

Multirobot Exploration of Communication-Restricted Environments: A Survey

Francesco Amigoni¹, Jacopo Banfi², and Nicola Basilico³

¹ Politecnico di Milano, Italy, francesco.amigoni@polimi.it

² Politecnico di Milano, Italy, jacopo.banfi@polimi.it

³ Università degli Studi di Milano, Italy, nicola.basilico@unimi.it

Abstract. Exploration of initially unknown environments is an online task in which autonomous mobile robots coordinate themselves in order to efficiently discover free spaces and obstacles. Several efforts have been devoted to study coordinated multirobot exploration assuming that communication is possible between any two locations. The problem of developing multirobot systems for effective exploration in presence of communication constraints, despite its remarkable practical relevance, is comparably much less studied. We provide a taxonomy of the field of communication-restricted multirobot exploration, we survey recent work in this field, and we outline some promising research directions.

1 Introduction

Adopting teams of autonomous mobile robots can provide significant advantages, like improved efficiency, reliability, and robustness, in accomplishing information-gathering tasks, like exploration, surveillance, and inspection [1]. Such advantages are obtained by employing some form of *coordination* between the members of the team, which is often developed assuming the possibility to communicate without limitations. However, real-world missions often require to deal with communication-restricted environments. In these settings, robots can share information only with teammates in a (local) communication range depending both on their transmission capabilities and on the environment itself (e.g., presence of obstacles or disturbances). As a result, achieving a satisfactory level of coordination may become problematic.

In this article, we survey the most significant methods that have been proposed to address the task of exploring initially unknown environments with a team of autonomous mobile robots when unlimited communication between all the robots is not available. To keep our illustration on a concrete ground, we focus on exploration in the context of map building [2]. Indeed, the construction of maps is one of the most basic and, at the same time, challenging tasks for a team of autonomous mobile robots. Its inherently *online* nature and the relevant impact that communication might have on some of its critical subtasks (like robot coordination or map merging) make it a highly representative domain for our

literature analysis. Applications in the fields of search and rescue, surveillance, and the like often include map building as part of their problem definitions and resolution techniques. As a consequence, many of the considerations we draw in what follows can be easily generalized to other problems.

From an abstract stance, the problem of exploration can be represented as follows. Assume to have an environment E and assume that the time is discretized in steps. There is a set of m robots, $R = \{r_1, r_2, \dots, r_m\}$, each one able to move in E . Call $p_i^t \subseteq E$ the portion of E perceived by robot i at time t . For a moment, assume also that all the robots can always communicate with each other. At a given time \bar{t} , the *map* representing the portion of the environment perceived by robots is given by:

$$M^{\bar{t}} = \bigcup_{i=1, \dots, m} \bigcup_{t=0, \dots, \bar{t}} p_i^t$$

The exploration is *complete* when there is a time \bar{t} such that $M^{\bar{t}} = E$ (or such that $M^{\bar{t}}$ is equal to the free space of E). The problem of multirobot exploration, introduced by [3], is that of choosing at each time step t towards which *frontiers* (namely, borders of the known portion of the environment M^t) the robots should move next, while optimizing a performance measure. Typical performance measures are the time \bar{t} to complete exploration (to be minimized), the amount of area mapped within a given time (to be maximized), the distance travelled by robots to complete exploration (to be minimized), or combinations of the above. The decisions about which frontier each robot should visit in order to optimize the performance measure are the product of the *exploration strategy*, which embeds the “intelligence” of the system.

While the unlimited communication setting has been widely studied (see, e.g., [4]), the case of multirobot exploration with restricted communication is much less investigated. Limiting communication has some consequences, though. A first critical issue is related to the joint knowledge of the environment during the exploration mission. With unlimited communication, such knowledge can be assumed to be accessible by each robot at any time. Indeed, map merging algorithms [5] (whether centralized or distributed) could rely on data sharing protocols to share the updated map. Clearly, if communication is restricted such a joint knowledge cannot be assumed to be always accessible, and the impact on exploration could become remarkable. In this case, the robots must move and coordinate to reconnect, exchange data, and, if needed, act as communication relays, considering that different robots can have different amount of knowledge. Obviously, this introduces extra costs. If the coordination mechanism is not carefully devised, significant degradations of the exploration performance could be experienced. An example scenario is depicted in Figure 1a and Figure 1b.

As second issue, consider that robots can predict, using a communication model, whether they will be in communication with a teammate at a particular location. Predictions, obviously, could be wrong and any coordination method relying on them can be deeply affected. For this reason, communication models

used in practice are typically very conservative (see Figure 1c for an illustrative example).

Third, additional connection requirements could be posed throughout the exploration mission. For instance, a human operator residing at a base station (BS) could be interested in receiving a video feed of the robots' cameras or a centralized map merging algorithm should receive the perceived data. Such cases can entail the need of constructing a multi-hop chain of locally-connected robots to route the video stream or the collected data towards the BS (see Figure 1d).

In the remainder of this article, we provide a critical analysis on how the problem of communication-constrained multirobot exploration has been tackled by different authors in different flavors. This is a recent research topic which, due to its practical relevance, is likely to become central for applications where intelligent exploration strategies (also including coordination of robots) must be developed under intermittent and unpredictable communication. To the best of our knowledge, we provide the first survey on it, trying to identify the main dimensions along which past contributions can be classified and discussing how relevant future research might stem from them.

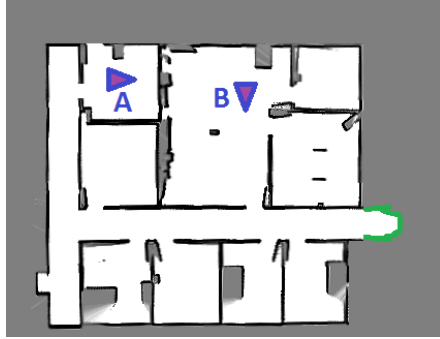
2 Classification dimensions

We start with a description of what emerged as the most significant classification dimensions from our literature review. Clearly, there are many alternative options for structuring a classification of the works. The dimensions that we present here are those that we found more suitable in terms of generality, recall, and adherence to the topic of communication-restricted exploration. Our two main classification dimensions are naturally related to the problem to be solved and to the solution adopted, respectively. The dimension related to the problem concerns the type of *connection requirements* imposed on mission objectives. The dimension related to the adopted solving methods is given by the *communication model* describing the contingencies under which robots are expected to connect to each other.

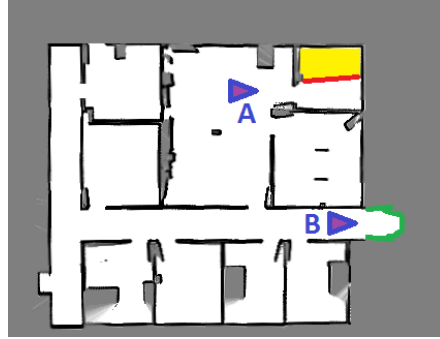
2.1 Connection requirements

The first dimension is related to the connection requirements to which a particular exploration mission should comply. We identify the following three main categories, presented in increasing levels of potential impact.

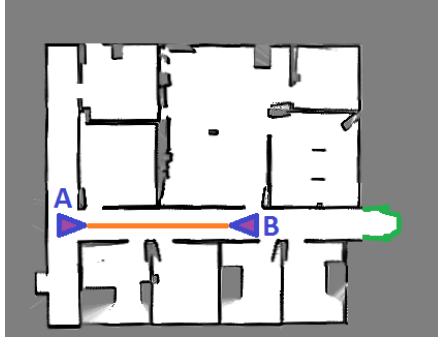
- **None**: robots are not required to communicate. This category covers the most straightforward scenarios for multirobot exploration in environments with limited communication.
- **Event-based connectivity**: robots must regain connection with some teammates according to a policy triggered by particular events, such as the discovery of new information about the environment or simply striking a given time. In the limit case, the whole team must regain global connection.



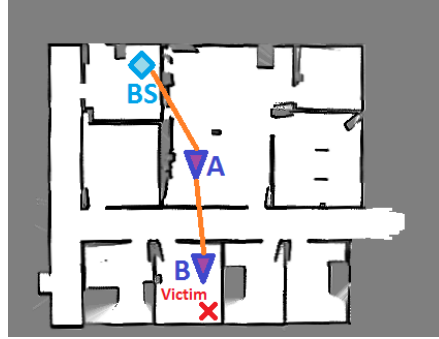
(a) Under unlimited communication, both robot A and robot B share the same map. They can decide who should explore the unique frontier (in green).



(b) Due to limited communication, the yellow portion of the map is known only by robot B. Therefore, robot A goes to re-explore the area beyond the red frontier.



(c) If the two robots have previously agreed on a line-of-sight meeting, they can merge their maps and achieve better coordination.



(d) A multi-hop chain allows to stream videos (or, more generally, the collected data) to a base station (BS) where human operators supervise the system.

Fig. 1: Example scenarios encountered in communication-restricted multirobot exploration. The environment is a portion of the “sdr_site.b” from the radish repository [6] (approximate size 40×30 m). Robots and communication links have been manually drawn for illustration purposes.

- **Continuous connectivity:** each robot must be always connected to any other teammate, either directly or in a multi-hop fashion. Custom backup procedures are often in charge of regaining global connection if lost.

Note that the above requirements have an impact on the exploration strategies that can be adopted (e.g., with continuous connectivity, the exploration strategy is “less free” to move robots around) and even on the possibility of completing the exploration task (e.g., requiring continuous connectivity could forbid to reach some points in the environment).

A final aspect, independent of the three listed above, is that the team could include a BS to which the gathered information must periodically be delivered. In this case, exploration is typically complete when there is a time \bar{t} for which map $M^{\bar{t}}$ known at the BS is such that $M^{\bar{t}} = E$.

2.2 Communication models

A typical mobile robot employed in an exploration mission is able to exchange information with its teammates (and, possibly, with a BS) over a radio channel (e.g., WiFi), whose quality is known to degrade with distance and presence of physical obstacles. Here, by “communication model” we intend the *prior knowledge* about communication capabilities that robots exploit during decision-making about where they should go next. For instance, the communication model may predict that a robot will be in communication with another robot only when they are mutually visible to each other within some range. Being hardly affected by false negatives, this is a conservative method that can be used to construct robust multi-hop chains between robots to achieve, for example, live video streaming between a location of interest and the BS. Sometimes no communication model is assumed. In such case, robots do not rely on the possibility of communicating when deciding which locations to reach. Communications can thus only occur opportunistically due to robots encounters, enabling explicit coordination only on an episodic basis. The communication models that are usually adopted are the following:

- **None:** robots do not make any assumption about the possibility of communicating between any two arbitrary locations.
- **Line-of-Sight (LoS):** a robot can communicate with another robot if the line segment connecting their positions (in the same two-dimensional plane) is entirely contained in the free space of E (and, usually, its length is within a maximum value d).
- **Circle:** a robot can communicate with any other robot within a fixed maximum distance d , independently of the presence of obstacles.
- **Signal:** a robot r_i can communicate with another robot r_j with a probability that depends on the estimated signal power between the positions of r_i and r_j (the higher the signal power, the higher the probability). A possible choice for estimating the signal power P (in dBm) at distance d from the source is to resort to the empirical relation [7]:

$$P = P_0 - 10n \log_{10}(d/d_0),$$

- where P_0 denotes the reference signal value at d_0 , and n is a model parameter.
- **Traces:** robots can communicate with other robots by leaving messages in the environment (e.g., by dropping beacons or tags). In this case, robots are usually not equipped with any radio communication device.

There are a few exceptions to the above scheme, that will be discussed in the next section. For instance, bandwidth limitations can also be modelled.

3 Survey

This section surveys the most relevant literature by first categorizing papers according to connection requirements. When presenting the work pertaining to a particular class, differences in the communication models will be highlighted.

Table 1 provides an overview of the surveyed papers organized according to the two categories introduced in Section 2. In addition, we use the superscript “*” to denote works that involve the presence of a BS to which relaying the gathered information, “+” to denote works in which an appropriate contingency plan or backup behavior is in charge of making the robots seeking connection with teammates when lost, “†” to denote works that explicitly take some form of bandwidth constraints into account, and “‡” to denote works containing experiments with real robots.

The majority of works that are present in the literature propose *distributed* approaches to coordinate robots in an exploration strategy. Decisions on where to move (selecting frontiers and assigning robots to them) are done at a local level, independently by single robots or by subsets of connected robots. In those works that instead rely on a *centralized* method, such tasks are typically the prerogative of a central entity acting as a collector of all the information and as a global planner. The works described in the following adopt a distributed approach if not stated otherwise. Table 1 shows centralized and distributed approaches in blue and black, respectively.

3.1 Exploration without any connection requirement

In this case, connections (either direct or multi-hop) are episodic, not planned in advance, events which robots can discover only when they actually happen and exploit to improve decision-making through opportunistic coordination. The literature includes several works following this approach, despite its well-known inefficiencies that could impact on the mission execution due to the difficulty of operating team coordination throughout the whole mission. A typical example is the occurrence of repeated explorations of the same area by different robots, which could result in an increase of the exploration time, as well as in useless travels and consequent greater energy consumption. Another example is given by the impossibility for human operators residing at a BS to enforce timely information updates, which may be crucial in some scenarios, like search and rescue.

One of the first papers belonging to this class is [8], where robots coordinate locally with teammates currently in communication in order to minimize the overlap in the information they collect when visiting new frontiers. Interestingly, in a small environment (18×14 m), experiments with a communication range equal to 30% of the environment diameter show performance comparable to that with global communication.

While the previous study assumes known robots' starting positions (and close to each other, as customarily done), the system presented in [9] is capable of coping with uncertainty in the robots' starting locations. In particular, the decision-theoretic framework of [8] is extended to bias the exploration towards areas allowing to verify hypotheses on the relative locations of robots.

Authors of [10] present a multirobot system that achieves coordination between teammates that are in communication through an auction mechanism in which each robot places bids relative to the attractiveness of candidate target locations. In this work, additional care is required to handle the inherent asynchronicity of the bidding mechanism. The proposed strategy works without the need of a central coordinator, but a central agent can be added to behave as an interface between the team and a human operator.

In the approach presented in [11], exploration proceeds by implicitly coordinating the robots by means of *potential fields* [12] which drive them towards unexplored areas far from those being explored by teammates.

In [13], multirobot exploration is cast into the framework of decentralized Markov Decision Processes (MDPs), and Distributed Value Functions (DVF) are employed to derive robots' exploration plans as optimal policies of the underlying MDPs. In particular, the classical methodology of DVFs is extended to cope with limited information sharing and communication breaks.

Generally, in multirobot exploration, it is assumed that robots operate in a static environment. This assumption is dropped in [14], which introduces the concept of *local collaborations* to denote physical actions that need to be performed collaboratively by more than one robot at a common position, and presents and evaluates some heuristics to decide *when* and with *whom* to collaborate. In selecting the robots for local collaborations, the proposed heuristics discard those assignments that can partition the network.

Contrarily to the previous works, authors of [15, 16, 17] consider instead the case where information must be relayed to a BS. The idea underlying the system presented by [15] is to divide robots in *explorers* and *relays*. The former ones are in charge of exploring new regions of the environment, while the latter ones travel back and forth to deliver the information collected by the explorers to the BS: explorers and relays coordinate themselves through the choice of appropriate *rendezvous* points. Notice that the periodic meetings at rendezvous points in some sense constrain the exploration process. However, this is a side effect of the particular exploration strategy, and it is not a requirement of the problem.

[16] presents a multirobot exploration system with a focus on the routing protocols required to share information between robots. At the decision-making level, this work investigates the behavior of two strategies: a role-based strategy,

inspired to that of [15], and a strategy in which the robots construct a multi-hop communication chain between frontiers and BS. Experimental results seem to suggest that the latter should be preferred only in the early phases of the exploration task. (A succinct presentation of the openly available source code used in [16] can also be found in [18].)

A role-based strategy is also proposed in [17] in the context of multi-UAV exploration. An interesting feature of this strategy is that UAVs are given the possibility of “sacrificing” themselves by continuing to explore even when they are about to run out of battery. When this actually happens, UAVs land and act as (fixed) communication relays for the remaining robots.

The Circle communication model is considered in [19, 20]. [19] takes inspiration from [10] to propose a market-based coordination mechanism. The assumption of a Circle communication model biases the robots to move towards frontiers where they will probably be in communication with other robots. A roadmap-based distributed exploration algorithm is proposed in [20], where the Circle communication model is exploited to derive a guarantee for the perfect functionality of the algorithm with respect to the robots’ perception range.

The Signal communication model is employed by [21, 22]. In [21], robots compute the expected quality of the signal between a fixed BS and the candidate target frontiers, and the exploration is biased towards those that are likely to allow the robots to communicate with the BS. The method requires global knowledge to assign frontiers to robots. [22] extends the role-based strategy of [15] with the possibility of arranging rendezvous assuming to be able to communicate through walls under the Signal communication model (actually, the strategy is agnostic about the particular communication model in use). Due to the possible presence of false positives, the execution of fallback plans in the form of “classical” rendezvous at physical locations can also be required. Interestingly, this work postulates the possibility of adopting a low bandwidth channel to send small broadcast messages to the whole team.

Finally, there are some works in which the Traces communication model is assumed. In [23], the authors present an exploration algorithm for graph-represented environments, providing strong bounds on the exploration time for the case in which the graph is a tree. Robots coordinate themselves by dropping “bookkeeping devices” when visiting graph vertices, whose state can be read and changed by subsequent visiting robots. The analysis is further extended in [24] for generic graphs.

Belonging to this last class, worth of a brief mention are also ant-inspired algorithms. Originally designed for distributed coverage of unknown environments [25], they have been recently proposed as an alternative to frontier-based exploration. [26] presents an ant-inspired algorithm where the robots are assumed to be able to “mark” explored parts of the environment, so that they are not accessed by their teammates. Similarly, [27] proposes an ant-inspired algorithm in which the robots deposit traces of “pheromone” (that can diffuse and evaporate as time evolves) and are guided in the exploration by very simple behaviors. It is worth noticing that in these last works experiments are per-

formed only in simulation, since a principled physical implementation of this kind of traces to be effectively used in multirobot exploration is still a research challenge.

3.2 Exploration under event-based connection requirements

A significant number of studies deal with multirobot exploration tasks in which connections between robots are subject to limitations. In particular, most of these works consider variants of the problem of connecting (a subset of) the robots at selected time instants, which depend on the occurrence of particular events. In a typical scenario, human operators reside at the BS with the aim of monitoring how the mission unfolds and so should be informed when something relevant happens, for example, when data is acquired from a frontier. It is not surprising that most of these works abandon the fully distributed coordination paradigm that dominates the previous class of works: indeed, the presence of a central entity that collects data can also be exploited to derive better exploration strategies, since it can exploit a “global view” of the current state of the mission. In particular, this planning scheme is adopted in [28, 29, 30, 31].

In [28], a set of robotic agents is assumed to be already present in an environment for an information-gathering mission, like an exploration task. Two problems are addressed: (i) find a deployment of relay robots which ensures global connection according to the LoS communication model between each agent and a BS and (ii) given the current deployment and the new locations that agents should reach (in our case, exploration frontiers), find the redeployment of relay robots which minimizes their traveling time. The former problem is reduced to the computation of a minimum Steiner tree where the agents’ locations constitute the terminal set, while the latter is solved by means of a (generally sub-optimal) dynamic programming algorithm.

In [29, 30], the Circle communication model is assumed. In [29], the authors propose a method for multirobot exploration that ensures, besides full connectivity from the frontiers to a BS each time the robots visit a new frontier, a sufficient bandwidth for the transmission of data on the relay chain. This is achieved by splitting the problem in three sub-problems (explorers placement, relays placement, and path generation) which are then solved as variations of known combinatorial optimization problems. The same setting is considered in [30], but in presence of a heterogeneous team composed by (a) robots that can be allocated to frontiers, (b) static relays, and (c) a single robot in charge of deploying the relays. Here, an additional challenge is given by the need to compute a minimum cost path for the relay-deploying robot.

While in the previous works new plans are computed once the whole explorers-relays network has been formed, the work presented in [31] moves along a complementary direction. Bandwidth constraints are not considered, but an exploration strategy is devised that does not depend on any specific communication model, assumed that the one employed is enough conservative (typical choices are LoS and Circle, but assuming a small range d), and does not require synchronous coordination among robots. In particular, new plans can be submitted to arbitrary

groups of robots, as soon as they have relayed to the BS the information from the frontiers to which they were assigned, in order to speed-up exploration.

The system proposed in [32] considers a generic information-gathering mission scenario in which the robots must synchronously regain global connection with a mobile BS after a fixed time interval. The authors prove the inapproximability of the problem and propose a heuristic algorithm based on planning robots' paths in turns, choosing the best path from a pool of samples according to a utility function which, in an exploration context, may be related to the information gain of the path. The proposed algorithm can be implemented in a centralized or distributed fashion, and can be easily extended to handle the presence of a fixed BS.

While all the above studies propose a planning scheme that is fundamentally centralized, the system presented in [33] is fully distributed. In presence of a fixed BS to which the gathered information should be relayed, the behavior of the robots is regulated by a utility function which considers the amount of information a robot has not yet delivered to the BS and the estimated amount of information known by the BS. Tuning a parameter, the mission planner is able to specify strategies ranging from a completely greedy exploration, with no returns to the BS, to an exploration ensuring the maximum update frequency at the BS. Here, a form of periodic reconnection is distributedly obtained as a collective emerging behavior, with no guarantees on the times the data are delivered to the BS. Robots do not assume anything about the communication model, but when two (or more) robots returning to the BS are in mutual communication range, they locally coordinate themselves so that only one remains in charge of delivering the information collected by both to the BS.

3.3 Exploration under continuous connection requirements

The most restrictive class of works imposes a continuous connection between all the robots (and, possibly, with a BS), either directly or in a multi-hop fashion. This could be needed, for instance, in situations where real-time image streaming must be available to human operators (e.g., in search and rescue), or to ensure a high degree of coordination, since new plans can be computed always assuming a global shared knowledge between all robots.

The systems in [34] and [35] feature a “mild” form of continuous connection requirement to which robots are subject while they try to visit the whole environment. Both works leverage on a small set of behaviors to perform exploration, with one behavior in charge of regaining connection with other team members as soon as it is lost. However, the disconnection period can be arbitrary long. In particular, the system presented in [34] considers the LoS communication model, and the exploration strategy is tested in scenarios with increasing amount of prior information about the environment. The exploration strategies proposed in [35] are guaranteed to achieve full exploration of an unknown environment with an architecture that uses few behaviors, messages locally exchanged between robots (that try to remain always connected with each other), and dropped beacons.

The majority of works belonging to this class adopts instead a “hard” form of continuous connection, in which robots’ movements are carefully planned to always maintain global connection according to a conservative LoS or Circle communication model.

[36] devises a centralized exploration strategy in which a local search method is in charge of guiding the team. In particular, the utility of the team is computed in terms of distances from the closest frontiers: a configuration that does not allow global connection is highly penalized and is never selected by the algorithm. A distributed version of this strategy is presented in [37] and is tested in an underwater mapping scenario with a high-fidelity simulator.

The centralized algorithm proposed in [38] considers a fixed BS and constructs a connected exploration tree in which the robots are organized in explorers and link stations: explorers are placed at the leaves of the tree, while the link stations are the inner nodes and ensure the connectivity between a BS (the root of the tree) and the explorers. These ideas are further elaborated in [39] to cope with the constraints given by a moving BS.

More recently, [40] models multirobot exploration under continuous connection requirements as a constrained optimization problem and tests the proposed (centralized) exploration strategy in an underwater map construction task.

Finally, [41] has recently addressed the problem of online exploration of generic targets by means of a team of robots that must remain continuously connected. This work proposes a role-based distributed algorithm capable of guaranteeing the exploration of all the targets appearing online throughout the mission. This algorithm can be applied to the multirobot exploration case, where targets could be associated to new frontiers.

4 Research directions

Some promising research directions can be derived from the analysis presented above. Here we just report some of them.

- Most of the presented works (except those of the first column of Table 1) rely on *a priori* knowledge about the possibility of communicating between any two points of the environment in order to decide where to move. However, the communication models expressing such a knowledge are chosen conservatively in order to avoid wrong predictions (i.e., LoS or Circle with a sufficiently small range d) or are based on some model of signal propagation (i.e., Signal). It would be interesting to have at disposal a method for building, as the exploration mission unfolds, less conservative and more realistic *communication maps* of the environment representing, for instance, the actual signal strength measurements between locations [42]. Note that this task is potentially more complex than mapping the signal strength of a fixed source (like in WiFi mapping [43]), because it would involve the estimation of the possibility of communicating between *any* two points in the environment.

Connection requirement	Communication model				
	None	LoS	Circle	Signal	Traces
None	Burgard <i>et al.</i> , 2005 [8] [†] Fox <i>et al.</i> , 2006 [9] [‡] Zlot <i>et al.</i> , 2002 [10] [†] Lin and Lyons, 2015 [11] [†] Matignon <i>et al.</i> , 2012 [13] [†] Andre and Betusterer, 2016 [14] de Hoog <i>et al.</i> , 2011 [15] [†] Andre, 2015 [16] [†] [‡] Cesare <i>et al.</i> , 2015 [17] [†] [‡]	Spirin and Cameron, 2014 [22] [†] ⁺	Sheng <i>et al.</i> , 2006 [19] Franchi <i>et al.</i> , 2009 [20] [†] Spirin and Cameron, 2014 [22] [†] ⁺	Visser and Slanec, 2008 [21] [†] Spirin and Cameron, 2014 [22] [†] ⁺	Brass <i>et al.</i> , 2011 [23] Brass <i>et al.</i> , 2014 [24] Andries and Charpillet, 2015 [26] Kuyucu <i>et al.</i> , 2015 [27]
Event-based	Spirin <i>et al.</i> , 2013 [33] [†]	Stump <i>et al.</i> , 2011 [28] [†] [‡] Banfi <i>et al.</i> , 2016 [31] [†] [‡] Hollinger and Singh, 2012 [32] [†]	Pei <i>et al.</i> , 2013 [29] [†] [‡] Pei and Mutka, 2012 [30] Banfi <i>et al.</i> , 2016 [31] [†] [‡] Hollinger and Singh, 2012 [32] [†]		
Continuous		Arkin and Diaz, 2002 [34] [†] Nestmeyer <i>et al.</i> , 2016 [41] [†]	Rooker and Birk, 2007 [36] Rathmann and Birk, 2011 [37] ⁺ Mukhija <i>et al.</i> , 2010 [38] [†] Pandey <i>et al.</i> , 2012 [39] [†] Kapoutsis <i>et al.</i> , 2016 [40] [†] [‡] Nestmeyer <i>et al.</i> , 2016 [41] [†]		Jensen <i>et al.</i> , 2016 [35] [†] [‡]

Table 1: Summary of literature survey. *: presence of a BS to which relaying the gathered information; †: an appropriate contingency plan or backup behavior is used to regain connection when lost; ‡: bandwidth constraints are explicitly taken into account; ‡: experiments with real robots. Centralized and distributed approaches are shown in blue and black, respectively.

- The majority of current methods for multirobot exploration under connection requirements usually do not consider any contingency plan when a communication link is expected between two points but, when the robots are there, they cannot actually communicate (see works in Table 1 that are not marked with “+”). In this case, some sort of pre-computed reconnection strategies could be employed, but this topic is largely uncovered by the literature.
- The relationships between communication models, connection requirements, and exploration strategies should be studied more carefully. From our survey, it emerges that centralized exploration strategies are increasingly preferred over distributed ones when roughly moving from the top left corner of Table 1 (no connection requirement and no communication model) to the bottom right corner (continuous connectivity and Circle).
- Issues relative to theoretical guarantees (for instance, in terms of worst-case bounds on reconnection time), to benchmarking different solutions, and to experimental protocols to validate exploration strategies with real robots constitute three avenues for future research. For instance, it would be interesting to apply the framework of [44] to communication-restricted settings of multirobot exploration.
- Extensions of the current framework that cover dynamic connection requirements are mostly unaddressed. For example, a system operating in a mission without any connection requirement could be required to switch to a continuous connection requirement to set up a video stream with the BS in case a robot discovers something important. Other extensions could consider the use of an existing communication infrastructure (like a cellular network) instead, or besides, the use of the *ad hoc* communication networks created by robots [45].
- It would be interesting to evaluate if the same methods used to cope with communication constraints in multirobot exploration could be generalized to other multirobot information gathering tasks, like patrolling and monitoring. It seems that many of our considerations about the connection requirements imposed on the mission, about the communication models, and about the centralized vs. distributed nature of exploration strategies can be easily generalized to other settings.

5 Conclusion

In this paper, we have surveyed the current state of the art of multirobot exploration in presence of restrictions on communication. We have identified two main dimensions for our analysis, namely the requirements posed on connections by the mission the robots are accomplishing and the models that robots use to predict from where they can communicate in partially explored environments. In addition, other dimensions, like the centralized or distributed nature of exploration strategies, have been proposed. Finally, we have illustrated some promising directions for future research in the field of communication-restricted multirobot exploration.

Bibliography

- [1] A. Farinelli, L. Iocchi, and D. Nardi. Multirobot systems: A classification focused on coordination. *IEEE T Syst Man Cyb*, 34(5):2015–2028, 2004.
- [2] S. Thrun. Robotic mapping: A survey. In *Exploring Artificial Intelligence in the New Millenium*, pages 1–35. Morgan Kaufmann, 2002.
- [3] B. Yamauchi. Frontier-based exploration using multiple robots. In *Proc. AGENTS*, pages 47–53, 1998.
- [4] M. Julia, A. Gil, and O. Reinoso. A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. *Auton Robot*, 33:427–444, 2012.
- [5] A. Birk and S. Carpin. Merging occupancy grid maps from multiple robots. *P IEEE*, 94(7):1384–1397, 2006.
- [6] A. Howard and N. Roy. The robotics data set repository (Radish). <http://radish.sourceforge.net/>, 2003.
- [7] P. Bahl and V. Padmanabhan. Radar: An in-building RF-based user location and tracking system. In *Proc. INFOCOM*, volume 2, pages 775–784, 2000.
- [8] W. Burgard, M. Moors, C. Stachniss, and F. Schneider. Coordinated multi-robot exploration. *IEEE T Robot*, 21(3):376–386, 2005.
- [9] D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schulz, and B. Stewart. Distributed multirobot exploration and mapping. *P IEEE*, 94(7):1325–1339, 2006.
- [10] R. Zlot, A. Stentz, M. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. In *Proc. ICRA*, pages 3016–3023, 2002.
- [11] T. Liu and D. Lyons. Leveraging area bounds information for autonomous decentralized multi-robot exploration. *Robot Auton Syst*, 74:66–78, 2015.
- [12] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. In *Autonomous robot vehicles*, pages 396–404. Springer, 1986.
- [13] L. Matignon, L. Jeanpierre, and A. Mouaddib. Coordinated multi-robot exploration under communication constraints using decentralized Markov decision processes. In *Proc. AAAI*, pages 2017–2023, 2012.
- [14] T. Andre and C. Bettstetter. Collaboration in multi-robot exploration: To meet or not to meet? *J Intell Robot Syst*, 82:325–337, 2016.
- [15] J. de Hoog. *Role-based multirobot exploration*. PhD thesis, University of Oxford, Oxford, 2011.
- [16] T. Andre. *Autonomous exploration by robot teams: coordination, communication, and collaboration*. PhD thesis, Alpen-Adria-University, Klagenfurt, 2015.
- [17] K. Cesare, R. Skeelee, S. Yoo, Y. Zhang, and G. Hollinger. Multi-UAV exploration with limited communication and battery. In *Proc. ICRA*, pages 2230–2235, 2015.
- [18] T. Andre, D. Neuhold, and C. Bettstetter. Coordinated multi-robot exploration: Out of the box packages for ROS. In *Proc. IEEE GLOBECOM*, pages 1457–1462, 2014.

- [19] W. Sheng, Q. Yang, J. Tan, and N. Xi. Distributed multi-robot coordination in area exploration. *Robot Auton Syst*, 54(12):945–955, 2006.
- [20] A. Franchi, L. Freda, G. Oriolo, and M. Vendittelli. The sensor-based random graph method for cooperative robot exploration. *IEEE-ASME T Mech*, 14(2):163–175, 2009.
- [21] A. Visser and B. Slamet. Including communication success in the estimation of information gain for multi-robot exploration. In *Proc. WiOPT*, pages 680–687, 2008.
- [22] V. Spirin and S. Cameron. Rendezvous through obstacles in multi-agent exploration. In *Proc. SSRR*, pages 1–6, 2014.
- [23] P. Brass, F. Cabrera-Mora, A. Gasparri, and J. Xiao. Multirobot tree and graph exploration. *IEEE T Robot*, 27(4):707–717, 2011.
- [24] P. Brass, I. Vigan, and N. Xu. Improved analysis of a multirobot graph exploration strategy. In *Proc. ICARCV*, pages 1906–1910, 2014.
- [25] S. Koenig, B. Szymanski, and Y. Liu. Efficient and inefficient ant coverage methods. *Ann Math Artif Intel*, 31(1-4):41–76, 2001.
- [26] M. Andries and F. Charpillet. Multi-robot taboo-list exploration of unknown structured environments. In *Proc. IROS*, pages 5195–5201, 2015.
- [27] T. Kuyucu, I. Tanev, and K. Shimohara. Superadditive effect of multi-robot coordination in the exploration of unknown environments via stigmergy. *Neurocomputing*, 148:83–90, 2015.
- [28] E. Stump, N. Michal, V. Kumar, and V. Isler. Visibility-based deployment of robot formations for communication maintenance. In *Proc. ICRA*, pages 4498–4505, 2011.
- [29] Y. Pei, M. Mutka, and N. Xi. Connectivity and bandwidth-aware real-time exploration in mobile robot networks. *Wirel Commun Mob Comput*, 13(9):847–863, 2013.
- [30] Y. Pei and M. Mutka. Steiner traveler: Relay deployment for remote sensing in heterogeneous multi-robot exploration. In *Proc. ICRA*, pages 1551–1556, 2012.
- [31] J. Banfi, A. Quattrini Li, N. Basilico, I. Rekleitis, and F. Amigoni. Asynchronous multirobot exploration under recurrent connectivity constraints. In *Proc. ICRA*, pages 5491–5498, 2016.
- [32] G. Hollinger and S. Singh. Multirobot coordination with periodic connectivity: Theory and experiments. *IEEE T Robot*, 28(4):967–973, 2012.
- [33] V. Spirin, S. Cameron, and J. de Hoog. Time preference for information in multiagent exploration with limited communication. In *Proc. TAROS*, pages 34–45, 2013.
- [34] R. Arkin and J. Diaz. Line-of-sight constrained exploration for reactive multiagent robotic teams. In *Proc. AMC*, pages 455–461, 2002.
- [35] E. Jensen, L. Lowmanstone, and M. Gini. Communication-restricted exploration for search teams. In *Proc. DARS*, 2016. To appear.
- [36] M. Rooker and A. Birk. Multi-robot exploration under the constraints of wireless networking. *Control Eng Pract*, 15(4):435–445, 2007.
- [37] R. Rathnam and A. Birk. Distributed communicative exploration under underwater communication constraints. In *Proc. SSRR*, pages 339–344, 2011.

- [38] P. Mukhija, K. Krishna, and V. Krishna. A two phase recursive tree propagation based multi-robotic exploration framework with fixed base station constraint. In *Proc. IROS*, pages 4806–4811, 2010.
- [39] R. Pandey, A. Singh, and K. Krishna. Multi-robot exploration with communication requirement to a moving base station. In *Proc. CASE*, pages 823–828, 2012.
- [40] A. Kapoutsis, S. Chatzichristofis, L. Doitsidis, J. de Sousa, J. Pinto, J. Braga, and E. Kosmatopoulos. Real-time adaptive multi-robot exploration with application to underwater map construction. *Auton Robot*, 40(6):987–1015, 2016.
- [41] T. Nestmeyer, P. Robuffo Giordano, H. Bühlhoff, and A. Franchi. Decentralized multi-target exploration using a connected network of multiple robots. *Auton Robot*, 41(4):989–1001, 2017.
- [42] J. Banfi, A. Quattrini Li, N. Basilico, I. Rekleitis, and F. Amigoni. Multi-robot online construction of communication maps. In *Proc. ICRA*, pages 2577–2583, 2017.
- [43] J. Fink and V. Kumar. Online methods for radio signal mapping with mobile robots. In *Proc. ICRA*, pages 1940–1945, 2010.
- [44] J. Faigl and M. Kulich. On benchmarking of frontier-based multi-robot exploration strategies. In *Proc. ECMR*, pages 1–8, 2015.
- [45] J. Banfi, N. Basilico, and F. Amigoni. Minimizing communication latency in multi-robot situation-aware patrolling. In *Proc. IROS*, pages 616–622, 2015.