

Discussion of Multi-Robot Exploration in Communication-Limited Environments

Arnoud Visser¹ Julian de Hoog² Adrian Jiménez-González³ José Ramiro Martínez de-Dios³

Abstract—Many robotics applications benefit from cooperative exploration by multiple robots. Often the environments that they operate in contain significant communication challenges due to their size or complexity. There has been much work in keeping teams of autonomously exploring robots connected to one another, but in certain environments, distant locations can only be reached if robots relay information via teammates or explore autonomously beyond team communication range.

I. INTRODUCTION

Mobile robots are already being used for a variety of tasks, from reconnaissance to search and rescue in disaster zones to inspection of hazardous areas. A recent high-profile use of robots involved the inspection of radioactive areas at the Fukushima nuclear reactor in the aftermath of the Sendai earthquake in Japan.

Today, most robots in such situations are controlled remotely by human operators, and control of a robot typically requires the operator's full concentration. However, the amount of incoming information can be overwhelming, there may be many other factors competing for the operator's attention, and in the case of very remote environments, control delay can be significant. As a result, there is much incentive to offload some of the work to the robots, and partial or full autonomy is a desirable capability in many scenarios.

A. Research Goals

Given the existing applications that robots are used for and given the work that has been done to date in multi-robot exploration (see for instance [1]), it is evident that further development of methods for multi-robot exploration are necessary that take communication drop-out into account. The method described in this article aims to take early steps towards filling this void.

The specific research questions that this article tackles are: (i) How can a team of robots be coordinated to explore a previously unknown and communication-limited environment as efficiently as possible; and (ii) how can new information obtained by this team be gathered at a single location as quickly and as reliably as possible?

¹Arnoud Visser is with the Intelligent Systems Laboratory, Universiteit van Amsterdam, The Netherlands

²Julian de Hoog is with the Department of Mechanical Engineering, University of Melbourne, Australia

³Adrian Jiménez-González and José Ramiro Martínez de-Dios are with the Robotics, Vision and Control Group, University of Sevilla, Spain

II. ROLE-BASED EXPLORATION

Role-based exploration was introduced in [2].

In this discussion we concentrate on the most surprising and perhaps the most interesting result of this paper; the emergent behavior. Once the role swap rule and a dynamic hierarchy had been implemented and verified in experiments, attention was turned to the effect of various hierarchy structures on the exploration effort. For example, given a team of six robots, is it better to have two branches of length 3, or three branches of length 2 (*i.e.* two Relays per Explorer, or one)? It turned out that longer chains of relays do not lead to an improvement, and can actually introduce difficulties due to the increased number of required rendezvous. Additionally, a smaller number of exploring robots means that exploration proceeds slower.

However, analysis of the runs involving many short branches, *i.e.* only a single Relay for each Explorer, revealed that when the Role Swap Rule is implemented and the hierarchy changes dynamically, *many short chains of robots can behave much like one long chain of robots*. This is an important result for coordination of any large team of robots: a simple hierarchy and a simple role exchange mechanism can lead to highly coordinated behaviour that adapts dynamically both to communication availability and to the shape of the environment.

III. EXPERIMENTAL RESULTS

To examine and evaluate the behaviour of Role-Based Exploration, and to compare it to some existing methods, a series of experiments were performed, initially in simulation [3], and subsequently as part of a practical implementation on a team of Pioneer robots [2]. Several commonly used multi-robot exploration algorithms were implemented in order to conduct a comprehensive comparison.

There are many variables when using a team of robots for exploration (such as sensor range, communication range, number of robots in the team, and type of environment, to name but a few). Only four representative data sets have been presented in [2], [3], but a comprehensive set of results can be found in [4]. Based on these comprehensive results, the behaviour of the competing exploration algorithms can be summarised as follows:

- **Speed of exploration:** Greedy exploration leads to the fastest discovery of new information in almost all cases. The Role-Based methods approach the speed of G only when the sensor or communication ranges are very large relative to the size of the environment, when

there are many robots exploring, or when environments contain long hallways (or fewer frontiers than robots in the team). However, it must be remembered that this speed of exploration is not necessarily of use in many situations, as it might not reach human responders. Greedy-Periodic exploration typically lags behind the Role-Based approaches, and by far the slowest approach is Leader-Follower exploration.

- **Relay of information to BaseStation:** Leader-Follower exploration demonstrates perfect relaying of new information to the BaseStation (since all team members are connected), but being quite slow still lags behind the other approaches in terms of the amount of information brought to the BaseStation per unit of time. The Role-Based approaches demonstrate the most regular, reliable relaying of new information to the BaseStation in almost all cases.
- **Information sharing within the team:** The Role-Based approaches also provide major advantages regarding information sharing within the team. In this regard they always perform better than Greedy exploration. In many experiments, Greedy-Periodic exploration demonstrated better information sharing than the Role-Based approaches, but it must be remembered that most of the time the Role-Based approaches had covered more area, and thus had longer paths and longer intervals outside of one another's ranges. Leader-Follower exploration of course demonstrates perfect information sharing, since all members are always connected.
- **Responsiveness to BaseStation commands:** If an operator had to make sudden changes to the exploration effort, such as designating parts of the environment as being of greater interest or pulling the whole team out, then again the Role-Based approaches would be of much greater use than Greedy exploration. Responsiveness is better for the Role-Based approaches over Greedy in almost all cases, and generally outperforms Greedy-Periodic as well when taking size of explored environment into consideration. Again, Leader-Follower exploration shows the greatest responsiveness of all, which follows from the fact that the team is fully connected at all times.

The simulation experiments reveal that Role-Based Exploration provides a compromise between those approaches that maintain perfect communication (but explore slowly), and those that explore quickly (but ignore communication limitations). In many robotics applications, such a compromise can be important.

IV. ALTERNATIVE TEAM STRUCTURES

A. Hierarchies with Higher Branching Factors

The team structure described and analysed in this article only considers chains of robots in a hierarchy having a branching factor of 1 (except for the root). In other words, a robot may have at most a single child and a single parent. What if this limitation were to be relaxed – for example, what

if one relay were to serve two (or three, or four) explorers? Multi-robot rendezvous is a well studied problem, and has been approached by multiple authors [5], [6], [7].

There is no inherent reason not to attempt such a hierarchy, and certainly the results would be interesting to compare. However, such an implementation would significantly increase the complexity of team coordination. A relay serving two explorers, for example, could perform rendezvous in two ways: (i) rendezvousing with both explorers at the same time; (ii) rendezvousing first with one and then with the other in a single trip before returning to the BaseStation; or (iii) rendezvousing with the first explorer, returning to BaseStation, rendezvousing with the second explorer, returning to BaseStation, etc.

Each of these situations introduces additional challenges. In case (i), rendezvous by three (or more) robots requires very careful timing in order for them all to meet at the same time, and risks one robot delaying two others if it is late. In case (ii), it is possible that the relay would have to travel long distances before it could return to the BaseStation. In case (iii), explorers would be on their own for long periods, and the intervals in which new information would reach the BaseStation would be much greater.

As a result, it seems that changing the branching factor of the hierarchy may not be worth the additional complications introduced. In some rare cases environment structure may suggest a higher branching factor (for example, a long hallway with two large rooms at the end would suggest using a relay with two explorers). But even then, the emergent behaviour that arises with the use of the Role Swap Rule (see section II) typically deals well with any shape of environment.

B. Dynamic Hierarchy Structures

In the current team structure of Role-Based Exploration with role swaps, robots may jump around within the team hierarchy and assume different roles. However, the structure of the hierarchy is fixed from the start, and the structure itself never changes. If the exploration effort begins with 6 relays and 3 explorers, these numbers of each role will stay the same throughout the entire effort, even if robots are swapping roles. Why not allow for a dynamic structure, where branches of the hierarchy can lengthen or shorten as required?

A dynamic hierarchy structure would certainly introduce additional challenges: any robot in a branch affected by a change in hierarchy structure would need to be informed of the change. Given that robots are likely to wander in and out of range, this would need to be managed carefully. Furthermore, the emergent behaviour resulting from the Role Swap Rule already provides a reasonably good solution. As demonstrated in section II and Fig. 1, the nature of Role-Based Exploration with role swaps is such that the team adjusts dynamically to the shape of the environment.

There is however at least one situation where a change in hierarchy would be quite beneficial, namely the very early stages of the exploration effort. Greedy methods tend

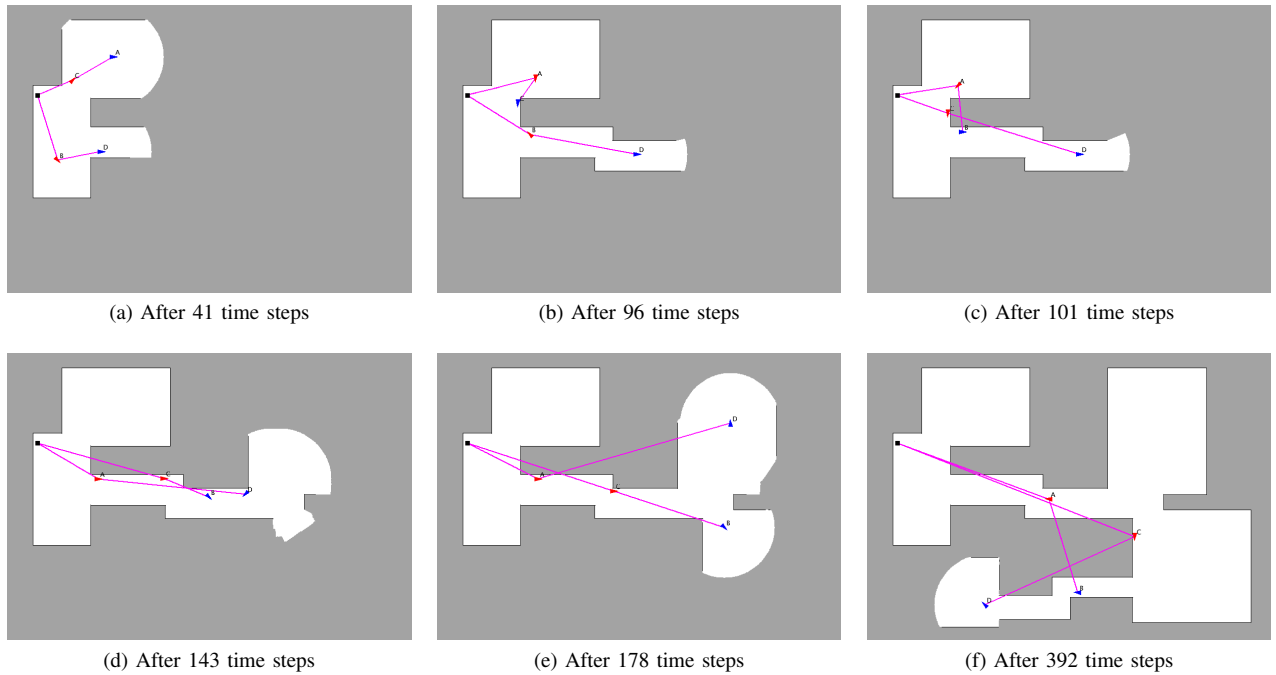


Fig. 1: Screenshots from a simulation run demonstrating the interesting behaviour that emerges from teams having hierarchies of short branches, using the role swap rule. Magenta lines indicate team hierarchy. Initially exploration proceeds as normal, with A and D as explorers, and C and B their respective relays (1a). After 96 time steps, A completes exploration of the top room and switches roles with C (1b). Soon after, C switches roles with B (1c). The exploration phase now involves traversal of a long hallway to the only open frontiers. The team dynamically adjusts, via role switches, in such a manner that A , C , and B act as a multi-hop link for explorer D (1d). As new frontiers open after 178 time steps, the team again restructures so that two robots actively explore (1e). When another long hall is found much later, again the team behaves like a 4-robot multi-hop chain (1f).

to outperform role-based methods significantly at the start, since there are twice as many robots exploring. A useful adjustment to Role-Based exploration could be to begin the exploration effort with all robots acting as explorers, and certain robots only taking on the roles of relays once the limits of communication ranges are reached. This as well would need to be managed carefully: robots would likely not be able to predict the moment that they lose contact, and agreements on which robots become explorers and which robots become relays would need to be formed *before* communication is lost.

An ability to dynamically form a hierarchy could also be useful in situations where robots do not start with a common frame of reference. For example, if two groups of robots enter a mine from different entrances and meet in the middle, it would be useful for them to be able to agree on some sort of hierarchy “on-the-fly”. This would be a very interesting and useful behaviour to develop in future work.

C. Heterogeneous Teams

The current team structure does not take into account potential heterogeneity in the team. It is possible that different types of robots with different capabilities and different sensor loads may be involved in the same effort.

For example, it is possible that the robot team could be composed of fast, simple robots (ideal for relaying) and more advanced robots with more sophisticated sensors (ideal for exploring). In such a scenario, the Role Swap Rule would

need to be adjusted to take robot types and their ideal roles into account.

As another example, the potential use of aerial robots for creation of a communication infrastructure has received more and more attention recently. Using UAVs as mobile relays that can either land in strategic positions or ferry information between robots without taking ground-based obstacles into account could be a very interesting extension to Role-Based Exploration.

D. Teammate and Environmental Prediction

Gains in exploration efficiency could likely be achieved with better prediction of where teammates may be when they are out of range. This is particularly true for explorers returning to rendezvous. A scenario observed often in simulation was that an explorer took a long path back to a rendezvous point, when it would have been faster to predict the relay’s location and intercept it at an earlier point. This is particularly true in environments with loops and multiple passageways. The route of a relay should be highly predictable to its child explorer, and this knowledge could be used to greater effect.

It is entirely possible that in some robotics applications prior knowledge of the environment exists. For example, if a building needs to be inspected, floor plans of the building may be available. In such cases, plans performed in advance could help to steer explorers into areas of greatest interest. One possible way to do this would be to incorporate

desirability into the utility calculations of frontier polygons. Polygons close to desired areas would thus be chosen more likely than those farther away.

Another problem that was observed in simulation is that an explorer doesn't prioritise "finishing the job". The point at which an explorer turns to return to rendezvous is timed exactly, and currently there is no system to slightly adjust the timing of the state change even when it is of advantage. This means that in some scenarios an explorer will explore 90% of a room before it turns to rendezvous, and must then subsequently come all the way back into the room to finish the job. This is expensive, and it would make more sense to leave the relay waiting for a short time, in order to complete exploration of the room and not have to visit it again.

This is certainly not an easy problem, as it involves knowledge of concepts such as what shape a room may have. Furthermore it cannot be solved perfectly when there is no prior knowledge of the environment; in the final unexplored 10% of a room, there may be a further passage leading to another room. Nevertheless, some degree of environmental prediction could lead to significant gains in exploration efficiency.

E. Extending to Three Dimensions

Finally, the work on Role-Based Exploration to date has been limited to flat environments. Even the experiments conducted with real robots involved a mostly planar floor and range sensing in a plane. Clearly most real applications would involve information on multiple planes or in three dimensions. The move to three dimensions is a big challenge in robotics today, and much work still needs to be done on every aspect, from mapping, to navigation, to planning. Nevertheless it is useful to examine how an approach such as Role-Based Exploration might hold up if it were to be employed in a 3D environment at some point in the future.

Consider the case of a team of robots travelling in three dimensions (*e.g.* a team of UAVs [8] or underwater robots, or a team of ground robots [9] in a pile of rubble). Roles could be assigned in the same manner, and robots could either explore, or relay. Localisation and mapping would be very challenging, but there has already been much work in this direction by the robotics community. Frontier polygons would become frontier polyhedra, and certainly the calculation of such polyhedra would be more complex and possibly involve some manner of contour surfaces instead of contour lines. Path planning would require significantly more resources, and consequently frontier-to-robot assignment would require significantly more calculation. Considerably greater memory would be required for storage of maps, and thus sharing of maps would require greater bandwidth. Skeletonisation would become a much more expensive operation.

In short, computation and memory requirements would soar in the move to three dimensions. It is likely that any early approach applying Role-Based Exploration (or any multi-robot exploration algorithm) would need to apply some clever compression and optimisation techniques while hardware development catches up with requirements. All of

that said, there is no inherent reason for some variant of Role-Based Exploration not to be applied to a team of robots jointly exploring in three dimensions. But there is certainly much future work to be done before that happens.

V. FURTHER ADVANTAGES

The strengths of Role-Based Exploration as compared with some competing approaches have already been discussed: Role-Based Exploration is generally better than greedy approaches at relaying new information to the BaseStation at regular intervals, demonstrates better inter-teammate connectivity and information sharing, and allows for tighter, more responsive control of the team, all while still providing faster exploration than leader-follower approaches. However, there are some further possible advantages worth mentioning here.

A. Cooperative Localisation

Given that rendezvous is an essential element of the approach, Role-Based Exploration leads to frequent meetings of teammates. Meetings between robots have been used for a long time to improve mutual localisation, and cooperative localisation is a well-studied problem [10], [11], [12], [13]. Thus, an approach such as Role-Based Exploration that plans on repeated mutual observation can lead to cleaner maps and better localisation.

B. Acyclic Communication

In the implementations of Role-Based Exploration presented in this article, information was gained in a monotonic manner, adding to or overwriting previously gained information. However, many multi-robot systems elsewhere use probabilistic information gathering methods, where each information update received through the network contributes to the degree of confidence in a given estimate of the map [14], [15], [16]. In a multi-robot team where messages are passed over multiple hops, there is a risk that a single information update will be received more than once via separate paths, incorrectly increasing the confidence in that estimate. This can be solved by data tagging messages and maintaining a record of their history, but in large networks this could lead to unnecessarily large messages. As a result, *acyclic* communication networks are of great interest in some multi-robot systems, and are assumed or explicitly created in several cases (*e.g.* [17], [15]). Therefore, when such additive probabilistic information gathering methods are employed, Role-Based Exploration could potentially provide a solution: an acyclic team hierarchy already exists, and can be used to prevent double counting of information updates.

VI. CONCLUSIONS

Extensive experiments using Role-Based Exploration have demonstrated the approach's strengths and weaknesses, and the following guide could be used to determine when Role-Based Exploration is most applicable.

Role-Based Exploration *should* be used when:

- 1) the highest priority is to receive quick and regular information updates at a central BaseStation.

- 2) the communication range of the robots is small compared with the size or complexity of the environment. The more difficult it is for the robots to communicate, the greater the benefit is in having mobile relays improving connectivity and information sharing within the team.
- 3) there are a large numbers of robots in the team (four or more). This is particularly true in environments with fewer frontiers. The more robots there are in the team, the more important it is for them to coordinate their effort well.
- 4) environments contain many obstacles or long hallways and passages.
- 5) the SLAM system is of a type that benefits from mutual localisation between the robots, or an acyclic communication protocol is desirable.

Role-Based Exploration should *not* be used when:

- 1) quick exploration of the full environment is the highest priority. Greedy methods perform this faster in most cases.
- 2) full connectivity of the team is required at all times.
- 3) availability of communication is ubiquitous throughout the environment.
- 4) the team is composed of robots that have difficulty turning on the spot, or cannot retrace their paths due to environmental factors.
- 5) the desired application is not primarily exploration-oriented. Role-Based Exploration was designed for exploration of unknown environments, and better algorithms are likely to exist for other applications, such as coverage of open space.

Extensive experiments in simulation and on a team of Pioneer robots compared Role-Based Exploration to some competing algorithms, notably greedy and leader-follower approaches. Greedy approaches can lead to quicker exploration of the full environment, and leader-follower approaches can maintain better connectivity and information sharing within the team. However, Role-Based Exploration provides a useful trade-off between the two, particularly as communication becomes less and less reliable, as more robots are used in the team, and as environments become more complex. In such situations, Role-Based Exploration provides regular updates to the BaseStation (which greedy approaches do not) while still exploring the farthest reaches of the environment (which leader-follower approaches do not).

REFERENCES

- [1] E. Olson *et al.*, "Progress towards multi-robot reconnaissance and the MAGIC 2010 competition," *Journal of Field Robotics*, vol. 29, no. 5, pp. 762–792, September 2012.
- [2] J. de Hoog, A. Jiménez-González, S. Cameron, J. R. M. de Dios, and A. Ollero, "Using mobile relays in multi-robot exploration," in *Proceedings of ACRA (Australian Conference on Robotics and Automation)*, 2011, submitted.
- [3] J. de Hoog, S. Cameron, and A. Visser, "Dynamic team hierarchies in communication-limited multi-robot exploration," in *Proceedings of the International Workshop on Safety, Security and Rescue Robotics (SSRR)*, July 2010.
- [4] J. de Hoog, "Role-based multi-robot exploration," Ph.D. dissertation, Department of Computer Science, University of Oxford, September 2011.
- [5] H. Ando, Y. Oasa, I. Suzuki, and M. Yamashita, "Distributed memoryless point convergence algorithm for mobile robots with limited visibility," *IEEE Transactions on Robotics and Automation*, vol. 15, no. 5, pp. 818–828, Oct 1999.
- [6] J. Lin, A. Morse, and B. Anderson, "The multi-agent rendezvous problem," in *Proceedings of the IEEE Conference on Decision and Control*, vol. 2, 2003, pp. 1508 – 1513 Vol.2.
- [7] J. Cortes, S. Martínez, and F. Bullo, "Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions," *IEEE Transactions on Automatic Control*, vol. 51, no. 8, pp. 1289–1298, 2006.
- [8] A. Visser, "The international micro air vehicle flight competition as autonomy benchmark," in *Robotics Competition: Benchmarking, Technology, Transfer and Education Workshop - European Robotics Forum*, March 2013.
- [9] H. Akin, N. Ito, A. Kleiner, J. Pellenz, and A. Visser, "Robocup rescue robot and simulation leagues," *AI Magazine*, vol. 34, 2013.
- [10] S. I. Roumeliotis and G. A. Bekey, "Distributed multirobot localization," *IEEE Transactions on Robotics and Automation*, vol. 18, pp. 781–795, 2002.
- [11] D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schulz, and B. Stewart, "Distributed multirobot exploration and mapping," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1325–1339, July 2006.
- [12] X. Zhou and S. Roumeliotis, "Multi-robot SLAM with unknown initial correspondence: The robot rendezvous case," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2006, pp. 1785–1792.
- [13] A. Franchi, G. Oriolo, and P. Stegagno, "On the solvability of the mutual localization problem with anonymous position measures," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 3193–3199.
- [14] N. Sorensen and W. Ren, "Rendezvous problem in multi-vehicle systems: Information relay and local information based strategies," in *Proceedings of the IEEE Mountain Workshop on Adaptive and Learning Systems*, 2006.
- [15] A. Chapman and S. Sukkarieh, "A protocol for decentralized multi-vehicle mapping with limited communication connectivity," in *IEEE International Conference on Robotics and Automation (ICRA)*, may 2009, pp. 357–362.
- [16] G. M. Mathews, H. Durrant-Whyte, and M. Prokopenko, "Decentralised decision making in heterogeneous teams using anonymous optimisation," *Robotics and Autonomous Systems*, vol. 57, no. 3, pp. 310–320, 2009.
- [17] G. Mathews and H. Durrant-Whyte, "Decentralised optimal control for reconnaissance," in *Proceedings of the IEEE Conference on Information, Decision and Control (IDC)*, 2007, pp. 314–319.