Combining Multi-Robot Exploration and Rendezvous

Malika Meghjani and Gregory Dudek
Center for Intelligent Machines, McGill University,
3480, University Street, Montreal, H3A 2A7.
{malika,dudek}@cim.mcgill.ca

***Abstract*—We consider the problem of exploring an unknown environment with a pair of mobile robots. The goal is to make the robots meet (or rendezvous) in minimum time such that there is a maximum speed gain of the exploration task. The key challenge in achieving this goal is to rendezvous with the least possible dependency on communication. This single constraint involves several sub-problems: finding unique potential rendezvous locations in the environment, ranking these locations based on their uniqueness and synchronizing with the other robot to meet at one of the locations at a scheduled time. In addition, these tasks are to be performed simultaneously while exploring and mapping the environment. We propose an approach for efficiently combining the exploration and rendezvous tasks by considering the cost of reaching a rendezvous location and the reward of its uniqueness. This cost and reward model is combined with a set of deterministic and probabilistic rendezvous strategies for the robots to meet during exploration. Experimental results suggest that the joint tasks of exploration and rendezvous are substantially improved by ranking the potential rendezvous locations based on the combined cost-reward criterion when compared to the ranking solely based on the uniqueness of the location.**

I. Introduction

In this paper we consider the problem of multi-robot *rendezvous*, that is, the problem of getting two robots to find and meet each other using local sensing information only. This problem is significant in applications where robots must meet to exchange information or material, for example multi-robot exploration or "bucket brigade" material transport. Our work considers a particular case where the robots are not equipped with long-range communication, and thus must meet by physically visiting some kind of potential meeting point. The analogy in everyday life is to try and meet a friend at the top of a hill or some place from a set of favorite "hang outs". In practice, such "no communication" rendezvous can arise underwater, in energy-minimal planetary exploration, or in radio-silence environments.

Our work is motivated by [1] which explicitly studies the rendezvous problem for two exploring robots. The authors propose three classes of rendezvous strategies, including deterministic, probabilistic and random, which are validated based on analytical and experimental results. A rendezvous strategy provides an order for visiting a set of ranked locations in the environment where the rendezvous is most likely to occur. These potential rendezvous locations are referred to *landmarks* and are selected based on their respective *distinctiveness measure*. The key aspect of any rendezvous strategy is to identify an appropriate distinctiveness measure and rank the landmarks accordingly. In this paper, we consider a laser range sensor to gauge the distinctiveness measure however any local appearance sensor can be used for this purpose.

In contrast to prior work, we propose a ranking criterion which considers the distinctiveness and the accessibility of the potential rendezvous locations. This combination accounts for both the rendezvous and the exploration tasks respectively. The distinctiveness measure signifies the relative importance of rendezvous locations which, in our work, is measured based on the area of coverage of the robot's sensor at that point. The accessibility of the rendezvous location is measured based on the distance from the current location to the corresponding rendezvous point. We perform a set of experiments based on different rendezvous strategies to evaluate the proposed ranking criterion. The evaluation is based on the comparison of the rendezvous time using the aforementioned ranking criterion against the pure distinctiveness-based ranking. The cost-reward based ranking criterion reduces the exploration time as presented in Section V.

II. Related Work

The multi-robot exploration task is a well-addressed problem in the field of robotics. Existing solutions to this problem differ primarily by the type of coordination that exists between the robots. The level of coordination depends on the type of communication the robots are expected to share. A majority of multi-robot systems [2], [3], [4] assume the robots to be in constant communication with each other or with a central control unit to help them explore the environment efficiently.

Some of these coordination methods are derived from the field of optimization where each target location is associated with a cost and a reward function which helps a team of robots to distribute the exploration task or get assigned for the task by a central agent [4]. For instance, the target locations generally belong to the frontier regions, which form the boundary between the free-space and the unexplored area [3]. The cost associated with the target is the distance of a particular robot to the target location, and reward is the estimated amount of information gain. The final decision of assigning the exploration task to a particular robot depends on whether the robot would remain within the communication range at the target location [5]. A variant of this task allocation method is implemented based on market principles where the target locations are auctioned and the robots bid for these locations based on their cost and reward values [6], [7].

The aforementioned approaches restrict robot exploration within the maximum communication range. One of the approaches to overcome this limitation is to form small groups of

978-0-7695-4362-8/11 $26.00 © 2011 IEEE
DOI 10.1109/CRV.2011.18

robots beginning from different locations in an unknown environment to autonomously explore the environment without any communication constrains [8]. The teams of robots coordinate with each other by using shared maps and providing updates about relative locations to efficiently perform the exploration task. Another interesting approach for short-range autonomous exploration is to use a team of hierarchical robots [9]. Exploration is performed using a role-based approach where the explorer robots split intermittently from the relay robots for exploring new regions while the relay robots convey the information to the command center. The relay robots merge again with the explorers by estimating their current location with the help of a topological map. This form of robot coordination without any communication, which purely depends on the environmental cues to meet with the other teammates, is similar to the problem analyzed in this paper.

There has been relatively less attention given to the problem where the robots are initially distributed in the environment without any communication until a line-of-sight communication is established at a pre-arranged time. This is the topic of concern in our work, which demonstrates a realistic search and rescue application with a requirement of exploring remote regions of interest. This particular problem of robots meeting in an unknown environment without any communication is referred to as the *rendezvous* problem and was first introduced by Dudek et al. in [1] for multi-robot exploration. The rendezvous problem originates from the field of game theory as a search problem and is well-known for its application in the mobile hider game, *princess and monster* [10]. Notably this "no communication" context is relevant to many practical scenarios including surveillance and energy-minimal rugged-terrain exploration.

The abstract rendezvous problem has been addressed by many authors such as in [11] and [12] for different contexts where the goal is to gather multiple agents with limited visibility at a common meeting point. The present work focuses on the rendezvous problem in the context of multi-robot exploration task which is discussed in detail in the following sections.

III. Proposed Approach

Our current work analyzes the rendezvous problem for a pair of robots placed randomly at separate locations in an unknown environment with no communication. The rendezvous problem can be defined in several ways. The instance we are concerned with, what we will refer to as the "topological rendezvous" problem where the mobile agents in the the environment can only meet one another at a specific set of locations which are implicit in the environment (although their location may also be a function of a particular sensor as well). In this paper the definition of such locations is extended slightly to include a finite region, rather than a single point.

Each robot is required to individually explore and create a map of the environment to find potential rendezvous locations or finite regions which they can visit during the rendezvous time in an attempt to find and meet the other robot. The exploration and rendezvous are performed intermittently until the robots meet or the exploration is completed. If the robots meet then a joint exploration of the robot pair can be more efficiently scheduled. We propose and implement a combination of different ranking criteria and rendezvous strategies, each of which are discussed in detail in the following sub-sections.

A. Mapping

The robots maintain individual maps of the environment using the grid-based approach [13]. An example of a partial map created during exploration is presented in Figure 1.

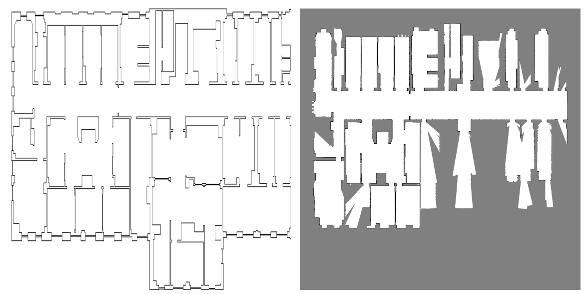

(a) (b)

Fig. 1. A partial grid-based map 1(b) is created during exploration. The original floor–plan is shown in Figure 1(a).

B. Exploration

The grid-based map is transformed into a skeleton structure of the free-space for exploration and rendezvous. The process of obtaining a skeleton by reducing the shape of the free-space to a fully-connected and centered structure is known as thinning. In our work, we implemented Hilditch's algorithm [14][1] to obtain the skeleton structure as illustrated in Figure 2.

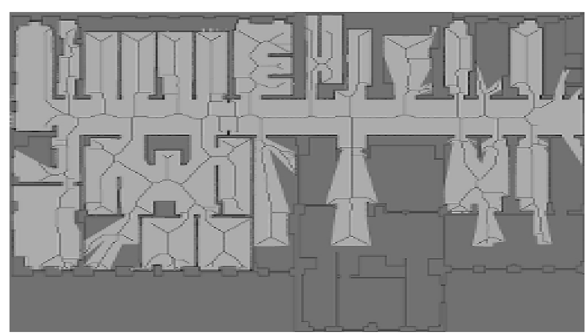

Fig. 2. The skeleton structure overlaid on the grid-based map 1(b) and the floor-plan 1(a)

This structure has several useful properties such as symmetric and compact representation of the free-space and fast calculations of the locations to be explored. It is similar in structure to a Generalized Voronoi Graph (GVG) [15]. However, they differ in their respective constructions. The

[1]For a useful tutorial on Hilditch's algorithm, refer: http://cgm.cs.mcgill.ca/ godfried/teaching/projects97/azar/skeleton.html

GVG is constructed by explicitly visiting the edge ends in order to complete the graph. On the other hand, a skeleton structure of the free-space of a maximum radius equivalent to the sensor range of the robot can be obtained by taking a scan from the current location and then planning if a particular node is worth visiting for exploration or rendezvous. This information is crucial as our primary objective is to minimize the combined exploration and the rendezvous time.

The skeleton structure is finally stored in form of an undirected graph data-structure such that the exploration task is reduced to that of traversing through the nodes of the graph, in particular, we use depth-first traversal to visit these nodes. The adjacency information for visiting the nodes is obtained by using a flood-fill algorithm. This algorithm determines the area connected to a given node which is also useful for ranking the node for rendezvous. Path-planning for visiting a particular node is implemented using the A* search algorithm [16].

C. Rendezvous

Once the environment is partially explored, mapped and stored as a graph, the robots need a strategy to meet in order to exchange information and speed-up the remaining exploration task. The rendezvous process is characterized by a sequence of attempts r_i of the robots $R \in \{1,2\}$ to meet with one another at specific times $t(r_i)$. These pre-scheduled attempts are interleaved with some other activity such as incremental exploration of the environment to discover more potential rendezvous locations p_i.

The rendezvous process for two robots is defined below which can easily be scaled for multiple robots.

- At the *rendezvous time* $t(r_i)$ the robots attempt to meet.
- Each robot R selects a location p^R to visit on this attempt, thus defining a sequence of attempted rendezvous locations p_i^R.
- Each robot travels to its appointed location p_i^R and, if $p_i^1 = p_i^2$ (they are in the same place) it finds the other robot there and the rendezvous process is completed.
- In the event of a failed rendezvous, the robots continue their background activity until the next rendezvous time.

The key challenge is to select matching points p_i^R to maximize the odds of finding the other robot, despite a lack of knowledge of the full set of locations or the location of the other robot. Moreover, following [1] we assume the locations are defined based on some sensor attribute that has been determined in advance, but due to sensor noise their sensor signatures may not be consistent between the two robots (for example they might agree to meet on the highest hill or the biggest open area, but that measurement may not be reliable and consistent). Complications to the basic scenario can also arise if the robots fail to arrive at a destination on time.

For this purpose, we implement three different rendezvous strategies which are derived from [1]. They are asymmetric sequential, symmetric sequential and exponential. The asymmetric and symmetric sequential strategies are deterministic in nature and their names represent the pre-defined roles of the robots during the rendezvous period. According to the asymmetric sequential strategy one of the robots is stationary while the other robot searches for it. This strategy is popularly known as "wait for mommy" strategy [10]. The symmetric sequential strategy requires both the robots to search simultaneously. The third strategy is probabilistic in nature. It assigns an exponential probability density function for visiting the ranked rendezvous locations.

The rendezvous locations are ranked using three different ranking criteria, namely, area-based ranking, linear distance-based ranking and sigmoid distance-based ranking. The area-based ranking as given in Equation 1 orders the potential rendezvous locations based on the area of coverage of the robot's sensor at the corresponding location. Intuitively, the larger the area of coverage at a given location the higher is the probability of spotting the other robot.

$$rank(p_i) = area(p_i) \tag{1}$$

The area-based ranking is modified to include the cost of traveling from the robot's current location to the potential rendezvous location. The cost of traveling is determined by considering the node-to-node distance from the source to the destination location. The inverse of the distance measure is multiplied with the area of coverage to obtain the linear-distance based ranking as in Equation 2. The sigmoid distance-based ranking criterion is obtained by weighing the distance measure by a sigmoid function as given in Equation 3. This function provides a threshold for scaling the closer landmarks with large weights and relatively smaller weights to the farther locations.

$$rank(p_i) = \frac{area(p_i)}{distance(p_i)} \tag{2}$$

$$rank(p_i) = \frac{area(p_i)}{sigmoid(p_i)} \tag{3}$$

The aforementioned rendezvous strategies are pairwise combined with the ranking criterion to evaluate the proposed objective: that the use of distance measure for ranking the rendezvous locations improves the exploration time.

IV. Implementation

We validate our approach in three different simulated environments using Player/Stage [17] [2]. Player is used to control the simulated hardware of the robot whereas Stage interprets the control commands and turns it into a simulation of the robot. An instance of the simulation is presented in Figure 3. The robots are equipped with laser scanners and GPS to sense and map the environment. Each robot also has a head-mounted camera to identify the other robot at the time of rendezvous. The robots detect each other using a color blob detection technique since each robot is represented by a distinct color.

[2] For a useful tutorial on how to use Player/Stage software for simulating the robots and the environment, refer: http://www-users.cs.york.ac.uk/ jowen/player/playerstage-manual.html

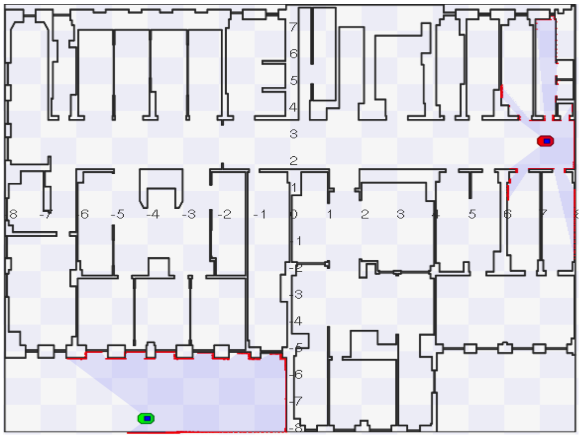

Fig. 3. Player/Stage simulation for "hospital section" environment.

The robots navigate through the environment using the skeleton structure as described in Section III-B. The nodes of the skeleton are considered as the potential rendezvous locations. This choice of using the nodes of the skeleton as the landmarks is inspired by human experience, where people intend to meet at junctions which are easily accessible (e.g. entrance or exits of buildings). The area covered by the nodes in the skeleton structure is used as a measure for ranking the node for rendezvous.

V. Simulation Results

We evaluated nine algorithmic variations produced by pairwise combining each of the rendezvous strategies with the ranking criteria for three different simulated environments given in Figure 4. These environments were specifically selected because each is an exemplar of a different canonical type of closed world. The environment in Figure 4(a) is obtained from the standard Player/Stage distribution whereas the ones in Figures 4(b) and 4(c) were obtained from [1].

The simulation results reported in this section are based on an average of 10 trials in each environment for all the algorithmic variants. In each experiment, a pair of robots is deposited in the environment at different extremal locations near the boundary and far from one another. They execute the rendezvous algorithm (one of the variants being evaluated) at pre-scheduled rendezvous intervals which are combined with the exploration task. A summary of the results obtained for the environments in Figures 4(a), 4(b) and 4(c) is illustrated in Figures 5, 6 and 7 respectively.

It can be observed from the aforementioned summaries that the algorithms exhibit slightly different results in each environment with a variation on the actual time-to rendezvous. This variability is due to substantially different locations that were visited and the extra time that was required if a rendezvous was missed. As indicated in the prior work of [1] there is a very large gap between the analytic best case and worst case times for any rendezvous strategy, and as a result of stochastic factors these extremal times can actually be achieved in practice.

The results suggest that the proposed linear distance-based ranking criteria consistently performed better (8 out of 9

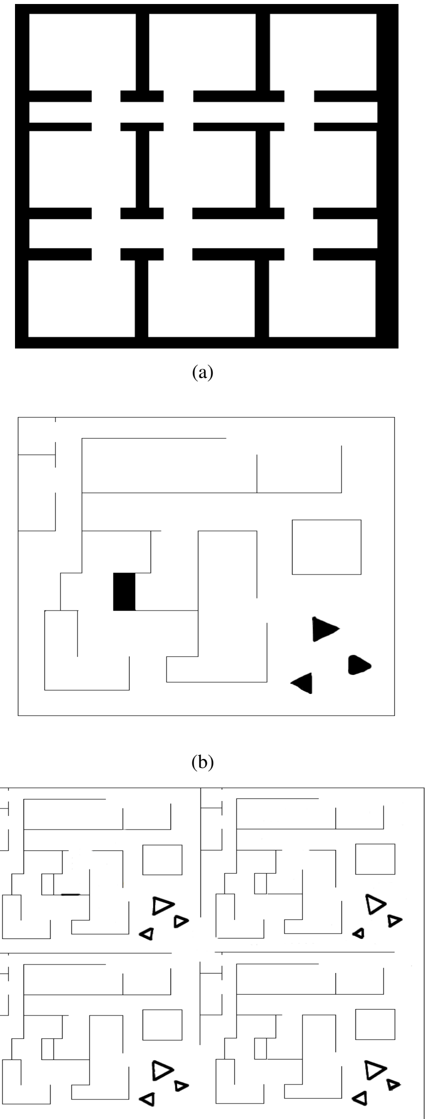

(a)

(b)

(c)

Fig. 4. Simulated environments used for the experiments

runs) than area and sigmoid distance-based ranking as per the total time to rendezvous. The results for the rendezvous strategy reflect that the symmetric strategy was convenient for environment 4(b) and 4(c) whereas the asymmetric strategy did better in 4(a). The reason asymmetric strategy performed better in environment 4(a) is due to the fact that the environment is symmetric and do not have many distinct features.

An additional set of experiment was performed to analyze the effect of variation in the initial distance between two agents on the number of rendezvous attempts, for different ranking criteria. The experimental setup was such that the agents were placed at random locations in the environment by fixing the initial distance between them. These distances ranged from quarter to a full factor of the environment diameter d i.e. $[d/4, d]$. The results for environment 4(a) recorded a total of

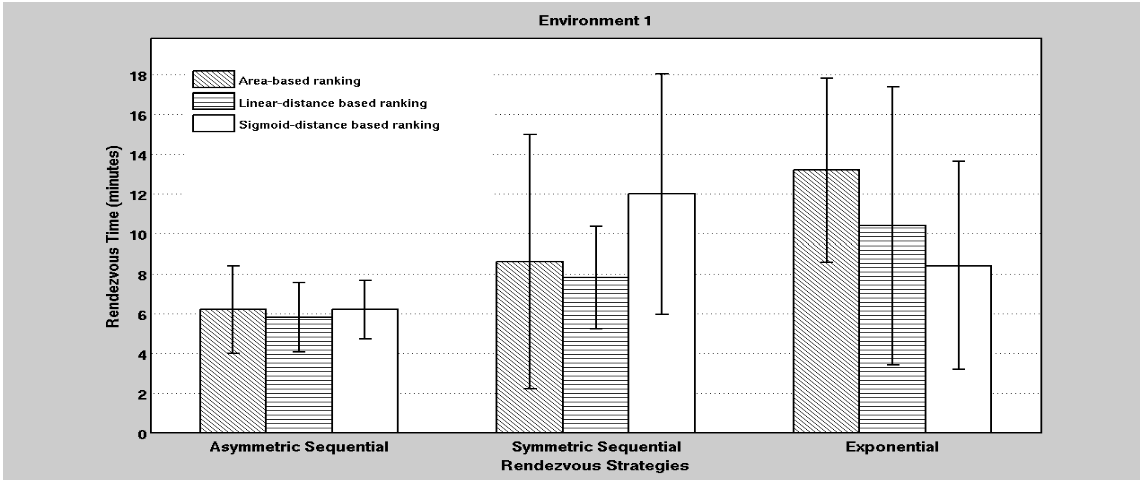

Fig. 5. Comparison of rendezvous strategies based on different ranking criterion for 4(a)

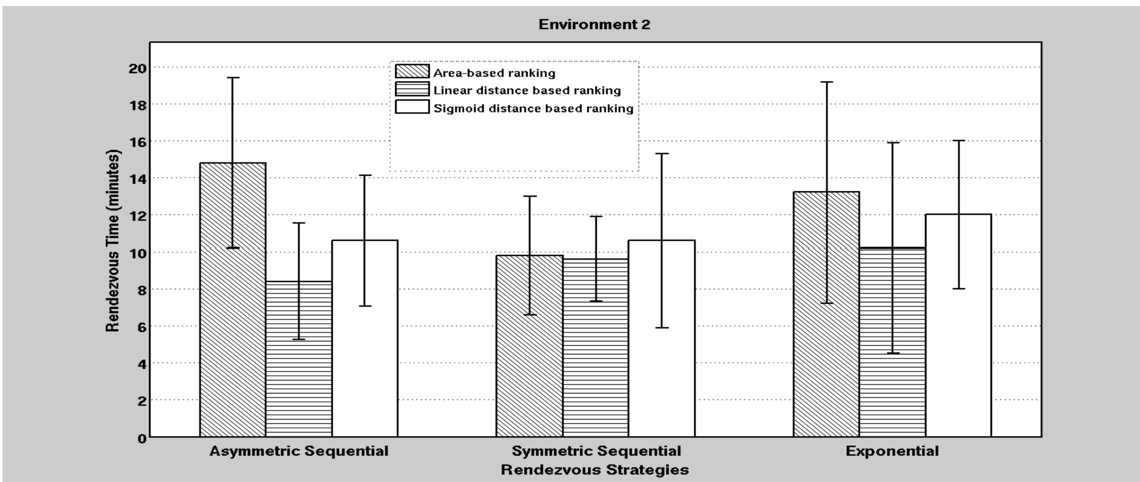

Fig. 6. Comparison of rendezvous strategies based on different ranking criterion for 4(b)

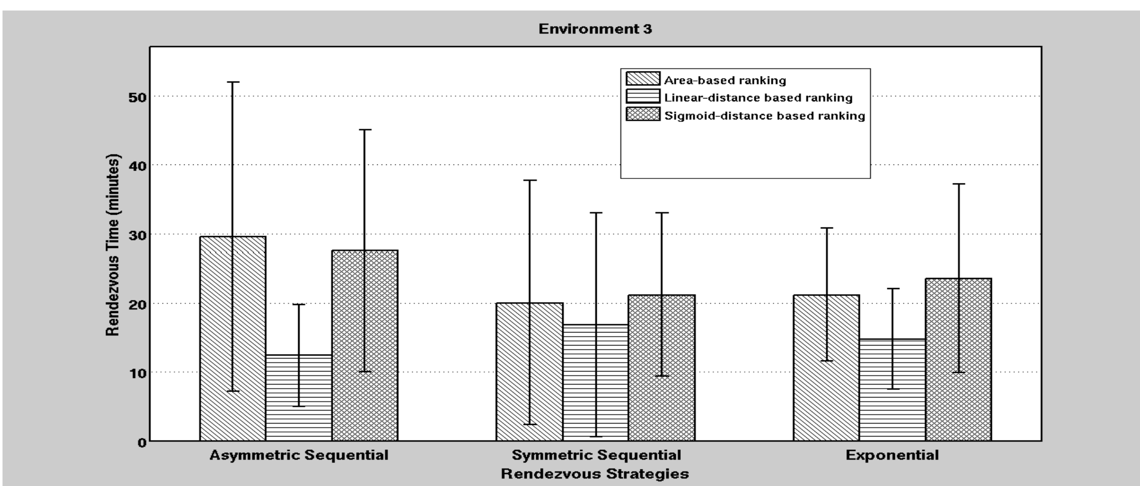

Fig. 7. Comparison of rendezvous strategies based on different ranking criterion for 4(c)

10, 12 and 9 rendezvous attempts for area, linear-distance and sigmoid-distance based ranking criteria respectively, as a function of initial distance. The sigmoid-distance based ranking required the least number of total rendezvous attempts and hence the minimum total rendezvous time. This is explained by the fact that sigmoid distance-based ranking penalizes the distant landmarks and advantages closer ones, but only to a limited extent so that irrelevant variations are not accentuated. It was also observed from Figure 8 that for smaller distance of separation between the robots, sigmoid distance-based ranking performed the best and when the distance between the agents was far off the area-based ranking out-performed the other ranking criterion. The reason for this trend is that sigmoid distance-based ranking initializes the rendezvous search from closer landmarks whereas area-based ranking looks for the best view-point at a given time.

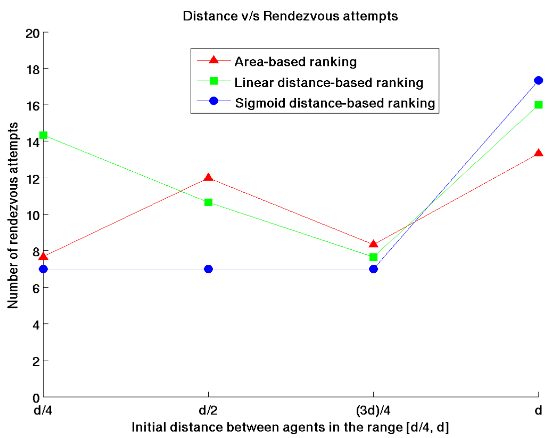

Fig. 8. The effect of variation in initial distance between two robot on the number of rendezvous attempts

VI. Conclusion and Future Work

The proposed approach validates our hypothesis that the use of distance-based ranking criteria can be effective in selecting locations of potential rendezvous. Our results indicate that such a consideration can reduce the exploration and rendezvous times across the range of environment types we examined. Naturally, only selecting to rendezvous in nearby locations could have negative consequences if used inappropriately, but the family of deterministic algorithms precluded the possibility of completely ignoring distant landmarks. This, however, may be the explanation for why the exponentially-weighted stochastic algorithm performed poorly.

The key conclusion is to take account of the accessibility of a rendezvous point while selecting where to go such that attempts to rendezvous can be interleaved with other activities including exploration. This result was, surprisingly, largely overlooked in prior work on this problem.

These results suggest two avenues for future work: one is the question of efficient rendezvous with more than two robots, and the related question of task partitioning in unknown environments. These problems are closely related since even after an initial rendezvous of two robots as part of a larger group, the attempt to make further rendezvous plans is itself a task partitioning problem.

References

[1] N. Roy and G. Dudek. Collaborative exploration and rendezvous: algorithms, performance bounds and observations. *Autonomous Robots*, 11(2):117–136, Sept. 2001.

[2] I. Rekleitis, G. Dudek, and E. Milios. Multi-robot collaboration for robust exploration. In *Proceedings of International Conference in Robotics and Automation*, pages 3164–3169, San Francisco, USA, April 2000.

[3] W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *IEEE International Conference on Robotics and Automation*, volume 1, pages 476–481, 2000.

[4] N. Atay and B. Bayazit. Emergent task allocation for mobile robots. In *Proceedings of Robotics: Science and Systems*, 2007.

[5] A. Visser and B. Slamet. Including communication success in the estimation of information gain for multi-robot exploration. In *6th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks and Workshops, 2008*, pages 680–687, 2008.

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

[7] Ioannis Rekleitis, Ai Peng New, Edward Samuel Rankin, and Howie Choset. Efficient boustrophedon multi-robot coverage: an algorithmic approach. *Annals of Mathematics and Artificial Intelligence*, 52(2–4):109–142, April 2008.

[8] R. Vincent, D. Fox, J. Ko, K. Konolige, B. Limketkai, B. Morisset, C. Ortiz, D. Schulz, and B. Stewart. Distributed multirobot exploration, mapping, and task allocation. *Annals of Mathematics and Artificial Intelligence*, 52(2):229–255, 2008.

[9] 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*, 2010.

[10] S. Alpern and S. Gal. *The theory of search games and rendezvous*. Springer, 2003.

[11] P. Flocchini, G. Prencipe, N. Santoro, and P. Widmayer. Gathering of asynchronous oblivious robots with limited visibility. *Symposium on Theoretical Aspects of Computer Science, 2001*, pages 247–258, 2001.

[12] J. Lin, AS Morse, and BDO Anderson. The multi-agent rendezvous problem. In *Proceedings of the 42nd IEEE Conference on Decision and Control*, volume 2, pages 1508–1513, 2003.

[13] A. Elfes and L. Matthies. Sensor integration for robot navigation: combining sonar and stereo range data in a grid-based representataion. In *26th IEEE Conference on Decision and Control, 1987*, volume 26, 1987.

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

[15] H. Choset and J. Burdick. Sensor based planning, part ii: incremental construction of the generalized voronoi graph. In *IEEE International Conference on Robotics and Automation*, pages 1643–1643. Citeseer, 1995.

[16] P.E. Hart, N.J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

[17] Brian P. Gerkey and Richard T. Vaughan and Andrew Howard. The player/stage project: tools for multi-robot and distributed sensor systems. *11th International Conference on Advanced Robotics*, pages 317–323, 2003.