



Search and rescue with sparsely connected swarms

Ulrich Dah-Achinanon¹ · Seyed Ehsan Marjani Bajestani¹ · Pierre-Yves Lajoie¹ · Giovanni Beltrame¹

Received: 1 December 2021 / Accepted: 27 November 2022 / Published online: 18 January 2023
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

Designing and deploying autonomous swarms capable of performing collective tasks in real-world is extremely challenging. One drawback of getting out of the lab is that realistic tasks involve long distances with limited numbers of robots, leading to sparse and intermittent connectivity. As an example, search and rescue requires robots to coordinate in their search, and relay the information of found targets. The search's effectiveness is greatly reduced if robots must stay close to maintain connectivity. This paper proposes a decentralized search system that only requires sporadic connectivity and allows information diffusion through the swarm whenever possible. Our robots share and update a distributed belief map, to coordinate the search. Once a target is detected, the robots form a communication relay between a base station and the target's position. We show the applicability of our system both in simulation and with real-world experiments with a small swarm of drones.

Keywords Swarm robotics · Drone · Sparse connection · Search-and-rescue · Real-world deployment · Rendezvous point

1 Introduction

Drones or unmanned aerial vehicles (UAVs) have been experiencing steady growth for the past few years in terms of their popularity, availability, and potential. This enables a wide range of applications in several fields going from surveillance to search and rescue missions (Apvrille et al., 2014). The emergence of swarms systems, made of multiple UAVs collaborating towards a common goal, further improves the efficiency of those solutions. Swarm robotics has been defined by Brambilla et al. (2013), as an approach to

collective robotics inspired by the self-organized behaviors of social animals, where a large group of simple robots aims to accomplish a complex task through simple rules and local interactions.

Using teams of robots, instead of a single one, brings robustness, scalability, survivability, and it increases the speed of execution (Hentati and Fourati, 2020). Despite those practical benefits, many unresolved challenges prevent swarm robotics from being used in commercial products. In particular, swarms of UAVs are confronted with multiple challenges such as: in-flight coordination, swarm layout reconfiguration, handling losses of swarms elements, and data relaying optimization among others (Wubben et al., 2020). Tarapore et al. (2020) outline one of those challenges and formalizes the notion of sparse swarms in which it can be prohibitively expensive for the robots to maintain close proximity. For example, during swarm-based search and rescue (SAR) operations, preserving close proximity among the robots would certainly restrain the area that can be covered. Therefore, a robust ad-hoc communication system, resilient to disconnections between the swarm agents, is essential to deploy such systems in realistic scenarios.

In this article, we attempt to bridge the gap between theoretical approaches and practical applications by proposing a SAR algorithm based on ad-hoc networks accepting sporadic connectivity. This algorithm leverages a search pattern accepting sparsely connected swarms by scheduling a ren-

Ulrich Dah-Achinanon, Seyed Ehsan Marjani Bajestani and Pierre-Yves Lajoie have contributed equally.

✉ Ulrich Dah-Achinanon
ulrich.dah-achinanon@polymtl.ca

✉ Seyed Ehsan Marjani Bajestani
ehsan.marjani@polymtl.ca

Pierre-Yves Lajoie
pierre-yves.lajoie@polymtl.ca

Giovanni Beltrame
giovanni.beltrame@polymtl.ca

¹ MIST Laboratory, Department of Computer Engineering and Software Engineering, Polytechnique Montréal, Chemin de Polytechnique, Montréal, Québec H3T 1J4, Canada

dezhvous point for periodic meetings and target discovery report. It then allows the swarm to adopt a relay tree formation connecting the meeting point (or base) to all the detected targets. Thus, the connectivity between the ground operators and the robots is restored only when targets are found. The search method is based on belief space exploration in order to incorporate crucial priors from the authorities, such as the last known locations of the targets.

This paper argues that ad-hoc networks accepting sporadic connectivity are the key to real-world deployments of swarms of UAVs in large areas. Such networks should handle adaptive topology and routing while assuring reliable data exchange within the swarm. Existing works exploit either a rendezvous point or a relay chain formation to improve the communication links, but, to the best of our knowledge, none combine these for SAR applications. For instance, Nickerson (2004), Hourani et al. (2013), and Belkadi et al. (2016) leverage a rendezvous point for exploration and SAR, but they do not use communication relays. Other works present relay chain architectures for communication enhancement (Li, 2019; Varadharajan et al., 2020), without considering the search pattern. This article aims to bridge these two principles with the mentioned algorithm. To summarize, the main contributions presented in this paper are:

- A search pattern based on rendezvous point, using ad-hoc networks with adaptive topology and routing, and accepting sparsely connected swarms;
- An algorithm to create an adaptive relay tree structure design to maintain the communication between the ground operators and the discovered targets;
- An implementation of a search method based on dynamic, distributed belief maps;
- An experimental robotic system for real-world deployment of swarms of drones;
- Simulation and real-world deployments of the proposed system.

We validate our approach through tests in simulation and real-world experiments. In simulation, we test our dynamic belief map search algorithm on different area sizes and number of drones. We also performed real-world field tests with three drones to confirm our findings. Figure 1 shows the real-world experiments setup with three DJI M300 quad-copters.

The rest of the paper is organized as follows: in Sect. 2 we present related works outlining similarities and differences with our approach. In Sect. 3 we describe our algorithm and its components. Sections 4 and 5 present our simulations and experiments setup and results. Finally, Sect. 7 draws concluding remarks while presenting possible future works.



Fig. 1 The real-world experiments setup with three DJI M300 quad-copters

2 Related work

As mentioned by Hentati and Fourati (2020), a reliable communication structure is essential to share information among group of neighbors in swarm applications. Sharing the information becomes essential in applications, such as SAR, where a group of robots is trying to find a target in a large area. As indicated by Alotaibi et al. (2019), finding the target could be faster using more robots, however, the lack of reliable communication could separate a group of robots or make the whole mission fail, especially in centralized applications. For instance, Alotaibi et al. (2019) proposed a layered search and rescue (LSAR) centralized “partitioning” algorithm, that needs reliable communication with a cloud server. Although their simulation results indicate that a better success rate could be achieved by increasing the number of robots, it is inherently limited by the central communication bottleneck to the server. Recent work (Ruetten et al., 2020) introduced an optimized self-organized mesh network to cover large areas, but do not consider disconnections.

To solve these communication issues in a team of robots during exploration missions, Hourani et al. (2013) considered a periodic rendezvous strategy in order to overlap the communication ranges of the robots. It also presented an approach to mitigate the negative impact of these meetings on the time efficiency of the overall mission. Our approach is similar, but we drop the connectivity maintenance requirement during the searching phase. Another benefit of our technique is that the drones continue to search for the targets while going to the periodic meetings. Andries and Charpillet (2013) and Andries and Charpillet (2015) also used a meeting point, but the robots only meet at the rendezvous when the exploration is completed. Adopting a different approach, Belkadi et al. (2016) plans the exploration in order to converge to a predefined spatial configuration around the rendezvous point.

Belief maps have been studied for a long time as a tool for multi-robot exploration (Kobayashi et al., 2002, 2003). Similar to our work, Khan et al. (2014) updates the belief map with local observations and merges data from multiple UAVs.

Our distributed belief map implementation is an adaptation of the work proposed by Vielfaure et al. (2021), in which the authors stored the shared belief map in a distributed database called virtual stigmergy (Pinciroli et al., 2016).

During the rescue phase of SAR missions, it is crucial to maintain the connection between the base station and the UAVs following the target. To this end, we propose to maintain a relay chain from the rendezvous position to the targets once these latter are found. Using a heuristic optimization method, Kim et al. (2020) increased the communication performance metric and determined the optimal positions for the communication relay robots. To keep the connection between a heterogeneous group of robots (on-ground and flying robots), Varadharajan et al. (2020) introduced a fully decentralized algorithm to create and keep a chain of robots from the ground station to the target. Li (2019) and Zhang et al. (2021) also used drones as relays and virtual potentials to create a stable link between a group of robots (or a survey drone) and a base station. Instead of using virtual potentials, Yamaguchi et al. (2017) measured the communication quality and expanded the drone relays when needed. While the above-mentioned works focused on the formation of relay chains during the search phase, this could make the search process very slow, which would be critical for SAR operations. Therefore, we propose to create the communication relay chains only in the rescue phase. Our communication relay approach is similar to the approach proposed by Majcherczyk et al. (2018); Çeltek et al. (2018), in which the creation and expansion of tree/chain topologies between drones and target(s) have been evaluated. We use a tree topology for the relay connection when rescuing more than one target.

Sperati et al. (2011) use robots to explore an unknown environment, find two distant target locations, and navigate between them. Their approach could technically be used in our context, but it has some important differences: it does not allow branching like in our method, and it does not exploit the fact that GPS localization is available for our robots. Similarly, Nouyan and Dorigo (2006) show two distributed mechanisms that use visually connected robot structures to form a path between two objects. Unlike our approach, they rely on visual connection between robots and do not explicitly present a branching strategy for multiple targets.

Some existing works use planing based methods to find an exploration path maximizing the information gain. For instance, Zhou et al. (2021) presents a hierarchical framework design for fast UAVs exploration of unknown environments. The system is a frontier based approach that leverage a frontier information system (FIS) that is incrementally maintained to provide the exploration planning with essential information. To detect frontiers and plan efficiently the exploration, the approach uses sensory data from the environment. These data aren't used in our algorithm. To limit

the exploration space and incrementally explore bigger space as in the FUEL algorithm, we used a search_speed parameter that allow the UAV to sample new positions in a limited range around the current position.

Stirling et al. (2010) proposed an energy efficient strategy to coordinate a swarm of UAVs for indoor exploration. Due to the attenuated signals in that environment, the use of GPS is impossible. To overcome the challenge of localization and positioning without global information, local sensing and low-bandwidth communication is used to create a sensor and communication network. In fact, robots are assumed here to be able to switch from active to passive surveillance and attach to the ceiling. That passive state reduces the energy consumption of the system, allowing a longer mission time. For the exploration, searcher UAVs are guided by beacons attached to the ceiling. In comparison to our approach, the work in this article considers indoor exploration when we assume an outdoor use case. That environment stop them from simply using GPS for the localization. Also, in this article the network formed by the beacons is necessary for the search phase, when our work presents a search pattern accepting sparsely connected swarm and therefore do not need connectivity during the search.

McGuire et al. (2019) presented a minimal navigation solution for unknown environment exploration with tiny flying robots. The proposed approach is called swarm gradient bug algorithm and allows the robots to come back to the starting point after the exploration is completed. To prove the potential of the algorithm, it was used for a proof of concept in a search and rescue application. Unlike our system, this algorithm is designed for indoor exploration and would not be usable outdoor since barriers such as walls used for the wall following would not be available.

The work presented by Rouček et al. (2021) a field report of the system developed by the CTU-CRAS-NORLAB team for the subterranean challenge. The challenge requires a team of mobile robots (including UAVs) to search an unknown environment to locate and report the positions and types of specific artifacts, with limited human interaction. Different exploration strategies were used depending on the robots. The UAVs, for example, aim at the exploration of further areas of the environment instead of a thorough exploration of nearby locations, and use a generated 3D Lidar map. As communication issues arise in such environments, the team had to propose an approach to maintain connectivity. Several redundant communication systems were then implemented to deliver at least some data to the human supervisor of the system. A Wi-Fi module is used for usual usage of the system. The mid-range link (mobicom) is used during the actual mission, with an approach similar to the one presented in our paper: using robots as retransmission nodes to maintain communication. The last system implemented is a long-range

link, called Motes. Robots carry these modules and can drop them to create a network of relays.

De Hoog et al. (2009) proposes a role-based approach to multi-robot exploration, that is robust to communication limitations. In fact, robots in the system can assume one of two roles. Searchers explore the environment and meet periodically with relays to share knowledge. Relays in turn carry information back to a central entity. The exploration approach presented here is similar to ours, as it exploits period meeting for information exchange. Although, the meeting point changes dynamically during the exploration and differs depending on the searcher, when we propose a static, predefined rendezvous point. It is also worth mentioning that the goal in the work is complete exploration, when this metric is not of interest in our work.

Wellman et al. (2011) presented an approach to unknown environment exploration, using sector search. In fact, robots in the system explore independently different sectors of the available area and use scheduled rendezvous to share new information. Studying the performance of their system in comparison to other communication paradigm, they conclude it is comparable to when robots communicate only with other robots in proximity. This approach uses a search approach similar to ours, but they assume robot teams to be small, allowing for a limited number of messages at rendezvous, when our approach is scalable to relatively bigger swarms.

Meghjani and Dudek (2012) proposed a method to have a dynamic rendezvous meeting point selection based on the shared cost of visiting the location. Since for the rescue part we need to have access to most robots, the dynamic rendezvous method is not preferred. Because in terms of robot failure or lost messages, some agents could meet the root after a long time in that case.

Instead of considering the network periodically connected same as Hollinger and Singh (2012), we performed a periodic connection at the rendezvous meeting point. This makes it possible to reconnect an agent that is far from the other robot.

Spirin et al. (2013) propose a scheme in which agents choose their actions based on the time preference of the base station for information, to minimize the rate of information update at a base station.

Pei et al. (2013) proposes another approach to solve the connectivity problem in multi robots exploration. The approach is bandwidth aware and allows the robots to realize a better exploration time with enhanced connectivity when compared to recent works. It aims to keep the connectivity, not provide a system robust to sparsely connected swarms.

The work by Spirin and Cameron (2014) is an extension of the role-based approach. It shows an algorithm that allows mobile robots to plan how to transfer information by setting rendezvous points either side of walls in the unknown environment exploration application. This approach can be

of use when considering obstacles during the relay formation in future works.

In comparison to the method introduced by Cesare et al. (2015), we start to create the relay chain after finding the target, and since the environment of the test could be open water, unlike the indoor environment, it is not possible to land and act as a relay or networker. As a result, in case of having a low battery, the only option is to go back to the station/meeting point to recharge the battery.

Banfi et al. (2018) address the problem of multi robot exploration missions with communication constraints. Considering the situation in which a recurrent connectivity is required in the system, two planning techniques are developed and extensively tested. Results from the carried out experiments show that the approaches are effective, providing good results with a better degree of freedom for the exploration.

The work presented by Shirsat et al. (2020) is a probabilistic consensus-based multi-robot search strategy that is robust to communication link failures.

3 Search and rescue with sparsely connected swarms

In this paper, we consider the scenario in which a swarm of drones needs to be deployed in an unknown environment to search for one or more targets, and track them as rescuers are dispatched to the target locations. The drones explore the area autonomously and in a decentralized manner, searching for targets. Communication links are needed between the swarm members either to inform the others when a target is found and to share the target positions, propagating this information to a base station so that the targets can be rescued.

In realistic scenarios, the search area is likely to be larger than the combined communication coverage of the robots in the swarm. Therefore, the searching robots need to disconnect from their neighbors to explore enough space to find the desired targets, creating a sparse swarm.

Let us consider a swarm \mathcal{S} of n robots $\mathcal{S} = \{1, 2, \dots, n\}$. At a given time step $t_s \geq 0$ during the mission, \mathcal{S} is considered a sparse swarm if a robot $r \in \mathcal{S}$ satisfies:

$$\begin{aligned} \text{cost}_r(\text{"move to nearest neighbor"}, t_s) &\gg \\ \text{cost}_r(\text{"perform typical operation"}, t_s) & \end{aligned} \quad (1)$$

where \gg is defined as "at least one order of magnitude greater than" and cost_r is a function defining the cost for robot r to perform a given task at a given time (Tarapore et al., 2020). In order to ensure coordination and efficient searching in such a swarm, we designed an algorithm inspired from typical search parties in rescue operations, shown in Fig. 2.

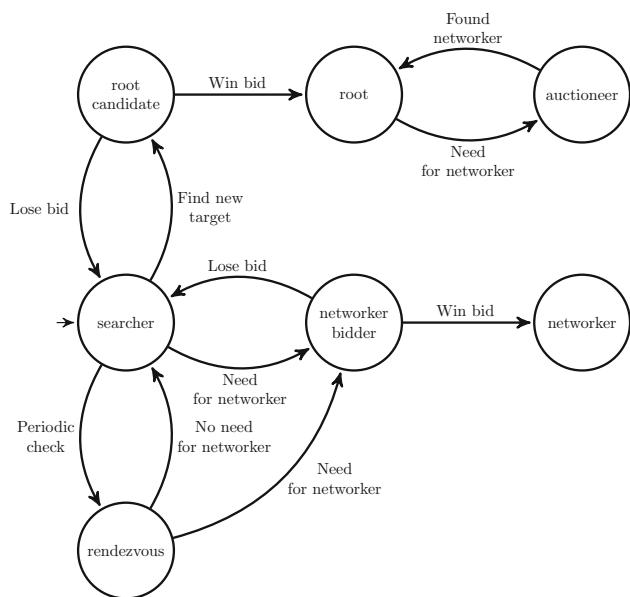


Fig. 2 Coordination algorithm for target searching

The overall idea is that robots perform their search for a target, regularly reporting at a fixed rendezvous location or meeting point. If a robot finds a target, it immediately goes to the meeting point to share the location of the target with the operator and the rest of the swarm. Assuming a robot is the first to go to the meeting point after finding a target, it becomes a *root* robot, and it coordinates the formation of a relay chain towards the target. The chain construction starts as the root robot broadcasts a call for *networkers* (i.e., relay robots) that other members of the swarm respond to with bids based on their distance from the required position of the relay, and the root robot assigns roles based on the received bids (Gerkey and Mataric, 2002). As robots find more targets, the root adds branches to the relay formation, and should robots find targets simultaneously, they elect a root through a basic consensus mechanism.

3.1 Agent roles

The overall strategy is based on a state machine that assigns roles to the robots in the swarm, with each role is associated with a task executed by the robot.

Searcher: When a robot is a searcher, it looks for targets using a predefined search method (a belief-based search in our case). During its operation, a searcher listens and responds to *calls for bids* from other agents. These calls offer networker (i.e., relay) roles, and a searcher bids based on its distance to the requested relay position.

Root candidate: Upon finding a new target, an agent will go to the rendezvous point to become a root robot. In case of a root being already present, the agent shares its target information with the existing root and go back to being a searcher. If multiple robots are heading to the rendezvous point, the first robot to arrive proclaims itself the winner and shares that information with every incoming robots. Should multiple agents arrive to the rendezvous at the same time, we use a conflict management method based on robot ID (Pinciroli et al., 2016).

Root: Winning the bid for the root node, an agent becomes the root: it listens for new target information from the swarm, it computes the number of networkers needed per target and their positions, and calls for robots to fill the networker roles. The root is the only agent that recruits the networkers. It keeps searching for networkers until the target is reached. Note that this strategy does not make the system centralized: the root is easily replaced in case of failure with a new election.

Auctioneer: When the root needs networkers to cover a target, it switches to the auctioneer state, for a typical market-based task allocation strategy (Gerkey and Mataric, 2002). The root/auctioneer broadcasts the relay position, opens the auction, listens for bids, and closes the auction after a predefined period of time, remaining in the same state until a winner is found. The auctioneer then broadcasts the winning bid and goes back to the root state.

Networker bidder: When a searcher receives a call for bids, it stops moving and bids for a networker role. Its bid value is inversely proportional to its distance to the assigned relay position.

Networker: The networker bidder that is the closest to a relay position wins the bid and becomes a networker. A network relays information between a target location and a base station at the meeting point, connecting the target with the operator and providing constant communication coverage in the area of the target.

Rendezvous: Robots regularly switch to the rendezvous state and go back to the meeting point to check for any new information (new root

node, request for a networker, or updated found targets list). Note that robots keep searching for targets on their way back to the meeting point. If there are no networker calls for bids happening, the robots in rendezvous state go directly back to being searchers.

3.2 Algorithm

All the robots execute their search for the targets based on the available belief information. The belief information, represented as a map, is updated and distributed to the neighbors during the search to avoid searching the same area multiple times. This distributed belief map based search is inspired from the work in Vielfaure et al. (2021). To distribute a belief map, we use the virtual stigmergy (VS), a system that allows a swarm of robots to agree on a set of (key,value) pairs through gossip communication (Pinciroli et al., 2016). Thus, at each step, a searcher will decrease the belief value for its current position if no target is detected there. The new value is then put into the VS, sharing the information with all neighbors in communication range. A searcher sampling a new position to navigate to, will then opportunistically get the most recent belief for the position (from the initial map if no updated version is available in the VS). It is worth noting that the propagation of the VS is strictly best-effort, and therefore tolerant to disconnections and communication delays. Based on the value of the belief map, the robots decide to either move to the sampled position or sample a new one (if the belief is below a certain threshold).

The idea of the search pattern is to realize multiple runs of a user-defined duration and come back to a rendezvous point between each of them. During the search, if a robot finds a target, it goes back to the initially fixed rally point and checks for the existence of a root node. The first drone to come back to the rally point after finding a target will be the root. When multiple drones arrive to the rendezvous at the same time, a conflict management routine selects one of the drones as the root. This robot becomes the first link of the communication and tracking relay between the meeting point and the targets. The root stays at the rendezvous point and broadcasts relevant information (updated found targets list, root id, networking positions, etc.) to all the robots in its communication range. As explained previously, this node computes the networkers' positions and manages an auction every time it needs a new networker by switching to the auctioneer state. The networkers' positions depend on the *communication range* of the robots. Let N be the set of networkers positions in the system and r the root's position. To find the networkers positions for a target located at t , we choose a branching node at

a position b such that:

$$\|t - b\| = \min(\|t - n\|) \quad (2)$$

for all $n \in N \cup r$.

The networkers are then placed one after the other on the line connecting b and t . They are spaced at a maximum distance of *communication range* to ensure connectivity in the relay.

Other robots that did not find any target, go back periodically to the meeting point to check if another robot found a target or if a networker is needed. Since the battery life of flying robots is quite limited, this periodic check is an opportunity for recharging or battery swapping. When reaching the meeting point, if the robot receives a message from the root for a networking position, it immediately sends its bid for the auction and waits for the results announcement. When a robot wins the bid it acknowledges the auctioneer that it received the message and goes to its assigned position, becoming part of the communication and tracking relay. The robots in relay positions form a tree from the meeting point, allowing the operator at a base station to constantly and simultaneously to "see" and monitor the state of all the detected targets. The chain starts from the meeting point and proceeds towards the targets. It is worth mentioning that new branches are added to the first chain as new targets are discovered.

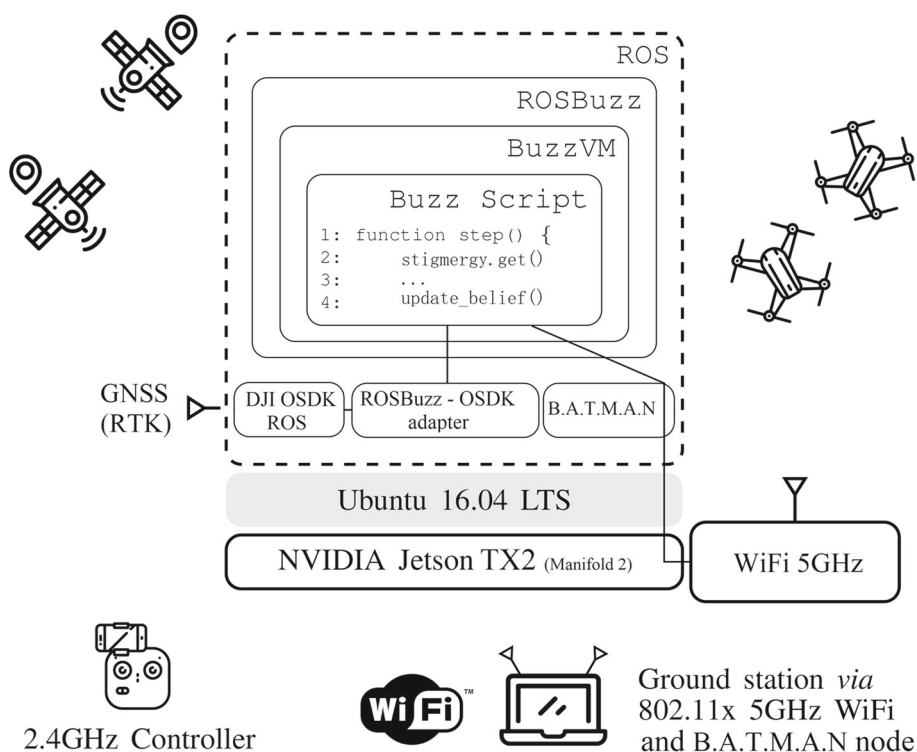
With a sufficient number of searcher robots, the relay should allow connectivity maintenance and a live camera stream of the target to the ground operators. In the case of a moving target, the closest robot to target has the responsibility of tracking its motion and sharing the updated position along the relay. This way, the relay can adapt itself and follow the target up until its rescue.

3.3 Belief map search

The belief map assigns a probability to find a target to each cell in the map. Each drone samples a new exploration target from cells in the map that are above a certain probability threshold, and are located in a square centered on the drone's current location. The size of the square is a parameter.

As they search the area, the UAVs update the belief value at their current position (i.e. reducing the probability if they did not find the target) and share it in the virtual stigmergy. Thanks to this information, UAVs with the updated value will avoid exploring areas that were already explored by other drones. It is worth noting that the virtual stigmergy is updated opportunistically, meaning that updates are propagated as soon as robots enter each other's communication range. This means that, despite the lack of an explicit task allocation algorithm, drones going towards the same or neighboring cells will update each other's stigmergies, ultimately preventing effort duplication.

Fig. 3 The control architecture. The global positioning system and (back-up) remote controllers joint with the DJIs OSDK and flight controller to perform the decentralized behavioural Buzz script. The entire fleet runs the same script, interfacing with the Flight Control Unit (FCU) through DJI OSDK ROS and communication device (WiFi) through the Robot Operating System (ROS). The communication between swarm has been achieved by creating a B.A.T.M.A.N ad-hoc mesh network and using the DJI-Manifold 2 WiFi as the network hardware (St-Onge et al., 2019)



3.4 Real-world deployment

For the real-world deployment of our solution, we used DJI Matrice 300 RTK (M300 RTK) drones equipped with a Manifold 2 onboard computer which was connected to the drone through a serial connection. The M300 RTK is a powerful UAV platform offering an adaptive onboard software development kit (OSDK) for autonomous control of the aircraft. It uses an advanced flight controller system, a 6 directional sensing and positioning system and FPV camera. Thanks to these features, the drones were able to perform a basic collision avoidance routine during the flights.

The decentralized control of the drones is achieved using Buzz, a domain-specific language designed for programming multi-robot teams and swarms behaviors (Pincioli and Beltrame, 2016). The software consists of three main layers: the Buzz control layer taking care of the algorithm logic, the ROSBuzz (St-Onge et al., 2017) layer responsible for the integration of the swarm-oriented programming language and its virtual machine (BVM) into the ROS environment, and the DJI OSDK layer that manages the flight controller and other UAV related features. The Buzz control layer is responsible for the system’s behavior, using a Buzz script to implement the proposed algorithm, while sending hardware specific commands to the lower layers.

Our whole experimental system is based on ROS (Robot Operating System), an open-source and now standard software system for robotic development (Quigley et al., 2009). To link the Buzz control layer to ROS, we use ROSBuzz,

an existing implementation of the BVM as a ROS node. ROSBuzz encapsulates all the BVM logic, publishes the Buzz script commands, and subscribes to external data such as sensor readings. A main feature of Buzz is the implementation of gossip-based situated communication among neighbors in a swarm. This feature is implemented in our system with batman-adv (Better Approach to Mobile Ad-hoc Networking), using a ROS node that manages the neighbors of each robots and broadcasts messages as needed by the Buzz script. Batman-adv is a layer 2-based protocol leveraging adaptive topology and routing to offer a robust ad-hoc networking solution (Kiran et al., 2018). As our algorithm assumes sporadic connectivity among the robots in the swarm, batman-adv can easily handle such a pattern, ensuring a reliable link between the robots when in communication range. The physical device supporting batman-adv is the 5 GHz WiFi antenna on the Manifold 2, set to communicate on an common IBSS ad-hoc wireless network.

The DJI OSDK layer is a DJI proprietary API that allows a developer to control the aircraft with a program. DJI proposes a version of the OSDK integrate with ROS that we used in our setup, and the UAVs are actuated via service calls. The interaction between the ROSBuzz ecosystem and the OSDK layer is taken care of by a custom adapter node. This node receives actuation commands from the Buzz script (in the form of topics) and sends flight controller-specific commands to the OSDK, as well as publishing the required data for ROSBuzz’s operation (e.g., GNSS readings). Such an architecture, allows for a easily portable code base. In fact,

since the adapter is the only M300 dependent node, using the same algorithm with a different robotic platform only requires to write a similar adapter. Figure 3 synthesizes the described software architecture, which is completely open source (see github.com/mistlab).

4 Experimental results: simulations

To validate the presented algorithm and evaluate the necessary time required to find targets and obtain the final relay chain, we performed a series of tests in a simulated environment. We used 15, 20, and 25 robots while randomly varying their starting positions and the position of the targets. We also compare the performance of a random walk search to our dynamic belief map search, while using in both cases the presented search pattern.

4.1 Simulation setup

We use ARGoS, a multi-physics robot simulator (Pincioli et al., 2012). The experiments were performed on an ARGoS model of the Spiri Mu quadrotor from Spiri Robotics. Three foot-bot mobile robots (Dorigo et al., 2013) were randomly spawned and used as targets to be detected by the Spiris. To perform the detection, we simulate a basic sensing mechanism on the drones with a downward facing camera using blob detection. Figure 4 presents an example of a starting state of the simulation.

We first used a random walk exploration pattern (Dimidov et al., 2016) where each drone sequentially samples a 2D position and autonomously navigates to it (keeping the height constant). The second search method is the proposed dynamic belief space exploration. In both cases, we simulate the system until all the targets are detected and the relay network is fully formed. The total time needed to find the targets is the metric used for our evaluation.

4.2 Results

We perform 6 test cases with three arena sizes. For these tests, 2 parameters were considered: the number of drones (15, 20 or 25 UAVs) and the search method (random walk or belief map search). We used three arena sizes for each test: (a) 20m × 20m, (b) 30m × 30m, and (c) 40m × 40m. To obtain statistically relevant results, each test case was executed 30 times, randomly assigning the initial positions for the drones, the targets' positions and the belief map if applicable. The other parameters, including the meeting point, were maintained constant through the experiments: 20 search steps, three targets and a communication range of 10 meters (not realistic, but distances can be simply scaled to realistic values). An output sample for a test with with 15 drones and

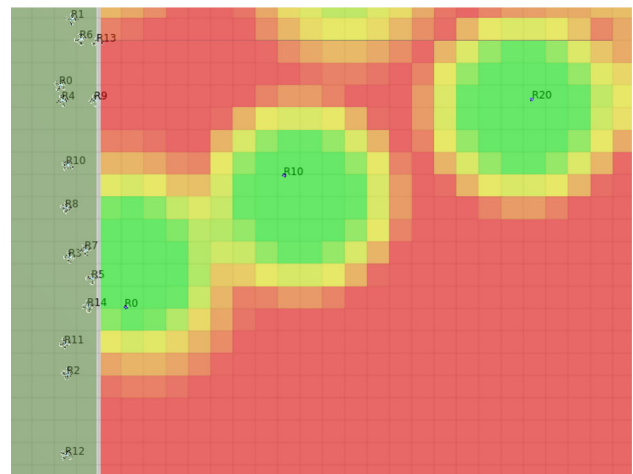


Fig. 4 Initial simulation setup sample using 15 drones (R0 to R19 on the left side) with 3 targets (R0, R10, and R20 in the searching area). The searching area is represented as a belief 2D map where red represent a probability close to 0 to find the target and green represent a probability close to 1




Fig. 5 Possible final state for a test with 15 drones and 3 targets and a communication range of 4 meters. Each drone logging its state: root, networker or searching. The searching area is represented as a belief 2D map where red represent a probability close to 0 to find the target and green represent a probability close to 1

3 targets is presented in Fig. 5 where we can see the relay tree connecting the meeting point to the three targets. The figure also show the remaining drones (not used in the relay) searching for any additional target.

For every experiment, we report the number of Buzz timesteps necessary to find the three targets. For reference, a Buzz timestep is 0.1s by default but it is fully configurable depending on the capabilities of the robots and communication system. Figures 6, 7 and 8 present the number of timesteps necessary to find the targets for different arena sizes and different number of drones. They show the results obtained with both the random search and the belief space

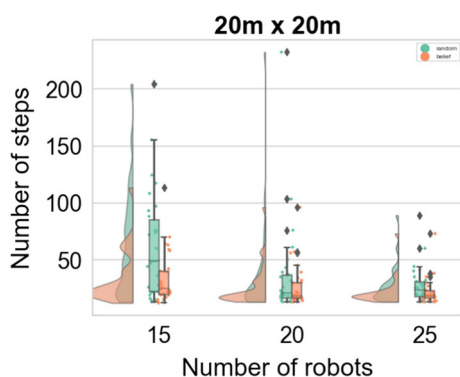


Fig. 6 Number of timesteps to find 3 targets in a 20 m × 20 m arena. The results for the random search are in green and the results for the belief space search are in orange. The BF and the error of the Bayesian paired samples t-test for each setup are respectively 92.404, and 8.273e-5 for 15 robots, 2.886 and 0.001 for 20 robots, and 70.606, and 1.043e-4 for 25 robots

search while increasing the number of drones and the arena size. From those results, we can see as expected for a swarm based SAR algorithm that the time to find the target decreases when the number of drones increases, for all the arena sizes (Figs. 6, 7 and 8). This observation is due to the fact that a bigger search area can be covered with more drones, allowing them to find the targets faster.

Also, the time decreases in all cases when we use a belief space search in comparison to a random walk search. In fact, by reducing step after step the searching area (thanks to the distributed belief information), our search method allows drones to converge faster towards the targets.

Finally, when the arena size increases (from Figs. 6, 7 and 8), we can see that the time needed to find the targets increases as well (the sub graphs have different vertical scales). This can be explained by the sampling space which gets bigger with the arena. Thus, there are more options to explore before we can find the target wherever it might be.

To confirm the hypothesis that the belief space search would be faster than the random walk, in any case, we did a Bayesian paired samples t-test on the results by using JASP Team (2021) software. The results show that all data accept the hypothesis with the Bayesian Factor (BF) greater than 1 (BF min = 2.886, BF max = 66550.284).

Another hypothesis we postulated was that the search pattern that was designed should create the relay formation in nearly constant time. To verify that assertion, we considered for each experiment the total time spent by the root node in the auctioneer state. Figures 9, 10, and 11 summarize the results obtained from that metric. It shows in most cases a slightly higher formation time for the belief based search, with a higher standard deviation. Those results hints that the relay formation might take more time for the belief based search, refuting our hypothesis. To get a clearer answer, we once again did a Bayesian paired samples t-test on the statistical




Fig. 7 Number of timesteps to find 3 targets in a 30 m × 30 m arena. The results for the random search are in green and the results for the belief space search are in orange. The BF of the Bayesian paired samples t-test for each setup are respectively 66550.284, for 15 robots, 914.282 for 20 robots, and 25.268 for 25 robots. No error estimate can be given for 15 robots, however the error for 20 and 25 are respectively 4.609e-8 and 2.954e-5




Fig. 8 Number of timesteps to find 3 targets in a 40 m × 40 m arena. The results for the random search are in green and the results for the belief space search are in orange. The BF and the error of the Bayesian paired samples t-test for each setup are respectively 210.765 and 1.358e-6, for 15 robots, 203.165 and 3.369e-7 for 20 robots, 3577.307 for 25 robots. No error estimate can be given for 25 robots

results considering the null hypothesis (i.e. the relay formation time is the same for both random, and belief search). The results show that 89% of the data reject the null hypothesis with the Bayesian Factor lower than 1 (BF min = 2.539e-6, BF max = 3.862).

We can also notice a constant, but slight increase of the formation time when the arena gets bigger. That fact is expected as the searching area increases. In fact, we have less chances of finding a drone in communication range of the root since they will now search further from the root. Another interesting result is that the formation time do not change much when we compare for the same arena size, different number of drones. Here, the number of agents do not impact the results as the auction time is based on whether a winner is found or not. Thus, more agents do not change the ability for the auctioneer to find a winner, assuming they are in range.




Fig. 9 Number of timesteps needed per experiment to obtain the whole relay structure in a $20\text{ m} \times 20\text{ m}$ arena. The results for the random search are in red and the results for the belief space search are in blue. The BF and the error of the Bayesian paired samples t-test for each setup are respectively 0.115 and $2.281\text{e-}6$, for 15 robots, $8.070\text{e-}4$ and $1.277\text{e-}6$ for 20 robots, $2.539\text{e-}6$ and $2.511\text{e-}8$ for 25 robots




Fig. 10 Number of timesteps needed per experiment to obtain the whole relay structure in a $30\text{ m} \times 30\text{ m}$ arena. The results for the random search are in red and the results for the belief space search are in blue. The BF and the error of the Bayesian paired samples t-test for each setup are respectively 0.264 and $3.022\text{e-}6$, for 15 robots, $3.327\text{e-}6$ and $2.817\text{e-}8$ for 20 robots, 0.022 and $1.896\text{e-}4$ for 25 robots




Fig. 11 Number of timesteps needed per experiment to obtain the whole relay structure in a $40\text{ m} \times 40\text{ m}$ arena. The results for the random search are in red and the results for the belief space search are in blue. The BF and the error of the Bayesian paired samples t-test for each setup are respectively 3.862 and 0.001, for 15 robots, 0.017 and $1.133\text{e-}4$ for 20 robots, 0.469 and $3.806\text{e-}6$ for 25 robots

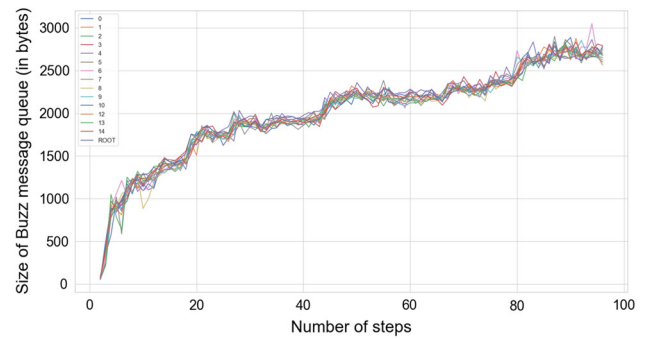


Fig. 12 Bandwidth usage for 15 drones using a belief search in a 20×20 arena (single experiment)




Fig. 13 Bandwidth usage for 15 drones using a random search in a 20×20 arena (single experiment)

To explain the relay formation time difference between the two search methods, we made the hypothesis that the network usage is at the root of this observation. In fact, because of the dynamic update of the belief map, the network is heavily used in the belief based search. The update messages could delay the handling of the auction-related ones and therefore extend the formation time. To confirm this hypothesis, we compiled the bandwidth usage for the belief based search and the random walk search. For the sake of brevity, we present in Figs. 12 and 13 the results of a randomly selected experiment's configuration. We can see on the graph that for the belief-based search, the bandwidth is constantly increasing (reaching a maximum of 3000 bytes in the message queue) and we observe an homogeneous network usage by all the agents. In some rare occasions, when a drone would be out of range, we can observe a decrease of its network usage which increases drastically when it is back in range. In contrast, Fig. 13 show a maximum of 380 bytes for the message queue size, with a relatively constant bandwidth usage after the election of the root node (around step 15). We can observe some high usage periods that correspond to the auction time in the system, with the root using more bandwidth in comparison to the other drones. Those observations confirmed our hypothesis by showing a fairly high network usage in the belief search.

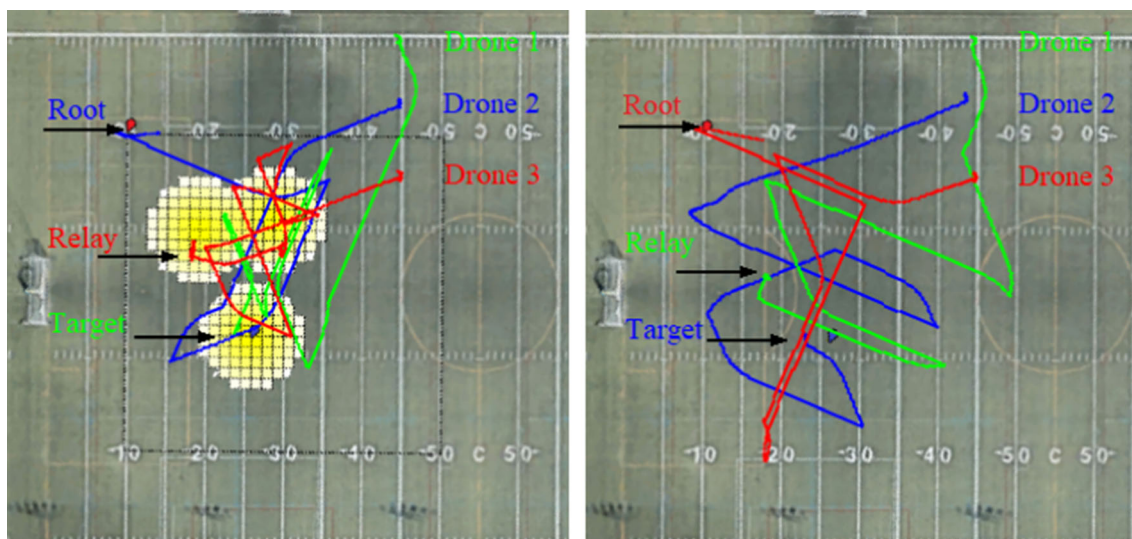


Fig. 14 GPS trajectory during tests on a football field. Right: random, Left: belief map, the yellow cells indicate the probability of having a target

It should also be noted that the number of search steps (set to 20 for our tests) could impact the time needed to have the relay chain, showing the importance for parameters tuning. For example, if the target is found by a drone at step 1, the first relays will not come to the meeting point before step 20 (unless they also find the target or are in *communication range* of the root). That fact can make the system stabilization unnecessarily long.

5 Experiments

The real-world experiments were performed in an outdoor football field with three DJI M300 quadcopters. The search area considered for the tests is a 36×36 meters field. The experimental platform was mainly presented in Sect. 3.4. The decentralized control of the drones is achieved using the same Buzz scripts used in the simulation, with some minor changes related to the auction duration, and the M300 RTK's flight controller.

The tests were performed with both the random walk search and the distributed belief map search. Each of the experiments, were performed three times to confirm the proper functioning of the algorithm. Figure 15 confirms the feasibility of our method. After performing the previously explained search pattern, the drones adopt, as seen in the picture, a relay formation. Given that we only have one target and three drones the relay was a line from the predefined rendezvous point to the target position. Sending a message from the agent at the target position, the information can now be relayed back to the root for analysis.

To give a better idea of the searching behavior during the experiments, we present in Fig. 14 the path taken by the

drones during two randomly selected runs (Right: random, Left: belief). That figure shows for the random search sparse lines, scattered randomly over the field, indicating a lack of pattern during the search. That lack of pattern can give fast discovery of the target under certain circumstances and a very long time in other cases. On the other hand, for a belief search, the lines are concentrated in areas with high belief.

Table 1 presents, for information, the recorded metrics for the belief map search. It shows the number of timesteps needed to find the target for each of the experiments. We obtained an average 1105 timesteps with the belief search, a value that is bigger than the ones obtained in simulation. In fact, the difference with field experiments is that the number of steps required to get from one position to another is higher because of the velocity of the drones set to a relatively low value. The table also contains the necessary time to form a relay chain (as the one shown in Fig. 15). The mean value here is 295 timesteps.

Unlike the simulations, these experiments were only performed 3 times with a fixed target and a fixed belief map. That configuration can easily produce some biased results. Therefore, more real-world tests would be necessary to have statistically relevant data.

6 Discussion of real-world use cases

Since the presented system was designed for a real-world deployment of UAVs swarm, a discussion in terms of practical use cases is essential to better understand when this system could be used. This section will present that discussion with some ideas to optimize the current solution.

Table 1 Real-world tests results

Experiment #	Target discovery (timesteps)	Relay formation (timesteps)
1	1159	301
2	863	324
3	1295	249

**Fig. 15** Relay chain formation during real-world experiment

Considering our approach during the rescuing phase, with the relay structure formation, some questions about the benefits of keeping the targets in sight instead of simply reporting the position might come in mind. For a search and rescue application, establishing a real-time video connection between the base station and the found targets provides multiple advantages in comparison to simply flying back to report the location. Among those we have:

- Allow human rescuers to monitor the state of the target at the detected position and assess their priority during rescue in case of multiple targets;
- Track the target in the likely case of movement (e.g. at sea);
- Communicate with the target till rescuer can reach it (either via radio or speakers).

It is also worth mentioning that in the proposed algorithm, the first UAV that finds a target, flies back to report the location, that information is then used to create the relay structure. Also, if we consider applications other than SAR, this approach can be used to stream live video of a monitored region to the ground station (e.g. fire monitoring).

If we consider the use of fixed wing aircraft (to which our method can easily apply using a circling, loitering motion), the flight range for the UAVs can be in the hundreds of kilometers, and even a recent DJI M300 could fly 20–25 km if we do not consider the range of the remote controller. It is

also worth noting that the horizon is at about 25 km distance when flying at 50 m altitude, which also limits communication range, and 4G or 5G communication is generally not available farther than a few km off the coast. Overall, we believe that relay chains have practical use with the current available communication systems, but in any case, our systems adapts and scales based on the communication range of the UAVs, using fewer or no drones for relay as the range expands.

Communication in the system occurs during bidding or the belief map update (via virtual stigmergy). If a message is lost during the bidding it could take more time to find a winner, creating some delays in the execution or even put the system in an undefined state. Also, if messages are lost during the belief map update, the search time will be affected because the agents will search at positions previously visited by their neighbors. To prevent this from happening, acknowledgement messages are used for the communication during the bidding to ensure that we don't get stuck in that state if some messages get lost. As indicated previously, the belief map update is best effort. Therefore, no specific measures were considered to mitigate message loss impact there. The experiments results show that the method is robust to message loss.

The current algorithm does not take into account obstacles to form the relay structure. Our previous work (Varadharajan et al., 2020) does consider relays in a cluttered environment and could be integrated. In general, the presence of obstacles

between the member of the relay might decrease the communication range and require the UAVs to get closer to one another. In our specific case, we assume an environment with obstacles that lie lower than the altitude of the UAVs. That assumption removes the need for obstacle management during the relay formation phase while staying relevant for most SAR operations.

7 Conclusion and future works

In this paper, we presented a novel swarm robotic system for search and rescue operations in realistic scenarios. The proposed system is fully decentralized and robust to sporadic disconnections. It was also coupled with an implementation of a distributed belief map search algorithm leveraging prior knowledge on the area to realize a faster search. Based on the results obtained in simulation, we were able to confirm that our search method performs better than a random walk. The results obtained during our simulations and the real-world experiments, confirmed the feasibility of our approach. The deployed architecture also provide a modular, easily portable and scalable system, that could be used in other swarm deployments. In future work, we will leverage the distributed belief map in the search algorithm to perform dynamic updates on the target location belief. Indeed, we could model the target motion (e.g. due to the flow of a river) to update the prior over time.

Funding This work was supported by the Fonds de recherche du Québec—Nature et technologies (FRQNT) under Grant No. 296737 and by the National Research Council Canada (NRC).

Declarations

Conflict of interest The authors declare no potential conflict of interest with respect to the research, authorship, and/or publication of this article.

References

- Alotaibi, E. T., Alqefari, S. S., & Koubaa, A. (2019). LSAR: Multi-UAV collaboration for search and rescue missions. *IEEE Access*, 7, 55817–55832.
- Andries, M., & Charpillet, F. (2013). Multi-robot exploration of unknown environments with identification of exploration completion and post-exploration rendezvous using ant algorithms. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 5571–5578). IEEE.
- Andries, M., & Charpillet, F. (2015). Multi-robot taboo-list exploration of unknown structured environments. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 5195–5201). IEEE.
- Aprville, L., Tanzi, T., & Dugelay, J. L. (2014). Autonomous drones for assisting rescue services within the context of natural disasters. In *2014 XXXIth URSI General Assembly and Scientific Symposium (URSI GASS)* (pp. 1–4). IEEE.
- Banfi, J., Quattrini Li, A., Rekleitis, I., et al. (2018). Strategies for coordinated multirobot exploration with recurrent connectivity constraints. *Autonomous Robots*, 42(4), 875–894.
- Belkadi, A., Ciarletta, L., & Theilliol, D. (2016). UAVS fleet control design using distributed particle swarm optimization: A leaderless approach. In *2016 International Conference on Unmanned Aircraft Systems (ICUAS)* (pp. 364–371). IEEE.
- Brambilla, M., Ferrante, E., Birtattari, M., et al. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(1), 1–41.
- Çeltek, S. A., Durdu, A., & Kurnaz, E. (2018). Design and simulation of the hierarchical tree topology based wireless drone networks. In *2018 International Conference on Artificial Intelligence and Data Processing (IDAP)* (pp. 1–5). IEEE, <https://doi.org/10.1109/IDAP.2018.8620755>
- Cesare, K., Skeeel, R., Yoo, S. H., et al. (2015). Multi-uav exploration with limited communication and battery. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2230–2235). IEEE.
- De Hoog, J., Cameron, S., & Visser, A. (2009). Role-based autonomous multi-robot exploration. In *2009 Computation World: Future Computing, Service Computation, Cognitive, Adaptive, Content, Patterns* (pp. 482–487). IEEE.
- Dimidov, C., Oriolo, G., & Trianni, V. (2016). Random walks in swarm robotics: an experiment with Kilobots. In *International Conference on Swarm Intelligence* (pp. 185–196). Springer.
- Dorigo, M., Floreano, D., Gambardella, L. M., et al. (2013). Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics & Automation Magazine*, 20(4), 60–71.
- Gerkey, B. P., & Mataric, M. J. (2002). Sold!: Auction methods for multirobot coordination. *IEEE Transactions on Robotics and Automation*, 18(5), 758–768.
- Hentati, A. I., & Fourati, L. C. (2020). Comprehensive survey of UAVS communication networks. *Computer Standards & Interfaces*, 72, 103,451.
- Hollinger, G. A., & Singh, S. (2012). Multirobot coordination with periodic connectivity: Theory and experiments. *IEEE Transactions on Robotics*, 28(4), 967–973.
- Hourani, H., Hauck, E., & Jeschke, S. (2013). Serendipity rendezvous as a mitigation of exploration's interruptibility for a team of robots. In *2013 IEEE International Conference on Robotics and Automation* (pp. 2984–2991). IEEE.
- JASP Team (2021). JASP (Version) [Computer software]. URL <https://jasp-stats.org/>.
- Khan, A., Yanmaz, E., & Rinner, B. (2014). Information merging in multi-uav cooperative search. In *014 IEEE international conference on robotics and automation (ICRA)* (pp. 3122–3129). IEEE.
- Kim, J., Ladosz, P., & Oh, H. (2020). Optimal communication relay positioning in mobile multi-node networks. *Robotics and Autonomous Systems*, 129, 103517. <https://doi.org/10.1016/j.robot.2020.103517>, URL www.sciencedirect.com/science/article/pii/S0921889019309145.
- Kiran, K., Kaushik, N., Sharath, S., et al. (2018). Experimental evaluation of batman and batman-adv routing protocols in a mobile testbed. In *TENCON 2018-2018 IEEE Region 10 Conference* (pp. 1538–1543). IEEE.
- Kobayashi, F., Sakai, S., & Kojima, F. (2002). Sharing of exploring information using belief measure for multi robot exploration. In *2002 IEEE World Congress on Computational Intelligence. 2002 IEEE International Conference on Fuzzy Systems. FUZZ-IEEE'02. Proceedings (Cat. No. 02CH37291)* (pp. 1544–1549). IEEE.
- Kobayashi, F., Sakai, S., & Kojima, F. (2003). Determination of exploration target based on belief measure in multi-robot exploration.

- In *Proceedings 2003 IEEE International Symposium on Computational Intelligence in Robotics and Automation. Computational Intelligence in Robotics and Automation for the New Millennium (Cat. No. 03EX694)* (pp. 1545–1550). IEEE.
- Li, J. (2019). Throughput-aware flying communication relay network for disaster area search and rescue. In *Proceedings of the 2019 8th International Conference on Networks, Communication and Computing* (pp. 138–141). <https://doi.org/10.1145/3375998.3376038>.
- Majcherczyk, N., Jayabalan, A., Beltrame, G., et al. (2018). Decentralized connectivity-preserving deployment of large-scale robot swarms. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 4295–4302). IEEE.
- McGuire, K., De Wagter, C., Tuyls, K., et al. (2019). Minimal navigation solution for a swarm of tiny flying robots to explore an unknown environment. *Science Robotics*, 4(35), eaaw9710.
- Meghjani, M., & Dudek, G. (2012). Multi-robot exploration and rendezvous on graphs. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 5270–5276). IEEE.
- Nickerson, J. V. (2004). Robots and humans reconvening. In *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No. 04CH37583)* (pp. 2803–2808). IEEE.
- Nouyan, S., & Dorigo, M. (2006). Chain based path formation in swarms of robots. In *International Workshop on Ant Colony Optimization and Swarm Intelligence* (pp. 120–131). Springer.
- Pei, Y., Mutka, M. W., & Xi, N. (2013). Connectivity and bandwidth-aware real-time exploration in mobile robot networks. *Wireless Communications and Mobile Computing*, 13(9), 847–863.
- Pincirol, C., & Beltrame, G. (2016). Buzz: An extensible programming language for heterogeneous swarm robotics. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 3794–3800). IEEE.
- Pincirol, C., Trianni, V., O'Grady, R., et al. (2012). Argos: A modular, parallel, multi-engine simulator for multi-robot systems. *Swarm intelligence*, 6(4), 271–295.
- Pincirol, C., Lee-Brown, A., & Beltrame, G. (2016). A tuple space for data sharing in robot swarms. In *Proceedings of the 9th EAI International Conference on Bio-inspired Information and Communications Technologies (formerly BIONETICS)* (pp. 287–294).
- Quigley, M., Conley, K., Gerkey, B., et al. (2009). Ros: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, (pp. 5). Kobe, Japan.
- Rouček, T., Pecka, M., Čížek, P., et al. (2021). System for multi-robotic exploration of underground environments ctu-cras-norlab in the darpa subterranean challenge. arXiv preprint [arXiv:2110.05911](https://arxiv.org/abs/2110.05911).
- Ruetten, L., Regis, P. A., Feil-Seifer, D., et al. (2020). Area-optimized uav swarm network for search and rescue operations. In *2020 10th Annual Computing and Communication Workshop and Conference (CCWC)* (pp. 0613–0618). IEEE. <https://doi.org/10.1109/CCWC47524.2020.9031197>.
- Shirsat, A., Elamvazhuthi, K., & Berman, S. (2020). Multi-robot target search using probabilistic consensus on discrete markov chains. In *2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)* (pp. 108–115). IEEE.
- Sperati, V., Trianni, V., & Nolfi, S. (2011). Self-organised path formation in a swarm of robots. *Swarm Intelligence*, 5(2), 97–119.
- Spirin, V., & Cameron, S. (2014). Rendezvous through obstacles in multi-agent exploration. In *2014 IEEE International Symposium on Safety, Security, and Rescue Robotics (2014)* (pp. 1–6). IEEE.
- Spirin, V., Cameron, S., & Hoog, J. d. (2013). Time preference for information in multi-agent exploration with limited communication. In *Conference Towards Autonomous Robotic Systems* (pp. 34–45). Springer.
- St-Onge, D., Varadharajan, V. S., Li, G., et al. (2017). Ros and buzz: consensus-based behaviors for heterogeneous teams. arXiv preprint [arXiv:1710.08843](https://arxiv.org/abs/1710.08843).
- St-Onge, D., Kaufmann, M., Panerati, J., et al. (2019). Planetary exploration with robot teams: Implementing higher autonomy with swarm intelligence. *IEEE Robotics & Automation Magazine*, 27(2), 159–168. <https://doi.org/10.1109/MRA.2019.2940413>
- Stirling, T., Wischmann, S., & Floreano, D. (2010). Energy-efficient indoor search by swarms of simulated flying robots without global information. *Swarm Intelligence*, 4(2), 117–143.
- Tarapore, D., Groß, R., & Zauner, K. P. (2020). Sparse robot swarms: moving swarms to real-world applications. *Frontiers in Robotics and AI*, 7, 83.
- Varadharajan, V. S., St-Onge, D., Adams, B., et al. (2020). Swarm relays: Distributed self-healing ground-and-air connectivity chains. *IEEE Robotics and Automation Letters*, 5(4), 5347–5354.
- Vielfaure, D., Arseneault, S., Lajoie, P. Y., et al. (2021). Dora: Distributed online risk-aware explorer. arXiv preprint [arXiv:2109.14551](https://arxiv.org/abs/2109.14551).
- Wellman, B. L., Dawson, S., de Hoog, J., et al. (2011). Using rendezvous to overcome communication limitations in multirobot exploration. In *2011 IEEE International Conference on Systems, Man, and Cybernetics* (pp. 2401–2406). IEEE.
- Wubben, J., Aznar, P., Fabra, F., et al. (2020). Toward secure, efficient, and seamless reconfiguration of uav swarm formations. In *2020 IEEE/ACM 24th International Symposium on Distributed Simulation and Real Time Applications (DS-RT)* (pp. 1–7). IEEE.
- Yamaguchi, S. P., Karolonek, F., Emaru, T., et al. (2017). Autonomous position control of multi-unmanned aerial vehicle network designed for long range wireless data transmission. In *2017 IEEE/SICE International Symposium on System Integration (SII)* (pp. 127–132). IEEE. <https://doi.org/10.1109/SII.2017.8279200>
- Zhang, H. G., Jin, G. Y., Qu, Y. X., et al. (2021). Servo relays as distributed controllable-mobility network to maintain long-term stable links for mobile robot swarms. *Ad Hoc Networks*, 117, 102497. <https://doi.org/10.1016/j.adhoc.2021.102497>, URL www.sciencedirect.com/science/article/pii/S1570870521000597.
- Zhou, B., Zhang, Y., Chen, X., et al. (2021). Fuel: Fast UAV exploration using incremental frontier structure and hierarchical planning. *IEEE Robotics and Automation Letters*, 6(2), 779–786.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.



Ulrich Dah-Achinanon is a master's student at Polytechnique Montreal. He received a bachelor's degree in computer engineering from Polytechnique Montréal, QC, CA. His research interests include swarm robotics systems design more specifically UAV control.



Seyed Ehsan Marjani Bajestani is currently a Ph.D. Candidate in Computer Engineering with the MIST Lab at Polytechnique Montreal. His research is focusing on event-based multi-robot perception. He is an Executive Committee member of the International RoboCup Federation in the soccer robots MSL since 2016.



Pierre-Yves Lajoie is a PhD student at Polytechnique Montreal. He is particularly interested in the challenges in artificial intelligence and/or computer vision including photogrammetry, SLAM, robust state estimation, distributed (multi-robot) perception, semantic segmentation and other applications of deep learning. He has extensive experience in computer vision, GPU programming and robotics and 2 journal publications as first author on robust Simultaneous Localization and Mapping systems.



Giovanni Beltrame obtained his Ph. D. in Computer Engineering from Politecnico di Milano in 2006. He worked as microelectronics engineer at the European Space Agency on a number of projects, spanning from radiationtolerant systems to computer-aided design. Since 2010 he is Professor at Polytechnique Montreal with the Computer and Software Engineering Department, where he directs the MIST Lab. His research interests include modelling and design of embedded systems, artificial intelligence, and robotics. He was awarded more than 25 Grants by government agencies and industry, and has published more than 100 papers in international journals and conferences.