Learning to Rendezvous during Multi-agent Exploration

Nicholas Roy and Gregory Dudek
Centre for Intelligent Machines and School of Computer Science
McGill University
3480 University St., Montreal, Québec, Canada H3A 2A7

Abstract. We consider the problem of rendezvous between two robots collaborating in learning the layout of an unknown environment. That is, how can two autonomous exploring agents that cannot communicate with one another over long distances meet if they start exploring at different locations in an unknown environment. The intended application is collaborative map exploration. Ours is the first work to formalize the characteristics of the rendezvous problem, and we approach it by proposing several alternative algorithms that the robots could use in attempting to rendezvous quickly while continuing to explore. The algorithms are based on the assumption that potential rendezvous locations, referred to as landmarks, can be determined by the robots as they explore; these locations are based on a distinctiveness measure computed with an arbitrary sensor. We consider the performance of our proposed algorithms analytically with respect to both expected- and worst-case behaviour. We then examine their behaviour under a wider set of conditions in simulation.

1 Introduction

A rendezvous is a meeting between two or more agents at an appointed place and time, for example, when two people meet at a familiar restaurant. The problem of rendezvous is ubiquitous in nature. Migratory animals must learn to meet to share information about food. Non-social animals must be able to find each other during mating season. Humans are equally familiar with the problem of rendezvous, as any family whose members become separated at a large zoo or mall well knows. Multi-agent robot systems also have an inherent need for inter-agent rendezvous. The ability to meet facilitates localisation, allows collaborative map learning and has a plethora of other advantages, but most importantly it allows communication.

In this paper, we consider how a pair of robots can jointly learn a map of an unknown environment. In particular, the key issue is how they can learn where to meet one another. Most existing hardware agents are only capable of communication over short distances. Environmental geometry, wireless transmission technology, power considerations and atmospheric conditions (or water conditions for underwater agents) all contribute to fairly short communication limits. In the absence of sophisticated satellite receivers or high power devices, a common constraint for successful communication is maintaining "line-of-sight" between agents, a constraint that is rarely satisfied in the real world. However, multi-agent robot systems for the majority of real-life applications enjoy substantial speed gains only with some level of communication, when compared with single-agent systems or multi-agent systems that do not communicate. (Balch & Arkin, 1994). Many distributed-agent algorithms, for instance dynamic path-planning, assume and rely upon instantaneous, infinite bandwidth communication between agents at all times in order to achieve promised performance levels (Brumitt &

Stentz, 1996).

In order to overcome these constraints on real communication, we have proposed that the agents use the information they have learned about the environment in order to find each other. However, many factors may inhibit the rendezvous process. Agents may not agree on rendezvous locations, or rendezvous times; domains of spatial knowledge may not overlap sufficiently. Consequently, the different agents must adapt to the different environmental conditions, and learn to overcome the factors that complicate the rendezvous process. Overcoming these factors can require that the agents learn each other's behaviours, and react accordingly.

Looking to biology, some simple algorithms are easily observed. Most animals rely upon established common meeting points, such as a beehive, or a watering hole. In an unknown environment, however, such an absolute reference point is almost impossible to define *a priori*. A common strategy has one agent (eg. a child lost at the zoo) wait to be found by other agents (eg. desperate parents). As we shall see, such simple strategies can perform poorly under many conditions.

Specifically, this paper discusses how to determine the best strategy for a successful rendezvous between two agents in optimal time, and ours is the first work to formalize the characteristics of this problem. We will consider how this rendezvous task can be efficiently accomplished under various assumptions about the environment and the perceptual abilities of the agents involved. In particular we are interested in the problem of rendezvous in the context of multi-robot exploration using video or sonar sensing. In practice, the particular sensing modality has numerous pragmatic implications, a major factor being the range at which the agents can either recognize one another or landmarks in the environment. In the context of a general rendezvous strategy we will, however, consider a generic "abstract" sensor that allows the agents to recognize one another when they are sufficiently close together and which allows them to evaluate any point in space as to its suitability as a rendezvous point.

2 Background

Several authors have considered interaction and coordination between multiple mobile robots. Arkin and his colleagues, in particular, have considered issues of co-ordinated motion as well as collaborative behaviour with minimal pre-planning (Arkin & Hobbs, 1992; Balch & Arkin, 1994). His work, however, has not focussed on exploration or the use of deliberative strategies to allow robots to meet.

There has been considerable work in addressing the problems of communication between agents in a multiple agent system, however the majority of work has been to maximise efficiency and minimise complexity (Mataric, 1992) (Hara, 1992). Mataric has looked at models of collaborative behaviour between mobile robots (Mataric, 1992), and observed that the form of communication plays an important role in how collaborative actions proceed. In this work, we deal with how to facilitate that communication by allowing the robots to meet.

Yoshida et al. addressed the problem of how to reduce a global communication network to local communication, in order to minimizing information complexity (Yoshida, 1995). However, there has not been much research in overcoming communication limitations, except by limiting the scope of the system to some area (such as a factory or a port) where communication between agents can be guaranteed by some global co-ordinator.

The selection of distinctive locations in a simple 2-D environment has been considered previously in the context of map-making (Kuipers & Byun, 1991), in which distinctive locations were determined by hill-climbing, that is, by local gradient ascent over some function of the sensor output. More generally, local maxima in some continuous property of the environment would seem to present an opportunity for converting a metric environment representation into a graph-like or topological one (Chatila & Laumond, 1985; Dudek *et al.*, 1991).

The problem of map generation from co-operative multi-agent exploration was discussed and implemented by Ishioka et al. (Ishioka, Hiraki, & Anzai, 1993). Their work is a canonical example of the potential applications of the technique presented in this paper, in which co-operative heterogeneous robots generated maps of unknown environments. They did not discuss the problem of rendezvous, but focussed only on how to merge maps once the rendezvous has occurred. It is worth noting that map fusion is also closely related to the generic image-registration problem.

3 The Rendezvous Problem

We have separated the rendezvous problem into two separate sub-problems. The first is learning what points in the environment are suitable for potential rendezvous. These points will be referred to as

landmarks. The second sub-problem that is the major focus of this work, is learning how to find which point, out of the set of potential rendezvous points, to visit at the assigned time. The context of the rendezvous is an unknown environment, with no shared spatial information between agents, and no communication until rendezvous. We are examining the rendezvous problem in terms of minimizing time to rendezvous under various conditions such as sensor noise, environment size, etc.. Determining when a rendezvous is necessary in the framework of another task is a task-dependent problem, and is outside the scope of this paper.

We start by briefly considering some properties of good distinctiveness measures. We will then present several rendezvous strategies, consider their statistical properties analytically and simulate their behaviour numerically. We conclude with a discussion of the results.

3.1 Landmark Selection

As an agent travels throughout the environment, every visited location is evaluated by the agent in terms of its uniqueness. The assumption is that distinctive locations (with respect to some sensor-based computation) serve as locations that both robots can independently select as good landmarks. This notion of a landmark also serves as the basis of the topological mapping strategy proposed by Kuipers (Kuipers & Byun, 1991). We refer to the scalar measure of suitability as a rendezvous point as **distinctiveness**: $D(x,y)$, or more generally, for a pose vector $\mathbf{q}$ we can define $D(\mathbf{q})$. This is implicitly a function of sensor data $\mathbf{f}(\mathbf{q})$, so we have $D(\mathbf{f}(\mathbf{q}))$. Although the agent's sensor may not return scalar values, some scalar suitability measure can be usually be computed from the sensor. Some intuitive examples of environmental attributes that might serve as distinctiveness measures are: symmetry, distance to the nearest obstacle, or altitude (for 3D surfaces – for example humans might select hill tops).

In order for two robots to agree on a good landmark, they must have similar perceptions of the environment or be able to convert their percepts into a common intermediate form. In the extreme case of two agents with dramatically different sensing modalities, there is essentially no way for them to rendezvous based on the recognition of environmental characteristics. Sensor noise can play a similar problematic role. We model this aspect of the problem by parameterizing the extent to which the two agents can reliably obtain the same measurement of distinctiveness at the same location. With full generality, we can consider one of the two agents as the reference perceiver (the arbiter of good taste) with a percept $D_1(x,y) = D(x,y)$ while the second robot obtains a sensor measurement which can be viewed as noisy with respect to that of the first robot:

$$D_2(x,y) = (1-\delta)D(x,y) + \delta\eta(x,y), \qquad (1)$$

where $\eta(x,y)$ is a noise process and δ specifies the extent to which both robots sense (or perceive) the same thing. If both robots have exactly the same perceptions of the environment we have $\delta = 0$. In the context of this formalism, $\eta(x,y)$ combines both intrinsic sensor noise and any differences in the type of sensor used.

The possible distinctiveness measures are heavily dependent on the types of sensor the robots have at their disposal. Because the robot learns the value at every point, a good modality is one that allows the distinctiveness to be defined at any location in the environment, and for which there exists some metric that can order the resulting landmarks in the environment in terms of distinctiveness. This ordering allows the landmarks to be ranked in terms of their likelihood to lead to a successful rendezvous.

Certain generic properties apply to suitable landmarks and the distinctiveness function $D(x,y)$ independent of the sensing modality. If the distinctiveness function is smooth and has few local extrema or inflection points, then it may be possible to define highly stable and mutually agreed-upon landmarks with great ease using gradient ascent. However, although this strategy is attractive in principle, we believe that in many real environments, sensor noise, occlusion and other factors may make the "distinctiveness surfaces" highly non-convex and thus complicate the process.

The distinctiveness function and the associated landmarks should be stable over time and should not depend on the trajectory or history of the robot. For example, the "Northern-most" point in the already-explored environment is a poor choice since if, for example, the explored area of each robot is circular, then two robots will only have the same "northern-most" point if the environment is highly constrained or if the explored regions are very similar.

In this paper, we will neglect issues of navigation and assume an agent can always accurately reach a desired goal in the environment. While our framework can accommodate navigational error, it is outside the scope of this paper. For concreteness, the reader can imagine a point robot capable of arbitrary motion within free space.

The distinctiveness measure typically used for research in this area, in the context of mobile robotics and sonar-based perception, is the mean distance returned by the sonar ring, which essentially uses the enclosed space as the value of the landmark - the bigger the room, the better a landmark. Note that errors due to specularity with respect to sonar make the physical interpretation of the measurement ambiguous.

3.2 Rendezvous Strategies

In order to estimate the effectiveness of alternative strategies for rendezvous, we have identified key attributes that must be formalised. Important attributes of the rendezvous problem are:

- Similarities — the reproducibility of the perceptions between agents (do they sense the same attributes, and do they even use the same sensors),
- Landmark Commonality — the extent of overlap between the spatial ranges of the agents (this may change with time),
- Synchronisation — the level of synchronisation between the agents (for example, can they agree to meet at high noon).
- Landmark Cardinality — the number n of landmarks selected by each agent.

Implicit in the description of these attributes are certain assumptions. It is assumed that all agents share some degree of synchronization - that is, all agents can agree on when rendezvous attempts should be made. However, this synchronisation may contain noise, which will be dealt with shortly. The second assumption is that all agents have the same landmark set cardinality - they all attempt rendezvous over the same number of landmarks (even if they are not using identically the same landmarks in their sets). Finally, it is assumed that all agents are performing the same task, and using the same rendezvous strategies.

We have developed several fundamental strategies for assuring a rendezvous. As we will see, the best strategy depends on several properties of the robots and of the environment. In the simplest, idealized, noise-free case, each robot should select the location in the environment that is the most distinctive. Given 100% landmark commonality and the absence of noise, they will select the same location. Each robot should navigate to this location and wait for the other robot(s) to arrive. At such a time, they could fuse their maps and suitably partition any remaining exploration to be done. The problem with this idealized scenario in practice is that due to sensor variations, or disjoint landmark sets, they may not agree on where the ideal landmark is situated.

Our formalization of the rendezvous problem takes the key attributes mentioned above into account as follows:

1. Sensor noise — the distinctiveness measure observed by the two robots is unlikely to match perfectly. This is expressed by the constant δ and leads to strategies that must effectively consider a larger number of candidate rendezvous landmarks since a single guaranteed candidate may not be determined reliably. The issue of non-repeatability of sensor readings due to noise is not relevant in the context of this abstract model and is not considered here.

2. Sensor dissimilarities — the two robots may not measure the landmarks the same as each other. As illustrated in Eq. 1, so long as we do not consider issues of repeatability, sensor differences can be modelled as a form of noise.

3. Asynchrony — when two robots are attempting to meet at the same landmark, the rendezvous may fail because one could not reach the landmark in a dynamic environment, or even more likely, one robot could not reach the landmark in time, and the other moved on. This asynchrony j is referred to in this paper as the probability that a given meet at a common landmark will fail, $j \epsilon [0, 1]$. This effect leads to a need for strategies that may re-visit the same landmarks repeatedly to compensate for missed meetings.

4. Non-identical Landmark sets — the robots may have explored different areas, and will have selected different landmarks that are not in the common region (assuming such a common region exists at all). This is modelled formally as the number d of landmarks out of a total set of n that are not common to the robots. The effect of the non-commonality is that both robots must consider a larger number of candidate landmarks, since any given subset of landmarks selected by one robot may not be known to the other robot.

We will show that the choice of an appropriate rendezvous strategy depends on the extent to which the robots have learned the same set of landmarks, the amount of sensor noise (or, equivalently, the similarity of the sensors) and the reliability of the robots being at a mutually selected rendezvous point and detecting one another.

In the presence of large amounts of sensor noise, the landmark selection will be essentially random, in which case the best strategy is simply to have one robot visit every landmark, and have the other robot sit and wait for it. However, this is also an unrealistically pessimistic scenario. If the robots have been constructed to facilitate rendezvous, they are likely to have a somewhat common perception of the environment and to have some commonality in their explored areas. In reality, the robots will probably experience some limited sensor noise, minimal dissimilarities, some asynchrony, and partial but not complete landmark commonality. So, the best strategy takes these factors into account, and chooses a series of landmarks to visit in some intelligent way.

We are interested in strategies that would permit a robot to interleave its learning of the environmental structure, and its rendezvous attempts so that if the rendezvous fails, the robots can continue their work and the strategies remain robust even in the face of a complete inability to find their associates. Below, we describe four alternative rendezvous strategies: these form exemplars of what we believe are two key representative algorithm classes.

1. Deterministic Algorithms – Given the same set of landmarks, these algorithms will always create the same ordering of landmarks.
    - **Sequential** – One robot picks a landmark and waits there for the other robot, which visits every landmark in turn. If the second robot has visited every landmark without encountering the first robot, the first robot moves to another landmark it has not yet visited.
    - **Smart-sequential** – Each pairwise combination of landmarks known to a robot is assigned a "goodness" value. This value is the product of the distinctiveness of the pair. The list of landmark pairs is sorted by this product, and one side of each pair is discarded, leaving an ordered list of n^2 landmarks from a set of n. The robot then visits the landmarks in this order.
2. **Probabilistic** – The landmarks are sorted with respect to their distinctiveness and then assigning a likelihood of visitation p_i for landmark i as a function of its rank in the sorted list i.e $p_i = f(i)$. The algorithm probabilistically selects a landmark to visit, using p_i for each landmark.
    - **Exponential** – The likelihood of visiting the $i-th$ best landmark is αe^i.
    - **Random** – On each attempted visit, each robot selects a landmark at random and goes there.

Each of these methods has particular advantages and disadvantages. The sequential method is simple, but makes no effort to account for relative likelihoods, or asynchrony. In view of the potential shortcomings of the sequential method, we have proposed an alternative method, the probabilistic method, that has an increased chance of compensating for a missed rendezvous and also attempts to compensate for small variations in the respective rankings of the landmarks selected by the two robots. For instance, the distinctiveness of each landmark could be the same, which would lead to a uniform random visitation strategy. The probability distribution $f()$ could be a linear function of value, or, if we assume that the amount of sensor noise is low, an exponential strategy. However, if sensor noise is high and the two agents do not share the same ordering of landmarks, then the agents may be forced into revisiting the incorrect landmarks much too often. A good compromise between these two methods is the smart-sequential method. The advantage of this method is that, if δ is low, landmark combinations with high values are explored before landmark combinations where one landmark has a very high value, and the other has a relatively low value. This leads to an increased probability of meeting even with substantial asynchrony. The smart-sequential method is tantamount to guessing where the other robot might be, given relatively similar but not identical landmark rankings.

4 Behaviour - Analytical results

We can make an analytical assessment of the performance of the deterministic rendezvous algorithms, compared to the random algorithm baseline. If there is no noise, no asynchrony, and 100% landmark commonality, then all of the algorithms which use the distinctiveness measure to sort landmarks

Algorithm	Simple	Async.	<100% Comm.
Random	$\frac{1}{log_2(\frac{n}{n-1})}$	$\frac{1}{log_2(\frac{n}{n-1+j})}$	$\frac{1}{log_2(\frac{n^2}{n^2-n+d})}$
Sequential	$n/2$	$\frac{n}{2}+j^{\frac{-1}{\log j}}$	$\frac{n}{2}+\frac{d}{n}^{\frac{-1}{\log \frac{d}{n}}}$
Smart-seq.	$\approx n$	$n+j^{\frac{-1}{\log j}}$	$n+\frac{d}{n}^{\frac{-1}{\log \frac{d}{n}}}$

Figure 1: Expected case behaviour. The columns denote the ideal case, the case where the asynchrony $j \neq 0$ and the case where the landmark sets are not identical, but each agent has d non-common landmarks.

will lead to a rendezvous after only one attempt (i.e., both robots will go straight to the mutually agreed upon best landmark.). The random algorithm can never assure a rendezvous but will have a small, equal probability of leading to a rendezvous on every attempt.

More interesting is the performance of the algorithms in the limit of high noise, $\delta = 1$, such that no common ordering between agents of the same landmarks can be reliably determined. The first assessment is the algorithmic time complexity, i.e., the expected time to rendezvous, for the three algorithms in the limit of $\delta = 1$. The expected time to rendezvous is the maximum number of unsuccessful rendezvous attempts, where the probability of no success on the next attempt is greater than or equal to 50%. For a landmark set of size n, the probability of any single, random rendezvous attempt being unsuccessful is:

$$P_{unsuccessful} = \frac{n-1}{n} \tag{2}$$

If the asynchrony rate is accounted for, then the probability of an attempt being unsuccessful rises to

$$P_{unsuccessful} = \frac{n-1+j}{n} \tag{3}$$

These equations give rise to table 1. The first column refers to both robots having the same set of landmarks. The second column considers the scenario where the robots may fail to get to the appointed landmark at the same time (or fail to notice one another). This probability is the asynchrony, j. The third column deals with the case where d of each robot's n landmarks are not in the other robot's landmark set.

In the deterministic sequential algorithm, the expected time of the simplest case (identical landmark sets, no asynchrony), is very straightforward. One agent sits at a landmark, and the other agent visits every landmark in turn until they meet - on average $n/2$ landmarks. However, in the presence of asynchrony, additional sweeps of all n landmarks will have to be performed. To find the expected number, k such additional sweeps, we use

$$0.5 = j^k \tag{4}$$

noting that each extra sweep i of k will reduce the probability of failure, and k such sweeps must reduce the probability of failure to 50%. Thus, on average $j^{\frac{-1}{\log j}}$ sweeps during the rendezvous will fail due to asynchrony. Similarly, for non-identical landmark sets, additional sweeps of n landmarks will have to be performed on average $\frac{d}{n}^{\frac{-1}{\log \frac{d}{n}}}$ times.

In the worst case, the performance time complexity is much more straightforward. For the probabilistic algorithms, such as the random strategy, or whenever asynchrony is an issue, the worst-case is always $\mathcal{O}(\infty)$, because a meet can never be guaranteed. Similarly, a rendezvous can never be guaranteed if any asynchrony is present, and so for $j \neq 0$, the worst case for all algorithms is $\mathcal{O}(\infty)$.

However, the deterministic algorithms are guaranteed to terminate when $j = 0$. In the worst case, the two algorithms terminate in n^2 iterations when they share no common landmarks. At this point, both agents can determine that they cannot meet, and continue exploration. If, however, the agents share identical landmark sets, the sequential algorithm has a much lower worst-case complexity than the smart-sequential strategy, because one agent is guaranteed to visit every landmark in the other agent's set in n iterations.

Algorithm	Simple	Async.	< 100% Comm.
Random	∞	∞	∞
Sequential	n	∞	nd
Smart Seq.	n^2-n	∞	$n^2-(n-d)$

Figure 2: Worst case behaviour. Cols. as in Fig. 1.

5 Numerical Simulation

In order to determine the behaviour of the various algorithms with increasing noise, the algorithms were tested in numerical simulation. Rather than simulating an actual exploration[1], two agents were modelled as having already explored an unknown area, and collected a set of landmarks. The distinctiveness values of the ordered landmarks were generated with a linear function, and then the random noise δ as developed in equation 1 was applied to the two sets.

[1]This was carried out as well, but is not reported here due to lack of space.

The visitation strategy was then executed on the two ordered sets of landmarks, creating a (potentially infinite) sequence of landmarks for each agent to visit. The sequence was terminated at the first instance where both agents had a landmark at the same position in the sequence, corresponding to a successful rendezvous. At $\delta = 0$, the two ordered sets were identical, and the deterministic algorithms (sequential and smart-sequential) were guaranteed to generate sequences of length 1. The length of the sub-sequences until rendezvous was used as a measure of time until successful rendezvous. The cardinality of the landmark set was 50 landmarks, unless otherwise specified.

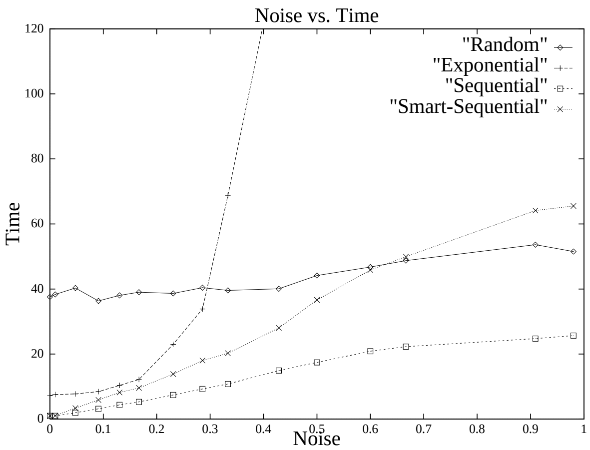

Figure 3: Baseline Performance - Time to Rendezvous as a function of Noise-level δ

The baseline simulation shows the performance of four algorithms in the face of increasing noise. The four algorithms are the deterministic sequential and smart-sequential algorithms, and weighted probabilistic distributions with exponential and linear probability functions. Recall that the exponential probabilistic function, for example, would have an exponentially higher probability of visiting the best landmark over any other. There is no asynchrony present, and the landmark sets were completely common. Figure 3 shows that the sequential algorithm is the best performer, especially in the face of high noise (i.e., $\delta > 0.2$) , which concurs with the analytical result. Clearly, exponential is a very fragile function, failing catastrophically with noise, $\delta > 0.2$.

Figure 4 shows the performance of the algorithms with a larger landmark set cardinality. Unsurprisingly, the performance of the algorithms scales appropriately with landmark set size.

In the face of asynchrony, however, the algorithms exhibit less intuitive behaviour. Asynchrony, again, is the probability that a particular ren-

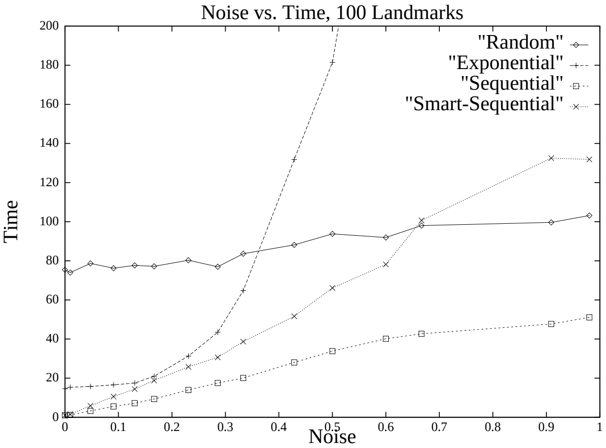

Figure 4: Performance with 100 Landmarks, ideal case

dezvous succeeded. The simulation (which creates landmark sequences) implemented asynchrony as the probability that a particular sequence element could be used. Even if the pair of landmark sequences contained the same landmark at identical positions, the sequence may not have terminated there, because the asynchrony probability prevented the first pair of matching landmarks in sequence from being compared, as if the robots had failed to rendezvous successfully despite attempting to do so at the same location at the same time.

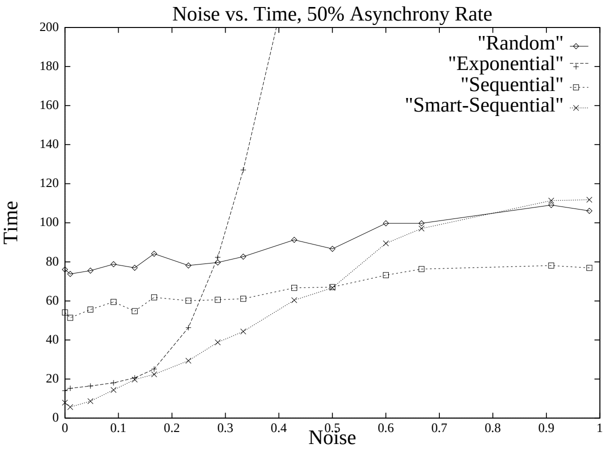

Figure 5: Performance with 50% Asynchrony rate

Figure 5 shows the performance of the algorithms given a 50% asynchrony rate, or a 50% probability of successfully making a rendezvous. In this case, the smart-sequential and exponential algorithms out-perform the sequential strategy, because the sequential form suffers from having to visit every other landmark before being able to return to the landmark that failed on a particular iteration, whereas the other two algorithms can return to landmarks relatively quickly. However, once noise dominates the values, ($\delta > 0.5$) the sequential algorithm becomes faster because it does not rely heavily on particular landmark values - it is not returning to the same landmark over and over again.

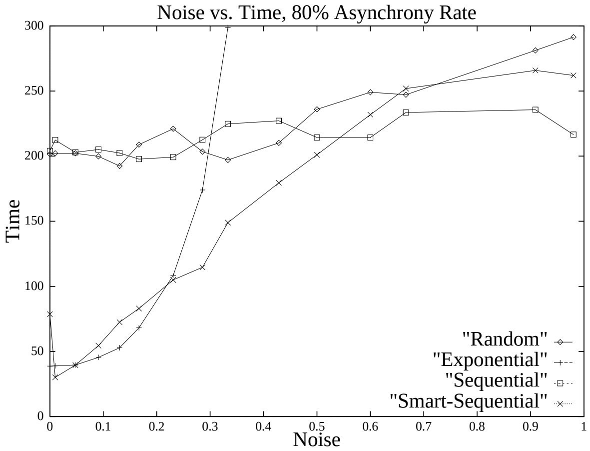

Figure 6: Performance with 80% Jitter rate

Even more interesting in the case of very high asynchrony, Figure 6 shows that the exponential probabilistic function outperforms the deterministic algorithms in the face of low noise ($0.5 < \delta < 0.25$), but again fails rapidly in the case of high noise ($\delta > 0.25$). The exponential algorithm essentially forces the robot to return to the same landmark over and over again, which is the correct strategy when asynchrony is high. However, when noise is high, the odds that the recurrent landmark is the wrong one increase, and the deterministic algorithms, which do not return to the same landmark as often, perform better.

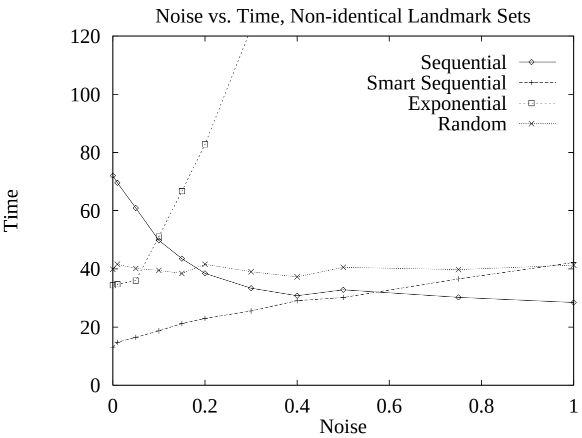

Figure 7: Performance with 25% non-commonality in landmark sets

Finally, Figures 7 and 8 show performance for maps with only 75% of the landmarks in com-

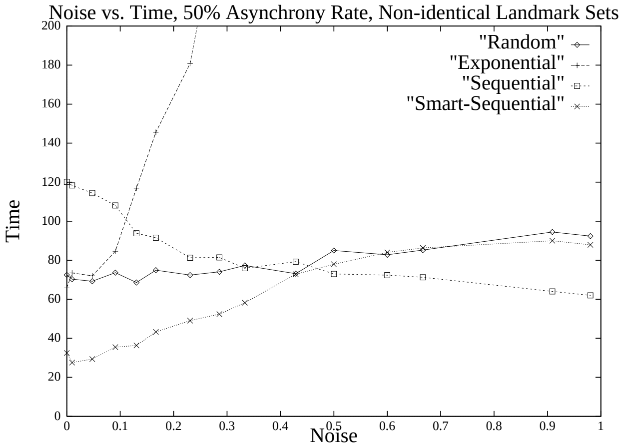

Figure 8: Performance with non-identical landmark sets, and 50% Asynchrony rate

mon. The performance with non-identical landmark sets (akin to non-isomorphic maps) is very similar to performance under low- to medium-asynchrony. The smart-sequential algorithm performs better with low noise because it can return to landmarks faster than sequential, but in the case of high noise ($\delta > 0.35$), returning to landmarks too frequently can be costly, and the sequential algorithm again dominates.

6 Experiments using Robots

While these abstract results have shown which algorithm the different robots must learn to use under different conditions, we would like to verify that the formalism is applicable to real robots. We first conducted the experiments in simulation, to show that the results under the range of noise conditions are upheld.

Using the map shown in figure 9, we tested a simulation of two robots learning the environment, and then learning to find each other. The two robots simulated were two cylindrical vehicles with holonomic motion constraints, using sonar transducers as their primary sensor.

The distinctiveness measure used the sonar sensors, and encapsulated a notion of both openness and symmetry; we summed the range from each sonar sensor, and then divided by the difference of opposing pairs, as equation 5 shows:

$$D = \frac{\sum_{n=1}^{64} R_i}{\sum_{n=1}^{32} \mid R_i - R_{i+32} \mid} \tag{5}$$

The experiment involved allowing the robots to learn the environment, determine where the best rendezvous locations were, and then learn to find

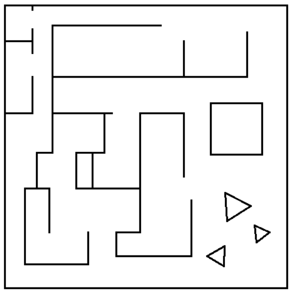

Figure 9: The map learned by the simulation of two robots.

the other agent. We performed 25 trials for 8 noise levels, shown in figure 10.

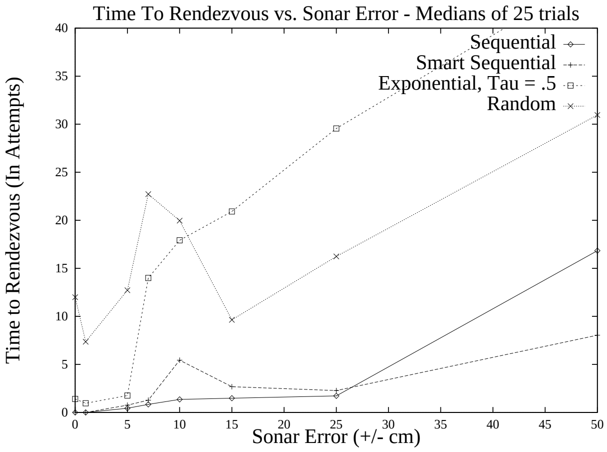

Figure 10: Simulated Robots - Time to Rendezvous as a function of Noise-level δ

As the figure indicates, the results were generally very similar to those of the numerical analysis, given in figure 3. Smart-sequential is a superior algorithm in the low-noise range, but once the noise begins to dominate the robots' perceptions of the environment, sequential is a preferable algorithm. The stochastic algorithm suffers the same dramatic failure with noise as was indicated by the numerical results.

Of more interest are the results concerning the ability of the agents to use rendezvous as a technique for overcoming communication deficits as they learn the environment. Figure 11 shows the increase in learning speed of the environmental structure available to the various algorithms. Notice that as the noise begins to dominate, the ability of the agents to find each other decreases, as does the rate at which the agents learn the environment. These figures were computed from the area learned by two agents in the time to explore added to the time to rendezvous, as compared to the area learned by one agent in the same time. Even though the agents must take the additional time to rendezvous, the increase in speed is still up to 50%.

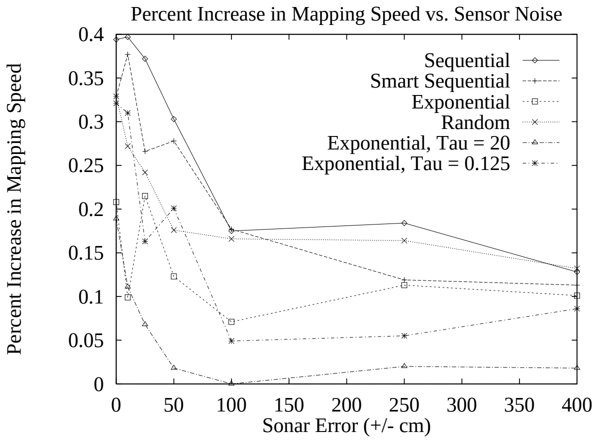

Figure 11: Increase in learning speed - Increase in speed as a function of Noise-level δ.

Finally, in figure 12, we show the result of an actual map learning and rendezvous experiment, implemented on a Nomad 200 and an RWI B-12 robot. This map is the result of the merging process, merging the information acquired by the two robots. This map is conclusive proof that the rendezvous process can be used successfully for multi-agent exploration and environmental learning.

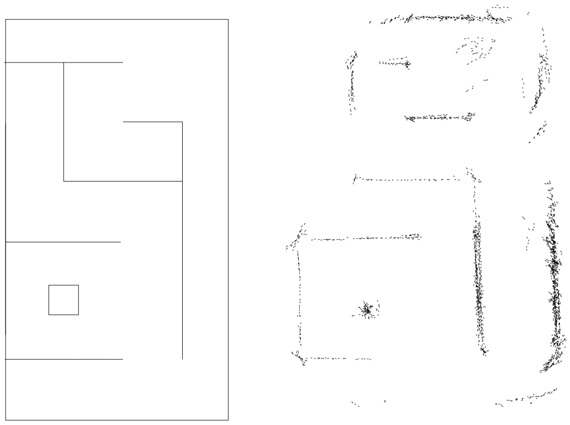

Figure 12: Map created from information learned by two robots and merged following rendezvous.

7 Conclusion

In this paper we have described the new problem of performing rendezvous between two exploring robots in an unknown environment. We are specifically interested in the ability of multi-robot systems to learn an environment, where communication is limited to short range. This is the first paper to describe a formalism for this problem, and we have presented an analytical and numerical analysis of some solutions to the rendezvous problem. Although we have primarily dealt with two agents, the algorithms we have presented could readily be adapted to larger collections of agents, or swarms.

We have shown that certain algorithms for performing a rendezvous in an unknown environment are especially good or bad under certain system and environment conditions. These factors include sensor noise, lack of commonality in the regions explored by the robots, and the possibility of the robots missing a scheduled rendezvous. An interesting result is that, depending on a combination of these confounding factors, no strategy is canonically a poor choice - under the correct circumstances, a heretofore poor choice of algorithm can outperform the erstwhile winner. These results are confirmed by both analytic closed-form solutions, and idealized numerical simulations.

Two major subtleties complicate the rendezvous problem. One is the possibility of missed rendezvous. This "asynchrony" factor may be a result of lack of synchronisation, a failure of the robots detecting each other, or navigation errors resulting in arriving at the wrong location. The other subtlety is whether both robots select landmarks from within commonly explored sub-regions of the environment. We refer to possible non-identical landmark sets as a lack of commonality. In the absence of these problems, several simple strategies are possible and rendezvous is a fairly simple problem. Expected time to rendezvous is between 1 and $n/2$ landmark visits depending on the strategy and the ability of the robots to agree on a consistent preference ordering of landmarks using noisy sensors.

In the presence of asynchrony and in the absence of 100% commonality of the explored areas, the algorithms referred to as sequential and smart-sequential each have their domains of superiority while surprisingly, a stochastic strategy based on an exponential probability of visit also has a small region of the parameter space in which it proves superior. That these small regions of parameter space exist indicates that the problem of rendezvous deserves further development. The smart-sequential strategy exploits the distinctiveness measure or preference ordering on landmarks to attempt to compensate for missed rendezvous and is superior under substantial levels of asynchrony and limited noise. The pure sequential strategy is preferable when asynchrony is low, since without asynchrony a meet is assured after n visits by avoiding visiting combinations that might otherwise compensate for asynchrony. With substantial levels of both asynchrony and moderate noise the stochastic search strategy is preferable to deterministic ones: a phenomenon we are investigating further.

Issues for future consideration are how the robots learn which algorithm to apply, and how to recognise the appropriate environmental conditions for each algorithm. Another interesting issue is how to refine the system parameters for subsequent rendezvous after an initial one has been achieved.

References

Arkin, R. C., and Hobbs, J. D. 1992. Dimensions of communication and social organization in multi-agent robotic systems. In *Proc. SJB 92*.

Balch, T., and Arkin, R. C. 1994. Communication in reactive multiagent robotic systems. *Autonomous Robots* 1(1):27–52.

Brumitt, B., and Stentz, A. 1996. Dynamic mission planning for multiple mobile robots. In *Proceedings of the IEEE/International Conference on Robotics and Automation*, 2396–2401.

Chatila, R., and Laumond, J. 1985. Position referencing and consistent world modelling for mobile robots. In *IEEE International Conference on Robotics and Automation*, 138–170.

Dudek, G.; Jenkin, M.; Milios, E.; and Wilkes, D. 1991. Robotic exploration as graph construction. *Transactions on Robotics and Automation* 7(6):859–865.

Hara et al., F. H. 1992. Effects of population size in multi-robots cooperative behaviors. In *Proceedings of the International Symposium on Distributed Autonomous Robotic Systems*, 3–9.

Yoshida et al., T. Y. 1995. A design method of local communication area in multiple mobile robots. In *Proceedings of the IEEE/International Conference on Robotics and Automation*, 2567–2572.

Ishioka, K.; Hiraki, K.; and Anzai, Y. 1993. Coorperative [sic] map generation by heterogeneous autonomous mobile robots. In *Proc. of the Workshop on Dynamically Interacting Robots*, 58–67. Chambery, France: IJCAI Press.

Kuipers, B., and Byun, Y.-T. 1991. A robot exploration and mapping strategy based on a semantic hierachy of spatial representations. *Robotics and Autonomous Systems* 8:46–63.

Mataric, M. 1992. Minimizing complexity in controlling a mobile robot population. In *Proc. IEEE Int. Conf. on Robotics and Automation*, 830–835.