Using Mobile Relays in Multi-Robot Exploration*

Julian de Hoog and **Stephen Cameron**
Department of Computer Science, University of Oxford, UK
{julian.dehoog, stephen.cameron} @cs.ox.ac.uk

Adrian Jiménez-González, J. Ramiro Martínez de-Dios and **Anibal Ollero**
Robotics, Vision and Control Group, University of Sevilla, Spain
{adrianjg, jdedios, aollero} @cartuja.us.es

Abstract

Many robotics tasks require autonomous exploration by teams of robots. In difficult or large environments, communication drop-out complicates this task. Several approaches exist that aim to keep the team connected, but even so there is an inherent limit to the range that can be explored. In this paper we describe and examine Role-Based Exploration, an approach that uses mobile relays to ferry information back and forth within the team, and compare it to methods that do not. There are significant advantages in the use of such relays, such as improved coordination and responsiveness, and adaptability to unexpected communication dropout. The approaches are implemented and validated on a team of real robots.

1 Introduction

Robots are already being used for exploration tasks, and it is likely that in the near future there will be an increased need for robots that can autonomously map and explore unknown environments. Recently there has been an emphasis on using teams, rather than single robots, for such tasks. While there are many advantages to having multiple robots, the use of a team introduces additional challenges of team coordination, communication and information sharing [Burgard *et al.*, 2005; Fox *et al.*, 2006].

For any multi-robot team there is an inherent limit to the range within which the team can stay connected. Lets assume that each robot has a large range, and communication is conducted in an ad-hoc fashion, over multiple hops when necessary. Even then, in large or complex environments there will still be areas that remain unaccessible unless one or more robots explore beyond the team's communication range. This is particularly true for environments with significant communication interference. In this research we are interested in comparing and developing multi-robot exploration algorithms that take this into account, *i.e.* methods which lead to full exploration of unknown environments, even beyond team communication range limits.

We assume that there is a central "BaseStation" which is immobile and can communicate with the robots. In exploration (*e.g.* of mines) or search-and-rescue, this corresponds to the human responders' point of entry; in reconnaissance or surveillance this corresponds to the command centre where information is gathered and analysed. Not all multi-robot exploration applications require a central basestation, but many do.

The exact problem we are interested in is: Given a central basestation, a team of robots, and an unknown environment, how can the robots be made to (i) explore the environment as efficiently as possible, while (ii) returning information to the basestation as frequently as possible and (iii) minimising the time that team members spend out of range of the basestation?

Previously, we have proposed "Role-Based Exploration" as a possible solution to this problem [de Hoog *et al.*, 2009] and compared its performance to some existing algorithms in simulation. In this paper we present the results of implementing Role-Based exploration on a team of real robots for the first time.

2 Related Work

Autonomous exploration of unknown environments by individual robots is a well studied problem. While the implementation depends on the sensors used (*e.g.* bearing-based vs. range-based) and the nature of the map (*e.g.* grid vs. topological), robots must always aim for new, unexplored areas of the environment. Often this is achieved using the concept of *frontiers* [Yamauchi *et al.*, 1998], which are boundaries between explored and unexplored space. Utilities can be assigned to such frontiers using a combination of path cost, expected information gain, and communication likelihood.

When a team of robots must be coordinated to cooperatively explore an unknown environment, such utilities can be used to assign individual team members to different frontiers. This has been achieved in the past using a

*This work was partially supported by CONET: Cooperating Objects Network of Excellence (INFSO-ICT-224053) funded by the European Commission under ICT, Framework 7.

central coordination mechanism [Simmons *et al.*, 2000; Burgard *et al.*, 2005]; more recently, the Hungarian method has been used to solve the robot-to-frontier assignment problem optimally [Wurm *et al.*, 2008]. Elsewhere, economic principles have inspired "market-based" approaches in which robots bid on points of interest [Dias and Stentz, 2000] or negotiate on desired goals [Zlot *et al.*, 2002]. Many of these types of methods take communication drop-out into account (typically in a reactive manner), but few explicitly plan to explore beyond communication range limits.

Several other approaches deal with limited communication intentionally, rather than reactively. An early attempt involves maintaining line of sight, in which robots only explore until the threshold of communication is reached [Arkin and Diaz, 2002]. Similarly, several authors have proposed "leader-follower" methods, where one robot explores and teammates follow, maintaining connectivity [Howard *et al.*, 2002; Nguyen *et al.*, 2004]. Methods from graph theory have been further employed to examine the numbers of robots required for such a task [Stump *et al.*, 2008]. Robot teams have also been made to explore unknown environments in "robot packs", using a heuristic that takes communication strength into account to guide their movement [Rooker and Birk, 2007].

A closely related field of research is that of communication coverage, *i.e.* arranging robots (or mobile sensor nodes) in such a manner that they are within communication range of the largest possible space. This has been achieved using potential field methods [Poduri and Sukhatme, 2004] or low level control laws [Esposito and Dunbar, 2006].

While many of the above approaches have shown success in maintaining team connectivity, there will still be situations where parts of the environment can only be reached by autonomous exploration beyond communication range boundaries. This is particularly true if information must be relayed back to a single location. Thus, there remains a need for robust multi-robot exploration algorithms that explicitly take this into account.

3 Proposed Algorithms

In this section we describe the two multi-robot exploration algorithms that we implemented on a team of real robots: Greedy Exploration and Role-Based Exploration. Our Role-Based approach has been described previously [de Hoog *et al.*, 2009]; here we only provide a short summary.

3.1 Greedy Exploration

The algorithm we call "Greedy Exploration" is closely related to the frontier and utility based approaches that have been used extensively elsewhere (*e.g.* [Simmons *et al.*, 2000; Burgard *et al.*, 2005; Fox *et al.*, 2006; Rooker and Birk, 2007]). All robots aim to explore open space as fast as possible, by steering towards new open spaces.

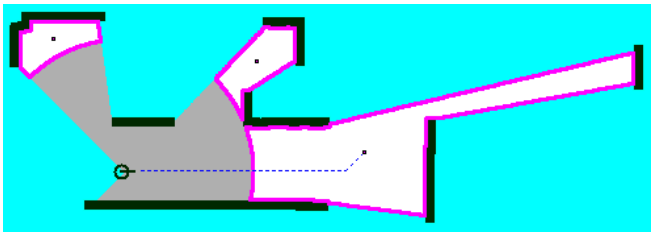

(a) Free space, safe space, and choosing a frontier polygon

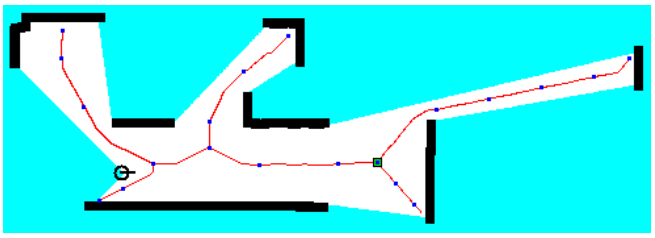

(b) Skeletonisation and rendezvous points

Figure 1: A demonstration of some of the principles used for autonomous exploration

When teammates are within range of one another, the effort is coordinated.

Localisation and Mapping

Each robot keeps track of its own pose (location in the x–y plane and yaw) and has a range sensor. In our research we focussed primarily on the exploration problem, and used existing tools for the typical localisation and mapping problems. Specifically, we used the *mricp* and *amcl* drivers of the Player-Stage framework [Gerkey *et al.*, 2001], which use scan matching and particle filter methods, respectively. Localisation does not need to be perfect, but robots need to be able to retrace their paths and get within communication range of specific rendezvous points. As for the robots' maps, we used occupancy grid based maps, in which each cell is either *free*, *occupied* or *unknown*. The approach could be tailored to topological maps if required.

Frontier Polygons

Whereas many approaches use as frontiers the *boundaries* between free and unknown space, we use *frontier polygons*: polygons of free space that exist beyond an artificially superimposed "safe space", typically half the range of the laser scanner. This is demonstrated in Fig. 1a: using a single range scan it is straightforward to determine obstacles (black) and free space (white); the safe space (grey) is artificially superimposed and the resulting areas of remaining free space (outlined in magenta) are the frontier polygons. For each of these polygons p_i we can then calculate a utility:

$$U(p_i) = A(p_i)/C^n(p_i)$$

where $A(p_i)$ is the area of p_i (the information gain), $C(p_i)$ is the length of the path to that frontier polygon's centre (the path cost), and exponent n determines the exploration behaviour. High values of n lead to lower

utilities for distant frontier polygons, resulting in exploration of nearby frontier polygons (such as rooms); low values of n mean that robots are more likely to pursue larger frontier polygons (such as hallways or open spaces) [Visser and Slamet, 2008]. For experiments reported later in this paper we used $n = 2$, since in practice this provided a good balance between trying to explore open space quickly, while not changing direction too frequently.

Communication and Map Sharing

Team members connect over an ad-hoc network. Whenever two robots are in range, they communicate their own pose estimate and map to one another. Knowing teammates' poses means that robots can avoid one another and erase erroneous obstacle measurements in their maps that are actually their teammates. Knowing teammates' maps means that robots can increase their knowledge of the environment and don't need to visit areas that have already been explored.

We assume that robots start the exploration with a common frame of reference, and are aware of one another's starting positions. This means that even if they have had a period without communication, they can still share maps when they reenter one another's range. This requires reliable localisation, but given today's solutions for this problem we considered this a reasonable expectation (and our experiences with the real robot system detailed in Section 4 validated this assumption).

In our maps, each cell in the occupancy grid maintained one of three values ($free$, *occupied* or *unknown*). To compress the maps we render them in the standard PNG[1] image format, which allows us to store and communicate environments measuring up to 300m×300m at a resolution of 5cm×5cm in sizes of tens of kilobytes or less.

Team Coordination

Regular communication with teammates means that robots can coordinate their exploration. After two robots have communicated, they will share the same map and know one another's poses. Each then calculates a frontier-to-robot assignment on its own, taking its teammates into account (similar to the methods described in [Burgard *et al.*, 2005; Wurm *et al.*, 2008]). Teammates thus choose different frontier polygons and minimise exploration overlap. If the number of unexplored frontier polygons is smaller than the number of robots, multiple robots may end up aiming for the same frontier polygon, but in practice this is a rare occurrence as new frontiers generally open quickly in most environments.

In Greedy exploration there is no explicit effort to relay new information back to the Basestation; knowledge at the Basestation is only extended when a robot happens to wander within the Basestation's communication range.

[1]The PNG format is lossless and raster-based, and uses a number of compression filters to significantly reduce image size. More information is available at `http://www.libpng.org/pub/png/`.

3.2 Role-based Exploration

In Role-Based Exploration, robots explore in much the same way as in Greedy Exploration: localisation, mapping and frontier polygon selection are performed in the same way; communication and map sharing are also conducted in the same way; and robots within range of one another coordinate their choice of frontier in the same way. However, role-based exploration uses several key extensions to the greedy approach.

Roles and Team Hierarchy

Each robot in the team is assigned one of two roles: (i) *Explorers* autonomously explore the environment, returning periodically to previously agreed rendezvous points to pass their knowledge to a Relay; (ii) *Relays* act as mobile links between Explorers and the Basestation, ferrying new knowledge from Explorer to Basestation and control commands from Basestation to Explorer. The team hierarchy corresponds to a tree, with the Basestation at the root and Explorers at the leaves.

Skeletonisation and Rendezvous Points

Each time an Explorer and Relay meet, they agree on the subsequent rendezvous location. The rendezvous location is a specific place on the map chosen by the Explorer and communicated to the Relay. We have shown previously that this choice of rendezvous location significantly affects the efficiency of exploration [de Hoog *et al.*, 2010b]. As can be expected, points in open space or at junctions lead to much more efficient rendezvous than points in dead-ends or near walls.

The Explorer calculates the next rendezvous location as follows: first, it performs thinning[2] on the free space in its map. Second, once a skeleton of the free space has been thus obtained, the Explorer chooses points at junctions, fills in extra points where distances are too far, and removes points where distances are too close. This leads to a small set of points that are suitable for rendezvous; the Explorer finally chooses the one that is closest to the frontier he intends to explore next. In this manner, each subsequent rendezvous is pushed deeper and deeper into the environment, leading to a full exploration of areas even far beyond team communication range limits.

An example of skeletonisation is provided in Fig. 1b. The skeleton is red, the possible rendezvous points are blue, and the eventual choice of rendezvous is the green square.

4 Implementation

For our implementation of these algorithms on a system of real robots, we used the CONET Integrated Testbed at the University of Seville, Spain [Jiménez-González *et*

[2]Thinning, also known as the medial axis transform, reduces a shape to its skeleton while keeping it connected and centred (thinning has much in common with Voronoi diagrams). Several simple algorithms for thinning exist; we use Hilditch's algorithm [Hilditch, 1969].

Figure 2: A team of Pioneer 3-AT robots exploring "Environment 2" in the CONET Integrated Testbed

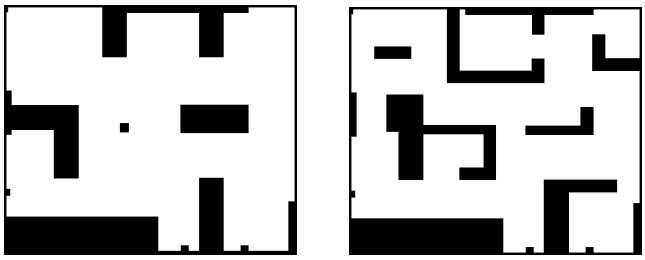

(a) Environment 1 (b) Environment 2

Figure 3: Floorplans of the two environments used: Environment 1 is simple and open (left); Environment 2 contains tighter spaces and has room-like elements (right)

al., 2010]. The Testbed is run in an open space measuring 23.9 m by 20.3 m (Fig. 2). We created two environments, the first being fairly simple and open, the second with tight spaces and room-like structures. Floorplans of these two environments are shown in Fig. 3.

The Testbed includes 5 *Pioneer 3-AT* robots, a skid-steered four-wheeled research platform. Each robot is mounted with a wireless a/b/g/n bridge, a *Hokuyo UTM-30LX* 2D laser range scanner having a range of 30 m and resolution of 0.25 degrees at 25 ms/scan, and an *Acer AspireOne 532h Netbook* (1.66GHz, 1GB Memory) running Ubuntu 10.04. To control the robot actuators and receive data from the laser scanner we used the Player-Stage framework [Gerkey *et al.*, 2001]. Each robot ran its own high level processes onboard (mapping, frontier selection, skeletonisation, etc.) and used Player to communicate with low level software.

Communication between robots and the basestation was performed over the wireless network using TCP/IP. Access to the wireless network within the testbed is ubiquitous – since the goal of our experiments was to examine the performance of exploration algorithms under *limited* communication, we artificially limited communication between all robots. This was performed using a central "Dispatch", a program that opened send and receive channels to each robot, and only forwarded messages between robots when they were within a designated range of one another. For most of the experiments detailed here, we used an 8m radius, since this was approximately one third the width of the lab and meant that a significant area remained beyond communication range.

5 Results & Discussion

5.1 Description of Runs

To properly examine the behaviours of both Greedy and Role-Based exploration, we ran a series of experiments for each of the approaches, in each of the two environments described in section 4, using teams of either 2 robots or 4 robots. The results were fairly consistent across all these tests; in other words, Greedy or Role-Based exploration show no obvious changes in relative performance in different environment types, nor does the use of more or fewer robots show significant relative performance gains. For purposes of space and simplicity we present in this paper only the results of two sets of runs: (i) A team of 4 robots exploring Environment 1 (Fig. 4); and (ii) a team of 2 robots exploring Environment 2 (Fig. 5). However, the results from these runs are representative of a large number of runs in both environments with different team sizes.

5.2 Performance Metrics

Our first performance metric (Figs 4a and 5a) measures the total area that has been discovered, *i.e.* the total amount of free space that has been sensed by all robots. Initially Greedy exploration leads to faster exploration, but Role-Based exploration soon catches up and performs as well or better. The reasons for this are simple: in Greedy exploration, twice as many robots are exploring as in Role-based exploration, so initial progress is faster. The Greedy teammates however do not make an effort to coordinate their exploration (unless wandering within range of one another by chance), and soon duplicate one another's efforts, re-exploring rooms. As expected, 4 robots explore the environment much faster than 2, for both approaches.

Our second performance metric (Figs 4b and 5b) measures how much is known at the basestation, *i.e.* the total free space discovered that has been communicated back to the base- station. Initially all robots are in range so the performance mirrors that of total exploration. Soon, however, robots leave the basestation's range. In Role-Based exploration, regular updates from the relaying robots mean that there are frequent increases in the basestation's knowledge. In Greedy exploration, updates are infrequent, happening only when a robot wanders back within range by chance.

Our third performance metric (Figs 4c and 5c) measures the responsiveness of the team, in terms of the average number of seconds (for all robots) since the last message from the basestation was received. In Greedy exploration, as robots wander out of range the responsiveness becomes poorer and poorer, with only occasional

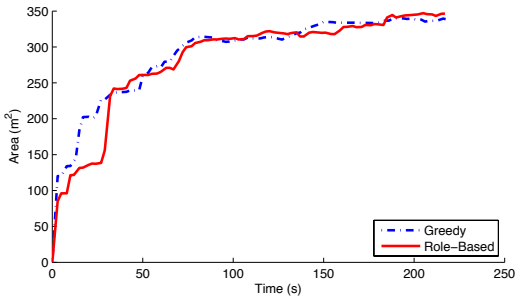

(a) Total free space discovered

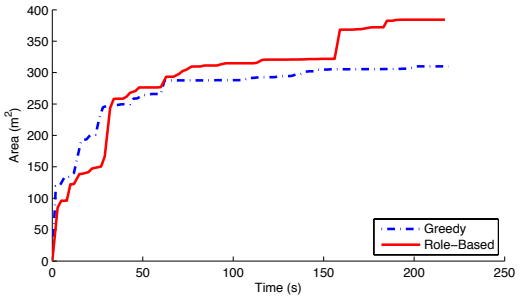

(b) Free space known at Basestation

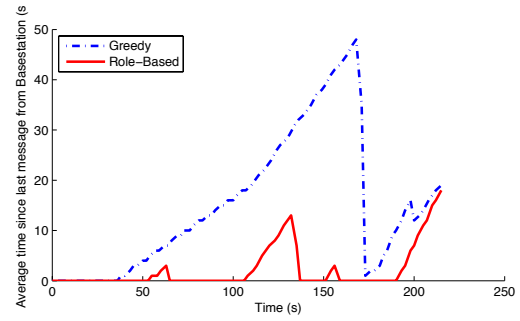

(c) Average time since last message from Basestation

Figure 4: Exploration by a team of 4 robots in Environment 1

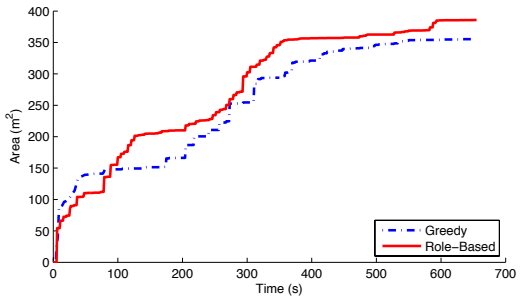

(a) Total free space discovered

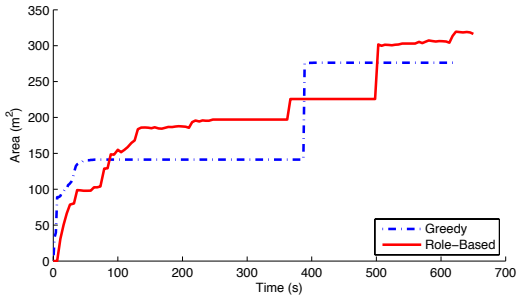

(b) Free space known at Basestation

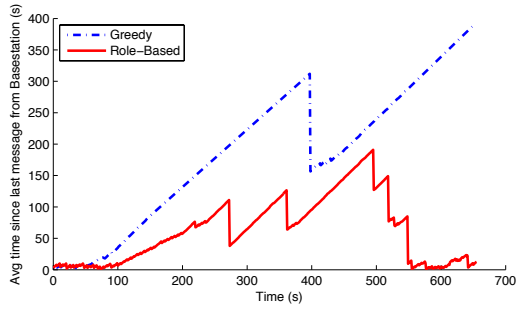

(c) Average time since last message from Basestation

Figure 5: Exploration by a team of 2 robots in Environment 2

improvements when a robot wanders back within range. In Role-Based exploration, responsiveness slowly deteriorates as the robots explore deeper and deeper within the environment, but is kept within manageable levels due to the relay's regular ferrying of messages back and forth between basestation and explorer.

Fig 6 presents screenshots demonstrating the relative performance of Greedy and Role-Based exploration for a set of runs involving 2 robots exploring Env. 2.

5.3 Comparison, Advantages, Points of Failure

As expected, Greedy exploration performs better in the very early stages of exploration, while Role-Based exploration performs better in the later stages. In Greedy exploration, there are simply more robots exploring, so for environments that are small or unlikely to contain communication difficulties, the Greedy approach makes more sense.

However, as environments become larger relative to the number of robots, and as communication becomes more of an issue, the advantages of a Role-Based approach quickly become clear. Updates are brought to the basestation much more frequently, and the team as a whole is considerably more responsive. Even total exploration proceeds as fast (or faster) than Greedy exploration since teammates share information better and are less likely to explore previously visited locations. This is true even for larger teams, where new information is fed up one branch of the team hierarchy tree and down another, and we have seen this as a consistent result for teams of variable size, in both simulation and in reality. The Role-Based approach is immune to communication limitations; it responds in an *adaptive* manner to com-

munication availability.

There are two further potential advantages to the Role-Based approach: First, several authors have suggested that repeated mutual observation by robots can lead to improved localisation (*e.g.* in [Fox *et al.*, 2006] robots actively seek one another to verify positions). Second, the Role-Based approach could be more suitable for heterogeneous teams: robots with more sophisticated sensor loads could be used for exploration, while smaller and faster robots with lighter loads could be used for information relay. Aerial relays could be a very interesting extension to the approach, for example.

It must be noted that both approaches are susceptible to robot failure (*e.g.* motor failure, laser failure, software crashes, etc.). In Greedy exploration this does not significantly affect the team, as all functioning teammates simply continue exploring. In Role-Based exploration robot failure must be handled with greater care as teammates depend on one another's actions. If a relay fails an explorer will end up stationary, waiting for it at the rendezvous location (and vice versa). Furthermore in dynamic environments, rendezvous points may become inaccessible. There are solutions to such problems: for example, we have introduced a time-out on rendezvous. If an explorer waits too long for a relay, it returns to the basestation itself; if a relay waits too long for an explorer, it gives up and becomes an explorer itself, returning to the basestation periodically.

6 Conclusions and Future Work

6.1 Conclusions

We have presented two algorithms for exploration of unknown environments. In Greedy exploration, an approach commonly used by multi-robot teams today, robots opportunistically seek to expand their knowledge of the world and coordinate with teammates when possible, but there is no effort to relay information to a basestation. In Role-Based exploration, the team conforms to a hierarchy with robots exploring the far reaches of the environment and relays acting as mobile messengers, ferrying information back and forth between basestation and explorers.

For applications where frequent updates at a central location are desirable (again we cite search-and-rescue as an example), the use of such mobile relays has significant advantages. Updates are received more frequently, the team is more responsive to central control commands, and the team exploration effort is coordinated better than in the Greedy approach, which can lead to overlap and repeated exploration of previously visited locations. In addition, the Role-Based approach responds adaptively to communication availability.

These results have been verified in simulation and on a team of real robots, and we have seen these results consistently on teams of 2 or 4 robots, and in different types of environments.

6.2 Future Work

We have examined scenarios involving failures of robots in simulation, and we have previously demonstrated in simulation the advantages of a dynamic hierarchy, *i.e.* robots swapping roles to explore more efficiently [de Hoog *et al.*, 2010a]. We intend to implement and examine these behaviours on the real robots.

Experiments presented in this paper were conducted in a controlled environment. This was useful for demonstrating relative advantages and disadvantages of exploration algorithms, but we will move into a fully realistic domain with our next set of experiments. These will involve no artificial constraints on communication; the environment we expect to use will contain very real communication challenges.

References

[Arkin and Diaz, 2002] R.C. Arkin and J. Diaz. Line-of-sight constrained exploration for reactive multiagent robotic teams. In *7th International Workshop on Advanced Motion Control*, 2002.

[Burgard *et al.*, 2005] W. Burgard, M. Moors, C. Stachniss, and F. Schneider. Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3):376 – 378, 2005.

[de Hoog *et al.*, 2009] Julian de Hoog, Stephen Cameron, and Arnoud Visser. Role-based autonomous multi-robot exploration. In *International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE)*, November 2009.

[de Hoog *et al.*, 2010a] Julian de Hoog, Stephen Cameron, and Arnoud Visser. Dynamic team hierarchies in communication-limited multi-robot exploration. In *International Workshop of Safety, Security and Rescue Robotics (SSRR)*, July 2010.

[de Hoog *et al.*, 2010b] Julian de Hoog, Stephen Cameron, and Arnoud 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.

[Dias and Stentz, 2000] M Bernardine Dias and Tony Stentz. A free market architecture for distributed control of a multirobot system. In *6th International Conference on Intelligent Autonomous Systems (IAS-6)*, pages 115–122, July 2000.

[Esposito and Dunbar, 2006] Joel M. Esposito and Thomas W. Dunbar. Maintaining wireless connectivity constraints for swarms in the presence of obstacles. In *International Conference on Robotics and Automation (ICRA)*, pages 946–951, 2006.

[Fox *et al.*, 2006] 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.

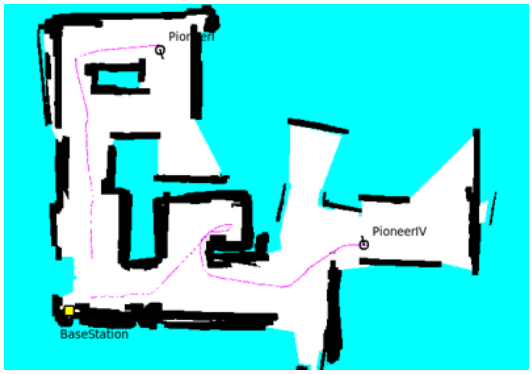

(a) Greedy Exploration after 226 seconds

(b) Role-based Exploration after 226 seconds

(c) Greedy Exploration after 605 seconds

(d) Role-based Exploration after 605 seconds

Figure 6: Screenshots of both the Greedy and Role-based efforts, using two robots in Environment 2. After 226 seconds the Greedy robots have explored more area, but are in different parts of the environment and out of range; the Role-Based teammates have explored less but maintained a chain of communication to the basestation. After 605 seconds both approaches have almost completely explored the environment, but in the Role-Based approach updates have been more frequent.

[Gerkey *et al.*, 2001] B.P. Gerkey, R.T. Vaughan, K. Stoy, A. Howard, G.S. Sukhatme, and M.J. Mataric. Most valuable player: a robot device server for distributed control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 3, pages 1226 –1231 vol.3, 2001.

[Hilditch, 1969] C.J. Hilditch. Linear skeletons from square cupboards. *Machine Intelligence*, 4:403–420, 1969.

[Howard *et al.*, 2002] A. Howard, M.J. Mataric, and G.S. Sukhatme. An incremental deployment algorithm for mobile robot teams. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.

[Jiménez-González *et al.*, 2010] Adrian Jiménez-González, J. R. Martínez-de Dios, and Anibal Ollero. An integrated testbed for heterogeneous mobile robots and other cooperating objects. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, October 2010.

[Nguyen *et al.*, 2004] Hoa G. Nguyen, Narek Pezeshkian, Anoop Gupta, and Nathan Farrington. Maintaining communication link for a robot operating in a hazardous environment. In *ANS 10th Int. Conf. on Robotics and Remote Systems for Hazardous Environments*, pages 28–31, 2004.

[Poduri and Sukhatme, 2004] Sameera Poduri and Gaurav S. Sukhatme. Constrained coverage for mobile sensor networks. In *International Conference on Robotics and Automation (ICRA)*, 2004.

[Rooker and Birk, 2007] Martijn N. Rooker and Andreas Birk. Multi-robot exploration under the constraints of wireless networking. *Control Engineering Practice*, 15(4):435–445, 2007.

[Simmons *et al.*, 2000] Reid G. Simmons, David Apfelbaum, Wolfram Burgard, Dieter Fox, Mark Moors, Sebastian Thrun, and Håkan L. S. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, pages 852–858. AAAI Press / The MIT Press, 2000.

[Stump *et al.*, 2008] E. Stump, A. Jadbabaie, and V. Kumar. Connectivity management in mobile robot

teams. In *International Conference on Robotics and Automation (ICRA)*, pages 1525 –1530, May 2008.

[Visser and Slamet, 2008] Arnoud Visser and Bayu A. Slamet. Balancing the Information Gain Against the Movement Cost for Multi-robot Frontier Exploration. In *European Robotics Symposium 2008*, Springer Tracts in Advanced Robotics, pages 43–52. Springer-Verlag, February 2008.

[Wurm *et al.*, 2008] K.M. Wurm, C. Stachniss, and W. Burgard. Coordinated multi-robot exploration using a segmentation of the environment. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1160–1165, Sept. 2008.

[Yamauchi *et al.*, 1998] B. Yamauchi, A. Schultz, and W. Adams. Mobile robot exploration and map-building with continuous localization. In *International Conference on Robotics and Automation (ICRA)*, 1998.

[Zlot *et al.*, 2002] R. Zlot, A. Stentz, M.B. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. In *International Conference on Robotics and Automation (ICRA)*, 2002.