Cooperative Localization During Exploration

IOANNIS M. REKLEITIS[1] GREGORY DUDEK[1]
EVANGELOS E. MILIOS[2]

[1]Centre for Intelligent Machines, McGill University,
3480 University St., Montreal, Québec, CANADA H3A 2A7
[2]Department of Computer Science, York University,
North York, Ontario, CANADA M3J 1P3

Abstract:- This paper presents a new approach on localization for a team of mobile agents. This approach is based on the cooperative use of each agent as a landmark by its partner(s). The proposed technique is applied to the exploration problem.

We present an algorithmic solution, simulation results, as well as a cost analysis and experimental data. In this approach a pair of robots observe one another's behaviour, thus greatly reducing odometry errors. We assume the robots can both directly sense nearby obstacles and see one another. We have implemented both these capabilities with actual robots in our lab. By exploiting the ability of the robots to see one another, we can detect opaque obstacles in the environment independent of their surface reflectance properties.

Key-Words:- Localization, Multi-Agent, Intelligent Collaboration, Exploration, Uncertainty Reduction.

1 Introduction

In this paper we discuss the benefits of cooperative localization during the exploration of a large environment. A new sensing modality is used in order to improve the accuracy of the position estimation of each robot. The robots explore the environment in teams of two; each robot is using the position of its partner in order to update the estimate of its position. Our approach is sufficiently robust to be able to cope with environments that may have uneven or slippery terrains, or whose surface reflectance properties are not well suited to conventional sensors.

This paper builds on previous work in which we have examined algorithms for multi-robot exploration in a theoretical context [16]. In this paper, we consider how to realistically implement a *low-cost* multi-robot position tracker and evaluate its performance. Moreover, we examine the empirical performance of multi-robot localization and exploration in a simulation context.

Observe that conventional approaches to robotic mapping and navigation are typically applied to test environments of rather limited size. Further, the sensing techniques used to both explore the environment and position the robot often make rather optimistic assumptions about the environment: diffuse visual reflectors, substantial reflectivity, etc. In practice, some surfaces may either be specular (mirror-like) reflectors or be hard to detect due to low reflectance, and some parts of the environment may have frictional properties that make large-scale odometry difficult.

We deal with these issues in two ways, both based on a polygonal approximation to the environment and the detection of convex (reflex) vertices. The presence of reflex vertices is critical since it is these reflex vertices that determine the occlusion of regions of the environment with respect to one another. We use a pair of robots observing one another to build a map and circumvent problems of object visibility. The exploration process is based on triangulation using an environment decomposition attached to reflex vertices.

In the next subsection, we will briefly discuss relevant background research. In Section 2 we discuss multi-robot localization and exploration including, in Subsection 2.1, an example of a visual "tracker" that we have used to implement the algorithm described in the paper. Section 3 contains a brief overview of the exploration algorithm. In Section 4 we present experimental results from simulations and from laboratory experiments.

1.1 Background

There are two major approaches regarding the localization of a mobile robot. The first approach to localization is to use landmarks in the environment in order to localize frequently and thus reduce the odometry error [4]. A common technique is to select a collection of landmarks in known positions and inform the robot beforehand [8, 10, 9]. Another technique is to let the robot select its own landmarks [18] according to a set of criteria that optimise its ability to localize, and then use that landmarks to correct its position [3]. The second approach to localization is to perform a matching of the sensor data collected at the current location to an existing model of the environment. Sonar, and laser range finder data have been matched to geometrical models [18, 11, 13, 14], and images have been matched to higher order configuration space models [1, 7] in order to extract the position of the robot.

The existence of clearly identifiable landmarks is a strong assumption for an unknown environment. Even in man-made environments the cost of maintaining labels in prearranged position has been prohibitive in the past. Moreover, in large scale explorations it is quite possible that the robot would have to travel a large distance (larger than its sensor range) before locating a distinct landmark.

Several authors have examined the issue of exploring space with one or more robots [15, 12, 5]. In general, multi-robot exploration techniques have tended to focus on models with limited coordination or communication between the robots [2]. In contrast, we consider a tight coupling between the exploring robots in the interest of greater accuracy of more efficient behaviour. Related work deals with exploring spaces large enough that the robots cannot see one another across the environment [16]. In this work, we consider the case where the robots do not lose visual contact as long as their view of one another is not occluded.

2 Cooperative Localization

Since sensing is being used to correct position estimation errors, the sole source of error in the selective localization of the robots is the inaccuracy of the "robot tracker" sensor that is used to update/correct the position of the moving robot relative to the position of the stationary one. Therefore, if the two robots start with one stationary robot in an initial position P_{origin}, then the moving robot could localize itself with respect to that position (see Fig. 1). Note that information from both sensing and odometry could be combined using either extended Kalman filtering or optimal multiscan alignment [11].

There are three potential sources of information for the localization of the moving robot. First, the odometry measurements $\hat{X}_{odom}(t)$ provide a base estimate of the moving robot's position (with high uncertainty σ_o. Second, the different objects in the environment, when sensed from different positions, could provide updates in the robot position [13, 18]. Finally, the robot tracker $\hat{X}_{track}(t)$ provides measurements relative to the position of the stationary robot $\hat{X}_{stat}(t)$. In practice over large scale environments, the position of different objects drifts over time and they cannot provide safe position updates. On the other hand the estimate of the robot tracker is influenced by the uncertainty in the position of the stationary robot σ_s plus the error of the tracker $\hat{X}_{track}(t)$. The accumulation of uncertainty on the position of the stationary robot depends only on the number of role exchanges the two robots had. Consequently, over large open spaces where the odometry error grows unbounded the moving robot could always reference back to a stationary landmark (played by the second robot).

$$\hat{X}(t) = \frac{\sigma_s(\hat{X}_{track}(t) + \hat{X}_{stat}(t))}{\sigma_s + \sigma_o} + \frac{\sigma_o \hat{X}_{odom}(t)}{\sigma_s + \sigma_o} \quad (1)$$

2.1 Tracker implementation

There are many sensors that could be used for the robot tracker. Our preliminary implementation is based on visual observation of a geometric target on the robot [6]. (Alternative possible implementations use retroreflectors or laser light striping – our actual robot is equipped with alternative such technologies.) Each robot is equipped with a camera that allows it to observe its partner. The robots are both marked with a special pattern for pose estimation. The first part of the pattern is a series of horizontal circles

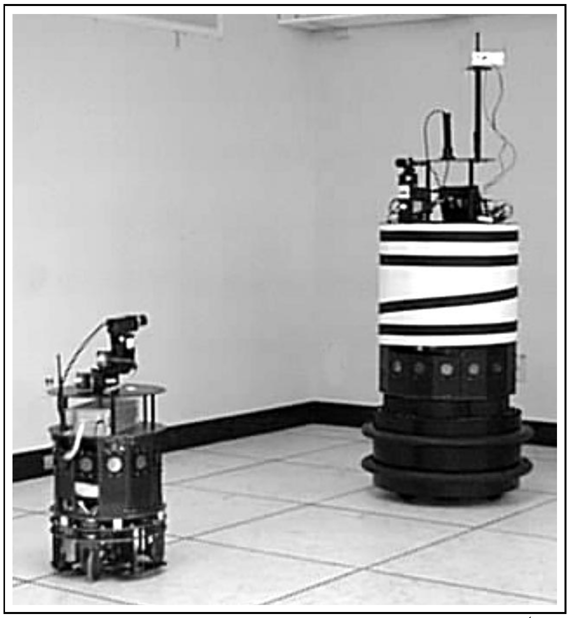

Figure 1: The visual robot tracker system (camera mounted on one robot, helix target pattern mounted on the second robot.

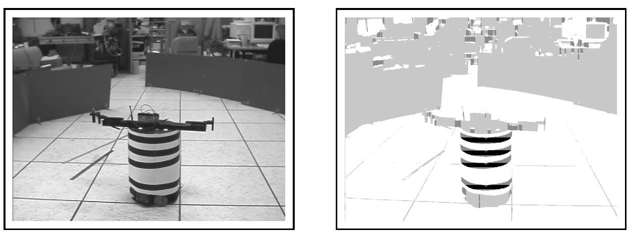

(a) (b)

Figure 2: Robot Tracker: (a) The raw image of the moving robot as observed by the robot tracker. (b) The helical and cylindrical pattern detected in the image.

(that project into an almost linear pattern in the image) which allows the robot to be discriminated from background objects: the ratio of spacing between the circles is extremely unlikely to occur in the background by chance. Thus, the presence of the robot is established by a set of lines (curves) with the appropriate length-to-width ratio and the appropriate inter-line ratios, as well as the correct position. The second component of the pattern is a helix that wraps once around the robot. The elevation of the center of the helix allows the relative orientation of the robot to be inferred (see Figure 2). In practice, this allows the robot's pose to be inferred with an accuracy of a few centimeters and 3 to 5 degrees.

3 Outline of the exploration algorithm

In [16] we presented an algorithm[1] for mapping the interior of an environment. The size of the area should be small enough to be covered by the range of the tracker sensor. Two mobile robots equipped with two different types of sensors are used in close cooperation to completely map the free space. Both robots use a traditional range finder in order to detect obstacles that are very close to them and, subsequently, to follow the object perimeter during the exploration. In addition, each robot has a robot tracker sensor that provides the pose of the other robot if the line of visual contact is uninterrupted, or a signal that an obstacle exists between the two robots.

The exploration algorithm is based on the following idea. At any single time one robot is positioned at a vertex (corner) of the environment operating as an intelligent landmark, while the other robot moves across the perimeter of the environment maintaining visual contact with the stationary robot. More precisely, as the moving robot follows one wall of the environment, it "sweeps" the line of visual contact across the triangle defined by the corner where the stationary robot is positioned and the two ends of the wall. Thus, the robot establishes the position of the wall and the occupancy of the swept free space inside the triangle.

Figure 3 illustrates the algorithm (simulation results). Figures 3 (a-i) present snapshots of the exploration as perceived by *Robot 0*, *Robot 1*, and the resulting map, respectively, at different instances of the exploration. The two robots exchange roles when the line of visual contact breaks. In the first row an early phase of the exploration is presented. The two robots have exchanged roles twice and *Robot 0* explores five new triangles. Consequently, in the second row *Robot 1* is exploring again, while *Robot 0* presents a portable landmark for localization. The third row illustrates the final stages of the exploration where *Robot 1* explores the final parts of the environment using *Robot 0* as a reference.

4 Experimental results

Different sets of experiments have been conducted in order to validate our approach. Experiments in a sim-

[1] Appendix A contains a formal description of the algorithm.

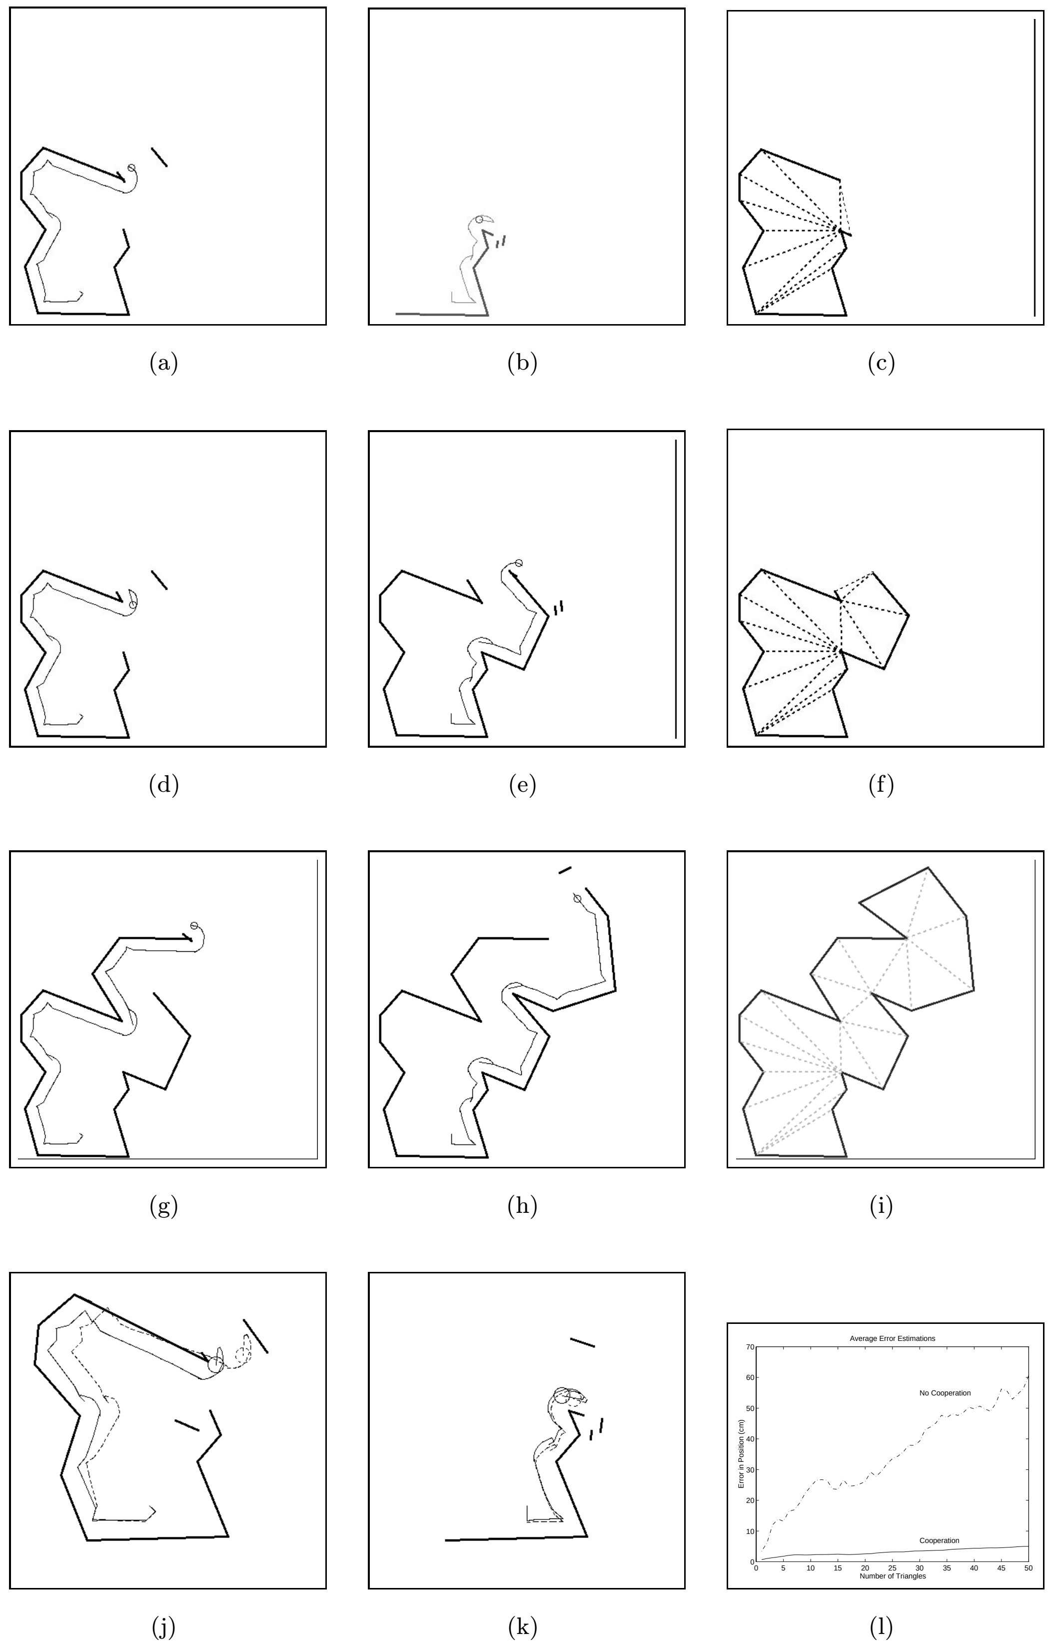

Figure 3: First three rows: Exploring an unknown environment, (a,d,g) The first column illustrates the trajectory of Robot 0. (b,e,h) The second column illustrates the trajectory of Robot 1. (c,f,i) Finally the third column presents the map up to that point. Last row: Close-up on the build up of the uncertainty when only odometry was used. The solid line is the odometry based estimation of the robots while the dashed line is the real position of the robots. (l) Taking the average error during the exploration of 50 triangle (over 100 experiments).

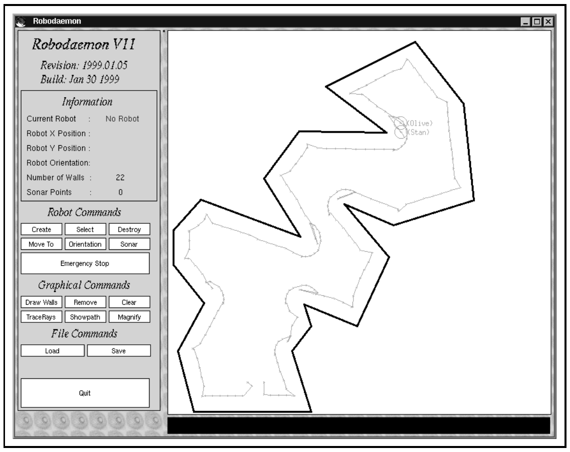

Figure 4: The paths of the two robots after the completion of the exploration.

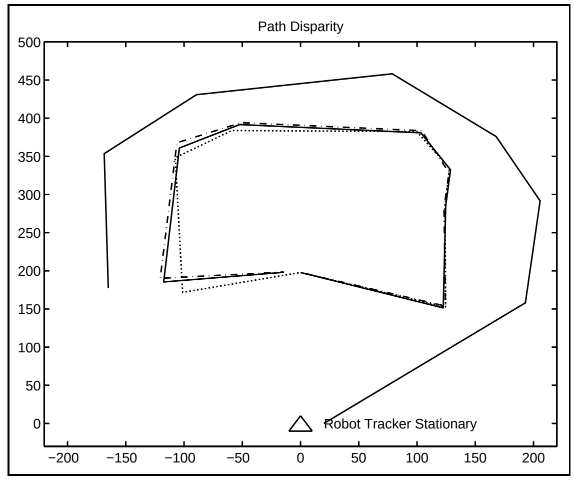

Figure 5: The different paths as they were perceived with the different localization methods.

ulated environment (using the RoboDaemon package, see Figure 4) provided verification in a variety of model worlds. In addition, laboratory experiments with the real robots helped us estimate realistic values for the uncertainty of the sensors and the odometry.

4.1 Simulation

Extensive experiments have been contacted using the robotic simulation package Robodaemon. The simulations allowed us to specify different parameters such as odometry error, robot-tracker uncertainty and the complexity of the explored environments. Figure 4 presents a typical environment used in the simulations and the path the two robots followed (144 m^2). As seen earlier (see Figure 3(a-i)) the two robots successfully mapped this model world. In Figure 3, in the last row, the early phase of the exploration is presented, using pure odometry for positioning. The dashed line depicts the real path of the robot and the solid line the odometry based paths. As can be seen in Figure 3(j,k) the accumulation of uncertainty, gradually distorts the map while maintaining local consistency. These distortions could lead over time to a map that is not even topologically sound.

The accumulation of uncertainty over time can be seen in Figure 3(l). The same experiment of exploring fifty triangles was performed one hundred times and the accumulated error was recorded. The dash-dotted line represents the average error when only odometry was used and the solid line when the tracker was used.

4.2 Physical Validation

In order to demonstrate the effectiveness of the proposed approach, several preliminary exploration tests were carried out in our laboratory in workspaces of roughly 16 m^2. This comparatively small testbed allowed us to control various factors such as inhomogeneities in the terrain as a function of trajectory and obtain ground truth data. Using this testbed we compared the time, accuracy, and robustness of different exploration strategies. In our experimental arrangement the role of the Nomad robot is played by a tripod mounted camera at the same height as the Nomad. This allowed us to more reliably and repeatably verify ground truth.

A laser pointer has been placed on top of the moving robot in order to accurately mark its current position on the floor. This setup allowed us to measure the displacement from the initial position after the completion of the tour.

In Fig. 5 the path of the robot is marked with a solid line, the odometry based estimates of position are marked by the dotted line, while the tracker estimates are illustrated with in a dashed-dotted line. The ouside solid lines mark the position of the walls the moving robot followed.

The final displacement from pure odometry estimates is approximately 15cm with an orientation error of 15°. The tracker estimate has approximately 1.3cm error. This corroborates our assumption that joint exploration and localization using a "tracker" can lead to much more robust modelling than odometry alone.

5 Conclusions

In this paper, we have described an approach to exploring and navigating in *large scale spaces* where positioning and sensing might be difficult. In fact, such difficulties are likely to arise in many real-world environments.

Our approach is based on exploiting a line-of-sight constraint between two robots to achieve exploration with reduced odometric error. This approach can also cope with obstacles with hard-to-sense reflectance characteristics.

We are currently planning large-scale experiments of this strategy in a real physical environment.

One issue in this context is that it is difficult to obtain accurate ground-truth to validate the performance of our approach over a large terrain. A standard practice is to simply observe the "clean-ness" of the resulting map and use this as a performance metric [18]. However we expect that the triangulation-based mapping we perform will yield results whose accuracy may be too great in polyhedral environment for such qualitative evaluation methods to be satisfactory.

In prior work, we have considered alternative strategies for environments where the distance is too large to permit reliable operation of the tracker across the workspace [16]. An open issue is how to automatically detect such situations *efficiently* during exploration and switch strategies, or switch back-and-forth between strategies based on local properties of the environment.

We are also considering combining this approach with more traditional localization methods (such as landmarks [17]) where they can be used effectively. Doing this efficiently appears feasible but remains unresolved.

References

[1] I. A. Bachelder and A. M. Waxman. A view-based neurocomputational system for relational map-making and navigation in visual environments. *Robotics & Auton. Systems*, 16:267–289, 1995.

[2] Tucker Balch and Ronald C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1994.

[3] M. Betke and L. Gurvits. Mobile robot localization using landmarks. *IEEE Trans. on Robotics & Automation*, 13(2):251–263, April 1997.

[4] J. Borenstein, H R. Everett, L. Feng, and D. Wehe. Mobile robot positioning: sensors and techniques. *J. of Robotic Systems*, 14(4):231–249, Apr 1997.

[5] Gregory Dudek, Paul Freedman, and Ioannis M. Rekleitis. Just-in-time sensing: efficiently combining sonar and laser range data for exploring unknown worlds. In *Int. Conf. in Robotics & Automation*, v. 1, pp 667–671, Apr. 1996.

[6] Gregory Dudek, Michael Jenkin, Evangelos Milios, and David Wilkes. Experiments in sensing and communication for robot convoy navigation. In *Int. Conf. on Intelligent Robots & Systems (IROS)*, v. 2, pp. 268–273, Aug. 1995.

[7] Gregory Dudek and Chi Zhang. Vision-based robot localization without explicit object models. In *Proc. Int. Conf. of Robotics & Automation*, 1996.

[8] T. Fukuda, et al. Navigation system based on ceiling landmark recognition for autonomous mobile robot - landmark detection based on fuzzy template matching (ftm). In *Int. Conf. on Intel. Robots & Systems*, v. 2, p. 150–155, 1995.

[9] F. Giuffrida, C. Massucco, P. Morasso, G. Vercelli, and R. Zaccaria. Multi-level navigation using active localization system. In *Int. Conf. on Intelligent Robots & Systems*, v. 1, p. 413–418, 1995.

[10] C. Lin and R. Tummala. Mobile robot navigation using artificial landmarks. *Journal of Robotic Systems*, 14(2):93–106, 1997.

[11] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.

[12] Vladimir J. Lumelsky, Snehasis Mukhopadhyay, and Kang Sun. Dynamic path planning in sensor-based terrain acquisition. *IEEE Trans. on Robotics & Automation*, 6(4):462–472, Aug. 1990.

[13] Paul MacKenzie and Gregory Dudek. Precise positioning using model-based maps. In *Int. Conf. on Robotics & Automation*, 1994.

[14] F. Nashashibi and M. Devy. Combining terrain maps and polyhedral models for robot navigation. In *1993 Int. Conf. on Intelligent Robots & Systems.*, p. 685–691, July 1993.

[15] B. Oommen, S. Iyegar, S. Nageswara, S. Rao, and R. Kashyap. Robot navigation in unknown terrains using learned visibility graphs, part i: The disjoint convex obstacle case. *IEEE J. of Robotics & Automation*, 3(6):672–681, 1987.

[16] Ioannis Rekleitis, Gregory Dudek, and Evangelos Milios. Multi-robot exploration of an unknown environment, efficiently reducing the odometry error. In *Int. Joint Conf. in Artificial Intelligence (IJCAI)*, v. 2, p. 1340–1345, Japan, Aug. 1997.

[17] Robert Sim and Gregory Dudek. Learning visual landmarks for pose estimation. In *Proc. of Int. Conf. on Robotics & Automation*, May 1999.

[18] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning and Autonomous Robots*, 31,5:29–53,253–271, 1998.

6 Appendix A

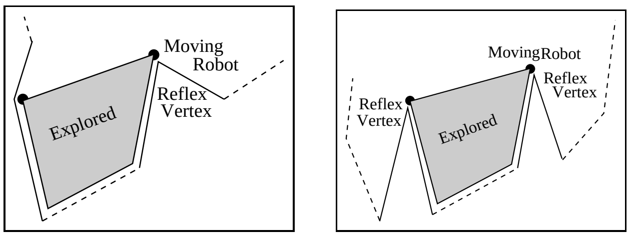

(a) (b)

Figure 6: (a) **Case 1:** The stationary robot is at a non-reflex vertex and the moving robot encounters a reflex vertex that would interrupt the line of visual contact (b) **Case 2:** Both robots are placed at reflex vertex such that any further exploration would break the line of visual contact.

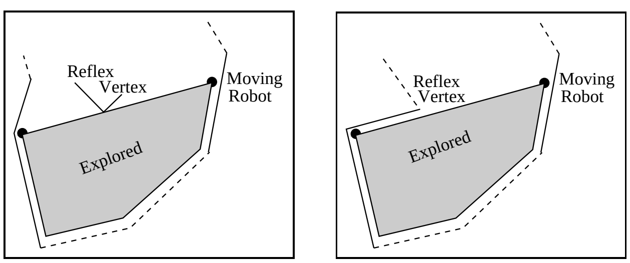

(a) (b)

Figure 7: (a) **Case 3:** Occluding Vertex between the two robots. (b) **Case 4:** Occluding Edge next to the stationary robot.

A sketch of the exploration algorithm is presented next. Both robots run the same exploration algorithm, taking turns moving thus mapping the free space and being stationary thus providing a fixed localization reference for the moving robot. In the following we assume no three points are co-linear if they are, it would involve a minor but tedious change to the algorithm. There are four different cases where the line of visual contact is interrupted (Fig. 6a,b and Fig. 7a,b), in these cases the moving robot can not continue its previous course and it has to make a decision where to move next in order to maintain visual contact with the stationary robot. The environment is explored in regions of free space composed by neighbouring triangles. The algorithm is summarised below.

```
While Unexplored Areas Do
 {
  Cover Nearest Unexplored Area
   {
    While No Occlusion Do
     Explore the next triangle of
     free space
    If Occlusion Then
     If Case 1 Then
      The two robots exchange roles.
     Else If Case 2 Then
      The Moving Robot goes to
      the Stationary Robot. Marking
      the reflex vertex as an opening to
      an Unexplored Area.
     Else If Case 3 Then
      The Moving Robot marks its position
      as a temporary vertex and moves
      towards the Stationary Robot until
      it encounters the occluding
      Reflex Vertex. The line
      between the occluding vertex and
      the temporary vertex is an
      opening to  an Unexplored Area.
     Else If Case 4 Then
      The two robots exchange roles
      The new Moving Robot follows
      the occluding edge to the next
      corner, then the two robots
      exchange roles again.
    Continue The Exploration.
   }
  If No Triangle of free space Then
  Move to the closest Unexplored
  Area.
 }
```