UvA-DARE (Digital Academic Repository)

Dynamic team hierarchies in communication-limited multi-robot exploration

de Hoog, J.; Cameron, S.; Visser, A.

DOI

[10.1109/SSRR.2010.5981573](https://doi.org/10.1109/SSRR.2010.5981573)

Publication date

2010

Document Version

Author accepted manuscript

Published in

Proceedings of the 8th IEEE International Workshop on Safety, Security & Rescue Robotics (SSRR 2010)

[Link to publication](#)

Citation for published version (APA):

de Hoog, J., Cameron, S., & Visser, A. (2010). Dynamic team hierarchies in communication-limited multi-robot exploration. In *Proceedings of the 8th IEEE International Workshop on Safety, Security & Rescue Robotics (SSRR 2010): 26-30 July 2010, Bremen* IEEE.
<https://doi.org/10.1109/SSRR.2010.5981573>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Dynamic Team Hierarchies in Communication-Limited Multi-Robot Exploration

Julian de Hoog and Stephen Cameron

Oxford University Computing Laboratory

Wolfson Building, Parks Road, OX13QD

Oxford, United Kingdom

{julian.dehoog, stephen.cameron} @comlab.ox.ac.uk

Arnoud Visser

Universiteit van Amsterdam

Science Park 107, NL 1098 XG

Amsterdam, Netherlands

a.visser@uva.nl

Abstract — In the near future, groups of autonomous robots using wireless communication will be used for a wide variety of tasks. In many such applications, communication may be unreliable and communication ranges difficult to predict. While most current approaches to this problem strive to keep team members within range of one another, we propose an approach in which navigation and exploration beyond range limits is explicitly planned for. Robots may either explore or relay known information, and the team hierarchy corresponds to a tree. As the exploration effort unfolds, robots swap roles within this tree to improve the efficiency of exploration. Since robots reactively adjust to communication availability, the resulting behaviour is robust to limited communication. This makes it particularly suitable for applications such as robotic search and rescue, where environments are likely to contain significant interference and unexpected communication ranges.

Keywords: *role-based exploration, limited communication, rendezvous, dynamic hierarchy, robotic search and rescue*

I. INTRODUCTION

Current trends in robotics and networking suggest that in the near future, teams of autonomous robots using wireless communication will be used for a wide variety of tasks, such as reconnaissance, surveillance, and search-and-rescue. For this to be possible, robot team members must be able to autonomously navigate in directions of interest, while maintaining a connection to the rest of the team. Taking communication drop-out and failure into account remains a major challenge in many applications.

While many multi-robot coordination strategies have ignored the communication problem in the past, recently there has been an increased effort to take limited communication into account. Line-of-sight approaches strive to keep team members within direct line of sight of one another [1], [7], [11]. Utility-based approaches factor communication likelihood into robots' decision of where to go next – this can be applied in terms of frontier utilities [2], [6], [12] or in terms of market-like bids [16], [14]. A closely related problem is that of environment coverage [3], [10], and decentralised control laws have been derived that lead to near-optimal sensing coverage of a given environment [13].

Many such approaches have proven effective at maintaining team connectivity. However, some environments may contain significant interference, or may extend beyond team communication limits even when multi-hop communication is applied. In such cases, full exploration of the environment is only possible if the whole team navigates through the environment together (*e.g.* like robot packs [12]), or if some members of the team explore beyond communication range limits. In this paper we are interested in the latter, and we present a multi-robot exploration approach that explicitly plans for exploration beyond communication range limits.

II. PROBLEM DESCRIPTION

The problem that we are particularly interested in is the consolidation of the knowledge of all robot team members at a single location. In a search-and-rescue scenario this corresponds to human responders' point of entry, while in reconnaissance or surveillance this corresponds to the base station where information is gathered and analysed.

We assume no prior knowledge of the environment. Given recent developments in robotic mapping and localisation, we also assume that robots are capable of maintaining maps of the environment and an estimate of their position. Recent approaches such as scan-matching [9] or particle filters [6] make this a realistic assumption. Localisation does not need to be perfect and there is some room for error. However, robots need to be able to find their way to within communication range of agreed rendezvous points. In our work we assume occupancy-grid based maps, but the approach could be tailored to topological maps as well.

Our main goals are to

- 1) explore the environment as efficiently as possible;
- 2) relay new information to the base station as quickly and as often as possible; and
- 3) minimise the time that team members spend out of range of the base station.

This must be achieved without placing an unrealistic burden on team communication systems.

† This conference paper is published in the following proceedings:

The 2010 IEEE International Workshop on Safety Security and Rescue Robotics (SSRR), pp.1-7, 26-30 July 2010

doi: 10.1109/SSRR.2010.5981573

The original publication is available at <http://ieeexplore.ieee.org/> © 2010 IEEE

In each of these two simple scenarios, a role swap leads to shorter paths, and consequently faster exploration. There are many other, similar situations where it is advantageous to have a dynamic team hierarchy.

C. Swapping Roles

Dynamism in the team hierarchy is achieved by a single, simple rule which we call the “Role Swap Rule”: Consider two robots A and B , each having destinations D_A and D_B , respectively. Let $\gamma(u, v)$ represent the path cost from location u to location v in a given map. When u and v are known, this value is easy to calculate using standard path planners (such as A*) on the map. Suppose A and B have encountered one another and established a communication link. If

$$\max\{\gamma(A, D_A), \gamma(B, D_B)\} > \max\{\gamma(A, D_B), \gamma(B, D_A)\}$$

then let A assume B ’s role, state, and location in the tree, and let B assume A ’s role, state, and location in the tree. This rule is applied equally to relays and explorers, both within the same branch and across branches.

For purposes of visualisation, this rule means that we are always eliminating the longest path among the four paths computed. The longest path is the deciding factor regarding how long a step in the exploration process will take, so by ensuring that the maximum path cost is always as small as possible, we are consistently reducing the bottleneck. Indeed we believe the choice of maximum path cost to be the best way to decide when to swap roles; we are interested in a formal proof but leave this as possible future work.

In our implementation, a role swap is achieved by exchange of a *role swap message* that contains all relevant information necessary for a role exchange. In a real implementation, a handshake process would likely be desirable. The role swap message contains the following information: ID, role, state, list of frontiers, child’s ID, parent’s ID, child rendezvous, parent rendezvous, and current goal. This information is sufficient for one robot involved in the swap to completely replace the role of the other robot involved in the swap.

Of particular note is that robots have unique ID numbers and recognise one another by these. When two robots swap roles, they also swap ID numbers. Thus, swaps outside a given robot’s communication range do not adversely affect that robot’s behaviour. For example, consider the case where robot A has to rendezvous with his parent B . However, unknown to A , B has swapped roles with D . D now has B ’s ID number. Since A is looking for the ID only, and not a specific robot, A finds D as his parent, and role-based exploration may proceed as expected.

Perhaps the best way to demonstrate this dynamic behaviour, and how it improves the exploration effort, is by example: see Figure 3 and the accompanying explanation in Textbox 1. In this simple demonstration, three separate applications of the role swap rule have been applied: the swap in Stage III involves a relay and an explorer in the same branch, the swap in Stage IV involves a relay and an explorer from separate branches, and the swap in Stage V involves two relays in separate branches. Additional possible applications of the rule exist, and the rule extends as well to larger hierarchies and longer branches.

Stage I: Four robots set out to explore an unknown environment. Initially, A and C are relays, B and D are explorers. Following initial range scans, two frontiers are discovered (F_1 and F_2). By joint utility maximisation, B chooses to explore F_1 and D chooses to explore F_2 .

Stage II: B explores F_1 and A follows B . The frontier F_1 opens up into a bigger frontier at F_3 . In the meantime D explores F_2 and C follows D . D reaches the end of the room, and decides to rendezvous with C to relay new information back to base station. The members of the team are still fully connected, although the connection to base station is lost.

Stage III: B and A reach the limits of the team communication chain, and break from it – B continues to explore, A continues to follow. Two frontiers open up (F_4 and F_5), and B chooses to explore F_4 . In the meantime, D and C have rendezvoused, and C must relay new information back to base station. However, D ’s only frontier of interest is at F_3 (D is not aware of A and B ’s latest exploration knowledge since these have been out of range). Since

$$\max\{\gamma(D, F_3), \gamma(C, Base)\} > \max\{\gamma(D, Base), \gamma(C, F_3)\}$$

the role swap rule is applied, and C and D trade positions in the tree. C is now an explorer with F_3 as its goal, and D is now its parent relay with the base station as its goal. C and D agree on R_1 as the next rendezvous point (calculation of rendezvous locations is performed as described in [5]).

Stage IV: Enough new information has been gained by B , so A and B rendezvous and A turns around to relay new information to the base station while B continues to explore. However, A encounters C , on its way to explore F_3 . Since

$$\max\{\gamma(C, F_3), \gamma(A, Base)\} > \max\{\gamma(C, Base), \gamma(A, F_3)\}$$

the role swap rule is applied, and C and A trade positions in the tree. C becomes a relay for B , while A becomes an explorer with D as parent relay.

Stage V: Now an explorer, A chooses the nearest frontier at F_6 and starts to explore. D , having relayed information to the base station, is on the way to rendezvous with its child (now A) at R_1 . C is on its way to the base station to relay new information. Since

$$\max\{\gamma(D, R_1), \gamma(C, Base)\} > \max\{\gamma(D, Base), \gamma(C, R_1)\}$$

the role swap rule is applied, and D and C trade positions in the tree. D becomes a relay for B , and C becomes a relay for A . D turns around to complete the task of relaying information, while C turns around to rendezvous with its child A at R_1 .

Stage VI: In the meantime, A and B have continued exploration of open frontiers and fully explored the environment. Eventually all robots return to the base station.

Textbox 1: A description of the events in Figure 3

Fig. 3: A demonstration of how a team hierarchy may change during an exploration effort. The team hierarchy is presented to the right of each stage of exploration. Dark parts of the map are unexplored, white parts have been sensed using range finders. The base station is the purple square on the left. Coloured circles indicate agents' respective communication ranges. A full explanation of the stages in this figure is provided in Textbox 1.

D. Performance and Considerations

1) *Completeness*: Since the actual exploration of new space is performed using frontier exploration, the approach is complete. Explorers will continue to explore as long as there are open frontiers, so the environment will be fully explored regardless of its size (given sufficient power).

2) *Deadlock, oscillations, and starvation*: Since a strict hierarchy exists and high-priority control commands may be issued top-down while low-priority information may be delivered bottom-up, we do not anticipate any deadlock issues, nor have we seen any arise in extensive simulations. In very rare cases, oscillation may occur, for example if two robots agree to exchange roles while moving, but it is no longer advantageous to have exchanged roles in their new locations once the exchange is complete. Such a scenario is easily circumvented by introducing a timeout on the rate of role swaps a robot may undergo.

Another common problem when large number of robots are involved is starvation, when there are more robots than open frontiers. Once all frontiers have been assigned to exploring robots we let remaining explorers choose any frontier they like. This means that some frontiers may be explored by multiple robots. In practice however, open frontiers often lead to multiple new open frontiers, so often this is a constructive approach. Starvation is a problem common to any frontier-based method and ours does not uniquely suffer from it.

3) *Path planner*: Since the role swap rule is heavily dependent on calculation of path costs in the map, an efficient path planner is essential. In our simulations a simple A* planner has proven sufficiently fast. In very large or three-dimensional environments, however, a different approach may be required. Our current rendezvous point calculation method provides an efficient calculation of the map skeleton – this could double as a topological estimate of the map for quick path cost estimation.

4) *Heterogeneous teams*: The current implementation does not take into account potential heterogeneity in the team. It is possible that different types of robots with different sensor loads may be involved in the same effort, in which case it may be desirable for certain types of robots to play particular roles (e.g. relays could be fast, simple robots while explorers could carry more intricate sensors). In such a scenario, the role swap rule would need to be adjusted to take robot types and their ideal roles into account.

IV. SIMULATION RESULTS

A. Comparison to other methods

To examine the behaviour of the role swap rule and to compare dynamic role-based exploration to other approaches, we ran several experiments in our custom-built MRESim simulator. Past work describes the simulator [4] and it is available upon request from the authors. Here we compare three exploration approaches:

- A) Greedy frontier-based exploration, where frontiers are chosen based on a utility function that takes into account

information gain and path cost [15]. This approach is similar to those used in [2], [6], [12].

- B) Role-based exploration with a static team hierarchy [5].
- C) Role-based exploration with a dynamic team hierarchy, using the role swap rule described in section III-C.

Experiments were conducted with a variety of team sizes and in a variety of environments. Here we present results that we believe to be representative of most of our experiments. As an environment, we used a slightly modified version of the `vasche_library_floor1` floor plan from the Radish data set¹. For each of the approaches, 10 robots were used. Both the static and dynamic role-based approaches used a hierarchy that contained 5 pairs of robots, *i.e.* one relay for each explorer. Figure 4 shows the full results of this run, and a screenshot is provided in Figure 5.

Dynamic role-based exploration leads to faster coverage of the environment than greedy frontier-based exploration, in spite of the fact that only half as many robots are actively exploring (the other half are relays). This is due to the fact that poor inter-team awareness in greedy frontier-based exploration means that robots are likely to cover areas that have already been explored.

Dynamic role-based exploration also outperforms static role-based exploration in every metric. It leads to faster exploration (figure 4a), greater awareness of the exploration at the base station (figure 4b), greater inter-teammate awareness (figure 4c) and quicker responsiveness to the base station (figure 4d). Only late in the experiment does connectivity to base station seem weaker, but this is mainly due to the fact that more of the environment has been discovered and robots must travel longer distances.

Overall, dynamic role-based exploration leads only to a small improvement in terms of speed of exploration, as compared with conventional frontier or utility-based approaches. The main gains, however, are inter-robot awareness and team responsiveness. For applications such as search and rescue, where instant control over the robots is highly desirable, this is an important characteristic.

B. Emergent behaviour

In a separate experiment we explored a variety of hierarchy structures and depths. For example, we compared two teams of six robots, the first having two chains of relay-relay-explorer, and the second having three chains of relay-explorer. In our experience longer chains of relays do not lead to an improvement, and can actually introduce difficulties due to the increased number of required rendezvous.

In fact, the dynamic role-swapping method introduced in this paper means that even short relay-explorer chains behave like much longer chains (as demonstrated in Figure 3). This was an unexpected but positive result for us: that the emergent behaviour of hierarchies having a depth of only two resembles what one would expect of much longer communication relay chains.

¹This data set was obtained from the Robotics Data Set Repository (Radish) [8]. Thanks go to Ashley Tews for providing this data.

(a) Percentage of environment explored. Dynamic role-based exploration and greedy frontier exploration perform similarly well.

(b) Percentage of the environment known at the base station. Dynamic role-based exploration leads to better relaying of information to the base station.

(c) Average percentage of currently known space known to all team members, in other words ability of the team to share information among teammates.

(d) Average time, over all agents, since last contact with the base station, either directly or via a mobile relay.

Fig. 4: Simulation results: a comparison of greedy frontier-based exploration, static role-based exploration, and dynamic role-based exploration.

V. DISCUSSION AND FUTURE WORK

While dynamic role-based exploration does not lead to vastly faster exploration than currently popular utility and frontier based approaches, it has other important advantages. Using the same number of robots, similar exploration can be achieved while maintaining considerably better team connectivity. In large and communication-challenged environments, this is very helpful. The main advantages and disadvantages can be summarised as follows:

Advantages: Explorers and relays adjust to size of the environment and communication availability reactively. Provided sufficient power is available, the approach leads to full exploration of environments regardless of how much interference or how short communication ranges are. Equal numbers of robots lead to similar exploration, but considerably better teammate awareness and team connectivity. We believe role-based exploration can be extended to three dimensions, making it suitable for UAVs in addition to ground-based robots.

Disadvantages: Since individual robots may be out of range of the base station, control over the full team may not be instantaneous. If a robot or a group of robots becomes lost or incapacitated, this information may not reach the base station (other than by lack of response). The approach is heavily reliant on reasonably accurate mapping and localisation.

In future work, we intend to consider integration of additional roles into the framework. In particular we are interested in the possibility of joint aerial and ground based teams, with airborne robots providing a communication infrastructure. This could lead to incorporation of both *static relay* and *dynamic relay* roles, in addition to standard *explorers*.

We are also intrigued by the emergent behaviour of short relay-explorer chains emulating much longer multi-hop chains. We hope to examine this property of dynamic role-based exploration in more depth.

Work to date has focussed on two-dimensional environments, but we do evaluate every aspect of the approach with an eye towards possible extension to three dimensions. Potential bottlenecks in the calculations include 3D path planning and 3D skeletonisation (for calculation of rendezvous points), and we intend to investigate possible solutions to these problems.

This work is in an early stage; nevertheless, we envision an extension of this approach to teams of tens of small ground or air-based robots jointly exploring unknown environments in three dimensions in the future. Since it is likely that future applications will require autonomous exploration beyond communication range limits, we hope that the ideas presented here provide an early step in that direction.

Fig. 5: A screenshot of our MRESim simulator showing dynamic role-based exploration in the `vasche_library_floor1` environment, after 1000 time steps.

Left: purple and yellow lines indicate team hierarchy, with yellow connections indicating current communication link. An example communication range has been presented for robot *A* (thin green polygon).

Top right: an example skeletonisation of the map, as performed by robot *H*. Green dots indicate possible rendezvous locations.

Bottom right: the team hierarchy as a tree.

REFERENCES

- [1] R. Arkin and J. Diaz. Line-of-sight constrained exploration for reactive multiagent robotic teams. *Advanced Motion Control, 2002. 7th International Workshop on*, pages 455–461, 2002.
- [2] W. Burgard, M. Moors, C. Stachniss, and F. Schneider. Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3):376 – 378, 2005.
- [3] J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.
- [4] J. de Hoog, S. Cameron, and A. Visser. Role-based autonomous multi-robot exploration. In *International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE)*, November 2009.
- [5] J. de Hoog, S. Cameron, and A. Visser. Selection of rendezvous points for multi-robot exploration in dynamic environments. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, May 2010.
- [6] D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schulz, and B. Stewart. Distributed multirobot exploration and mapping. *Proceedings of the IEEE*, 94(7):1325–1339, July 2006.
- [7] A. Howard, M. Mataric, and G. Sukhatme. An incremental deployment algorithm for mobile robot teams. *Intelligent Robots and System, 2002. IEEE/RSJ International Conference on*, 3:2849–2854 vol.3, 2002.
- [8] A. Howard and N. Roy. The robotics data set repository (radish), 2003.
- [9] M. Pfingsthorn, B. Slamet, and A. Visser. A Scalable Hybrid Multi-Robot SLAM method for Highly Detailed Maps. In *Proceedings of the 11th RoboCup International Symposium*, July 2007.
- [10] S. Poduri and G. S. Sukhatme. Constrained coverage for mobile sensor networks. In *International Conference on Robotics and Automation (ICRA)*, 2004.
- [11] M. Powers and T. Balch. Value-based communication preservation for mobile robots. In *7th International Symposium on Distributed Autonomous Robotic Systems*, 2004.
- [12] M. N. Rooker and A. Birk. Multi-robot exploration under the constraints of wireless networking. *Control Engineering Practice*, 15(4):435–445, 2007.
- [13] M. Schwager, J.-J. Slotine, and D. Rus. Decentralized, adaptive control for coverage with networked robots. In *International Conference on Robotics and Automation*, April 2007.
- [14] W. Sheng, Q. Yang, J. Tan, and N. Xi. Distributed multi-robot coordination in area exploration. *Robot. Auton. Syst.*, 54(12):945 – 955, 2006.
- [15] A. Visser and B. Slamet. Including communication success in the estimation of information gain for multi-robot exploration. *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks and Workshops, 2008. WiOPT 2008. 6th International Symposium on*, pages 680–687, April 2008.
- [16] R. Zlot, A. Stentz, M. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. *Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on*, 3:3016 – 3023, 2002.