Distributed Multirobot Exploration Maintaining a Mobile Network

Jose Vazquez, Chris Malcolm
School of Informatics
University of Edinburgh
Mayfield Road
Edinburgh EH9 3JZ, UK
J.M.Vazquez-Diosdado@sms.ed.ac.uk,Chris.Malcolm@ed.ac.uk

Abstract

We present a behaviour-based architecture for multirobot exploration and mapping. The architecture is designed to guide the exploration in a decentralized fashion constrained to maintain local short-range communication in a mobile adhoc network. Exploration with multiple robots has been extensively studied but communication constraints have been largely ignored. We focus on the integration of such constraints and develop a behaviour-based approach. The behaviours are designed to enhance global performance and are triggered based on local information. The robots are encouraged to move together forming a mobile network and sharing relevant information for the team. To this date there is no research addressing the effects of communication and sensor constraints on decentralized co-operative exploration and mapping in structured (office like) environments. Our experimental results show that the approach performs better than previous decentralized approaches.

1. INTRODUCTION

The problem of exploration of an unknown environment has been extensively studied, firstly using single robot systems with a variety of sensors and later using teams of robots (Shatkay and Kaelbling, 1997),(Yamauchi et al., 1998). The first implementations of multirobot exploration systems were simple extensions of the single robot implementations (Yamauchi, 1998). We use Yamauchi's approach as a basis for comparative evaluation of ours.

Multiple robot systems are more complex than other distributed systems because they have to deal with a real environment, which is more difficult to model since it is dynamic, unpredictable, noisy, etc.

Multirobot exploration systems are usually classified as centralized and decentralized. Centralized systems obtain solutions close to the optimal but are computationally intensive and have a single point of failure. On the other hand, decentralized systems are flexible and robust, but frequently achieve significantly suboptimal solutions.

Centralized exploration strategies implement the concept of dispersion: to explore an unknown environment efficiently robots should stay apart in order to minimize interference (Simmons et al., 2000). This concept by itself may be suitable for small teams when inter-robot communication is available but it's not scalable because as the number of robots in a team increases the communication bandwidth becomes a bottleneck for the system.

On the other hand, in the last three years there has been a growing interest in developing distributed sensing systems using local wireless networks with a local range of communication. The main characteristics of these devices are low power consumption, size, and cost which make them particularly suitable for larger numbers of cheap robots.

The research in this field has focused mainly on static arrangements of sensors to obtain maximal coverage (Howard et al., 2002b), the analysis of the delays on communications systems (Pereira et al., 2002), and the development of efficient techniques to route information.

The integration of these new technologies in multirobot tasks is a recent research area. Remarkable examples of robot platforms using such an approach are Robomote (Sibley et al., 2002) and Millibots (Navarro et al., 2002) which have the characteristics of low power consumption, basic sensing capabilities (infrared or sonars) and short range communication. Robots exploring an unknown environment under such constraints must work cooperatively in order to avoid overlapping and collisions. Using local short range communication has the advantage for small robots of low power requirements. While this can be exploited in a decentralized approach to reduce communication traffic and bottlenecks as the number of robot increases, it can't be used in centralized approaches precisely because the range is too short.

To meet these requirements, this paper introduces a behaviour based architecture that is designed to encourage a team of robots to explore an unknown space while maintaining a local communication network and being controlled in a decentralized fashion. We present experimental results conducted in simulated structured office like environments comparing the proposed architecture with Yamauchi's approach (Yamauchi, 1998). Our results show that the proposed approach performs better than previous approaches; moreover the approach performs incrementally better with larger teams of robots.

The present architecture implements the set of behaviours necessary to maintain robot network connection. Two strong assumptions are made, each robot has a global localization system ($\pm 5\%$ error in experiments) and there are no radio-opaque obstacles in the environment, a less important assumption is that information is shared among all the robot members only being affected in terms of the delay-time introduced by paths through the ad hoc network they form. These are only simplyfing assumptions for assesing the general promise of the approach. Future work will relax these assumptions.

2. RELATED WORK

2.1 Multirobot Exploration

In previous work using decentralized approaches there was no coordination between robots, instead the robots acted on their own and exchange their maps whenever they were within communication range (Shatkay and Kaelbling, 1997).

Yamauchi developed a technique in which robots build a common map (an occupancy grid) (Yamauchi et al., 1998). The work introduced the notion of frontier, which is a location near a unexplored part of the environment. Each robot heads for its closest frontier to acquire new information. The robots act independently without coordination. Yamauchi reports that robots tend to waste time exploring areas already mapped by other robots. Similar results had been obtained in research by Jennings (Shatkay and Kaelbling, 1997) using topological maps.

Simmons developed a coordination mechanism in multirobot exploration (Simmons et al., 2000). The approach is based on Yamauchi's concept of frontiers. The robots evaluate a set of frontiers and determine the travel costs and information gain, based on the estimated number of unknown map cells visible from the frontier. The robots then submit this information to a central entity which assigns an allocation to each robot based on their bids. The drawback is the fact that the system relies on a central entity and therefore the entire system will fail if the central agent fails.

We propose a decentralized approach in which robots are encouraged to maintain communication with at least one other robot. By maintaining communication with other robots they are able to coordinate by sharing their goals, minimizing the overlapping of tasks. The proposed approach has the fault tolerance and scalability of decentralized systems and avoids the drawback of highly suboptimal solutions by means of coordination behaviours.

2.2 Mobile sensor networks

Recent platforms such as Millibots (Navarro et al., 2002) and Robomote(Sibley et al., 2002) are seen as mobile sensor networks. These new platforms focus on the formation of cooperative large teams and are meant to be deployed in hazardous unknown scenarios.

Robomote is a test platform for mobile sensing based on communication and coordination in unplanned environments (Sibley et al., 2002). Based on this platform algorithms for incremental sensor deployment have been developed (Howard et al., 2002a). The issue of exploration has not been addressed yet, but it is one of the final goals of the project.

Research on the Millibots project has been oriented to cooperative localization using a beacon system (Navarro et al., 2002). Each robot is equipped with ultrasonic sensors. The robots perform trilateration through the use of distance measurements to other three Millibots. Millibots leap frog each other to maintain good position estimates as they transverse unknown terrain.

In Arkin's research algorithms have been developed to maintain a team of robots maintaining line of sight communications between robots while searching for a hazard with different degrees of a priori knowledge of the environment (Arkin and Diaz, 2002). Robots move one at a time while all the others maintain a position. The drawback of this approach is that since only one robot moves at a time having several robots to explore an area does not decreases significantly the time required to find the hazard. Based on this approach Balch developed an algorithm denominated VBCP (Value-Based Communication Preservation) for communication preservation in the context of other tasks (Powers and Balch, 2004).

Santos proposed an algorithm to explore unknown environments with small obstacles using multiple robots (Pimentel and Montenegro, 2002). In his approach robots have a limited range of communication and remain connected forming a mobile network. One of the robots remains static serving as a base for the rest of the robots. The exploration area is limited by the total communication range. In our case, we are interested in generating a map of the environment autonomously rather than maintaining a base station and explore the nearby areas.

Decentralized architectures have been implemented in these approaches due to their simplicity, scalability and fault tolerance. These approaches have been oriented to solve problems related to multirobot exploration such as cooperative localization systems and sensor coverage nevertheless a exploration system suitable for mobile sensor networks has not been proposed yet. Our approach aims to produce a suitable solution to the problem.

3. APPROACH

The behaviour based approach proposed is designed to reactively adapt to the dynamic conditions of the mobile network that the robots form while trying to explore the environment. Behaviours are designed to avoid collisions between the robots, encourage the exploration of unknown environments, and maintain the mobile robot network. The selection of behaviours is based on the current network conditions.

The exploration algorithm implemented is based on the multirobot architecture provided by Simmons (Simmons et al., 2000). The environment is represented by means of a global probabilistic grid map. The algorithm has been adapted to constraint the search space according to the current selected behaviour.

Robots broadcast their sensorial information, their estimated position, and their current heading (frontier selected) at regular time intervals. Robots receive this information from their neighbors directly and from the rest of the team by means of multicast. Based on this periodic information the robots are able to identify the topology of the communications network.

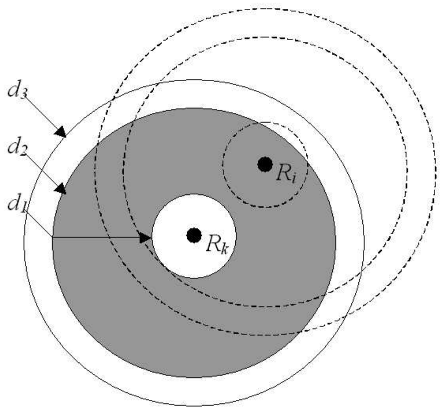

Figure 1: R_k induces a constraint in R_i The shadowed area defined as the comfort zone in the configuration space for R_i where collisions are avoided and connectivity is maintained.

3.1 Mobile network constraints

Robots forming a mobile network must constrain their motion to remain connected. A pair of robots impose mutual motion constraints only if they constitute a solitary bridge connection according to the current network topology and if the signal strength among the pair of robots falls below a threshold.

3.1.1 Bridge detection

Based on the topology of the network a robot its able to classify its connections with the rest of the robots by means of a graph. In the graph the robots are represented by vertices and the connections are represented by edges. From graph theory fundamentals (Weiss, 1999) it is known that an articulation is a vertex whose removal disconnects a graph and a bridge is an edge whose removal disconnects a graph. Using a bridge detection algorithm (Weiss, 1999) a robot is able to identify the bridge connections in a topology.

3.1.2 Constraint definition

Using measurements of communication signal strength the robots impose constraints on each others' movements in order to maintain the communication network. For these experiments we assume that there is no interference caused by the walls, and the signal strength decays by the square of distance.

In general, considering a generic robot R_k, that induces a constraint in R_i, given by $g(q_i, q_k) > d_3$ beyond which the communication between R_k and R_i, is broken. Then, q_i and q_k are the coordinates of the robots, and $g(q_i, q_k)$ is defined as $g(q_i, q_k) = (x_i - x_k)^2 + (y_i - y_k)^2$.

In addition to the above constraint and to avoid the loss of connections and collisions between robots the concept of comfort zone is introduced (Figure 1). The comfort zone is defined as $d_1 < g(q_i, q_k) < d_2$ where $d_1 < d_2 < d_3$ and is the area in which the robot can move safely. d1 is defined as a collision constraint and d2 is a preventive constraint. d_1 is designed to avoid possible collisions between a pair of robots and in our experimental set is equal to 3/2 of the diameter of the sensor range of the robot. d_2 is designed to be a precautionary constraint in order to avoid the total loss of communication between a pair of robots (in the case of a bridge) and for the experimental set is equal to $d_2 = 0.9d_3$.

Due to the dynamic nature of the network if the robots move out the comfort zone their behaviour changes in order to return. If two robots form a bridge connection in the network they will tend to explore new areas in the comfort one; otherwise their exploration is not constrained.

3.2 Exploration Algorithm

The exploration algorithm is based on the implementation of Simmons idea (Simmons et al., 2000)which introduces the concept of information gain. This section presents an overview of the algorithm and its adaptation to decentralized robots. For a detailed explanation consult (Simmons et al., 2000).

The map generated is a probabilistic grid map where a cell with a probability $p(x, y) > p_o$ is an obstacle, a cell with a $p_o > p(x, y) > p_f$ is an unknown cell and a cell with a probability $p(x, y) < p_f$ is free space. p_o and p_f are the thresholds for occupancy and free space respectively.

A Frontier cell is defined as a free space cell adjacent to at least one unknown cell. Each robot generates a list of possible destinations from various frontier cells.

Frontier cells are evaluated according to the estimated costs and the utility of the information. To estimate the cost of visiting a frontier cell, we compute the optimal path (assuming deterministic motion) from the robot's current position.

The costs are computed simultaneously, using a simple flood-fill algorithm (Latombe, 1991) to propagate minimum path costs through the map. The utility refers to the nearby unexplored area and is computed by counting the number of unexplored cells that lie inside the circumference of the sensor range of the robot. The utility expected by a robot considering moving to a particular frontier is lessened if there are any robots near that destination. In our approach each robot executes the algorithm in a decentralized fashion, while in Simmons approach each robot makes a bid for each frontier, the bid being based on information gain. A robot is assigned a frontier based on the best bid and the whole process repeated for unassigned robots.

In some occasions robots may abort the plan to explore a frontier due to constraints, in such case the frontier may be selected by another robot. Once a frontier is selected based on the map of costs generated by the flood-fill algorithm the path from the selected frontier to the robot's position is computed by iteratively finding the neighboring grid cell with minimal cost (Alexander and Yuta, 1993).

3.3 Selection of behaviours

Robot behaviours are selected according to the current conditions in the network topology. In Figure 2 its presented the schema to select the behaviours and their transitions.

The achieve-connectivity behaviour is triggered when a robot detects that the current connection with a robot is a bridge connection and such connection is bigger than d_2. Under the achieve connectivity behaviour the pair of robots that form the bridge connection search for frontier cells inside the comfort zone (if any) or move towards each other in order to return the comfort zone.

The maintain-connectivity behaviour is triggered when a robot is in the comfort zone with respect to their bridge connections (if any) or when the robot does not

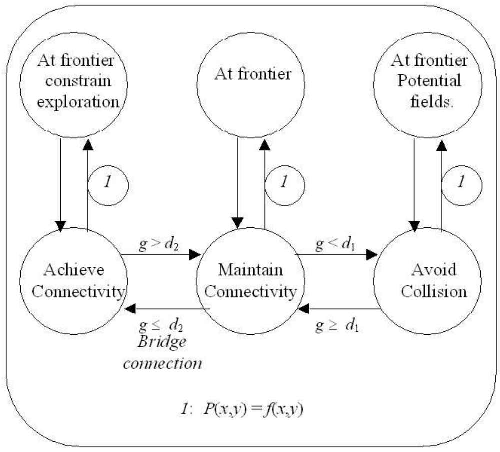

Figure 2: Schematic switching of behaviours. When a change in behaviour is detected the robot will reevaluate its current plan. If the robot arrives at the current frontier the exploration algorithm its executed according to the current network conditions.

have bridge connections and its distance to each one of his neighbors is bigger than the collision distance.

The avoid-collision behaviour is triggered when the robot detects a robot closer than the collision distance. In such case the robot recalculates its trajectory applying potential fields to avoid a possible collision.

When $P(x, y) = f(x, y)$ where p is the robot position and f is the current frontier the robot had arrived to the frontier. In this case the exploration algorithm its executed according to the constraints imposed by the current behaviour. As indicated in Figure 2 the exploration algorithm could constraint the search area or incorporate potential fields when finding a new frontier, otherwise the algorithm is executed as explained in Section 3.2. In the next subsections we provide details about the implementation of achieve-connectivity and avoid-collision behaviours.

3.3.1 Achieve-connectivity

When the connectivity behaviour is triggered the pair of robots that are out of the comfort zone in a bridge connection are encouraged to return. To achieve the goal the search of frontier cells is constrained to the comfort zone between the pair of robots. When no frontier cells are found the robot computes a path to the current position of the other robot in the bridge connection. Once the robot achieves the connection status it constrains future plans in order to avoid being caught in the same situation.

In Figure 3 an example of the execution of the behaviour is shown. The allowed search space for r_0 lies inside the big circle (blue) since r_0 and r_4 are mutually

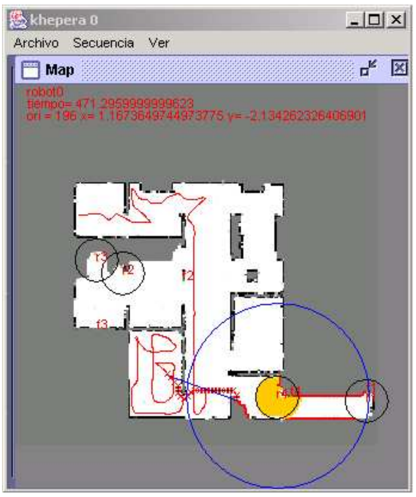

Figure 3: r_0 executing the achieve connectivity behaviour to return to the comfort zone with respect to r_4. The exploration area for frontier cells is constrained to the area that lies inside the comfort zone with respect to r_4.

constrained. As result of the search r_0 selected f_0.

3.3.2 Avoid-collision

The avoid-collision behaviour causes a re-evaluation of paths to the current goal to avoid collisions. This is reactive and avoids the necessity of deliberative plans that are expensive in terms of communication and computer resources due to searching in high dimensional spaces.

The current implementation is based on the potential fields implementation developed by Warren (Warren, 1990). In this implementation robots that are closer than the collision distance are considered as static obstacles in the planning space. In future implementations we will improve this by avoiding potential deadlocks. The advantage of the use of potential fields is that they offer a relatively fast and efficient way to solve for safe trajectories around stationary and moving obstacles. Although the current implementation is simple it has proven to be efficient in practice to avoid collisions between the robots.

Figure 4 shows an example of the implementation where r_1 is moving to f_1, the route generated has been modified to avoid a potential collision with r_0. As can be observed from the figure the original path is much shorter than the avoidance path.

4. EXPERIMENTS

The objective of the experiments is to asses and compare the performance of the proposed architecture with Yamauchi's approach (Yamauchi, 1998). The comparision is based on the time required to generate a complete

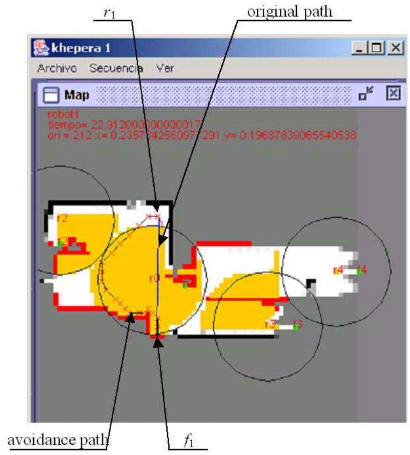

Figure 4: avoid collision behaviour for r_1 moving to f_1 to avoiding r_0.

map of the environment.

The experiments have been conducted in the Webots simulator (Michel, 1998). The simulated robots are khepera. In Figure 5 the experimental environment is presented. The environment as has been divided in different places to resemble different sizes.

In the current set of experiments the robots start off relative positions to a known location. The position of each robot its assumed to be known with an error margin of $\pm 5\%$. Additionally the robots know the topology of the network and each broadcasts its sensorial information every half second. The information arrives by multi hops to all the members of the team.

The robots have a diameter of 5 cm and their sensor range diameter is 45cm. The communication ranges compared were 1m, 1.5 m and 2m. The team sizes range from 2 to 10 robots. The total space to explore in the environments were $10.2m^2$ and $23.5m^2$.

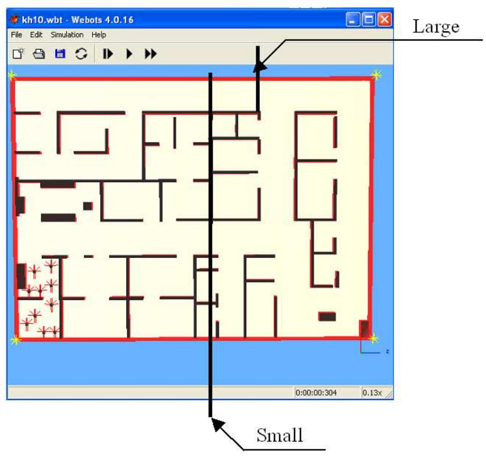

Figure 5: Structured environment used in the experiments.

4.1 Results

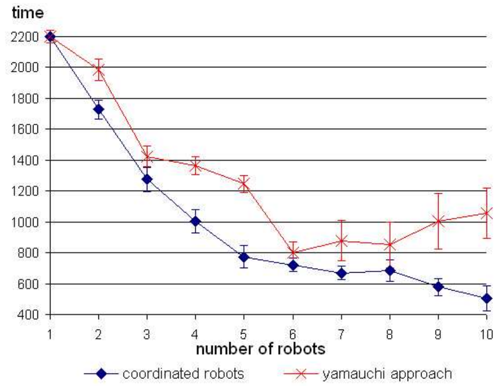

Figure 6: Small size environment with 1 meter communication range

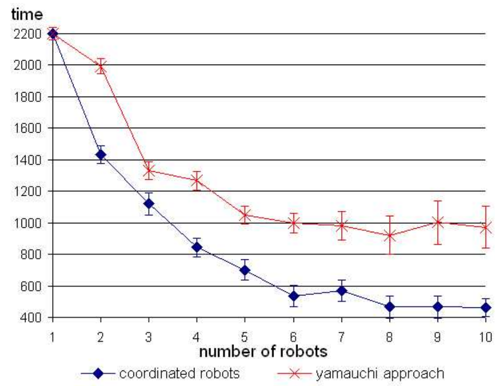

Figure 7: Small size environment with 1.5 meter communication range

In Figures 6-11 we present the results obtained in the two different size environments and three different ranges of communications. For each team of robots ten trials were run and averaged. The error bars represent the variance in the results for each particular configuration.

It is observed that the time of completition drops in a linear fashion with respect to the number of robots up to some maximum number of robots. Similar results in time of completion vs. number of robots have been found in Simmons (Simmons et al., 2000) and Maio (Maio and Rizzi, 1995). Others have found the same thing in various implementations (Howard et al., 2002a), and have argued that there is a size of the environment that can only afford up to a certain number of robots. Up to the threshold number of robots there is a linear performance improvement. After the threshold there are diminishing returns for more robots. This characteristic trend is confirmed in the graphs.

The robots maintaining the network perfom better than the robots in Yamauchi's approach, moreover in the case of the small environment when the number of robots is above seven robots in Yamauchis approach the robots perform worse. This behaviour is attributed to the lack of coordination in this approach where robots end up interfering with each other most of the time.

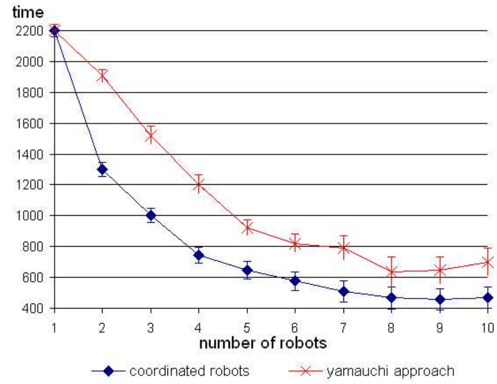

Figure 8: Small size environment with 2 meter communication range

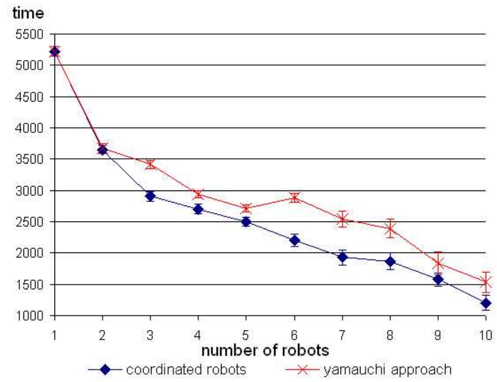

Figure 9: Large size environment with 1 meter communication range

As reported in Yamauchis research robots under his approach tend to move to previously explored areas constantly, this is more critical as the number of robot increases in small environments; this situation is reflected in the varianze of the results which is significantly higher than the varianze in the proposed approach, this is more notorious on Figure 7 when there are larger number of robots with the shortest communication range.

In Figure 12 we present a comparison of the proposed approach under different ranges of communication, it is observed that the effect of the communication range presents an acceptable trade-off with respect to the time of completion. Moreover as the number of robots increases the difference is diminished.

For the purpose of our research is particularly important to identify the trade-offs introduced by the commu-

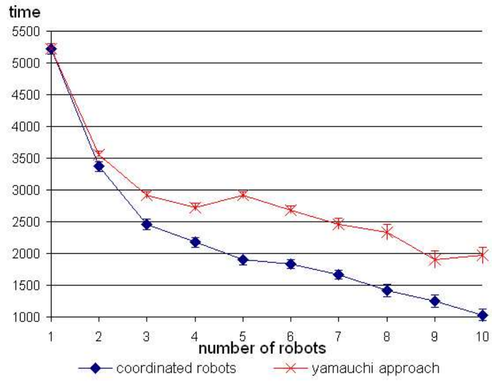

Figure 10: Large size environment with 1.5 meter communication range

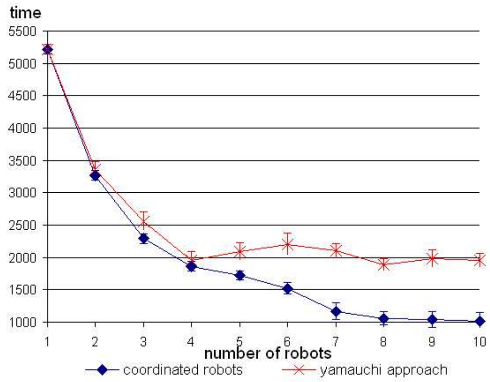

Figure 11: Large size environment with 2 meter communication range

nication constraints with respect to the time of completion of the exploration. The importance is due to the fact that in indoor environments the signal decay tends to be proportional to the cube of the distance rather than the square. This implies that even if the robots are equiped with powerful transmitters the consumption of energy will increase considerably. Moreover a smaller communication range gives benefits when the number of robots is large due to the decrement in interference. Similar ideas have been applied successfully in mobile technologies succesfully where the devices are able to moderate their power to minimize interference and energy consumption.

We argue that the constraint introduced by lessening the range of communication is a reasonable tradeoff since robots are able to maintain a similar performance. Nevertheless, we have to conduct more exhaustive experiments to prove this argument and compare the results with a non-cooperative multirobot mapping algorithm such as Yamauchis' (Yamauchi et al., 1998).

Although the environments used in this initial experiments resemble office like environments the inclusion of

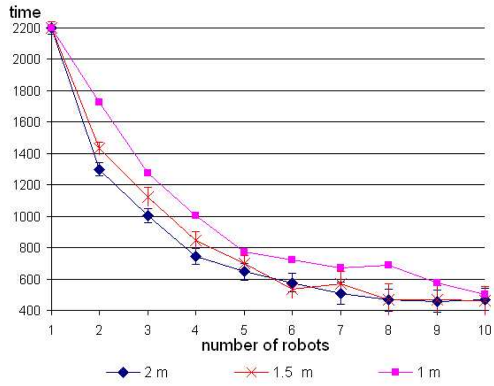

Figure 12: Comparison of the proposed approach under different ranges of communication

experiments in cluttered areas its one the future areas of interest.

5. FUTURE WORK

This paper gives the first experimental results of a PhD project investigating cooperative decentralized multi-robot mapping of structured (office-like) environments. In the present experiments two strong assumptions have been made, non accumulative error in the position of the robots and interference-free communication. Future work will relax these assumptions.

Additionally, we plan to conduct experiments to analyze the communication cost of maintaining the robot-network because research in the area of mobile networks has determined that the communication concerned with the administration and updating of the network connectivity takes about 50% of the bandwidth.

In the current implementation every robot has global knowledge of the map and the topology. This gives problems of scalability. To reduce these we plan to conduct experiments analyzing the tradeoffs of partial topology knowledge as well as the implementation of a hierarchical map representation to reduce communication overheads. Distant robots would communicate compressed map information occasionally, whereas local robots would communicate rawer local map details more frequently.

The use of grid maps to represent the environment has as disadvantage the large space required to store them. We plan to address this issue by means of a hierarchical representation of the environment.

Currently, we are developing a small team of cheap robots to implement the algorithms and identify problems omitted in the simulations.

6. CONCLUSIONS

Although multirobot exploration tasks have been relatively well studied integrating the constraints of these tasks into mobile communication networks is a relatively unexplored new area of research.

We have presented a behaviour based approach that takes into account the restrictions (and exploits the advantages) of these new short range communication technologies. The behaviours are designed to enhance global performance and are triggered based on local information.

Our main concern is the scalability to large numbers of robots and the restrictions imposed by the range of communication. The objective of the experiments is to assess and compare the proposed approach under different constraints of communication and different team sizes. Results so far encourage us pursue this research, and suggest that the constraints introduced by restricting communication range do not significantly harm performance. Moreover the proposed approach performs better than previous decentralized approaches.

For the purposes of our research it is of vital importance to evaluate the effect of restricting communication range in the architecture to identify the trade-offs between this and other variables in the system.

In decentralized systems robustness is a natural consequence. Scalability requires economy of communication bandwidth. In future research we plan to address this issue by using strictly local low power communication, combined with hierarchical map representations.

References

Alexander, Z. and Yuta, S. (1993). A unified approach to planning, sensing and navigation for mobile robots. In *third international symposium on experimental robotics*, kyoto, japan.

Howard, A., Mataric, M., and Sukhatme, G. (2002a). An incremental deployment algorithm for mobile robot teams. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2002)*, volume 3, pages 2849–2854.

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

Latombe, J.-C. (1991). *Robot Motion Planning*. Kluwer Academic Press.

Maio, D. and Rizzi, S. (1995). Unsupervised multi-agent exploration of structured environments. In Lesser, V., (Ed.), *Proceedings of the First International Conference on Multi–Agent Systems*, pages 269–275, San Francisco, CA. MIT Press.

Michel, O. (1998). Webots: a powerful realistic mobile robots simulator. In *Second International Workshop on RoboCup*.

Navarro, L., Grabowski, R., Paredis, C., and Khosla, P. (2002). Millibots. *IEEE Robotics and Automation Magazine*, pages 31 – 40.

Pereira, G., Pimentel, B., Chaimowicz, L., and Campos, M. (2002). Coordination of multiple mobile robots in an object carrying task using implicit communication. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA'02)*, pages 281–286, Washington, DC.

Pimentel, S. and Montenegro, M. (2002). Multi-robot exploration with limited-range communication. In *congreso brasileiro de automatica*.

Powers, M. and Balch, T. (2004). Value-based communication preservation for mobile robots. In *International Conference on Robotics and Automation*.

Arkin, R. and Diaz, J. (2002). Line-of-sight constrained exploration for reactive multiagent robotic teams. In *AMC'02 7th International Workshop on Advanced Motion Control*.

Shatkay, H. and Kaelbling, L. P. (1997). Learning topological maps with weak local odometric information. In *IJCAI (2)*, pages 920–929.

Sibley, G., Rahimi, M., and Sukhatme, G. (2002). Robomote: a tiny mobile robot platform for large-scale ad-hoc sensor networks. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1143 –1148.

Simmons, R., Apfelbaum, D., Burgard, W., Fox, D., Moors, M.and Thrun, S., and H., Y. (2000). Coordination for multi-robot exploration and mapping. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*.

Warren, C. W. (1990). Multiple robot path coordination using artificial potential fields. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 500–505.

Weiss, M. A. (1999). *Data structures and algorithm analysis in Java*. Harlow.

Yamauchi, B. (1998). Frontier-based exploration using multiple robots. In *Proceedings of the second international conference on Autonomous agents*, pages 47–53. ACM Press.

Yamauchi, B., Schultz, A. C., and Adams, W. (1998). Mobile robot exploration and map-building with continuous localization. In *1998 IEEE International Conference on Robotics and Automation*, pages 3715–3720.