Adapting Frontier-Based Exploration for Multi-Robot Rendezvous in Unknown Environments

Mauro Tellaroli[†], Michele Antonazzi, Matteo Luperto[*,†] and Nicola Basilico

Department of Computer Science, University of Milan, via Celoria 18, 20133, Milan, Italy.

Abstract

Multi-robot rendezvous and exploration are fundamental challenges in the domain of mobile robotic systems. This paper addresses the coordination problem of enabling multiple robots to efficiently converge on a common location within an initially unknown and communication-restricted indoor environment. The robots start from distinct locations with no prior knowledge of the environment's layout or any pre-established meeting point. Communication is only possible after rendezvous, adding an extra layer of complexity. In this context, we propose a novel approach that integrates exploration and rendezvous into a unified strategy, hence bridging the gap between the two. Traditionally, exploration has been focused on rapidly mapping the environment, often leading to suboptimal rendezvous performance in later stages. Our method adapts standard frontier-based exploration techniques, prioritizing frontiers not only for map expansion but also for rendezvous opportunities. We introduce a mechanism that allows robots to "forget" previously explored regions, redirecting their attention to unexplored areas and enhancing rendezvous likelihood. To evaluate our approach, we conduct experiments in realistic 3D simulations using ROS, showcasing its effectiveness in achieving faster rendezvous times compared to conventional exploration strategies while maintaining a comparable level of environmental coverage.

Keywords

Multirobot rendezvous, Frontier-based exploration

1. Introduction

Multi-robot rendezvous is a coordination problem characterizing many multi-robot systems (MRS) where mobile robots must efficiently travel to or in the immediate vicinity of a common location in a shared environment. A rendezvous strategy is typically evaluated by the total time or distance taken for all robots to reach the meeting point, the smaller the better. These metrics are often interpreted as proxies for both energy consumption and quality of service.

Computing and executing efficient rendezvous strategies represent key components in MRS application domains where robots need to physically meet to share collected information or collaborate on some localized task. Distributed data-gathering offers many settings where multi-robot rendezvous is required or can play a fundamental role.

10th Italian Workshop on Artificial Intelligence and Robotics (AIRO 2023)
mauro.tellaroli@studenti.unimi.it (M. Tellaroli); michele.antonazzi@unimi.it (M. Antonazzi); matteo.luperto@unimi.it (M. Luperto); nicola.basilico@unimi.it (N. Basilico)
0000-0001-6396-7567 (M. Antonazzi); 0000-0002-8976-2073 (M. Luperto); 0000-0002-4512-3480 (N. Basilico)

Examples include autonomous exploration [1], persistent surveillance or monitoring [2, 3], and search and rescue [4]. The distributed nature of these tasks reflects the lack of a centralized infrastructure covering the whole environment and enabling global communication and coordination among robots. Instead, robots have to rely on peer-to-peer interactions which, being subject to minimum-range constraints, require physical proximity. Meeting with teammates enables sharing partial maps in exploration or exchanging findings and collected data in monitoring and search, hence allowing robots to compute more informed plans for their common task. Communication is not the only domain where rendezvous might play a role. Multi-robot task allocation scenarios [5] often feature tasks that, to be executed, require concurrent cooperation by more robots (perhaps also heterogeneous ones). In such cases, meeting at the same (task) location is the pre-condition to complete joint task assignments. In this work, we consider the multi-robot rendezvous problem in the challenging setting of a communication-restricted [6] and initially unknown indoor environment. We shall assume that the robots start from a set of given different 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 fact that the environment is communication-constrained implies that communications are possible only after a rendezvous. An example of a situation where this problem happens is when two or more people need to gather in an environment that they don't know, such as in a shopping mall, a hospital, or an airport.

In this scenario, the rendezvous problem is augmented by the difficulty of online exploration, namely the problem of building a complete and accurate map of an initially unknown environment. Exploration is customarily performed by steps. (i) The robot identifies a set of promising candidate locations in the already mapped portion of the environment. (ii) Among them, the robot selects the most promising one, computing utilities and costs of taking perceptions along the way to that location. (iii) Upon selection, the location is reached and the obtained perceptions are integrated into the partial map of the environment. The process then repeats from (i) until a termination condition is met (often a threshold over the percentage of the explored area). When a multi-robot team is involved, each robot can carry out its exploration independently, or they can coordinate by communicating with each other, and exchanging maps and plans [7]. A popular choice in step (i) is to use frontiers [8], namely the boundaries between the explored and unexplored part of the environment, as candidate locations, and to select the most promising one according to an exploration strategy.

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 ones happen when the meeting is already scheduled among agents following some kind of coordination. The latter ones take place without a pre-arrangement. 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 [9].

Multi-robot rendezvous and exploration have been studied together, mainly as alternating phases that robots undergo in the scope of a given task. Often the two phases of exploration and rendezvous are mutually exclusive; when the robots are exploring, their interest is to acquire new knowledge; when the robots decide to rendezvous, they travel to a location in the mapped area. This decoupling between the two phases can introduce suboptimal performance, as robots are forced to alternate between the two behaviors in a coordinated way. The work of [10] tries to reduce such a limitation by introducing the concept of serendipity in exploration, i.e., to create a robot behavior that tries to facilitate an unplanned encounter with other robots while those are still in an exploration phase, by adopting a behavior that is a mix between episodic and planned rendezvous.

In this work, we study how to address rendezvous and exploration simultaneously in a multi-robot system. Importantly, we investigate how standard frontier-based techniques for exploration can be extended to tackle both problems. Frontier–based exploration strategies are widely used as they are simple yet effective, allowing robust exploration in heterogeneous contexts. The inherent greedy nature of frontier-based exploration (see (i)-(iii) above) results in quickly mapping those regions where the most free space lies, leaving behind portions of the environment that are less informative. As a consequence of this fact, robots quickly explore the larger portion of the map; after that, they spend a considerable amount of time filling the gaps of the not-selected frontiers. As noted in [11, 12], in certain scenarios, up to 71% of the total exploration time is spent covering the last 10% of space. This kind of profile is a direct consequence of the typical criterion that drives exploration, namely to obtain the largest map of the environment in the shortest time. 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. In our method, we propose an exploration strategy that is biased towards rendezvous. In such a strategy, frontiers are assigned a selection priority that not only is based on their potential contribution to the map's expansion (classical frontier-based approach) but also on the opportunity to meet a teammate. To do so, we introduce a mechanism that allows each robot to forget about parts of the environment that have been mapped, thus putting those parts back into the portions of the environment still to be explored. Differently from other methods that evaluate the exploration and rendezvous problem on graphs or 2D simulated environments, we test our framework in 3D realistic simulations made with ROS. Preliminary results show how the proposed method allows the MRS to meet in less time required for a standard exploration run while, at the same time, they perform exploration.

2. Our Method

2.1. Problem Formulation

We consider a team of m homogeneous robots equipped with the same perception, navigation, and communication capabilities. We assume a strongly limited communication model: each robot has a communication range of radius d and, to exchange any information, two robots must be inside each other's range and have an unobstructed line of sight. Each robot starts from a random location, unknown to the others, in a given environment whose map is initially unknown. Our objective is to have the robots form a connected group so that, from that moment on, they can collectively plan and navigate in the environment while keeping a formation.

Formally, we define the rendezvous condition by requiring that the m locations occupied by the team members satisfy a hard and a soft constraint. The hard constraint requires the robots to form a connected configuration as per the communication model described above (so each robot should, in principle, be able to communicate with another one in a multi-hop fashion). The soft constraint requires that robots should be close to each other, to navigate in a formation. Ideally, this requirement should enforce a maximum distance between the locations of any pair of robots. However, setting it as a hard constraint to reach a rendezvous would in many cases prevent the problem's feasibility. Indeed, with many robots and in environments that do not offer open-area regions, connected groups of robots might not find a way to maneuver into a joint configuration that, without collisions, realizes the required mutual uniform-distance formation. For this reason, in our method, we seek this condition as much as possible while allowing for its violation when no other option is available.

2.2. Robot Exploration for Rendezvous

Our proposed method is based on extending a standard frontier-based exploration strategy [8], allowing the robot to perform episodic rendezvous. A frontier $f \in F$ is the boundary between a part of the map that is explored and is free (i.e., without obstacles) and an unknown part of the environment (i.e., not mapped yet). Each robot implements a frontier-based exploration strategy, which is executed independently of the other robots. The common frontier-based exploration process is an iterative process consisting of the following steps:

1. extracting from the map a list of frontiers $\mathscr{F}$; from each frontier, select a candidate location to reach (e.g., as in our case, the centroid of the frontier);
2. ranking the frontiers according to an exploration strategy, and select the next best location;
3. planning and executing a path towards the frontier, integrating the perceptions acquired during the path into the map;
4. once the frontier is reached, restart from (1) until no frontier is left.

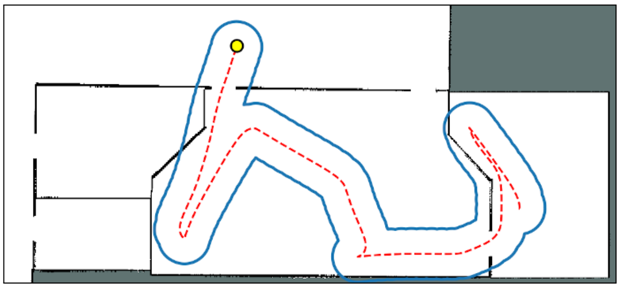

Figure 1: In blue is represented the *exploration trace* of the robot that performed the trajectory highlighted in red.

Using this framework, the robot aims to reach the boundaries of its current map, ideally without going back on its steps but continuously moving forward, until the whole environment is mapped. In our method, we bias exploration by introducing an information decay mechanism on the mapping process so that the robot is also driven to go back on its steps, following the intuition that this backtracking mechanism will promote accidental rendezvous among robots.

To do so, we keep track of the exploration trace E of the robot, i.e., the area around the trajectory of the robot during the exploration run. An example of that area is shown in Figure 1. The exploration trace is obtained by taking track of an ordered set of robot poses P from its trajectory, sampled at a given frequency from the global robot trajectory. We use the communication range d as the width of each pose, centered around the robot E, thus creating a circular pattern of radius d around each pose.

We create a new set of frontiers $\bar{f} \in \bar{F}$ that is obtained by using an information decay mechanism on the exploration trace. More precisely, we retain a subset of the robot exploration trace containing the $N \leq P$ most recent robot poses. Each pose is characterized by a timestamp indicating when it was obtained; poses are removed after a time t is passed after their acquisition. When a pose is removed we artificially create a new frontier $\bar{f} \in \bar{F}$ as the contour of the difference of the area before and after the removal of the pose. An example of this mechanism is shown in Figure 2, where such artificial frontiers are highlighted in dashed red.

As a result of this, while in a standard frontier-based exploration the robot selects the next most promising location from a set of frontiers $\mathcal{F} = F$, in our method the robot selects the most promising location from $\mathcal{F} = F \cup \bar{F}$. In this way, the robot is pushed to partially backtrack on its steps to revisit parts of the environments that have already been explored.

In this paper, we rank frontiers according to a linear combination of their distance from the current position of the robot, and their length, similar to what is proposed in [8]. We consider as the most promising frontiers those that are large frontiers and are close to the robot. More precisely, indicating with $dist(f)$ the distance between the current position of the robot and the centroid of the frontier, and with $len(f)$ its length, the cost of a frontier is indicated as $\theta(f) = \alpha * len(f) - (1 - \alpha) * dist(f)$. Note that, as our method

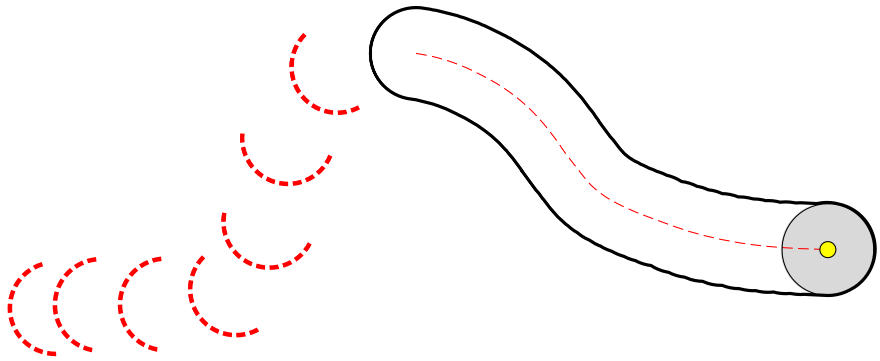

Figure 2: Frontiers in dashed red are created using our *information decay* method.

adds new frontiers to $\mathscr{F}$, our method can be used with other exploration strategies.

2.3. Multi-robot information sharing

When two or more robots are within their sensing range d, they join a common cluster. In each cluster, one robot is elected as a leader, while the other robots are termed followers. The robots' leader in a cluster continues the frontier-based exploration, while the others are tasked with following it. Each cluster C formed by robots $\{R_i, R_j, ...\}$ has a set of frontiers $\mathscr{F}_C = \{\mathscr{F}_1, \mathscr{F}_2, ...\}$ and a shared exploration trace $E_C = \{E_1, E_2, ...\}$ that is created by sharing the frontiers $\mathscr{F}_i$ and the exploration trace E_i for all robots $R_i \in C$.

The process of frontier and exploration-trace sharing is performed once each time a robot i joins a cluster; after that, the knowledge is retained by the leader of the cluster. During this process, we remove also frontiers $\bar{F}_C$ that overlap with the cluster exploration trace E_C, as E_C represents the area of the environment that was recently jointly explored by the cluster. An example of the process of merging two exploration traces E_i and E_j when the two robots merge in a single cluster C is shown in Figure 3. After a cluster is formed, only the leader of the cluster adds frontiers to $\mathscr{F}_C$. The leader of the cluster is selected using a predefined agreement (e.g., the lexicographic order).

The robots belonging to the same cluster move together in a formation, following the cluster leader. If a robot loses connection with the remainder of a cluster, as an example because it was stuck after hitting an obstacle, the robot is removed from the cluster and resumes the frontier-based exploration task.

3. Preliminary Results

We implemented our method using ROS [13]. Our multi-robot configuration is based on the namespace system of ROS, which allows us to logically separate the control stack

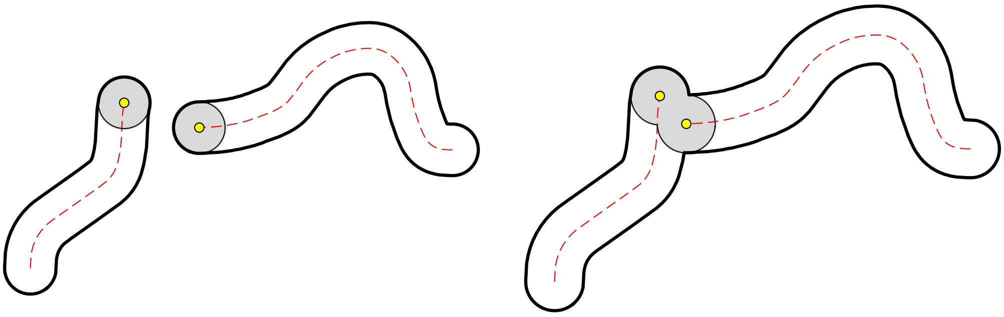

Figure 3: Example of two robots that, after they meet (on the right) share their exploration trace.

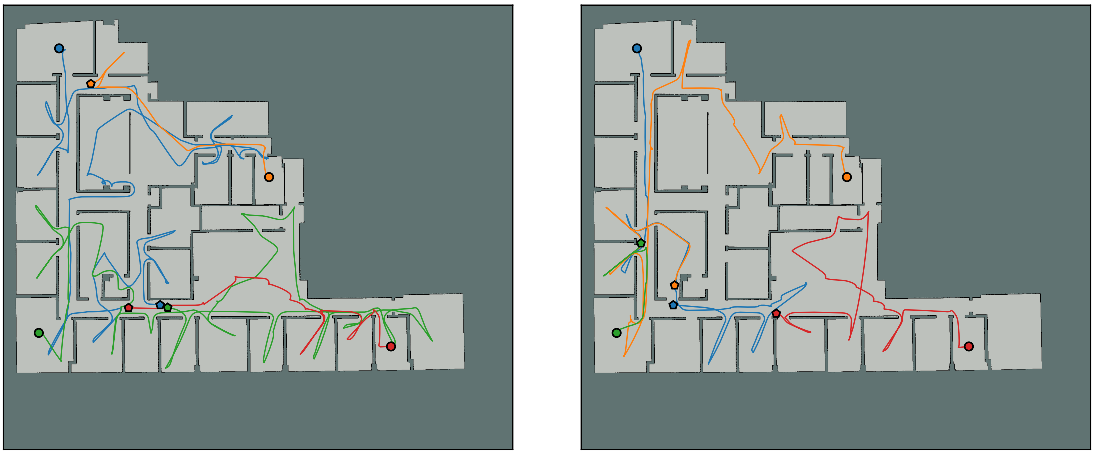

Figure 4: The paths followed by the robots using the classical frontier exploration strategy (FE, left) compared with our method (`Ours`, right). Each robot has a different color. The circles indicate the starting position of the robot, while the stars indicate the location when a robot joins a cluster. 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.

associated with each robot. We compare the results obtained by our method (label `Ours`) against those obtained with a standard frontier-based exploration strategy as in [8] (label FE).

The experiments are performed in a complex indoor environment simulated in Gazebo, using 4 robots, and performing 6 random runs for each method and averaging the results. During each exploration run, we record the time t required to perform a rendezvous among all robots, the size $\max |C|$ of the largest cluster of robots, and the combined area explored by the robots A. A run is ended when $\max |C|$ is equal to the number of robots.

We used a Turtlebot3-burger robot with a speed of $0.3\,\mathrm{m\,s^{-1}}$ and with a 2D lidar with a range of $10\,\mathrm{m}$; each robot performs SLAM using Gmapping [14], and is equipped with its

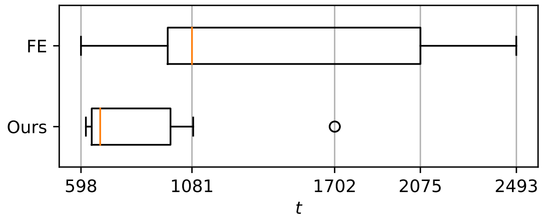

Figure 5: The time to complete rendezvous using the frontier exploration strategy (FE) against our method over 6 runs.

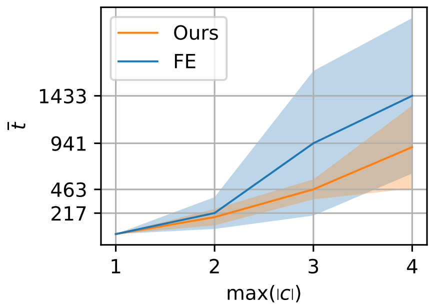

Figure 6: The average time *bart* (and standard deviation) to complete the rendezvous changing the maximum size of the cluster.

own navigation stack implemented using `move base`[1]. We simulate sensors and actuation errors of the robot, to have a setting that poses challenges similar to those of a real-world experimental run. We set a parameter of 333 s for the information decay of a robot pose in E, and we consider a sensing range for multi-robot clustering and communication of 2.7 m. A new pose is added to the exploration trace each 2 s, and we add a frontier to $\bar{F}$ after 9 poses have been collected (18 s). These parameters are experimentally set to have a set of frontiers $\bar{F}$ that cover the robot path, without adding too many frontiers. We set $\alpha = 1/4$.

Figure 4 shows the trajectories performed in two of the runs of FE (left) against our method (right). While with FE multiple robots cover the same parts of the environment, our method allows the robot to perform a shorter trajectory before meeting altogether.

Figure 5 shows the average time required by the four robots to perform a rendezvous, while Figure 6 shows the average time and standard deviation for forming a cluster of sizes 2, 3, and 4, respectively. It can be appreciated how our method not only allows the robot to perform a rendezvous in less time but also how our performance is more stable across different runs, with a lower standard deviation.

[1] http://wiki.ros.org/move_base

4. Conclusion and future works

In this work, we have presented a framework for allowing an MRS to perform a rendezvous in a previously unknown environment while performing exploration. To do so, we have introduced a mechanism for information decay on top of a frontier-based exploration approach. Preliminary results are promising and clearly show how our framework allows the robot to perform a rendezvous in less time than a classic frontier-based exploration approach. Future results will involve a more exhaustive experimental evaluation.

References

[1] M. Corah, N. Michael, Efficient online multi-robot exploration via distributed sequential greedy assignment., in: Proc. RSS, volume 13, Cambridge, MA, 2017.

[2] L. Iocchi, L. Marchetti, D. Nardi, Multi-robot patrolling with coordinated behaviours in realistic environments, in: Proc. IROS, IEEE, 2011, pp. 2796–2801.

[3] B. Liu, X. Xiao, P. Stone, Team orienteering coverage planning with uncertain reward, in: Proc. IROS, IEEE, 2021, pp. 9728–9733.

[4] Y. Liu, G. Nejat, Multirobot cooperative learning for semiautonomous control in urban search and rescue applications, Journal of Field Robotics 33 (2016) 512–536.

[5] A. Khamis, A. Hussein, A. Elmogy, Multi-robot task allocation: A review of the state-of-the-art, Cooperative robots and sensor networks 2015 (2015) 31–51.

[6] F. Amigoni, J. Banfi, N. Basilico, Multirobot exploration of communication-restricted environments: A survey, IEEE Intelligent Systems 32 (2017) 48–57.

[7] J. G. Rogers, C. Nieto-Granda, H. I. Christensen, Coordination strategies for multi-robot exploration and mapping, in: Experimental Robotics: The 13th International Symposium on Experimental Robotics, Springer, 2013, pp. 231–243.

[8] 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.

[9] D. Ozsoyeller, A. Beveridge, V. Isler, Rendezvous in planar environments with obstacles and unknown initial distance, Artificial Intelligence 273 (2019) 19–36.

[10] H. Hourani, E. Hauck, S. Jeschke, Serendipity rendezvous as a mitigation of exploration's interruptibility for a team of robots, in: Proc. ICRA, IEEE, 2013, pp. 2984–2991.

[11] L. Ericson, D. Duberg, P. Jensfelt, Understanding greediness in map-predictive exploration planning, in: Proc. ECMR, IEEE, 2021, pp. 1–7.

[12] M. Luperto, M. Antonazzi, F. Amigoni, N. A. Borghese, Robot exploration of indoor environments using incomplete and inaccurate prior knowledge, Robotics and Autonomous Systems 133 (2020) 103622.

[13] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, A. Ng, Ros: an open-source robot operating system, in: Proc. ICRA, Workshop on Open Source Robotics, 2009.

[14] G. Grisetti, C. Stachniss, W. Burgard, Improved techniques for grid mapping with Rao-Blackwellized particle filters, IEEE Transactions on Robotics 23 (2007) 34–46.