from its sub-team must meet at its rendezvous location and share maps. Robots update their sub-team rendezvous location and, for simplicity, they can update the plan if the encounter is of a sub-team composed by all robots.

Our solution is illustrated in Fig. 2. In summary, the robots will first create an exploration plan. They will then follow their local policy to explore the frontiers based on predefined agreements, which specify the time they should spend exploring and which sub-team they should meet next. Finally, the robots will navigate to one of their sub-teams rendezvous locations in a specific order, share maps, define a new rendezvous location for the current sub-team, and resume exploring. We explain the specific details of the process in the following sections and provide a pseudo-code in Algorithm 2. In particular, in Section III-A and III-B, we present our method to generate the plans, and in Section III-D, we present the policy for exploration.

#### A. Representing the Rendezvous Plan

We propose an Agreement matrix  $K$  and a Steps matrix  $W$  to represent the plan efficiently. Each element of  $K$  is a binary number, and each element of  $W$  is a positive real. The rows of both matrices holds information about the agreements, and columns represent robots' index. For example, the following  $K$  and  $W$  instances represent the agreements from Fig. 3,

where, the first row of  $K$  tells that the first rendezvous of robot 1 should be with robot 2, whereas the second row tells that the first rendezvous of robot 3 should be with robot 2. Since agreements are fulfilled in order, this example implies that robot 3 is free to explore while robots 1 and 2 fulfill their first agreement. The  $W$  matrix tells how many time steps each robot should explore before fulfilling an agreement from  $K$ . For instance, the first row of  $W$  tells that robots 1 and 2 must explore for 50 time steps before meeting. The structure of matrices  $K$  and  $W$  also inform us the number of sub-teams in  $\mathcal{R}$ . In this example there are three sub-teams, where  $\mathcal{T}_1 = \{1, 2\}$ ,  $\mathcal{T}_2 = \{2, 3\}$ , and  $\mathcal{T}_3 = \{1, 3\}$  the rendezvous locations are associated with the sub-teams, that can meet at different moments and locations.

#### B. Generating $K$ and $W$

While  $K$  and  $W$  serve as compact representations of agreements, it is crucial to define plans that effectively coordinate the robots' activities, balancing exploration and travel to rendezvous locations. This coordination minimizes waiting times and prevents resource waste. Consequently, to generate the matrices, we propose transforming  $K$  and  $W$  into a JSSP instance  $\mathcal{I}$ , changing the instance parameters by optimizing metrics such as the *Makespan* (i.e., the total time to finish a set of sub-tasks) and obtaining the optimized  $K$  and  $W$  from  $\mathcal{I}$ . A JSSP instance  $\mathcal{I}$  represents the execution of the plan through time, and it is composed of machines (robots), where each machine has a list of jobs to execute. It is important to note that jobs are sub-tasks of exploring the environment associated with starting and ending times. In this regard, robots explore the environment when executing jobs and fulfill one agreement for each job they finish.

In this work, we particularly designed a *Genetic Algorithm* [18] solver with a population of  $(K, W)$  pairs, initialized

---

**Algorithm 1:** Job-shop Scheduling Transformation
 

---

**Input:** Number of robots  $|\mathcal{R}|$   
 Rendezvous agreement matrix  $K$   
 Rendezvous exploration steps matrix  $W$

**Output:** JSSP instance  $\mathcal{I}$

```

1 function JSSP( $|\mathcal{R}|, K, W$ )
2   Let  $\mathcal{I}$  be a JSSP instance with  $|\mathcal{R}|$  machines;
3   foreach agreement row in  $K$  do
4     foreach robot column in  $K[\text{agreement}]$  do
5       if  $K[\text{agreement}][\text{robot}] = 1$  then
6         create a new job  $job_{new}$  and allocate it to the
           $j$ th machine in  $\mathcal{I}$ ;
7         set the length of  $job_{new}$  as
           $W[\text{agreement}][\text{robot}]$  and update its starting
          and ending times according to what was
          already allocated to the  $j$ th machine;
8         mark  $job_{new}$  as belonging to  $\text{agreement}$  to
          help recovering  $K$  and  $W$  through the inverse
          process;
9   return  $\mathcal{I}$ ;
```

---

randomly, to generate the plan, because it is not trivial to find gradients to iterate the plan over the fitness landscape. The solver evaluates the population by converting each pair into a JSSP instance  $\mathcal{I}$  and extracting metrics we present in Section III-C. Algorithm 1 shows the transformation of each individual  $(K, W)$  into  $\mathcal{I}$ , which creates jobs by iterating over rows of  $K$ . At the same time, it uses the values from the steps matrix  $W$  and the length of the jobs already inserted into  $\mathcal{I}$  to adjust their starting and ending times. The algorithm marks jobs from the same agreement at line 8 to help recover  $K$  and  $W$  after the optimization.

### C. JSSP Instance Metrics

We designed the following objective function to evaluate a JSSP instance  $\mathcal{I}$  obtained from a  $(K, W)$  pair,

$$z(\mathcal{I}) = \sum_a \alpha_a h_a(\mathcal{I}) + \sum_b \beta_b g_b(\mathcal{I}) \quad (1)$$

where,  $h_a$  and  $g_b$  are metrics obtained from  $\mathcal{I}$ ,  $\alpha_a$  and  $\beta_b$  are weights. The intuition behind  $z(\mathcal{I})$  is to minimize the Makespan and waiting times during rendezvous encounters and it is represented by the following metrics:

- **Plan Error:**  $\sum \alpha_a h_a(\mathcal{I})$  is used to generate plans considering an estimate of how big is  $\Omega$  with minimum waiting times, a fulfillment frequency, and the Makespan.
- **Connectivity Metrics:**  $\sum \beta_b g_b(\mathcal{I})$  is used to check if the robots are going to be intermittently connected if they follow the plan.

1) *Plan Error:* From the JSSP instance  $\mathcal{I}$ , the algorithm calculates the exploration estimate error,

$$h_1(\mathcal{I}) = ((\Omega_A - \Omega_E) - \sum_{job_j \in \mathcal{J}_{\mathcal{I}}} job_j \mathcal{X})^2 \quad (2)$$

where,  $(\Omega_A - \Omega_E)$  is an estimate of how much area is left to explore, given its estimated size  $\Omega_A$ , and the area robots already explored  $\Omega_E$ ,  $job_j$  is the length of the  $j$ th job, and  $\mathcal{X}$  is a constant that estimates how much area robots can explore, given the sensor they use to acquire information from  $\Omega$  (e.g., lidar, sonar, camera) and traveling speed; the work done error,

$$h_2(\mathcal{I}) = \left( \sum_{job_j \in \mathcal{J}_{\mathcal{I}}} job_j - \sum_{job_j \in \mathcal{J}_{\mathcal{I}}} (job_j + \mathcal{Y}(job_j)) \right)^2 \quad (3)$$

where,  $job_j$  is the job length and  $\mathcal{Y}(job_j)$  is the number of time steps that should be added to reduce waiting times at rendezvous locations; and the deviation between rendezvous encounters,

$$h_3(\mathcal{I}) = \sqrt{\frac{1}{|\mathcal{J}_{\mathcal{I}}|} \sum_{job_j \in \mathcal{J}} (job_j - \bar{job})^2} \quad (4)$$

where,  $job_j$  is the job length,  $|\mathcal{J}_{\mathcal{I}}|$  is the number of jobs from  $\mathcal{I}$ , and  $\bar{job}$  is the average job length. Importantly, we estimate the size of the environment  $\Omega_A$  as the area of the bounding rectangle that encloses the entire mission zone.

2) *Connectivity Metrics:* Let  $\mathcal{H}_{\mathcal{I}}$  be a set of connectivity graphs obtained from a JSSP instance  $\mathcal{I}$ . We propose creating a new graph  $\mathcal{G}_{dense}$ , which is a condensed version of all graphs in  $\mathcal{H}_{\mathcal{I}}$  to calculate the connectivity metrics. Each vertex from  $\mathcal{G}_{dense}$  is a robot, and an edge exists between robots  $i$  and  $j$  from  $\mathcal{G}_{dense}$  if any of the connectivity graphs from  $\mathcal{H}(\mathcal{I})$  connects them. From  $\mathcal{G}_{dense}$ , we propose extracting the connectivity index,

$$g_1(\mathcal{I}) = 1 - \frac{1}{1 + \mathcal{P}(\mathcal{G}_{dense})} \quad (5)$$

where,  $\mathcal{P}(\mathcal{G}_{dense})$  is a function that returns the length of the biggest connected component of  $\mathcal{G}_{dense}$ . We also propose extracting,

$$g_2(\mathcal{I}) = |\mathcal{E}'| \quad (6)$$

where,  $|\mathcal{E}'|$  is the number of edges from  $\mathcal{G}_{dense}$ . The intuition behind the number of edges is to avoid over-connectivity.

During the planning phase, it is feasible to estimate the DEC-POMDP's return from the JSSP instance. This estimation relies on various factors derived from the JSSP, including the time allocated for robot exploration, rendezvous scheduling, and the extent of information exchange during map sharing from eq. 2. The optimization process in this context correlates to the horizon of the DEC-POMDP. Specifically, minimizing the Makespan implies in a more condensed exploration timeline, leading to a shorter planning horizon within the DEC-POMDP framework.

### D. Exploration Policy with Intermittent Rendezvous

After receiving a rendezvous plan (matrices  $K$  and  $W$ ), each robot  $i$  explores the environment with its policy  $\pi_i$ . Robots capture information from the environment  $\Omega$  with their sensors and receive observations  $o_i(t)$  that they use within their current state. Based on their state, they decide which action to take. An action can be exploring the environment or fulfilling an agreement from  $K$ . Importantly, when navigating, robots constantly try to share maps and their relative poses with others within communication range.

Algorithm 2 shows how robots gather information, assemble their state, and select their actions. At each time step, a

**Algorithm 2:** Exploration with Intermittent Rendezvous

---

**Input:** Agreement matrix  $K$   
Exploration steps matrix  $W$

```

1 function ExplorationIntermittentRendezvous( $K, W$ )
2   Let  $\mathcal{T}_s$  be the sub-team of the first agreement to fulfill in  $K$ ;
3   Let  $steps$  be the number of steps to explore in  $W$  from the
   agreement of  $\mathcal{T}_s$ ;
4   foreach time step  $t$  do
5     Receive relative poses of other robots and their maps if
     they are within communication range, update local map
     and merge with the ones received, and estimate my pose;
6      $\triangleright s_i(t)$  is defined as a tuple with the relative poses, the
     current map, and the robot's pose.
7     Detect the frontiers set  $\mathcal{F}_i(t)$  from  $s_i(t)$ ;
8     Assemble the connectivity graph  $\mathcal{G}_i(t)$  from  $s_i(t)$ ;
9     if The frontier set  $\mathcal{F}_i(t)$  is empty then
10      break;
11     if  $steps > 0$  then // Exploration Actions
12       if Not selected an action then
13         Select an exploration action  $a_i(t)$  with a greedy
         action selection mechanism, considering the
         single and multi-robot cases;
14       else
15         Navigate towards  $a_i(t)$  location;
16          $steps = steps - 1$ 
17       else // Agreement Fulfillment Action
18         Navigate to the sub-team  $\mathcal{T}_s$  rendezvous location;
19         if I am in charge of the update and
20          $\mathcal{G}_i(t)$  connects the sub-team  $\mathcal{T}_s$  then
21           Calculate the centroids set  $\mathcal{C}(t)$  with the KNN
           from the frontiers in  $\mathcal{F}_i(t)$ ;
22           if  $\mathcal{T}_s \neq \mathcal{R}$  then // Location update
23             Pick a new rendezvous location with a
             random sampling method from  $\mathcal{C}(t)$  and
             send it to all robots in the sub-team  $\mathcal{T}_s$ ;
24           else // Synchronization agreement
25             Generate the new rendezvous matrices  $K$ 
             and  $W$ ;
26             Pick centroids with maximum distance from
             each other from  $\mathcal{C}(t)$ , one for each
             sub-team in the new  $K$ ;
27             Send the new matrices and the sub-team
             rendezvous locations to all robots in  $\mathcal{T}_s$ ;
28           if received sub-team rendezvous locations or
           received plan or timeout then
29             Set  $\mathcal{T}_s$  to the sub-team specified by the next
             agreement this robot participates in  $K$ ;
30             Update  $steps$  with the nex value from  $W$  given
             the next agreement this robot participates in  $K$ ;
31         Go back to the initial location;

```

---

robot  $i$  gathers information from the environment to build and maintain a local map and estimates its pose. This process involves integrating the observations into the local map and correcting its pose based on sensory input. It also receives maps, that are merged with its local map, and the relative poses from robots within its communication range. Finally, a robot represents its state  $s_i(t)$  as a tuple containing the relative poses of nearby robots, its local map, and its pose. Through its state, it calculates the set of frontiers  $\mathcal{F}_i(t)$  with a flooding-fill algorithm and the connectivity graph  $\mathcal{G}_i(t)$ . Both the frontiers and the connectivity graph are used to navigate and verify the requirements to fulfill agreements.

The policy keeps track of the current state and which agreement to fulfill to select an appropriate action. The action space  $\mathcal{A}_i(t)$  at a time step  $t$  of a robot is composed of

two types of actions: I) *Exploration Actions*. II) *Agreement Fulfillment Action*.

- **Exploration Actions:** An Exploration Action refers to exploring a frontier, calculating and updating a path towards it given its merged map, and following this path. Robots perform Exploration Actions until they explore for their assigned time to fulfill their current agreement.
- **Agreement Fulfillment Action:** Differently, an Agreement Fulfillment Action refers to navigating towards the current rendezvous location, waiting for the assigned sub-team or until a timeout, sharing maps, and selecting the next agreement to fulfill. Robots execute the Agreement Fulfillment Action if they have explored the number of time steps in  $W$  specified by the current agreement. After fulfilling the meeting, the assigned sub-team updates its rendezvous location and can update the plan if necessary.

Robots decide which type of action to take by verifying the variable  $steps$  at line 11 of Algorithm 2. They also track which sub-team they should meet when performing an Agreement Fulfillment Action with the variable  $\mathcal{T}_s$ . After executing an action, a robot selects the next agreement it participates in by updating both variables.

1) *Exploration Actions:* For each frontier in  $\mathcal{F}_i(t)$ , there is an Exploration Action in  $\mathcal{A}_i(t)$ . Basically, a robot estimates the utility of an action as the utility of visiting a frontier  $f_j$  from the frontiers set  $\mathcal{F}_i(t)$ . It selects an Exploration Action by iterating over all frontiers from  $\mathcal{F}_i(t)$  seeking for the one that maximizes the following utility function,

$$u_i(f_j) = \frac{\mathcal{N}_{f_j}}{\Phi(f_j)} \quad (7)$$

where,  $\mathcal{N}_{f_j}$  is the area to be explored when navigating towards  $f_j$ , and  $\Phi(f_j)$  is its path cost.

The single and multi-robot action selection cases work as follows:

- **Single Robot:** If a robot can not communicate with others, then it is going to select an action  $a_i(t)$  associated with  $f_j^* = \underset{\forall f_j \in \mathcal{F}_i(t)}{\operatorname{argmax}} u_i(f_j)$ .
- **Multi-robot:** When within communication range, robots from a sub-team have the same merged-map and can select which frontier to explore. In this case, one of the robots selects a joint action using a greedy action selection heuristic that aims to locally maximize the sum of the utilities  $u_i(a_i(t))$  for each robot from the sub-team. This greedy selection also maximizes the DEC-POMDP reward for a joint action  $a(t)$ .

2) *Agreement Fulfillment Action:* A robot tries to fulfill its current agreement after it has explored the assigned number of steps from the plan. In this regard, it stops executing a previously selected Exploration Action, calculates a path towards the current agreement sub-team  $\mathcal{T}_s$  rendezvous location, and follows it, as shown at line 18 of Algorithm 2. After reaching the rendezvous location, robots share their maps and update the location associated with their sub-team. To select a new rendezvous location, one of the robots from the sub-team creates the locations set  $\mathcal{C}(t)$  by clustering frontiers from  $\mathcal{F}_i(t)$  with a *K-Nearest Neighbors* (KNN).

	NY	DC	Tokyo	Beijing	Dubai	Moscow	Berlin	London	Paris	Rio
BS	6992 ± 648	11849 ± 301	10120 ± 719	11538 ± 1279	14643 ± 1273	14613 ± 1117	11973 ± 1153	11818 ± 1354	11399 ± 735	13871 ± 1213
CRN1R	6060 ± 377	11540 ± 1277	9579 ± 856	11209 ± 1126	13813 ± 1621	13634 ± 1361	10798 ± 1005	11109 ± 821	10826 ± 829	12475 ± 1582
CRN2R	6398 ± 782	10889 ± 549	9326 ± 614	10341 ± 980	13167 ± 1888	12989 ± 830	11729 ± 684	10558 ± 827	10359 ± 706	11140 ± 481
Ours	4607 ± 108	8048 ± 188	7056 ± 205	7174 ± 309	8061 ± 249	8297 ± 183	7632 ± 204	7760 ± 345	7436 ± 130	8073 ± 282
OursSync	<b>4436 ± 117</b>	<b>7181 ± 218</b>	<b>6500 ± 399</b>	<b>6459 ± 274</b>	<b>7700 ± 413</b>	<b>7870 ± 316</b>	<b>6932 ± 117</b>	<b>6996 ± 116</b>	<b>7084 ± 170</b>	<b>7460 ± 295</b>
	prk01	prk02	prk03	air01	air02	air03	air04	cons01	cons02	cons03
BS	10078 ± 881	2980 ± 285	7967 ± 744	5005 ± 656	10053 ± 1740	6460 ± 547	6710 ± 718	10822 ± 1690	7003 ± 495	7186 ± 496
CRN1R	9000 ± 1098	3016 ± 235	7398 ± 597	4582 ± 329	10270 ± 1784	6306 ± 469	6120 ± 601	9262 ± 622	6379 ± 400	6379 ± 400
CRN2R	9172 ± 802	2886 ± 165	6778 ± 512	4634 ± 497	9039 ± 868	6230 ± 661	6283 ± 682	8764 ± 644	6451 ± 302	6388 ± 495
Ours	5833 ± 251	2242 ± 49	5054 ± 182	<b>4310 ± 294</b>	6985 ± 286	5197 ± 618	5427 ± 361	7034 ± 338	4954 ± 184	<b>5343 ± 216</b>
OursSync	<b>5458 ± 222</b>	<b>2226 ± 63</b>	<b>4895 ± 265</b>	4436 ± 215	<b>6538 ± 378</b>	<b>5023 ± 416</b>	<b>5263 ± 307</b>	<b>6553 ± 293</b>	<b>4717 ± 123</b>	5449 ± 260

TABLE I: Average number of time steps to finish the exploration for our dataset composed of cities, parking lots (prk), airport sections (air), and constructions sites (cons).



Fig. 4: Heuristic to pick new rendezvous locations when robots have common information during synchronization for 3 sub-teams from the example from Fig. 3 in a construction site. The orange circles are frontiers, and the blue circles are the center of masses obtained from the KNN algorithm. The center of masses  $c_1$  and  $c_3$  were chosen to the 1<sup>st</sup> and 2<sup>nd</sup> sub-teams. Differently,  $c_2$  was chosen to the 3<sup>rd</sup> sub-team by maximizing the distance  $\sqrt{(c_1 - c_2)^2 + (c_2 - c_3)^2}$ .

Next, it randomly samples a location from  $\mathcal{C}(t)$  and sends it to all other robots in the sub-team. Finally, the robots from the sub-team select their next agreement and continue with subsequent exploration.

Conversely, to update the rendezvous plan, robots must have common information. Otherwise, they might generate agreements that can not be fulfilled. They achieve this by fulfilling an agreement in which the sub-team has all robots (*i.e.*,  $\mathcal{T}_s = \mathcal{R}$ ) which we call a *synchronization agreement*. One robot updates the plan through a synchronization agreement at lines 25 – 27 of Algorithm 2 after a predetermined number of time steps. They define the rendezvous location of this agreement as the center of mass of all frontiers they had in common during the last synchronization.

The robot updating the plan generates new values for  $K$  and  $W$ , which are transmitted to the other robots. In particular, it exploits the fact that they have common information and assign new rendezvous locations to all sub-teams of the new plan. As shown in Fig. 4, it selects the locations from  $\mathcal{C}(t)$  with the maximum distance from each other as the new sub-teams rendezvous locations. This encounter is less harmful for exploration than having all-time connectivity or a base station because communication is majorly intermittent in the plan, which we show through our experiments.

#### IV. BENCHMARK SIMULATION RESULTS

To evaluate our method, we assembled a dataset composed of 20 urban environments from satellite images from Google Maps (*e.g.*, Fig. 5) and prepared their cell decomposition in advance. We executed our approach in our simulator, where robots execute the exploration task in a completely distributed

fashion as shown in Fig. 2, plan their paths in the decomposed maps we prepared, and measured the total number of time steps to finish the exploration as done in previous MRE benchmarks [19]. In our simulator, robots account only for paths that are executed in discrete time steps.

Our experiments are: I) *Benchmark*: we benchmark our method against our baselines to evaluate its performance. II) *Consistency*: we evaluate the spread of the results on 6 instances of the dataset. III) *Larger Mission Areas*: we got an instance (London) from our dataset and doubled its dimensions to evaluate the impact of larger areas. IV) *Return*: we evaluated the return of the DEC-POMDP, as defined in Section II, in two instances of the dataset (cons01 and London). We measured the average explored area per time step, and its standard deviation between 10 runs with 4 robots for 50000 steps each, where they have 5 cells as visibility range and 4 cells of communication diameter, which resembles low-spec robots. At each run, we let the robots start as a cluster in the same location on the map, picked randomly between runs.

We designed three baselines inspired by [2] and [1] to compare against our method: a Constrained Static Base Station (BS), a Constrained Relay Network One Relay (CRN1R), and a Constrained Relay Network Two Relays (CRN2R) method. The base station in the BS and the relays have a communication range of 12 cells in the cell decomposition. In the CRN1R, each robot can place 1 extra relay after 2000 simulation steps at the communication boundaries with the base station or other relays, while robots in the CRN2R can place 2 relays. The comparisons were made against our method (Ours) without updating the plan and our method (OursSync), which performs roughly 4 synchronizations per mission. In all baselines, robots must deliver information to the base station or relay network at a fixed frequency.

1) *Benchmark*: In Table I, we show the average number of time steps robots took to complete the exploration for all 20 environments in our dataset. For city sections, our method took a maximum of 7870 time steps and a minimum of 4436 time steps to explore Moscow and New York City, respectively. Robots achieved this performance because Moscow is more structured, while New York has many possible routes that are close to each other, facilitating the mission. When operating at airports, parking lots, and construction sites, our method took a maximum of 6553 time steps at cons01 and a minimum of 2226 time steps at prk02, where cons01 is more complex. There are no meaningful differences between *Ours* and *OursSync*, which indicates the feasibility of the




Fig. 5: Samples from our benchmark and simulator. (From left to right) Sections of London, Dubai, Washington DC, Beijing, Rio de Janeiro, construction site, and Dallas International Airport.

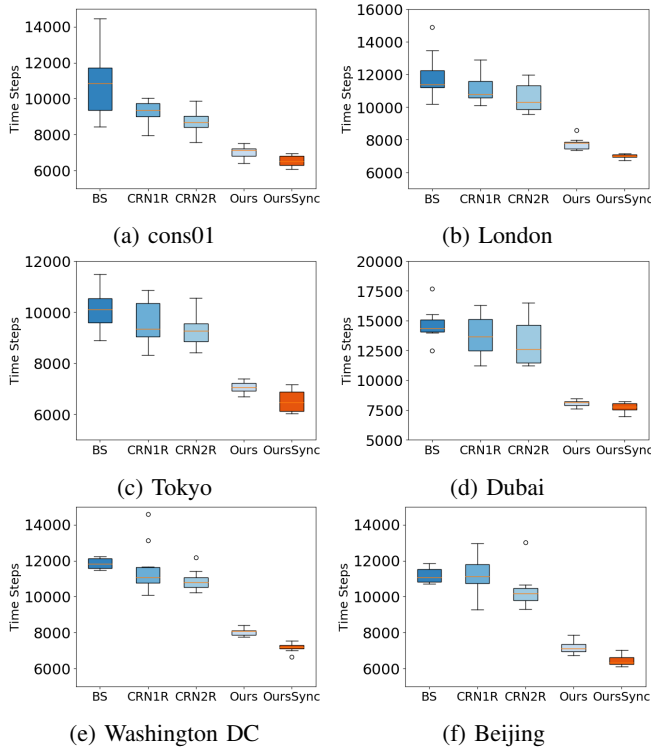


Fig. 6: Interquartile range of our baselines against our method. The Y-axis represents simulation steps.

synchronization without degrading the overall performance.

2) **Consistency:** We show the spread of the results for different scenarios (cons01, London, Tokyo, Dubai, Washington DC, and Beijing sections) in Fig. 6. The results suggest that our method is more stable with a smaller variation between runs. We can enhance the baselines by extending the base station range or adding more relays. However, extending this range towards infinity or keeping the robots connected might be impossible, and we expect that our method could achieve even better results for larger areas, as we present next.

3) **Larger Mission Area:** Fig. 7 compares our method against the CRN2R, which is the best performing baseline, in the London instance from our dataset. For this test, we doubled the area of the map in both dimensions. These results show that the other methods are constrained by their connectivity, because they have to deliver information to accomplish the mission. In contrast, our method obtained gains related to updating the rendezvous location and intermittently sharing maps across the workspace.

4) **Return:** Fig. 8 shows the average return of the DEC-POMDP for our methods in cons01 and London sections




Fig. 7: Percentage of the explored area for each simulation step. The X-axis represents simulation steps, and the Y-axis is the percentage of the area explored.




Fig. 8: Average return for the DEC-POMDP for each simulation step. The X-axis represents simulation steps, and the Y-axis is the return of the DEC-POMDP as defined in Section II. We captured the return from the baseline when robots explore frontiers or share maps with others and the relay network.

against the baseline CRN2R. Our method obtained average returns around 2000 for both scenarios, while the CRN2R obtained average returns below 1000. Our analysis indicates that our method can gather and share more information among robots about the workspace per unit of time while traveling less distances without maintaining a relay network.

## V. GAZEBO SIMULATION RESULTS

Our proof-of-concept addresses all aspects of the exploration mission using a mock communication model based on ROS network infrastructure as described by the robots' components in Fig. 2. We evaluated an exploration mission with 4 robots equipped with a 360-degree lidar for localization and mapping, and four RGB cameras short-range dynamic obstacle avoidance, each. The robots have a communication range of 3 meters. Each robot calculates its paths using a roadmap-based method, employs a local planner to handle exploration and navigation uncertainties, and has an average traveling speed of 0.5 m/s. For *Simultaneous Localization and Mapping* (SLAM), each robot uses the gmapping package to solve SLAM individually and further




Fig. 9: Rviz view of one robot (top) and the robots' trajectories (bottom). In the Rviz view, blue boxes represent frontiers and green boxes depict other robots. The trajectories plot below displays white circles indicating rendezvous locations, with colored lines representing the paths traversed by the robots.

shares their maps when appropriate. Unlike the benchmark experiments, robots receive predetermined  $K$  and  $W$  matrices in the proof of concept, which is not detrimental to the purpose of evaluating the agreement execution subsystems integrated with a decentralized multi-robot exploration stack. We evaluated the policy in two different trials and collected the robots' trajectories. The exploration stack source code is available at: <https://github.com/multirobotplayground/Noetic-Multi-Robot-Sandbox>.

Fig. 9 shows the execution of the method in two runs. We observed that our method can be deployed in a realistic multi-robot exploration stack with minor changes to existing controllers. With robust path and trajectory planners, the robots were able to fully explore the environment in the proposed testing area, executing agreements during the exploration process. Furthermore, we observed the necessity of having a communication protocol to synchronize rendezvous events. Despite this, our system was able to accomplish all posed challenges, highlighting the adaptability of our method in realistic scenarios. It successfully achieved the mission objectives of exploring the environment and maintaining intermittent connectivity under navigation and exploration uncertainties with communication constraints.

## VI. CONCLUSION

This paper proposes a multi-robot frontier-based exploration method with intermittent rendezvous. We model the problem as a DEC-POMDP to maximize the utility of visiting frontiers while maintaining intermittent communication across the assigned mission area. In our method, robots perform rendezvous through agreements at locations spread dynamically through the environment. We provide the means to generate the rendezvous plan automatically by transforming the MRE problem into an instance of the JSSP and optimizing it, given its metrics. Our results suggest that our method can explore faster with higher DEC-POMDP rewards without a base station or a relay network. Further research involves

explicitly incorporating uncertainties into the JSSP to reduce the time robots wait at rendezvous locations, incorporating a heterogeneous team of robots, and considering robots with energetic constraints and trajectories.

## REFERENCES

- [1] L. Bramblett, R. Peddi, and N. Bezzo, "Coordinated multi-agent exploration, rendezvous, task allocation in unknown environments with limited connectivity," in *2022 IEEE/RSJ IROS*, 2022, pp. 12 706–12 712.
- [2] E. A. Jensen, "Restricted communication in online multi-robot exploration," in *Proceedings of the Twenty-Seventh IJCAI-18. International Joint Conferences on Artificial Intelligence Organization*, 7 2018, pp. 5767–5768. [Online]. Available: <https://doi.org/10.24963/ijcai.2018/827>
- [3] P. Ong, B. Capelli, L. Sabattini, and J. Cortés, "Network connectivity maintenance via nonsmooth control barrier functions," in *2021 60th IEEE CDC*, 2021, pp. 4786–4791.
- [4] P. Ong, B. Capelli, L. Sabattini, and J. Cortés, "Nonsmooth control barrier function design of continuous constraints for network connectivity maintenance," *Automatica*, vol. 156, no. C, oct 2023. [Online]. Available: <https://doi.org/10.1016/j.automatica.2023.111209>
- [5] F. Pratisoli, B. Capelli, and L. Sabattini, "On coverage control for limited range multi-robot systems," in *2022 IEEE/RSJ IROS*, 2022, pp. 9957–9963.
- [6] Z. Zhang, J. Yu, J. Tang, Y. Xu, and Y. Wang, "Mr-topomap: Multi-robot exploration based on topological map in communication restricted environment," *IEEE RAL*, vol. 7, no. 4, pp. 10794–10 801, 2022.
- [7] H. Zhang, J. Cheng, L. Zhang, Y. Li, and W. Zhang, "H2gnn: Hierarchical-hops graph neural networks for multi-robot exploration in unknown environments," *IEEE RAL*, vol. 7, no. 2, pp. 3435–3442, 2022.
- [8] G. A. Hollinger and S. Singh, "Multirobot coordination with periodic connectivity: Theory and experiments," *IEEE Trans. on Robotics*, vol. 28, no. 4, pp. 967–973, 2012.
- [9] Y. Kantaros, M. Guo, and M. M. Zavlanos, "Temporal logic task planning and intermittent connectivity control of mobile robot networks," *IEEE Trans. on Automatic Control*, vol. 64, no. 10, pp. 4105–4120, 2019.
- [10] H. Rovina, T. Salam, Y. Kantaros, and M. Ani Hsieh, "Asynchronous adaptive sampling and reduced-order modeling of dynamic processes by robot teams via intermittently connected networks," in *2020 IEEE/RSJ IROS*, 2020, pp. 4798–4805.
- [11] F. Bullo, J. Cortes, and S. Martinez, *Distributed Control of Robotic Networks: A Mathematical Approach to Motion Coordination Algorithms*. USA: Princeton University Press, 2009.
- [12] L. Matignon, L. Jeanpierre, and A.-I. Mouaddib, "Coordinated multi-robot exploration under communication constraints using decentralized markov decision processes," in *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, ser. AAAI'12. AAAI Press, 2012, p. 2017–2023.
- [13] K. Gao, Z. Cao, L. Zhang, Z. Chen, Y. Han, and Q. Pan, "A review on swarm intelligence and evolutionary algorithms for solving flexible job shop scheduling problems," *IEEE/CAA Journal of Automatica Sinica*, vol. 6, no. 4, pp. 904–916, 2019.
- [14] Y. Fang, C. Peng, P. Lou, Z. Zhou, J. Hu, and J. Yan, "Digital-twin-based job shop scheduling toward smart manufacturing," *IEEE Trans. on Industrial Informatics*, vol. 15, no. 12, pp. 6425–6435, 2019.
- [15] W. Song, X. Chen, Q. Li, and Z. Cao, "Flexible job-shop scheduling via graph neural network and deep reinforcement learning," *IEEE Trans. on Industrial Informatics*, vol. 19, no. 2, pp. 1600–1610, 2023.
- [16] E. A. Jensen and M. Gini, "Rolling dispersion for robot teams," in *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, ser. IJCAI '13. AAAI Press, 2013, p. 2473–2479.
- [17] C. Amato, G. Chowdhary, A. Geramifard, N. K. Üre, and M. J. Kochenderfer, "Decentralized control of partially observable markov decision processes," in *52nd IEEE Conference on Decision and Control*, 2013, pp. 2398–2405.
- [18] Y. Zou, J.-K. Hao, and Q. Wu, "A two-individual evolutionary algorithm for cumulative capacitated vehicle routing with single and multiple depots," *IEEE Trans. on Evolutionary Computation*, pp. 1–1, 2024.