Robotics and Autonomous Systems 194 (2025) 105137

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Multi-robot rendezvous in communication-restricted unknown environments via backtracking and semantic frontier-based exploration

Matteo Luperto *, Mauro Tellaroli, Michele Antonazzi, Nicola Basilico

Department of Computer Science, University of Milan, Milano, Italy

ARTICLE INFO

Keywords:
Multi-robot rendezvous
Decentralized coordination
Frontier-based exploration
Unknown communication-restricted environments
Semantic-aware exploration and rendezvous strategies
Multi-robot systems in ROS

ABSTRACT

The multi-robot rendezvous problem requires coordinating a team of mobile robots to converge at a common location. Efficient decentralized execution with limited communication and in unknown environments can be essential in several applications, but such real-world robotic features are often not captured by theoretical rendezvous frameworks. In this work, we present an approach to this problem that extends traditional frontier-based exploration strategies to facilitate efficient rendezvous in such conditions. Our method allows robots to backtrack their exploration of the environment and exploit semantic knowledge on the map by prioritizing high-connectivity areas like corridors and hallways. We define and evaluate different variants of our method to study a trade-off between the time taken to perform a rendezvous and the amount of discovered area in the environment. Extensive experimental evaluation in ROS using 3D simulations demonstrates the feasibility of our method and its performance improvements over baselines.

1. Introduction

Multi-robot rendezvous involves coordinating a system of mobile robots, or a Multi-Robot System (MRS), to efficiently converge at or near a shared location. This is a fundamental coordination problem which is relevant in a variety of MRS missions. Multi-robot rendezvous can be beneficial in exploration and search missions to facilitate data sharing and to enhance team-level planning [1], and in autonomous surveillance to allow an improved coverage of the surveilled area [2, 3]. In search and rescue missions, MRS can better allocate different tasks according to the shared knowledge after they rendezvous, thus improving their performance [4]. In heterogeneous teams of robots, a rendezvous is a necessary precondition for collaborate in synergy; as an example, a drone can recharge its batteries at a mobile docking station placed on a wheeled robot, or a robot with mobile manipulation capabilities can allocate its tasks according to the informed knowledge acquired by other team members. It is also a prerequisite in situations where the robots have to collaborate at a specific location of the environment, a problem often found in multi-robot task allocation setups [5].

To be efficient, a rendezvous strategy must typically minimize the total time or distance for all robots to reach the meeting point, as these costs are proxies for energy consumption and task-execution efficiency. Achieving such an efficiency in real-world robotic setups poses additional challenges with respect to the abstract mathematical frameworks in which this problem is traditionally studied. One major difficulty is the decentralized, communication-limited nature of these settings. Unlike environments with infrastructure that supports global communication among robots, real robots must rely on peer-to-peer interactions constrained by physical proximity. Another significant challenge is the need to perform rendezvous in environments with initially unknown maps, requiring exploration while searching for team members. Despite these difficulties, rendezvous in such scenarios can be critically beneficial. Take for example the work of [6] that investigates how low-powered drones can rendezvous with a base-station ground robot for recharging. Similarly, [7] shows how rendezvous between different robots can be used to merge the (different and separated) maps obtained by each robot in the team into a single one, which can later be used for strategic decision-making.

In this work, we consider the rendezvous problem for a team of autonomous mobile robots in the challenging setting of a communication-restricted [8] and an initially unknown indoor environment. We assume that the robots start from different arbitrary locations, but no map of the environment is available to any of them and no pre-determined meeting location or coordination strategy has been agreed upon. The communication in the environment is restricted to occur only after a rendezvous. An example of this problem is when two or more people need to gather in an environment they do not know, such as in a shopping mall, a hospital, or an airport.

In this scenario, the difficulty of the rendezvous problem is augmented by the fact that the MRS needs to simultaneously perform an

* Corresponding author.
E-mail address: matteo.luperto@unimi.it (M. Luperto).

https://doi.org/10.1016/j.robot.2025.105137

Available online 31 July 2025

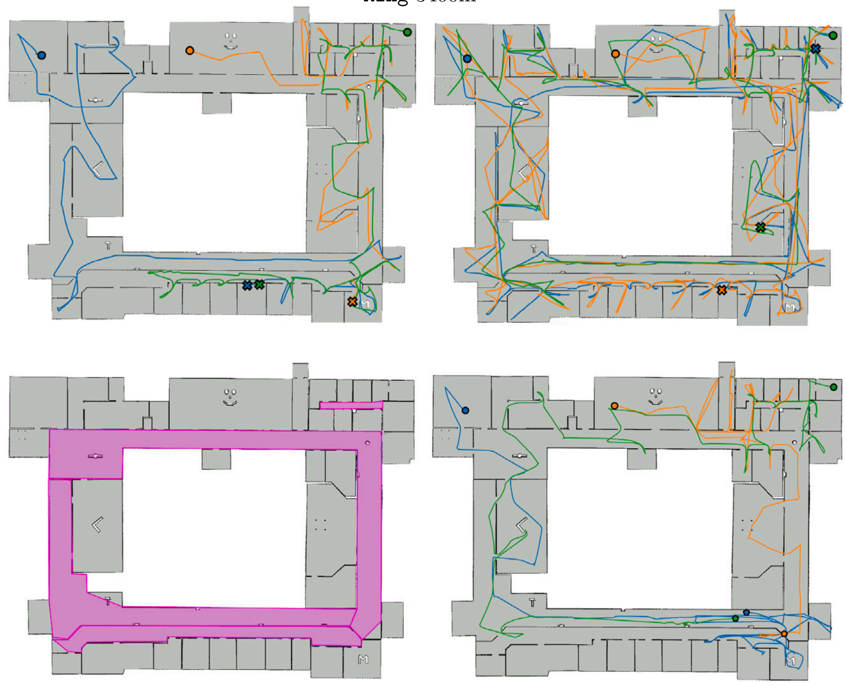

Fig. 1. Rendezvous paths of 3 robots in an initially unknown, communication-restricted environment, using our method (Frontier-Based Rendezvous, FBE-BT, top-left) and the classical frontier exploration strategy (Frontier-Based Exploration, FBE, top-right) in the map Ring, 3400 m^2. The second row shows how the knowledge of the high-connectivity areas of the environment (highlighted in pink, bottom-left) allows our method to bias the robot exploration strategy (Frontier-Based Exploration with utility scaling - FBE-BT-U(0.5) in this case, bottom-right) to travel across corridors and hallways, further facilitating rendezvous. Circles indicate the starting positions of each robot, crosses indicate the locations where a robot joins a cluster. When robots are in the same cluster, we indicate only the trajectory of the leader. Each robot has its own map and its own SLAM module; we use the full map of the environment for visualization purposes.

online exploration task, as the environment is unknown. Exploration is the problem of navigating an initially unknown environment to build its accurate map, and it is customarily performed by iteratively (i) identifying a set of promising candidate locations in the already mapped portion of the environment, (ii) selecting the most promising one according to a set of criteria, and (iii) reaching the target location while integrating new perceptions into the partial map of the environment. A popular choice in step (i) is to use *frontiers* [9], namely the boundaries between the explored and unexplored part of the environment, to determine candidate locations to reach, and to select the most promising one according to an *exploration strategy*. When an MRS is involved, each robot can carry out its exploration independently, or they can coordinate by communicating and exchanging maps and plans [10].

Frontier-based exploration strategies are widely used as they are simple yet effective, allowing robust exploration in heterogeneous contexts. However, the inherent greedy nature of frontier-based exploration (see (i)-(iii) above), which aims to obtain the largest map of the environment in the shortest time, results in quickly mapping those regions where the most free space lies, leaving behind frontiers in peripherical and less informative portions of the environment, that are eventually explored in later stages of the exploration run [11]. Such a rationale might be in contrast with what is required by an efficient rendezvous, which, in principle, does not strictly require building a complete map.

Motivated by real-world concrete mobile-robot settings, we propose a definition of rendezvous that enforces connectivity (to ensure communication) and promotes (without enforcing) physical proximity. Such a practical definition departs from the classical one typically studied in theoretical frameworks [12,13], where exploration and rendezvous are addressed on discrete representations like graphs or 2D polygonal environments (neglecting perceptions and navigation errors).

The method we propose is an exploration strategy that is biased towards rendezvous, where frontiers are created not only for their potential contribution to the map's expansion (classical frontier-based approach) but also to increase the opportunity of meeting a teammate. To do so, we introduce a mechanism that allows each robot to *forget* about parts of the environment that have been explored and mapped, thus putting those parts back into the portions of the environment still to be explored. In this way, the robot is encouraged to backtrack on previously explored areas, often passing through high-connectivity zones as corridors, facilitating accidental rendezvous. Our goals are minimizing the rendezvous time and maximizing the chances of rendezvous. At the same time, as the environment is unknown, we favour strategies that allow the total joint area covered by the MRS to be as high as possible; in this way, after the rendezvous, the robots can have access to more knowledge to coordinate and jointly plan their next actions.

We present two different methods; in the first one, each robot is biased to backtrack across its entire past trajectory. In the second one, we assume the robot is equipped with some sensing capabilities that allow it to distinguish high-connectivity areas of the environment from peripheric ones, by exploiting the fact that high-connectivity areas of indoor environments are easy to detect [14]. Using this knowledge, the robot is biased to backtrack to high-connectivity parts of the environment, thus increasing the chances of performing an accidental rendezvous in such areas. To detect high-connectivity areas, we refer to methods used for place categorization [15,16] and we assume that the robot can detect semantic features of the environment as the presence of corridors or hallways. In this way, the robot can exploit the fact that indoor environments have easily detectable features that can be used to identify hallways and corridors [17], and also from the fact that the topological structure of indoor environments, by design, is characterised by a few high-connectivity nodes that are placed in hallways and corridors [18]. An example of this can be seen in Fig. 1.

We evaluate this method in ROS with 3D realistic simulations of complex large-scale environments. The results obtained from our extensive experimental campaign show how our contribution extends existing frontier-based exploration frameworks to achieve rendezvous with different teams of robots efficiently.

This paper extends our previous contribution of [19,20], where we introduced our method to bias a frontier-based exploration strategy to backtrack on previously-visited areas of the environment to facilitate accidental rendezvous. In this work, we significantly expand previous contributions by investigating how using semantic knowledge about high-connectivity areas of the environment can further facilitate accidental rendezvous. Moreover, we investigate how different exploration strategies impact the combined explored area covered by the MRS before rendezvous, showing that our method allows also for implicit coordination of the MRS: not only the robot can perform a rendezvous earlier, but also the total combined area explored by the MRS is similar to the total area explored by an MRS team that is not tasked to rendezvous, but it just exploring the environment. Finally, we have significantly extended the experimental evaluation of our method, by performing hundreds of runs in a realistic simulation setting using ROS.

2. State of the art

Methods for multi-robot rendezvous typically leverage pre-determined coordination strategies among robots (pre-established rendezvous areas, biases applied to strategies to search for others, etc.) or execute online search strategies in a commonly known map to meet others. One distinction among existing approaches can be made between symmetric and asymmetric rendezvous. In symmetric rendezvous, all robots have the same role in seeking a meeting with others. In the asymmetric case, instead, some robots can be explorers, while others are relays, i.e., they have to meet and transfer knowledge acquired by other robots to each other or to a base station. Another distinction is between intentional rendezvous and accidental rendezvous. The former type happens when the meeting is already scheduled among agents following some kind of coordination. In the latter type, the robots know that they need to meet but no time and place is pre-arranged. Other approaches undertake an offline study of the problem, by deriving theoretical properties of the optimal solution from the specific geometrical features of the environment [21]. Finally, when communication is considered, some works [22] distinguish between opportunistic, continuous, and recurrent connectivity. Opportunistic connectivity is when the meeting between robots happens by chance (as in the setting we consider), continuous connectivity is when the MRS is always, directly or indirectly, connected to all other members of the team, while recurrent connectivity is when the robots perform periodic rendezvous for coordination. Several of these approaches are not directly applicable to our scenario, where the target environment and starting locations of the MRS are unknown, and communication is restricted. Ozsoyeller and Tokekar present in [23] a symmetric strategy to solve the rendezvous problem in a setting similar to ours. In such a work, the MRS has to rendezvous in a communication-restricted unknown linear environment; to do so, they alternate forward and backward movements with an increasing radius. As the environment considered is a line, robots estimate if they are the first, the last, or one of the central robots in the line, and act accordingly. Interestingly, the work computes the competitive ratio to rendezvous following a theoretical analysis. However, the proposed strategy is difficult to adapt to settings more complex than linear environments. Multi-robot rendezvous and exploration have been studied together, but are often dealt with as mutually-exclusive alternating phases: when the robots are exploring, their interest is to acquire new knowledge; when the robots decide to rendezvous, they travel to a meeting location in the mapped area. The decoupling of the two phases can lead to suboptimal performance, as robots must switch between behaviours in a coordinated way. This issue is investigated in [24] by introducing the concept of serendipity in exploration, i.e., to create a robot behaviour that tries to facilitate an unplanned encounter with other robots while those are still in an exploration phase, by adopting a behaviour that is a mix between episodic and planned rendezvous.

The work of [25] investigates how to perform arranged rendezvous in an asymmetric system where a robot is a relay of information, the others are *explorers*, and a rendezvous is set up in locations that may not be directly connected but are in communication range (e.g., are separated by a wall).

In [26] the robots, after performing a full exploration, rendezvous at a location on the final map of each robot, using an ant-based strategy.

The work of [13] presents a framework for the simultaneous multi-robot exploration and rendezvous on graphs. The MRS, without any communication, decides where to perform a rendezvous using three different strategies, namely asymmetric, symmetric, and exponential. In the asymmetric strategy, one of the robots is stationary at a node while the others explore the environment. The roles of the robots are pre-determined. In the symmetric strategy, all robots are exploring the environment. In the exponential strategy, the MRS uses an exponentially decreasing function to estimate rendezvous locations. In the beginning, the robots perform exploration using a breadth-first strategy. After a time budget is reached, the robots try to attempt a rendezvous by going to high-likelihood rendezvous locations. This approach is also used in [12] to minimize rendezvous time while having a maximum speed gain for the exploration task, with the least possible dependency on communication. A rendezvous strategy is used to rank and visit promising rendezvous locations.

The work of [27] presents a framework, similar to ours, to perform exploration, search, and rendezvous in a real-world deployment with line-of-sight communication, relying on the content of virtual stigmergy [28]. However, the goal of such a work is significantly different from ours, as it aims to perform periodic rendezvous at a pre-arranged time to share the information acquired during single-robot search tasks. After that, the MRS sets up a relay tree to connect the team to a base, sharing the position of the targets identified.

In [29], similarly to our work, new sets of frontiers are added to those used by a frontier-based exploration method to facilitate loop closures, that improve map quality. In our work, we go beyond hallucinating frontiers to facilitate rendezvous.

The work of [30] addresses a problem that shares some similarities with our setting, but with a slightly different flavour; an MRS is tasked to explore an environment with limited communication and performs intermittent rendezvous with subteams of robots to coordinate and share acquired knowledge. First, the MRS creates an exploration plan when all robots are within a communication range. In the exploration plan, frontiers are assigned to each robot; then each robot follows a local policy to explore its set of frontiers. Finally, robots perform intermittent rendezvous at pre-defined locations to share their newly acquired maps, define new rendezvous locations for each subteam, and restart the exploration process.

Recently, Song et al. presented in [31] a method for joint exploration and rendezvous in unknown environments. The strategy alternates different steps: first, robots explore the environment separately while building a topological map. When a shared area is observed by more than one robot, their topological map is merged, and their knowledge is shared. This is used to partition the environment. When all topological maps can be merged into a unique connected map, a rendezvous location is selected, and rendezvous is performed. Similarly to us, this work prioritises the exploration in areas with high connectivity, such as crossroads, following the insights that those areas are likely to be observed by multiple robots and can be used for rendezvous. A neural network using laser scans as input is used to detect high-connectivity areas. Differently from us, this method assumes that robots can communicate with each other, albeit with the availability of a limited bandwidth, and that they use communication to share their maps.

In [22], the problem of exploration and rendezvous for an MRS with intermittent connectivity is dealt with deep reinforcement learning using curriculum learning. The main task of the MRS, as in [30], is to explore the environment, and each rendezvous is functional to share information among robots to speed-up the exploration process. Each robot plans its actions using a hierarchical topological graph representation of the environment: a fine-grained local topological graph, centred in the robot's position, is used to determine the next location to reach for the robot. Each robot then builds a coarse global topological graph of the environment, that is shared and merged with another robot's global topological graph after a rendezvous. The policy of the robot is based on deep-reinforcement learning, where the reward is given both for explorative actions (i.e., local actions that allow the robot to observe new features of the environment) and for coordination one (i.e., local actions where the robot is biased to go to a location where it is more likely the communication with other members of the MRS).

In [32], a team of robots has to discover a set of targets, called *tasks*, in a previously unknown environment with limited communication. The robot periodically performs planned rendezvous to share the acquired knowledge. Robots agree on the rendezvous point when they are all connected, before starting an exploration step. After a time budget spent doing exploration, the MRS meets at the predefined location, shares knowledge, agrees to a new rendezvous location, and starts a new exploration step. If one or more robots fail to meet at the predefined rendezvous location, the MRS can search for their location, using the available knowledge about the other robots' statuses and exploration missions.

The work of [33] aims to perform a rendezvous under the following setting: each robot does not have global knowledge about the environment or the position of other robots, the MRS do not share a common reference frame, and they do not communicate; each robot can sense other robots where they are in line-of-sight. When two robots are in line-of-sight, an edge exists between them in a *visibility graph*. Differently from our work, it is assumed that the initial position of the robots creates a *connected* visibility graph. Thus, the proposed approach to solving the rendezvous problem involves two main ideas: first, the visibility graph should not lose connectivity during the evolution of the group; second, while preserving the connectivity of the graph, the robots must move closer to each other.

Hollinger and Singh introduced in [34] the concept of periodic connectivity for an MRS. The robots start in a connected configuration; after that, they split and their network becomes completely disconnected while exploring the environment. After a predefined time budget, the network of the robots regains connectivity by doing a periodic rendezvous, where they share the acquired knowledge and replan for the next meeting. Connectivity is ensured by an implicit coordination mechanism, where each robot plans its next position taking into account the possible connectivity with other members of the MRS.

The work of [35] presents a dual resolution scheme to perform autonomous exploration; each robot maintains a high-resolution local map of its vicinity, which is used to plan the next exploration steps locally, and a low-resolution global map of the whole environment, which is used to direct the robot to expand its map. Similar to other works, the rendezvous of robots is used to share information and coordination. For this, a pursuit strategy is presented, where robots can opportunistically go after each other to communicate if they believe that a communication action (that is, to perform a pursuit of another robot) is more cost-effective than performing an exploratory one (that is, to go to an unvisited location of the environment).

2.1. Semantic knowledge in exploration

The task of identifying the semantic properties of indoor environments is a well-investigated subject in autonomous mobile robotics [16, 36]; among different types of knowledge that can be acquired, several works have proposed methods to perform *place recognition* [37], that is to distinguish each room and to assign them a semantic label such as *corridor*, *office*, or *kitchen*. In literature, several heterogeneous methods to perform place recognition have been proposed. Some of them are based on vision, such as [37]. Recent methods focus on neural networks, with several architectures being proposed, ranging from CNNs [37] to less conventional approaches based on LLMs [38]. A popular dataset to train neural models for place recognition is Place365 [39], containing more than 10 million images from over 400 types of rooms.

Other works investigated how place recognition can be performed using LiDAR range-sensor data, such as in [15]. More recently, graph-neural networks are also used for this task, by leveraging on the topology of the environment to perform place recognition [14,40].

The use of knowledge about semantically-relevant features of the environment can be beneficial for several tasks performed by single robots or by multi-robot systems. The work of [41] shows how semantic knowledge can be embedded in LiDAR readings to have semantically-labeled LiDAR scans that can be used to improve path planning. Similarly, the work of [42] shows how the detection of semantic-relevant features such as doors, hallways, or openings, can be integrated with the SLAM process. The work of [43] uses semantic knowledge to improve the localization performances of indoor robots. Finally, the work of [44] shows how the performance of an MRS that is tasked to find victims in an unknown building improves when the robots have the availability of the semantic knowledge of the different types of rooms, i.e., if they are small, medium, or big rooms. In our work, we propose to exploit semantic knowledge about the environment to increase the chances of rendezvous.

3. Rendezvous-oriented exploration strategies

3.1. Problem formulation

We consider a team of m homogeneous robots equipped with the same perception, navigation, and communication capabilities. Analogously to a multi-robot exploration scenario, each robot starts randomly within the free, reachable areas, without prior knowledge of the environment's map or their teammates' initial positions. The capabilities that, in our scenario, robots exhibit in terms of mapping and communication are a key part of our problem formulation.

For mapping, we assume that robots are equipped with environmental sensors (e.g., LiDAR and/or cameras) that enable them to perceive the space in their surroundings and integrate these local perceptions into a metric map of the environment, which can be leveraged to carry out localization and navigation. Importantly, we assume that robots can annotate their map, distinguishing parts where meeting teammates might be more likely or desirable. We term these regions as *high-connectivity areas*. Examples might include hallways, corridors, or sections that are frequently traveled. In Fig. 2, we depict these areas in two environments that we shall consider for our experimental evaluation. For simplicity, we assume that the high-connectivity areas are predefined in the environment and that robots can recognize and flag them when they are added to the map. This assumption allows us to focus on the advantages that the availability of this semantic knowledge might introduce while abstracting away from the specific method required to acquire it. For such a purpose, common techniques based on semantic mapping are suitable solutions to augment an exploration robot's metric map with subregion categorization [16]. This allows selecting as high-connectivity areas those portions of the map that, for example, belong to specific categories.

We assume that robots can establish wireless communication channels, adhering to a connectivity model with specific limitations. Two robots can exchange data only if there is an unobstructed line of sight between their locations and each one resides within the communication range of the other. The communication range is determined by a

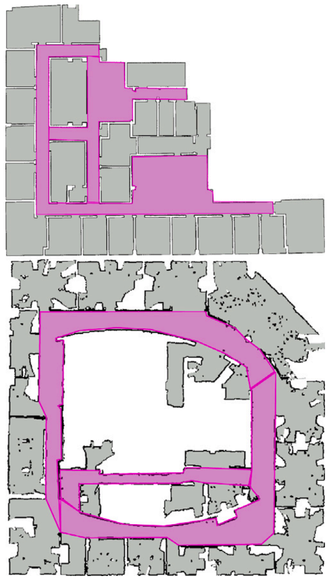

Fig. 2. Environments labelled as `Office` (top) and `Intel` (bottom) where high-connectivity areas are highlighted in pink.

circular area centred on the current location of the robot and with a radius d.

The task that robots have to solve is to perform a rendezvous. We define such an event as compliance with two conditions over the joint set of locations they occupy at a given time. The first condition (i) requires the robots to form a communication network according to the communication model described above. This configuration ensures that any pair of robots can communicate via a multi-hop route, assuming the existence of a multi-hop communication protocol. We enforce this condition as a hard constraint: we do not consider as successful rendezvous configurations where the communication network has more than one connected component among the m robots. The second condition (ii) requires robots to keep close to each other, similar to what they would do if navigating in a formation. This complements network connectivity with physical proximity, ensuring that the rendezvous involves not just forming a network, but also maintaining closeness, which can be important when robots need not only to exchange information but also to cooperate physically in some common region of the environment. The most natural way to enforce this condition would be to impose a maximum distance between the locations of any pair of robots. However, as the number of robots increases and open spaces in the environment are limited, achieving condition (i) may be easier, but reconfiguring into a joint formation without collisions while not exceeding a maximum distance between robots becomes challenging. Hence, we treat condition (ii) as a soft constraint: the team aims to adhere to the upper bound of mutual distance as much as possible, allowing its violation when no other option is available.

Ultimately, robots are required not only to execute a rendezvous but to do so efficiently. The cost of a rendezvous is generally assessed based on the time it takes the robots to accomplish it. The approach we present in this paper aims to optimize this cost and will be evaluated (in Section 4.2) using both quantitative metrics and qualitative assessments.

3.2. From frontier-based exploration to rendezvous

Given our problem formulation, an essential behaviour for the robots to achieve rendezvous involves exploring and discovering areas of the environment. For this reason, we develop our method as an extension of the Frontier-Based Exploration (FBE) approach [9], which is established as the simplest and most widely used strategy to enable one or more autonomous robots to explore and map an initially unknown environment.

A robot carrying out frontier-based exploration iteratively updates a 2D metric map with its sensory data, which encodes those areas of free space the robot has sensed and partially explored up to the present time. In this map, a *frontier* is an unsplit boundary separating the mapped free space from the unknown parts of the environment. Schematically, the robot follows these steps.

1. Extract the frontier set F_{exp} from the map. If F_{exp} is empty terminate, otherwise for each frontier $f \in F_{\text{exp}}$, identify a target location to approach, typically the frontier's centroid.
2. Select the best frontier f^* by maximizing a utility function. Following the approach of [9], for a frontier f, we define the utility as

   $$u_{\text{FBE}}(f) = \alpha \cdot len(f) - (1-\alpha) \cdot dist(f)$$

   where $len(f)$ is the frontier's boundary length, $dist(f)$ is its centroid's distance from the robot, and $\alpha \in [0, 1]$ is a parameter that balances these two factors. According to this function, FBE prioritizes nearby frontiers that are also long, as these are more likely to reveal large areas of free space.
3. Plan and execute a path to f^*'s centroid, updating the map with the new sensory data acquired while traveling.
4. Repeat from (1).

In our scenario, each of the m robots starts by executing an FBE strategy on their own in an asynchronous manner. As this process unfolds, we allow the formation of *clusters* upon meeting among robots. A *cluster* $C_t = \{R_i, R_j, \dots\}$ consists of one or more robots forming a network at time t according to condition (i). We assume that a cluster is formed whenever this condition is met among a set of robots. When a cluster forms one robot is elected as the *leader*, while others, if present, are *followers*. We assume that with network connectivity ensured by condition (i), roles can be assigned through a distributed consensus-reaching protocol. The leader will play the role of guiding the other robots by keeping on executing its FBE strategy after having included in its partial map the frontiers of the followers. Followers adhere to the leader's decisions, ensuring that they maintain proximity and do not exceed a specified maximum mutual distance to remain within the group. Splits can occur if one or more robots lose contact with the group, possibly due to lagging behind the leader or getting stuck in cluttered areas. In such cases, the isolated cluster (possibly a single robot) would elect a new leader (potentially itself) and continue the exploration process as described previously. For simplicity, in our simulations, we did not consider this type of events.

3.3. Actively seeking rendezvous

Collective execution of the FBE strategy with the addition of cluster formation might lead to serendipitous rendezvous. However, encounters among robots are sought neither directly nor indirectly. The FBE strategy prioritizes exploration, aiming to expand map boundaries while avoiding backtracking on previously visited areas. Our goal is to guide this process towards efficient rendezvous. To achieve this, we propose a set of techniques to steer exploration in that direction.

First, we enhance the above frontier-based exploration strategy by incorporating a mechanism that allows robots to backtrack to already explored areas. We achieve this by introducing *information decay* in the mapping process, inducing robots to revisit some of their paths. Each robot R_i maintains an *exploration trace* E_i, which outlines the area around its path up to the current exploration time. This region is defined by periodically recording the robot's timestamped location at a frequency f_{pose}. The footprint of each location is defined by a circular

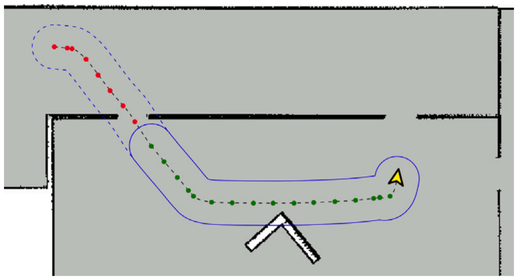

Fig. 3. Example of a robot's exploration trace. Green dots represent locations recorded within the last T time units; the area formed by merging their footprints is outlined by the solid blue line and represents the robot's current trace. Red dots denote locations that, with their associated area, have been discarded due to information decay.

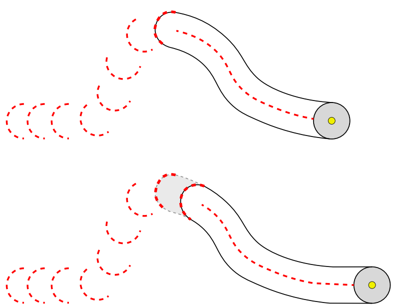

Fig. 4. Information decay and creation of backtracking frontiers for a single robot. The top and bottom images represent two subsequent steps in the trajectory of the same robot. When the robot moves forward (bottom image), the oldest location in the exploration trace and its footprint (dashed grey line) are removed, and a new backtracking frontier is created (dashed thick red line).

area centred on it with a radius r. The trace's area is then obtained by merging the footprints.

Each robot applies the information decay at a frequency of $f_{\text{decay}} \leq f_{\text{pose}}$, removing locations from its exploration trace that are older than T time units. This process effectively splits the trace into an active head and a vanishing tail. Fig. 3 illustrates an example of a robot's execution trace: the green locations are still in the trace, while the red ones have been discarded due to information decay. Whenever a location is removed from the trace, a new frontier f' is created by extracting the boundary of the differential area covered by the trace before and after removal. This new frontier and its centroid are completely surrounded by areas already explored by the robot. We call these frontiers, formed through information decay, *backtracking frontiers*, as they serve as potential locations for the robot to revisit on its path. We assume that each robot adds such frontiers to an initially empty set F_{bkt}. Fig. 4 provides a visualization of this process for a single robot.

In a cluster comprising multiple robots, the leader will inherit all the frontiers (whether exploration or backtracking ones) of the merged parties. At this stage, any segment of a frontier that overlaps with the unexpired trace of another robot in the cluster is removed. Consequently, the leader will get frontiers that are confined to regions not recently explored by any robot in the cluster. Notice how this may cause the frontiers to split into multiple sub-segments, also requiring updating the location of their centroid. From that point on the leader continues with its FBE strategy with the augmented set of frontiers and applies the information decay to its trace. An example of this process is illustrated in Fig. 5, where two robots encounter each other and form a new cluster.

In this setting, the leader robot will choose the next frontier to direct the cluster by following the FBE logic described in steps (1)-(4), while considering additional elements originating from our setting. First, in

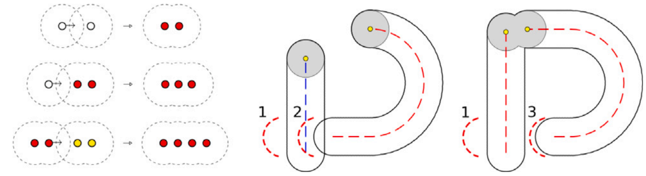

Fig. 5. Information decay in clusters. On the left: a new cluster is formed; On the right: two robots meet and share their exploration traces. The red robot's frontier (labelled as "2") is deleted as it overlaps with the exploration trace of the blue robot.

addition to traditional exploration frontiers F_{exp}, the leader can access F_{bkt}, which contains backtracking frontiers whose centroid would direct the cluster towards already mapped portions of the environment. We assume that step (2) of the FBE strategy is modified to keep track of this set too.

3.4. High-connectivity areas

The second element influencing this phase is the presence of high-connectivity areas in the cluster's current map of free space. Specifically, the robot can assess whether each frontier, whether exploration or backtracking, lies within a mapped high-connectivity area. We model this capability using the following binary function, computed by the leader. For any frontier f, the function is defined as:

$$m(f) = \begin{cases} 1 & \text{if } f\text{'s centroid is in a high-connectivity area} \\ 0 & \text{otherwise} \end{cases}$$

The elements are combined in an updated frontier selection method for step (2) so that f^*, the next frontier to reach, is computed as:

$$f^* = \arg \max_{f \in F_{\text{exp}} \cup F_{\text{bkt}}} \left(u_{\text{FBE}}(f) + \gamma m(f) \left| u_{\text{FBE}}(f) \right| \right) \quad (1)$$

Eq. (1) generalizes the standard functional form of the utility used in frontier-based exploration strategies to account for efficient rendezvous. One key extension is allowing backtracking: the next frontier is selected from $F_{\text{exp}} \cup F_{\text{bkt}}$, a set that includes both traditional exploration frontiers and those created through our information decay mechanism. Additionally, the utility of the frontiers within a high-connectivity area is increased by a scaling factor $\gamma \geq 0$, thus introducing a new parameter.

In this work, we focus on indoor environments and we consider corridors, hallways, and large open spaces as high-connectivity areas. These types of rooms have different functionalities and features [17] when compared with other categories of rooms; as an example, corridors and hallways are likely connected to multiple different rooms, while other room categories have one or two doors [18]. For these reasons corridors and hallways are easily identifiable and important for autonomous mobile robots; the choice of selecting these room categories for high-connectivity areas is thus motivated by the fact that their relevance is well known in the literature, with several works which have investigated how to distinguish between rooms and corridors directly from the robot's map or from perception data [15,40,44,45].

An approach similar to this has been recently proposed in [31]; in such a work, the MRS, when exploring an unknown environment, prioritizes high-connectivity areas as crossroads. This increases the chance that different robots travel across the same crossroad; this event is used to merge the maps of all robots and to set up a rendezvous location on the common map obtained by merging all maps of the MRS.

In indoor settings, selecting semantically relevant types of rooms as high-connectivity areas is straightforward, as some categories of rooms are central in the topology of the environment and their function is - by-design - to connect rooms [18]. In other settings, such as in urban outdoor environments, a different definition of high-connectivity area is needed. As an example, in a university campus or a shopping mall, parking lots, entrances, or large open spaces can be selected as high-connectivity areas. However, we deem that this is beyond the scope of this paper and is left for future work.

3.5. *Frontier-selection strategies for rendezvous*

In this work, we instantiate four different frontier selection methods to be evaluated and compared. Each one is obtained by applying specific choices to Eq. (1).

- FBE, or Frontier-Based Exploration, applies the standard frontier selection described in step (2) of the FBE strategy. This method exploits the utility function u_{FBE}. In Eq. (1) it amounts to disable backtracking by setting $F_{\mathrm{bkt}} \leftarrow \{\emptyset\}$ and canceling the utility scaling in the high-connectivity area with $\gamma = 0$. This method only achieves accidental rendezvous and will be considered as a baseline.
- FBE-BT, or FBE with backtracking, allows frontier selection to include also backtracking frontiers, which are still evaluated with the standard utility function u_{FBE}, without any utility scaling. It can be obtained by setting $\gamma = 0$. The method combines exploration of unknown regions with revisits to already mapped areas, thereby increasing the probability of encounters among robots.
- $\mathtt{FBE-BT-U}(\gamma)$, or FBE-BT with utility scaling, increases the utility of backtracking frontiers that lie within an environment's high-connectivity area. This is achieved by ensuring that $\gamma > 0$. This method prioritizes navigation within these known high-connectivity areas, resulting in more localized backtracking in regions considered suitable for encounters. FBE-BT-U represents our first approach to exploit the robot's capability of extracting semantic knowledge from the map in the form of high-connectivity areas.
- FBE-BT-A, or FBE-BT with backtracking inside the high-connectivity areas, restricts the generation of backtracking frontiers to parts of the map marked as high-connectivity areas. Then these frontiers are evaluated without any utility scaling using the standard utility function u_{FBE}. This approach can be obtained by overwriting F_{bkt} as $F_{\mathrm{bkt}} \leftarrow \{f \in F_{\mathrm{bkt}} \mid m(f) = 1\}$ and setting $\gamma = 0$. FBE-BT-A represents an alternative approach to FBE-BT-U for exploiting the knowledge of high-connectivity areas. Notice that, unlike FBE-BT-U, this method does not require setting the γ parameter.

In Fig. 6, we depict an example of a rendezvous between three robots. The robots ①, ②, and ③ start from three different corners of the environment. In Step A, none of the robot met, and the first backtracking frontiers were created in the initial part of each robot's path. In Step B, robots ① ② are backtracking, following their backtracking frontiers, and they rendezvous. The rendezvous location is highlighted with a purple square in Step C. Note how some of the backtracking frontiers of robot ② are shorter than others; this is because the robot's information trace, while backtracking, partially covers such frontiers, thus reducing their size. After rendezvous, robot ① is the leader of the cluster, and ② is the follower. Step C shows how the backtracking frontiers of robot ② are assigned to the cluster leader ①. Also in this step, some backtracking frontiers are reduced in size due to an overlap with the current exploration trace of the robot. Eventually, robot ① and ③ rendezvous in the central corridor of the map.

4. Experimental results

In this Section, we present an extensive experimental campaign to evaluate our methods, where more than 400 real-time exploration runs have been performed in large-scale realistic simulations of indoor environments. Sections 4.1 and 4.2 describe the experimental setting

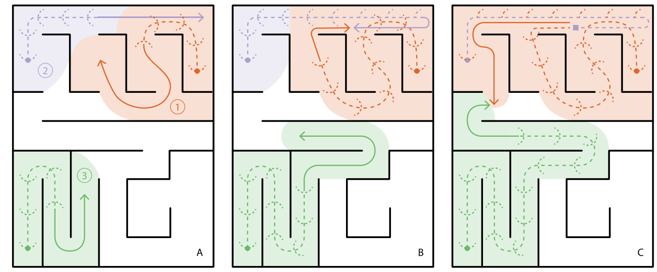

Fig. 6. Example of a rendezvous with three robots. The robots' ①, ②, and ③ acquired map and paths are represented using the colour orange, purple, and green, respectively. We highlight with a solid line the current path within the exploration trace, with a dashed line each robot's paths outside the exploration trace, and with a dashed line segments the backtracking frontiers of each robot. When two robots rendezvous, the frontiers of all robots in the clusters are assigned to the cluster's leader, and change their colour accordingly. Frontiers in F_{exp} are not shown for simplicity.

and the metrics, respectively, while Section 4.3 describes the main results of our method. Sections 4.4 and 4.5 investigate how different components of our method affect performance. Sections 4.6 and 4.7 investigate the performance of our method when an increasing number of robots are used and the impact of having different communication ranges. Finally, Section 4.8 evaluates performance in different and complex environments.

4.1. *Experimental setting*

We evaluate our method by implementing it in a real multi-robot control setup based on ROS. Thanks to the *namespace* system provided by ROS, we can concurrently run different and logically separated control stacks, one for each robot, within the same ROS environment. With such a control setup, we perform runs on environments simulated with Gazebo. We used the Turtlebot3-burger robots with a speed of $0.3\,\mathrm{m\,s^{-1}}$ and a 2D LiDAR with a range of $10\,\mathrm{m}$; each robot performs SLAM using Gmapping [46], and runs a navigation stack with movebase. We simulate sensors and actuation errors, to pose challenges comparable to those of a real-world run. After preliminary experiments, we set the information decay so that poses are forgotten each $\approx 5\,\mathrm{min}$, a communication range of $d = 2.7\,\mathrm{m}$, and an exploration trace radius $r = 1.5\,\mathrm{m}$. The exploration trace radius is empirically set to match the average width of rooms. A new pose is added to a trace each $2\,\mathrm{s}$, while a frontier is added to F_{bkt} after 9 poses have been collected ($18\,\mathrm{s}$). For the frontier evaluation function, we set $\alpha = 1/4$. Initially, we implemented leader–follower dynamics with the method from [47] which results in several robot failures (e.g., when trying to follow the leader through a narrow passage). Since the development of a robust formation control mechanism is outside the scope of this paper, in our experiments we assumed perfect formation control dynamics (see Section 3.1).

We use 5 different environments for the experimental evaluation, whose maps are shown in Figs. 1 and 13. The environments have been selected to have different features, allowing for the evaluation of the generalizability of our approaches. Not all experiments have been performed in all environments. Ring and Office are large-scale environments with ≈ 40 rooms and were used also to assess performance with a different number of robots, specifically from 3 to 8. We then evaluated, for teams of 3 robots, particularly challenging situations. First, we use the Intel environment [48], which is cluttered and with a less regular profile of obstacles. Then, we use Campus-1 and Campus-2 [49], two large-scale buildings with ≈ 100 rooms each. The starting location of the robots is chosen randomly and then fixed for each environment. For each configuration, we performed 10 runs, providing the mean μ and the standard deviation σ of the metrics, defined in Section 4.2.

We evaluate different variants of our method against a baseline running a standard frontier-based exploration as in [9] (label FBE).

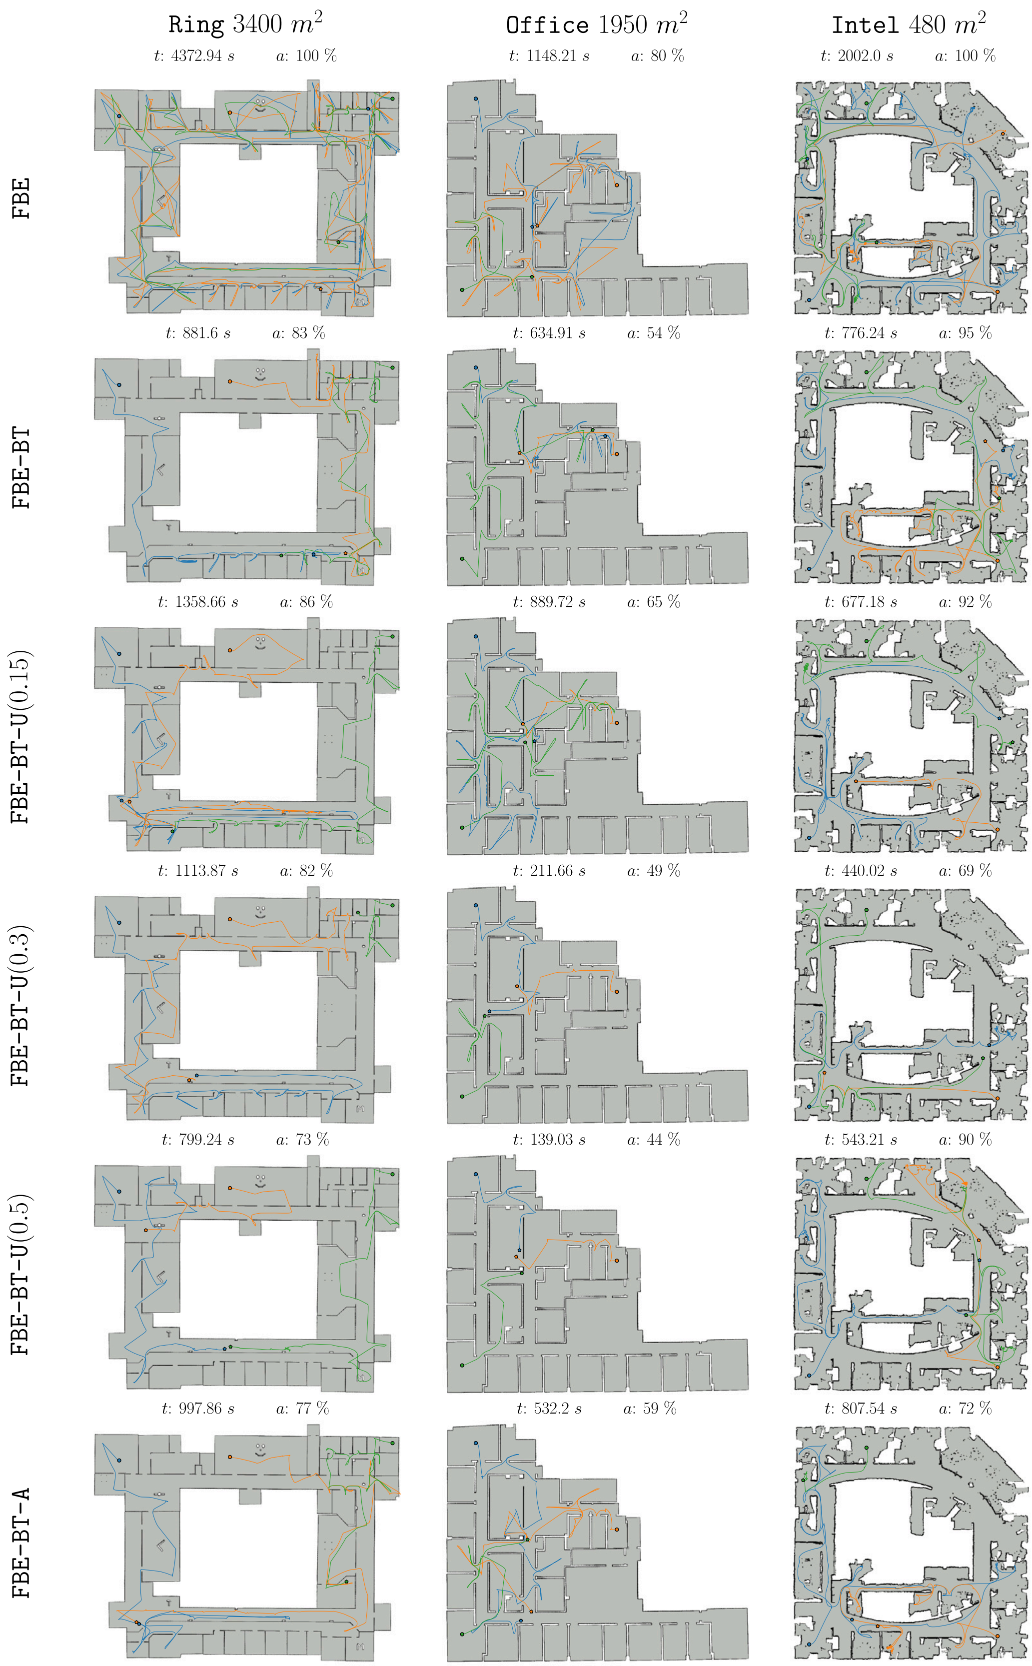

Fig. 7. Examples of trajectories of the MRS in three environments when 3 robots are used. We report on top of each map the total time t for rendezvous and the combined area explored a.

We label FBE-BT the Frontier-Based Rendezvous strategy, which is explained in Section 3.3 and does not use semantic knowledge. We then test the two strategies that use semantic knowledge, labelled FBE-BT-U and FBE-BT-A. For FBE-BT-U, we test different variants of the parameter γ, which increases the bias of the strategy to revisit backtracking frontiers in high-connectivity areas. More precisely, we test 3 configurations with $\gamma \in \{0.15, 0.30, 0.5\}$, labelled respectively as FBE-BT-U(0.15), FBE-BT-U(0.3), FBE-BT-U(0.5). Those results are reported in Section 4.4.

The methods FBE-BT-U and FBE-BT-A require that the robots be equipped with some sensing methods to distinguish between high-connectivity areas such as corridors, hallways, and other rooms. This semantic knowledge can be obtained by categorizing LiDAR point clouds directly [14,15,40] or using the semantic appearance and vision data [37,39]. As the task of semantic labeling of indoor environments is beyond the scope of this work, we assume that the robot knows whether each frontier is within a semantically-relevant high-connectivity area, or not. The high-connectivity areas of the Ring, Office, and Intel cover, respectively, are the 37, 32 and 30 % of the total area of those environments, and are shown in Figs. 1 and 2. To do so, we assume that the robot can query an oracle that, for each location in its known map, returns its status as part of a high-connectivity area or not.

Our experimental setting is a realistic benchmark, where robots experience similar challenges as those encountered in a real-world deployment, and that is significantly more realistic and difficult than the experimental framework of several works described in the state of the art, that are evaluated in more abstract settings (e.g., by considering a graph of the environment). Our code and the full set of experimental results are available in a public repository.[1]

4.2. Metrics

For each exploration run, we record the time t to perform a rendezvous among all robots, the progression with time of $\max |C|$ (the size of the largest cluster), and the combined area a explored by the robots obtained by computing the union of all maps aligned to the same reference system. While time t is in seconds, area $a \in [0, 1]$ is the portion of the total area of the environment covered by the combined maps of the MRS, after merging them. We indicate as success rate R the percentage of runs concluded with a rendezvous, and as t_i the time required for $i \in \{1, \dots, N\}$ robots to do a rendezvous. A run ends at time t when $\max |C|$ equals m, or when all robots explore the entire environment (no frontiers left). This last case can occur only with the baseline since our method, by definition, keeps adding virtual frontiers. In principle, our method could fail to terminate if the robots neither complete the exploration nor perform a rendezvous. While this situation can be easily handled by introducing a time limit, it never occurred in our experiments. To register a performance in this degenerate case, we use the time t required by the last robot to complete the exploration of the full environment. We thus assume that, when all robots have completed the full exploration of the environment without rendezvous, they can immediately travel to a default fallback location at the center of the map where they are guaranteed to rendezvous. The time required to reach such a location is not included in the total; in this way, the time used to measure performance when robots are not able to rendezvous is a pessimistic baseline for our method, even in the setting when there is no rendezvous.

We compute $\Delta t = (t_{\text{FBE}} - t_{\mathcal{E}})/t_{\text{FBE}} \cdot 100$, where $t_{\mathcal{E}}$ is the time required using one of the proposed multi-robot rendezvous exploration strategies $\mathcal{E}$ (i.e., FBE-BT, FBE-BT-U, FBE-BT-A), while t_{FBE} is the time required by FBE as a reference value. The higher the value of Δt, the better, as it would indicate that our method allows the MRS to perform a rendezvous faster.

[1] https://aislab.di.unimi.it/research/fberendezvous/.

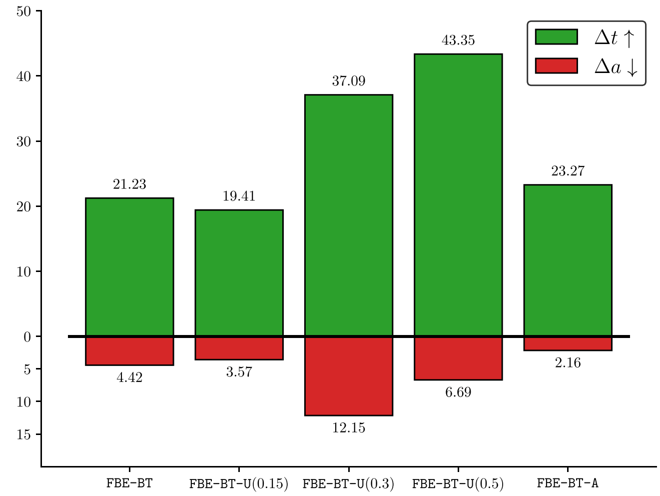

Fig. 8. Average values of Δt (green) and Δa (red) for the different exploration strategies over three environments when 3 robots are used. Full results are available in Tables 1 and 2.

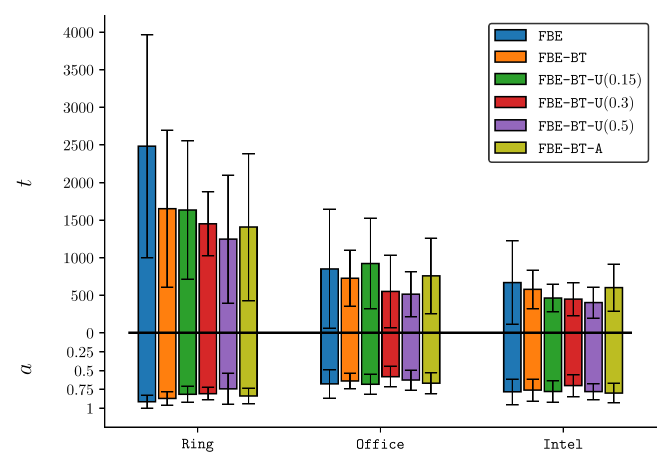

Fig. 9. Average time t and total area a for different exploration strategies to perform an accidental rendezvous among all robots of the MRS. Experiments are run in three environments, with 3 robots. Full results are in Tables 1 and 2.

Similarly, we compute $\Delta a = (a_{\text{FBE}} - a_{\mathcal{E}})/a_{\text{FBE}} \cdot 100$, where $a_{\mathcal{E}}$ is the percentage of the combined area explored by the team of robots when they perform a rendezvous over the total area of the environment, while a_{FBE} is the portion of the total area covered by FBE as a reference value. Differently, we aim to have a low value of Δa, as it would indicate that the total area explored by the MRS is similar. Note that after robots have performed a rendezvous, they can also jointly conclude the exploration by coordinating, as shown in [30].

Our average results do not include runs where one or more robots experienced a permanent failure during navigation (e.g., due to a collision or because they were incapable of maneuvering away from a too close obstacle). In our realistic setup, these kinds of events are fairly frequent, but in the majority of cases, robots managed to resolve the impasse by triggering a recovery behaviour on the navigation stack.

4.3. Results

Our results are shown in Fig. 9, which shows the average time t and the average combined area explored a when a rendezvous happens when 3 robots are used. These values are reported in Tables 1 and 2. We start our analysis from Table 1 and Fig. 8, which report, respectively, the rendezvous time required by the different methods, and Δt against

Table 1
Average time μ and variance σ required by the different exploration strategies to rendezvous among all robots of the MRS, in seconds. Data reported here are also visualized in Figs. 8 and 9.

Environment	FBE		FBE-BT		FBE-BT-U(0.15)		FBE-BT-U(0.3)		FBE-BT-U(0.5)		FBE-BT-A	
	μ_t	σ_t	μ_t	σ_t	μ_t	σ_t	μ_t	σ_t	μ_t	σ_t	μ_t	σ_t
Ring	2483.54	1484.68	1652.30	1041.26	1634.76	921.29	1451.54	427.22	1257.26	847.71	1408.72	976.98
Office	851.11	791.46	727.46	372.16	923.25	601.19	552.28	481.85	515.79	298.15	759.51	500.91
Intel	688.19	710	580.13	256.63	464.42	183.87	450.00	219.97	404.16	205.37	579.69	304.24

Table 2
Average portion of the area μ and variance σ covered during the exploration before the MRS rendezvous among all robots of the MRS. The value is in $[0, 1]$. Data reported here are also visualized in Figs. 8 and 9.

Environment	FBE		FBE-BT		FBE-BT-U(0.15)		FBE-BT-U(0.3)		FBE-BT-U(0.5)		FBE-BT-A	
	μ_a	σ_a	μ_a	σ_a	μ_a	σ_a	μ_a	σ_a	μ_a	σ_a	μ_a	σ_a
Ring	0.914	0.090	0.871	0.091	0.815	0.108	0.807	0.083	0.804	0.092	0.838	0.103
Office	0.678	0.187	0.639	0.101	0.682	0.132	0.581	0.136	0.626	0.133	0.667	0.139
Intel	0.782	0.169	0.760	0.147	0.778	0.144	0.700	0.148	0.780	0.106	0.809	0.125

FBE used as a baseline, in three different environments, when 3 robots are used. From this, several interesting remarks can be drawn. First, robots that are exploring the environment using FBE (so without trying to perform a rendezvous, but with the goal of covering the largest area possible, as quickly as possible) still have a good chance to do a rendezvous with each other. This is due to the fact that, especially in later stages of the exploration, the robots backtrack to previously explored zones to reach far away frontiers that have left behind during the exploration's early stages [11]. While doing so they travel mostly through connecting areas of the map as corridors, thus facilitating the event of accidental rendezvous. An example of this can be seen in Fig. 7, where trajectories that are performed by the MRS in three environments, while using different strategies, are shown.

This backtracking behaviour is considered an undesirable side effect of several exploration strategies, as during late stages of the exploration process robots are often driven to move across previously visited locations of the environment to reach far away frontiers not explored at early moments [50]. When this happens, the robot spends significant time traveling across the map, eventually adding only a limited knowledge to the map, while the robot's goal is to quickly observe new unmapped parts of the environment. However, we have empirically observed how robots that are backtracking in later stage of the exploration have higher chances of doing a rendezvous.

Our work encourages this backtracking behaviour from the early stages of exploration. The result of this is that our proposed methods significantly increase the chances of episodic rendezvous, as can be seen by the gain in terms of Δt in Fig. 8. The FBE-BT strategy almost reduced by one third the total time t required by the robot to perform a rendezvous, with $\Delta t = 21.23$. The use of semantic knowledge to bias the robot to re-visit high-connectivity areas such as corridors and hallways further increases Δt. When the semantic knowledge is used, it can be seen how the biggest gain in terms of $\Delta t = 43.35$ is seen when the bias of the robot to revisit semantic areas is higher, that is when FBE-BT-U(0.5) is used.

When we consider FBE-BT-A, we can see how using only high-connectivity areas for backtracking has interesting results. Fig. 8 shows how FBE-BT-A performances Pareto-dominate those of FBE-BT and FBE-BT-U(0.15) for both Δ_t and Δ_a; at the same time, using a higher value of γ improves Δ_t, at the expense of a higher Δ_a. A deployment of a MRS that is interested into fast rendezvous can thus select higher values of γ, while FBE-BT-A can guarantee low rendezvous time and high explored areas.

In all the runs performed with our method, the MRS were always able to perform a rendezvous ($R = 1$). This holds for all of the 150 runs considered in Table 1, and for both our two proposed methods, with and without using semantic knowledge. When FBE is used, the MRE sometimes fails to perform an accidental rendezvous before ending the exploration run, with $R = \{0.7, 0.9, 0.9\}$ for Ring, Office, and Intel, as shown in Table 4. At the same time, the performance is far more stable than those of FBE, with a lower standard deviation in t, a, and R.

Table 2 reports combined area a explored by the whole MRS at rendezvous time. Despite the fact that our method allows the robots to rendezvous in a significantly shorter time, the total area explored by the MRS is only slightly smaller than the area explored by robots performing exploration independently using FBE. This indicates how our method encourages the robot to perform a rendezvous faster, also by reducing overlapping exploration of the same portions of the environment. To highlight this, Fig. 8 shows the values of Δt and Δa for the different strategies. In particular, using FBE-BT and FBE-BT-A allows the robot to perform a rendezvous in $\approx 25\%$ less time ($\Delta t \approx 25$), while the total area explored is only $< 5\%$ less of those explored by FBE. It is interesting also to observe what happens when the bias of backtrack to semantically meaningful portions of the environment is higher, i.e., when FBE-BT-U(0.5) is used: the MRS time to perform a rendezvous is reduced half of the time required by robots using FBE ($\Delta t = 43.35$), but the total area explored by the MRS is only $\Delta a = 6.89$ less.

While, overall, our method improve the performances across different environments, the performance gain for Office is less evident, as shown in Fig. 9. This is because the topology of this environment, shown in Fig. 2, facilitates the robot to meet, as two open areas connect most of the rooms; after the beginning of exploration, robots have high chances of going into one of these two areas, and thus to meet, even when FBE is used. Nevertheless, our method, by explicitly encouraging the robots to backtrack to promising areas, has still some noticeable positive effect: it both improves performance and reduces the variance.

Fig. 7 shows some examples of the different trajectories of the MRS using different strategies. It can be seen how robots that follow a pure exploration strategy as FBE are prone to explore more the environment. At the same time, their trajectories are more overlapped, as they have visited the same parts of the environment at different moments, without meeting and thus without performing an accidental rendezvous. When FBE-BT is used, the robots are driven to backtrack on previously explored portions of the map; the trajectories are shorter, as the robots meet earlier, and also there is less overlap between trajectories. The use of semantic knowledge further biases the robot not only to backtrack to previously visited portions of the maps, but also to prefer to (re)visit high-connectivity areas as corridors. This can be seen in the trajectories of the robot using FBE-BT-U(0.3) on Ring, where all robots while they are moving across the two corridors at the bottom of the environment. This result is particularly interesting: the MRS using our method not only can perform a rendezvous faster than a MRS that is using FBE, but also it has a similar knowledge at rendezvous time to plan the next actions that the MRS has to perform. The backtracking mechanism implicitly allows the MRS to be more efficient in exploring different parts of the environment before rendezvous. For further examples of these results, with the full visualization of the trajectories of the MRS, please refer to our repository.

4.4. Semantic knowledge ablation study

Tables 1 and 2 present the results when different values of γ are used, i.e., when more/less importance is given to the strategy followed by the MRS to backtrack to high connectivity areas, while average results are shown in Fig. 8. It can be seen how giving more importance to high-connectivity areas increases the chances of the MRS to perform a rendezvous, with significant shorter rendezvous times associated to higher values of γ. However, these positive effects can also be linked with some drawbacks. The first is that the total area explored by the MRS decreases. Then, robots are more subject to repetitive patterns where they continue traveling across high-connectivity areas, without exploring anymore the environment. An example of this behaviour is in Fig. 14. The blue robot, starting from the top-left corner of the map, quickly moves to two parallel corridors at the bottom of the map. After exploring them, the robot starts to do several *laps* by repetitively going across the two corridors clockwise. This is because the robot prioritizes the backtracking frontiers that are in the corridors from the previous lap. During this time, the orange robot is doing a tour of the entire environment. Eventually, the two robots meet after three *laps* of the corridor circuit made by the blue robot. The blue robot and the green meet at the end of the first lap of the blue robot, at the bottom-right corner of the map. While these repetitive behaviours are undesirable, the fact that robots are selecting frontiers also in rooms creates slightly different paths at each lap of the robot, allowing them to eventually rendezvous. While it is true that these unwanted repetitive patterns are more frequent with a high parameter of $\gamma = 0.5$, we deem that choosing a high value of γ is a preferred choice, due to benefits in terms of saved t.

4.5. Backtracking ablation study

In this section we evaluate the impact of backtracking on different exploration strategies that consider - or not - high priority areas. We perform experiments with 3 robots, with the same setting described in previous sections, and using two environments, `Ring` and `Office`. We compare four different strategies:

– `FBE`, which does not use either backtracking or connectivity areas;
– `FBE-BT`, which uses backtracking, but does not prefer high connectivity areas;
– `FBE-U(0.5)`, which prefers high connectivity areas, but does not use a backtracking mechanism;
– `FBE-BT-U(0.5)`, which uses both backtracking and high connectivity areas.

Fig. 10 shows the results. It can be seen how both the backtracking mechanism and the use of high-connectivity areas have an effect in decreasing rendezvous time; among the two, the backtracking mechanism is the one which has larger benefits.

While the total time t to explore the environments changes by using or not backtracking, the total area explored by the MRS is somehow constant, showing that the extra time required by the robot to rendezvous is spent by mapping areas that has been already observed by other members of the team.

Interestingly, robots are always able to rendezvous when backtracking is used ($R = 1$), while sometimes they are not able to rendezvous both when `FBE` and `FBE-U(0.5)` is used.

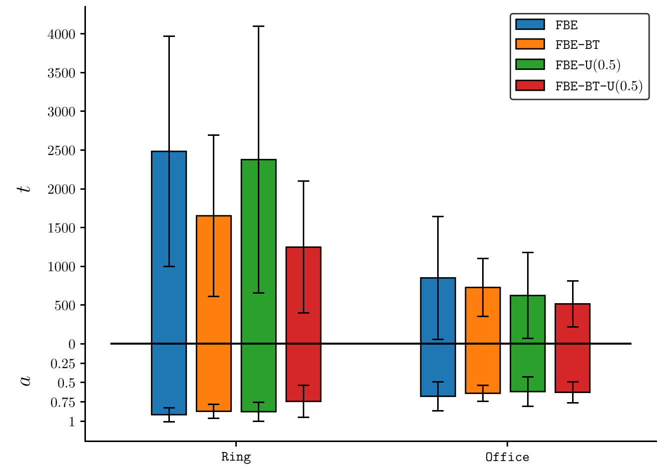

Fig. 10. Average time t and area a for the exploration strategies using and not using backtracking (`-BT`) over two environments when 3 robots are used.

4.6. Evaluation with different numbers of robots

Table 3 reports the results obtained in `Ring` and `Office`, using 3, 5, and 8 robots. To better illustrate the different performances when different numbers of robots are used, we focus on two strategies: `FBE` as a baseline, and `FBE-BT` as a comparison. The results show that our method facilitates faster accidental rendezvous for the robots, with the improvement being particularly pronounced when there are fewer robots distributed across a large environment.

Furthermore, our method ensured that all robots successfully performed a rendezvous in all 60 runs ($R = 1$). In contrast, there is a significant chance that the robots may not rendezvous when using the standard frontier-based exploration method `FBE`. This typically occurs when, by chance, the robots initially choose different directions during the exploration phase (e.g., one moving left, another moving upward, and another moving right). As a result, they may end up changing locations without encountering the other members of the Multi-Robot System (MRS).

These findings are illustrated in Fig. 11, which depicts the time t taken for rendezvous with 3 to 8 robots in the `Ring` environment. As the number of robots increases, the likelihood of episodic rendezvous also rises, making the differences between the two methods less pronounced. With the MRS has a high number of robots, `FBE` is an effective strategy for multi-robot rendezvous.

4.7. Communication range ablation study

In this section, we evaluate how having a different communication range for robots changes performance. We considered two environments, `Ring` and `Office`, and four different variants of strategies: `FBE` as a baseline, `FBE-BT`, `FBE-BT-U(0.5)`, and `FBE-BT-A`. We tested three values for communication range d, namely d, $2d$, and $3d$, with $d = 2.7\,\mathrm{m}$. The results are reported in Fig. 12.

A first consideration is that having a larger communication radius increases the chances of rendezvous. When a short communication range is used, we have $R = 0.9$ for `Ring` and $R = 0.7$ for `Office`, when `FBE` is used (and $R = 1$ when a strategy using backtracking is employed). A larger communication range increases the chances of rendezvous even when `FBE` is used; with a communication range of $2d$ and $3d$, we have $R = 1$ in all settings except for `Office` using `FBE` and a $2d$ communication range, resulting into $R = 0.8$.

When `FBE` is used, increasing the communication range linearly decreases the rendezvous time.

When backtracking is used, a larger communication range of $3d$ reduces the rendezvous time, as robots can detect each other even when

Table 3
Results of FBE-BT compared with FBE when a team of different m robots are used; t is in seconds, a and R are in $[0,1]$. The last two columns report the Δs, comparing the result of FBE-BT compared with FBE.

Environment	m	FBE-BT					FBE					Δ	
		t		a		R	t		a		R	Δt	Δa
		μ_t	σ_t	μ_a	σ_a		μ_t	σ_t	μ_a	σ_a			
Ring	3	1652.30	1041.26	0.871	0.091	1	2483.54	1484.68	0.914	0.090	0.7	33.47	4.79
	5	2131.28	986.16	0.931	0.063	1	2684.84	1502.32	0.945	0.090	0.7	20.62	1.48
	8	2106.52	819.02	0.971	0.027	1	2016.90	815.46	0.978	0.013	1	−4.44	0.72
Office	3	727.46	372.16	0.639	0.101	1	851.11	791.46	0.678	0.187	0.9	14.53	5.75
	5	1358.55	718.19	0.947	0.052	1	1478.21	594.59	0.970	0.032	0.8	8.09	2.37
	8	1309.45	437.50	0.957	0.029	1	1251.54	762.47	0.966	0.018	1	−4.63	0.93

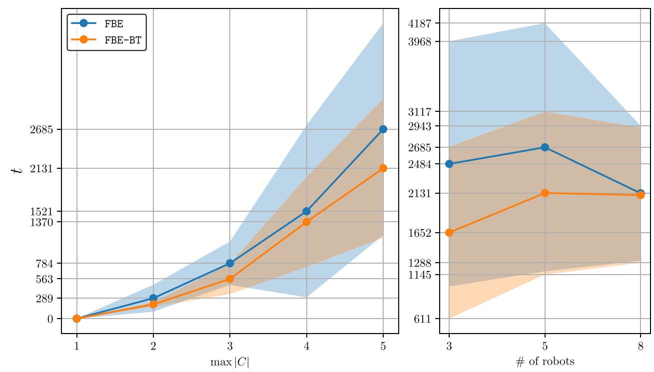

Fig. 11. (left) The average time t (and standard deviation) to complete the rendezvous as the maximum size of the cluster changes in the environment ring when an MRS of 5 robots is used; note how, in this environment the $R = 0.7$ for FBE. When the MRS does not perform rendezvous, the time required by the last robot to fully explore the map is used instead. (right) The average time t (and standard deviation) required for a run in Ring when a different number of robots is used.

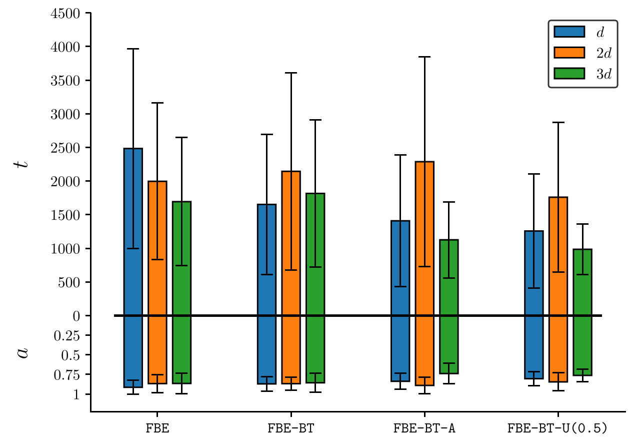

(a) Ring

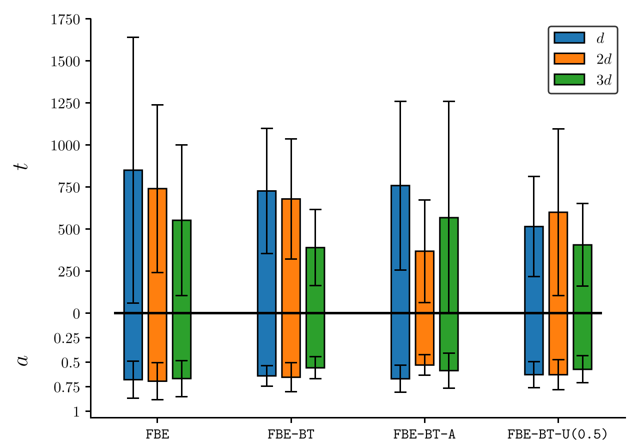

(b) Office

Fig. 12. Average time t and total area a for different exploration strategies with different values of communication range d. Experiments are performed with teams of 3 robots in 2 environments.

they are at a few meters distance. Interestingly, for an intermediate value for the communication range, $2d$, the trend in the performance is not clear. In some settings, less strict communication constraints reduce the rendezvous time; in others, it increases it. An example of this can be seen for the Ring office when FBE-BT-A is used in Fig. 12(a). By inspecting the robot's behaviour in these runs, we have seen how this is due to the following events: sometimes, two robots do a partial rendezvous very early in the exploration run. As they prefer to go towards high-connectivity areas, they may choose to explore the same corridor and do a partial rendezvous almost immediately. This has a positive effect of allowing the two robots to meet quickly; as a negative side effect, an early rendezvous results in the fact that their combined exploration trace covers only a portion of the map. Sharing the joint exploration trace has a great impact in facilitating rendezvous as it encourages the MRS to re-visit all the portions of the map already observed by any of the team members; in this case, because the robots have just started exploration before doing a partial rendezvous, this step adds a little knowledge to the team, thus reducing the effect of this process and increasing the time required for exploration. When larger communication ranges are used, the speed-up due to the reduced communication constraints is greater than the negative effect due to a less significant sharing of exploration traces. Nevertheless, in all cases, the use of a backtracking mechanism and biasing the exploration towards high connectivity areas has, on average, a benefit against using FBE.

4.8. Evaluation in more complex maps

We then evaluate FBE and FBE-BT in two additional complex environments, Campus-1 and Campus-2, using 3 robots. Results are reported in Table 4, while the maps and an example of a set of trajectories per method is shown in Fig. 13. While the two environments present different types of challenges, the results obtained confirm our findings and the robustness of our approach, as our method is always able to conclude the run with a rendezvous ($R = 1$), while requiring less time when compared with FBE. Despite the size of the environments, and their complex topology, results confirm those observed in the other maps, showing a trend between a strong decrease of the total time Δt with a minor decrease of the total explored area Δa.

5. Conclusion and future works

In this work, we presented a framework that enables a Multi-Robot System (MRS) to rendezvous in an unknown environment while simultaneously exploring it. To achieve this, we have introduced a mechanism for information decay that builds on a frontier-based exploration approach.

Table 4
Comparison of FBE and FBE-BT in 5 large-scale environments, when the MRS is of 3 robots. The last two columns report the Δs, comparing the result of FBE-BT compared with FBE.

Environment	FBE-BT					FBE					Δ	
	t		a		R	t		a		R	Δt	Δa
	μ_t	σ_t	μ_a	σ_a		μ_t	σ_t	μ_a	σ_a			
Ring	1652.30	1041.26	0.871	0.091	1	2483.54	1484.63	0.914	0.090	0.7	33.47	4.79
Office	727.46	372.16	0.639	0.101	1	851.11	791.46	0.678	0.187	0.9	14.53	5.75
Intel	580.13	256.63	0.760	0.147	1	688.19	553.28	0.782	0.169	0.9	15.70	2.81
Campus-1	1070.29	487.82	0.772	0.104	1	1591.70	1246.16	0.813	0.189	1	32.76	5.04
Campus-2	818.09	503.57	0.631	0.157	1	1370.68	959.97	0.745	0.171	1	40.31	15.30

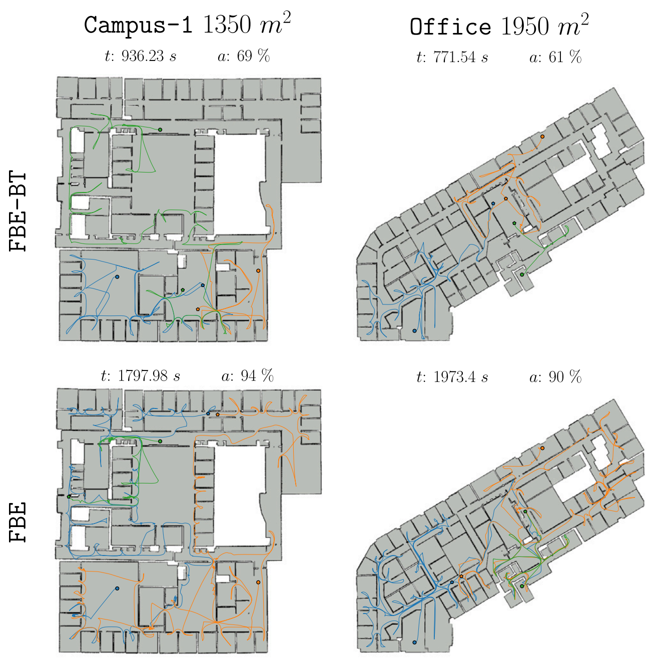

Fig. 13. Rendezvous paths obtained with our method (FBE-BT - first row) and the baseline (FBE - second row) two complex environments with multiple rooms.

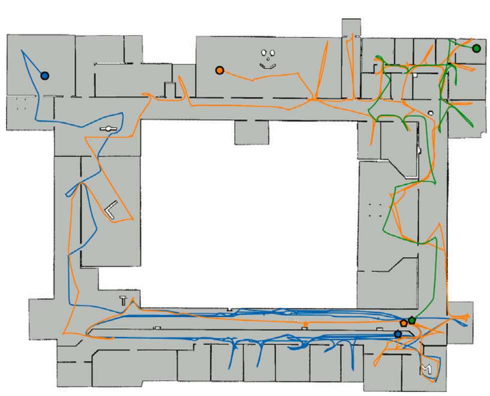

Fig. 14. An example of an undesirable behaviour obtained when FBE-BT-U(0.5) is used: as the robots are strongly biased to visit high-connectivity areas such as corridor, the robot blue robot engages a repetitive behavior where it keeps traversing the two corridors at the bottom of the map, clockwise, three times. Eventually, the orange robot meets the blue one, solving this deadlock.

This mechanism encourages the robots to backtrack to previously explored areas, an action which facilitates accidental rendezvous in highly connected parts of the environment, such as corridors and hallways. Our experimental evaluation clearly demonstrates that our framework allows the robots to complete a rendezvous in less time when compared to a traditional frontier-based exploration approach. We also note that standard frontier-based exploration methods can still be effective for rendezvous when there is a higher number of robots involved.

Understanding high-connectivity areas further increases the chances of accidental rendezvous by encouraging the robots to travel through these regions. However, we identify a trade-off between the total area explored by the MRS and the time required for a rendezvous. While our method can significantly reduce the time needed for rendezvous, it only slightly decreases the total area explored by the MRS.

Future work will focus on identifying high-connectivity areas directly from the partial map available to the MRS and perception data. Moreover, we will investigate how our method can be adapted to outdoor and unstructured environments. In those environments, high-connectivity areas are more difficult to identify than in indoor environments, as their definition is less straightforward. Nevertheless, domain-specific features may help identify locations where accidental rendezvous is more likely, for example large open fields in a forest or public squares in urban scenarios.

The concept of high-connectivity areas can also be further investigated by proposing asymmetric rendezvous strategies, where robots can decide to wait in high-connectivity areas to maximize the chance of a rendezvous, or when different types of heterogeneous robots with different capabilities and strategies are employed.

Our proposed framework can be eventually applied to settings where multi-robot rendezvous is required, such as exploration, search and rescue, surveillance or tasks involving multi-robot task allocation problems.

Finally, future work involves conducting empirical evaluations of our approach using real multi-robot systems.

CRediT authorship contribution statement

Matteo Luperto: Writing, Supervision, Conceptualization, Visualization, Methodology, Investigation, Validation . **Mauro Tellaroli:** Writing, Visualization, Validation, Software, Resources, Methodology, Investigation. **Michele Antonazzi:** Writing, Validation, Supervision, Data curation. **Nicola Basilico:** Writing, Supervision, Resources, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

[1] M. Corah, N. Michael, Efficient online multi-robot exploration via distributed sequential greedy assignment, in: Proc. RSS, 2017.
[2] L. Iocchi, L. Marchetti, D. Nardi, Multi-robot patrolling with coordinated behaviours in realistic environments, in: Proc. IROS, 2011, pp. 2796–2801.
[3] B. Liu, X. Xiao, P. Stone, Team orienteering coverage planning with uncertain reward, in: Proc. IROS, 2021, pp. 9728–9733.
[4] Y. Liu, G. Nejat, Multirobot cooperative learning for semiautonomous control in urban search and rescue applications, J. Field Robot. 33 (4) (2016) 512–536.
[5] A. Khamis, A. Hussein, A. Elmogy, Multi-robot task allocation: A review of the state-of-the-art, Coop. Robot. Sens. Netw. (2015) 31–51.
[6] N. Mathew, S.L. Smith, S.L. Waslander, Multirobot rendezvous planning for recharging in persistent tasks, IEEE Trans. Robot. 31 (1) (2015) 128–142.
[7] Y. Jang, C. Oh, Y. Lee, H.J. Kim, Multirobot collaborative monocular SLAM utilizing rendezvous, IEEE Trans. Robot. 37 (5) (2021) 1469–1486.
[8] F. Amigoni, J. Banfi, N. Basilico, Multirobot exploration of communication-restricted environments: A survey, IEEE Intell. Syst. 32 (6) (2017) 48–57.
[9] B. Yamauchi, A. Schultz, W. Adams, K. Graves, Integrating map learning, localization and planning in a mobile robot, in: Proc. ISIC, 1998, pp. 331–336.
[10] J.G. Rogers, C. Nieto-Granda, H.I. Christensen, Coordination strategies for multi-robot exploration and mapping, in: Proc. ISER, 2013, pp. 231–243.
[11] M. Luperto, M. Antonazzi, F. Amigoni, N.A. Borghese, Robot exploration of indoor environments using incomplete and inaccurate prior knowledge, Robot. Auton. Syst. 133 (2020) 103622.
[12] M. Meghjani, G. Dudek, Combining multi-robot exploration and rendezvous, in: Proc. CRV, 2011, pp. 80–85.
[13] M. Meghjani, G. Dudek, Multi-robot exploration and rendezvous on graphs, in: Proc. IROS, 2012, pp. 5270–5276.
[14] K. Zheng, A. Pronobis, R. Rao, Learning graph-structured sum-product networks for probabilistic semantic maps, in: Proc. AAAI, Vol. 32, 2018.
[15] H. Zender, O.M. Mozos, P. Jensfelt, G.-J. Kruijff, W. Burgard, Conceptual spatial representations for indoor mobile robots, Robot. Auton. Syst. 56 (6) (2008) 493–502.
[16] S. Garg, N. Sünderhauf, F. Dayoub, D. Morrison, A. Cosgun, G. Carneiro, Q. Wu, T.-J. Chin, I. Reid, S. Gould, et al., Semantics for robotic mapping, perception and interaction: A survey, Found. Trends Robot. 8 (1–2) (2020) 1–224.
[17] M. Luperto, A. Quattrini Li, F. Amigoni, A system for building semantic maps of indoor environments exploiting the concept of building typology, in: Proc. ROBOCUP, Springer, 2014, pp. 504–515.
[18] M. Luperto, F. Amigoni, Predicting the global structure of indoor environments: A constructive machine learning approach, Auton. Robots 43 (4) (2019) 813–835.
[19] M. Tellaroli, M. Luperto, M. Antonazzi, N. Basilico, Adapting frontier-based exploration for multi-robot rendezvous in unknown environments, in: Proc. of the Italian Workshop on Artificial Intelligence and Robotics, AIxIA, 2023.
[20] M. Tellaroli, M. Luperto, M. Antonazzi, N. Basilico, Frontier-based exploration for multi-robot rendezvous in communication-restricted unknown environments, in: Proc. IROS, 2024, pp. 5807–5812.
[21] D. Ozsoyeller, A. Beveridge, V. Isler, Rendezvous in planar environments with obstacles and unknown initial distance, Artificial Intelligence 273 (2019) 19–36.
[22] D.M.S. Tan, Y. Ma, J. Liang, Y.C. Chng, Y. Cao, G. Sartoretti, IR2: Implicit rendezvous for robotic exploration teams under sparse intermittent connectivity, in: Proc. IROS, 2024.
[23] D. Ozsoyeller, P. Tokekar, Multi-robot symmetric rendezvous search on the line, IEEE Robot. Autom. Lett. 7 (1) (2021) 334–341.
[24] H. Hourani, E. Hauck, S. Jeschke, Serendipity rendezvous as a mitigation of exploration's interruptibility for a team of robots, in: Proc ICRA, 2013, pp. 2984–2991.
[25] V. Spirin, S. Cameron, Rendezvous through obstacles in multi-agent exploration, in: Proc SSRR, 2014.
[26] M. Andries, F. Charpillet, Multi-robot exploration of unknown environments with identification of exploration completion and post-exploration rendezvous using ant algorithms, in: Proc IROS, 2013, pp. 5571–5578.
[27] U. Dah-Achinanon, S.E. Marjani Bajestani, P.-Y. Lajoie, G. Beltrame, Search and rescue with sparsely connected swarms, Auton. Robot. (2023) 1–15.
[28] C. Pinciroli, G. Beltrame, Buzz: An extensible programming language for heterogeneous swarm robotics, in: Proc. IROS, 2016, pp. 3794–3800.
[29] D. Pittol, M. Mantelli, R. Maffei, M. Kolberg, E. Prestes, Loop-aware exploration graph: A concise representation of environments for exploration and active loop-closure, Robot. Auton. Syst. 155 (2022).
[30] A.R. da Silva, L. Chaimowicz, T.C. Silva, A. Hsieh, Communication-constrained multi-robot exploration with intermittent rendezvous, in: Proc. IROS, 2024.
[31] K. Song, G. Chen, W. Liu, Z. Xiong, Multi-robot rendezvous in unknown environment with limited communication, IEEE Robot. Autom. Lett. (2024).
[32] L. Bramblett, R. Peddi, N. Bezzo, Coordinated multi-agent exploration, rendezvous, & task allocation in unknown environments with limited connectivity, in: Proc. IROS, 2022, pp. 12706–12712.
[33] A. Ganguli, J. Cortés, F. Bullo, Multirobot rendezvous with visibility sensors in nonconvex environments, IEEE Trans. Robot. 25 (2) (2009) 340–352.
[34] G.A. Hollinger, S. Singh, Multirobot coordination with periodic connectivity: Theory and experiments, IEEE Trans. Robot. 28 (4) (2012) 967–973.
[35] C. Cao, H. Zhu, Z. Ren, H. Choset, J. Zhang, Representation granularity enables time-efficient autonomous exploration in large, complex worlds, Sci. Robot. 8 (80) (2023).
[36] I. Kostavelis, A. Gasteratos, Semantic mapping for mobile robotics tasks: A survey, Robot. Auton. Syst. 66 (2015) 86–103.
[37] N. Sünderhauf, F. Dayoub, S. McMahon, B. Talbot, R. Schulz, P. Corke, G. Wyeth, B. Upcroft, M. Milford, Place categorization and semantic mapping on a mobile robot, in: Proc. ICRA, 2016, pp. 5729–5736.
[38] R. Bigazzi, L. Baraldi, S. Kousik, R. Cucchiara, M. Pavone, Mapping high-level semantic regions in indoor environments without object recognition, in: Proc. ICRA, 2024.
[39] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, A. Torralba, Places: A 10 million image database for scene recognition, IEEE Trans. Pattern Anal. Mach. Intell. 40 (6) (2017) 1452–1464.
[40] K. Zheng, A. Pronobis, From pixels to buildings: End-to-end probabilistic deep networks for large-scale semantic mapping, in: Proc. IROS, 2019, pp. 3511–3518.
[41] O. Mendez, S. Hadfield, N. Pugeault, R. Bowden, SeDAR: reading floorplans like a human—using deep learning to enable human-inspired localisation, Int. J. Comput. Vis. 128 (5) (2020) 1286–1310.
[42] N. Akai, SLAMER: Simultaneous localization and map-assisted environment recognition, in: Proc. ICRA, IEEE, 2023, pp. 6203–6209.
[43] N. Zimmerman, T. Guadagnino, X. Chen, J. Behley, C. Stachniss, Long-term localization using semantic cues in floor plan maps, IEEE Robot. Autom. Lett. 8 (1) (2022) 176–183.
[44] A. Quattrini Li, R. Cipolleschi, M. Giusto, F. Amigoni, A semantically-informed multirobot system for exploration of relevant areas in search and rescue settings, Auton. Robots 40 (2016) 581–597.
[45] A. Pronobis, P. Jensfelt, Large-scale semantic mapping and reasoning with heterogeneous modalities, in: Proc. ICRA, 2012, pp. 3515–3522.
[46] G. Grisetti, C. Stachniss, W. Burgard, Improved techniques for grid mapping with Rao-Blackwellized particle filters, IEEE Trans. Robot. 23 (2007) 34–46.
[47] A. Testa, A. Camisa, G. Notarstefano, ChoiRbot: A ROS 2 toolbox for cooperative robotics, IEEE Robot. Autom. Lett. 6 (2) (2021) 2714–2720.
[48] A. Howard, N. Roy, The robotics data set repository (Radish), 2003, URL http://radish.sourceforge.net/.
[49] F. Amigoni, V. Castelli, M. Luperto, Improving repeatability of experiments by automatic evaluation of SLAM algorithms, in: Proc. IROS, 2018, pp. 7237–7243.
[50] J.A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, J.A. Castellanos, A survey on active simultaneous localization and mapping: State of the art and new frontiers, IEEE Trans. Robot. 39 (3) (2023) 1686–1705.

Matteo Luperto is an assistant professor the Dipartimento di Informatica Giovanni Degli Antoni at the Università degli Studi di Milano(Italy). He received his Ph.D. in Information Technology from the Politecnico di Milano (Italy) in 2017. His main research interests lie within the field of semantic mapping for autonomous mobile robots in indoor environments, with a particular attention on the analysis of the structural properties of buildings, and long-term autonomy for social assistive mobile robots.

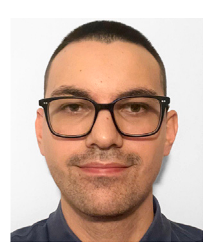

Mauro Tellaroli received his B.Sc. in Computer Science in 2023 from the University of Milan. He collaborates with AISLab at the Department of Computer Science, University of Milan. His research focuses on multi-robot systems.

Michele Antonazzi received his B.Sc. in Computer Science in 2018 from the University of Padova and his M.Sc. in Computer Science from the University of Milan in 2021. He is currently a PhD candidate in Robotics. His research investigates domain adaptation and privacy preservation for Robotic Vision, specifically focused to the paradigm of cloud robotics.

Nicola Basilico received an M.Sc. degree in Computer Science and Engineering in 2007 and a Ph.D. in Information Technology in 2011 from Politecnico di Milano (Italy). He has been a postdoctoral scholar at the Robotics Laboratory at the University of California, Merced, and worked as a research assistant at the Swiss AI Lab IDSIA. From 2014 to 2019, he was an assistant professor, and since 2020, he has been an associate professor at the Department of Computer Science of the University of Milan (Italy). His main research interests lie in the area of Artificial Intelligence, with a particular focus on Multi-Agent Systems and Autonomous Robotics.