Graph-Based Exploration using Multiple Robots

Ioannis M. Rekleitis[1], Gregory Dudek[1], and Evangelos E. Milios[2]

[1] Centre for Intelligent Machines, McGill University,
3480 University St., Montreal, Québec, CANADA H3A 2A7
[2] Faculty of Computer Science, Dalhousie University,
6050 University Avenue Halifax, Nova Scotia, Canada B3H 1W5

Abstract. We present an approach to multi-robot exploration of large environments. Our method is designed to be robust in the face of arbitrarily large odometry errors or objects with poor reflectance characteristics. The algorithm achieves its robustness by using a team of cooperating agents. The critical aspect of our method is the use of a vision system that sweeps areas of free space and generates a graph-based description of the environment. This graph is used to guide the exploration process and can also be used for subsequent tasks such as place recognition or path planning. As a result of the guidance provided by the dual graph of the triangulated environment, our system can guarantee complete exploration without any overlaps.

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 behavior, thus greatly reducing odometry errors. We assume the robots can both directly sense nearby obstacles and see one another (if their view is not obstructed). We have implemented both these capabilities with actual robots in our lab. [1]

Keywords: Multi-robot, exploration, mapping, cooperation, sensor fusion.

1 Introduction

We present an approach to the exploration and mapping of large environments that functions despite arbitrarily large odometry errors or objects with poor reflectance characteristics. We also present experimental results that validate this approach and discuss specific implementation issues for a suitable sensor.

The exploration is done by two robots which jointly assist one another. We utilize a vision sensor (**Robot Tracker**) with a dual purpose: localization and mapping. At any time, one of the robots is stationary while the other robot is moving. The stationary robot acts as an artificial landmark in order for the moving robot to recover its pose with respect to it. Therefore,

[1] Appeared in "5th International Symposium on Distributed Autonomous Robotic Systems (DARS) 2000", pages 241-250, Knoxville, Tennessee, USA, October 4-6, 2000.

a detectable landmark is provided without any modification of the environment. We call this approach *Cooperative Localization*. Moreover, when one robot is moving maintaining uninterrupted the line of visual contact with the stationary robot, it effectively maps the area covered by the line of visual contact. As the two robots move through the environment, they map areas of free space using the fact that they constantly are able to "*see*" each other. 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.

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 diverse 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 the dual graph of the triangulation using an environment decomposition attached to reflex vertices.

This paper builds on previous work in which we have examined algorithms for multi-robot exploration in a theoretical context [9]. Now we consider how to realistically implement a *low-cost* multi-robot position tracker and evaluate its performance, and we examine the empirical performance of multi-robot localization and exploration in a simulation context.

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

1.1 Background

There are two major approaches to localization while mapping for 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 [3]. A common technique is to select a collection of landmarks in known positions and inform the robot beforehand [5]. Another approach is to let the robot select its own landmarks [11] according to a set of criteria that optimize its ability to localize, and then use these landmarks to correct its position. A variation in this theme is to perform a matching of the sensor data collected at the current location to an existing model of the environment. In addition to vision, sonar and laser range finder data have been matched to geometrical models [11,6,8], and images have been matched to higher order configuration space models [1] in order to extract the position of the robot.

A related dichotomy is between natural and artificial landmarks; rather than exploit landmarks already occurring in the environment, whether learned or manually selected, it has sometimes been found preferable to physically modify the environment to ensure that good landmarks are present. Artificial landmarks that have been used include visually detectable ones, sonar beacons, and radio beacons. The key advantage of artificial landmarks is not only that they assure the presence of the landmarks, but that the landmarks are assured to be detectable and unambiguous.

For navigation, it is often desirable to combine information from multiple sensors rather than using a single modality. In particular, the results of odometry can be combined with external sensing. A common mechanism for this is the Extended Kalman filter [2].

Several authors have examined the issue of exploring space with one or more robots [7]. 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 behavior. Related work deals with exploring spaces large enough that the robots cannot see one another across the environment [9]. 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

Localization is achieved in our approach by tight collaboration of the robots. Since sensing the other robot is used in order to correct position estimation errors, the main source of error in the cooperative localization of the robots is the inaccuracy of the "robot tracker" sensor. This sensor is used to update/correct the position of the moving robot relative to the position of the stationary robot. At any given time one robot is stationary and the other robot is moving (later they exchange roles). After the robot moves it tracks the stationary robot and updates its position estimate according to the data from the robot tracker. Note that information from both sensing and odometry could be readily combined in practice using either extended Kalman filtering or optimal multiscan alignment [6].

[2] In this paper, while we acknowledge the desirability of such sensor fusion, we will not consider it in detail for reasons of brevity.

There are three potential sources of information for the localization of the moving robot. First, the odometry measurements provide a base estimate of the moving robot's position $\hat{X}_{odom}(t)$ (with high uncertainty σ_o). Second, the different objects in the environment, when sensed from different positions, provide updates in the robot position [8,11]. 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 estimates of the position of different objects drift over time and 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).

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

2.1 Tracker implementation

There are many sensors that could be used for the robot tracker. Our current implementation is based on visual observation of a geometric target on the robot [4]. (Alternative possible implementations use retroreflectors or laser light striping – our actual robot is equipped with 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.

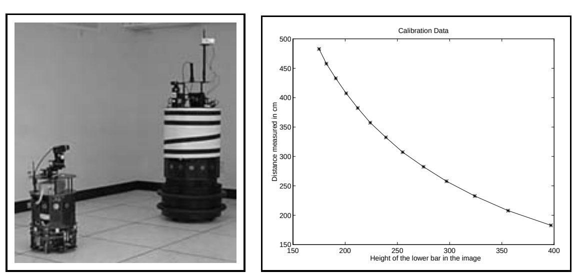

(a) (b)

Fig. 1. (a)The visual robot tracker system (camera mounted on one robot, helix target pattern mounted on the second robot. (b)Calibration data for the distance estimation relating observed image position to actual distances

By mounting the observing camera above (or below) the striped pattern, the distance from one robot to the other can be inferred from the height of the stripe pattern in the image, due to perspective projection (scaling of the pattern could also be used). The difference in heights between the observing camera and the target can be manipulated to provide a tradeoff between range of operation and accuracy (see Fig. 1b). One advantage of the helical target for orientation estimation is that it functions correctly even at very large distances (although with reduced accuracy, of course).

3 Outline of the exploration algorithm

In [9] we presented an algorithm for mapping the interior of an environment. The size of the area should be small enough to be covered by the range of the robot 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. This establishes a **line of visual contact**. If the view from one robot to the other is interrupted, we can determine the presence of an obstacle somewhere along this line.

For environments of limited size, we will show how this line of visual contact can be exploited to precisely explore the environment and determine its geometry. The essential operation in our strategy is triangulation: that is, the decomposition of the environment into triangular regions. The triangulation is accomplished by exploring one triangle of free space each time. One robot is stationary and acts as a landmark while the other robot moves mapping a triangle of free space. The stationary robot does not accrue odometry error and it can determine the position of the moving robot (while sweeping out free space with the line of visual contact).

3.1 Triangulation

The moving robot follows a wall of the environment while the stationary robot is positioned at a corner. If the two robots maintain visual contact during the exploration of a wall then the line of visual contact sweeps through a triangle of free space. These triangles of free space are the basic building blocks of the map. Any environment can be described using a polygonal approximation, a triangulation of a polygon is a well known spatial decomposition of polygons providing full coverage of the interior of the polygon. As can be seen in Fig. 2b, the moving robot explores a triangle at a time until it covers the complete triangulation.

In order to complete the exploration without any overlaps the dual graph of the triangulation is used. Every triangle corresponds to a vertex in the

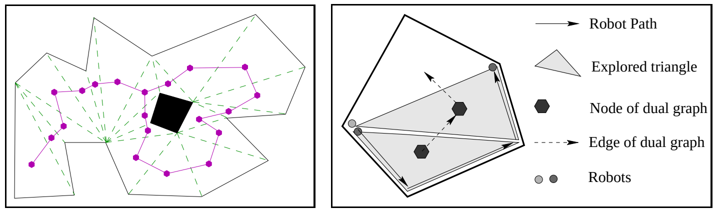

(a) (b)

Fig. 2. (a) A polygonal environment, the triangulation and the embedded graph. (b)Exploration of two triangles in a simple environment

dual graph. If two triangles share a common internal diagonal, then an edge connects the corresponding vertices. Fig. 2a presents a polygonal world with an obstacle and its triangulation with the dual graph embedded in it. During the exploration, the two robots traverse the dual graph. It is worth noting that the reflex vertices of the environment produce decision points due to the break of the line of visual contact. Therefore, when the moving robot reaches a reflex vertex that would interrupt the line of visual contact, the two robots exchange roles and the stationary robot starts moving. If the line of visual contact is interrupted by a reflex vertex at the middle of exploring an edge, then an internal triangle is constructed and the dual graph bifurcates. The moving robot proceeds to the occluding vertex, marking one triangle as a gate to an unexplored area, and maps the third triangle adjacent to the internal triangle (see Fig. 2a, the internal triangle attached to the left corner of the obstacle).

3.2 Illustrative Example

Fig. 3 illustrates the operation of the algorithm (simulation results). Figures 3(a-o) present snapshots of the exploration in three columns. The first column presents the world as perceived by *Robot 0*. The second column presents the world as perceived by *Robot 1*. Finally the last column presents the resulting map up to that point showing the triangulation and the dual graph. 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 once, *Robot 0* is stationary (Fig. 3a), while *Robot 1* explores the third triangle (Fig. 3b). Consequently, in the second row *Robot 1* is stationary after it reaches a reflex vertex that interrupts the line of visual contact. *Robot 0* is mapping the sixth triangle. The third row demonstrates the case of an occluding obstacle. *Robot 0*, after losing visual contact with the stationary *Robot 1*, moves towards it mapping the occluding vertex. The resulting map contains an internal triangle and a bifurcation is introduced to the underlying graph (see Fig. 3i). The guiding graph contains two possible paths, *Robot 0* explores the right path keeping the line of visual contact unbroken with the stationary *Robot 1*. The fourth row presents the exploration at a later stage. *Robot 0* is stationary at a reflex vertex while *Robot 1* explores the other side of the environment. The fifth 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

Several sets of experiments have been conducted in order to validate our approach. Experiments in a simulated environment (using the RoboDaemon package, see Fig. 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.

Simulation: Extensive experiments have been conducted using the robotic simulation package RoboDaemon. The simulations allowed us to examine the impact of different values of the different parameters such as odometry error, robot-tracker uncertainty and the complexity of the explored environments. Typical values for the standard deviation of the odometry errors are: 3cm/m in translation $3°/360°$ in rotation and $2°/m$ drift during translation. Fig. 4a presents a typical environment used in the simulations and the path the two robots followed. An obstacle is placed in the middle of the environment and the two robots construct the map by moving around it. As seen earlier (see Fig. 3(a-o)) the two robots successfully mapped this model world. Fig. 4b presents a complex environment with two occluding vertices that create bifurcations in the graph. In small worlds and/or cluttered environments, multiple observations of the same object could be used in order to correct the positioning of the moving robot.

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 several trials 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. In preliminary experiments using this arrangement, we have verified that our approach to collaborative exploration can improve pose estimation and map generation even over fairly short trajectories (of as little as 2m).

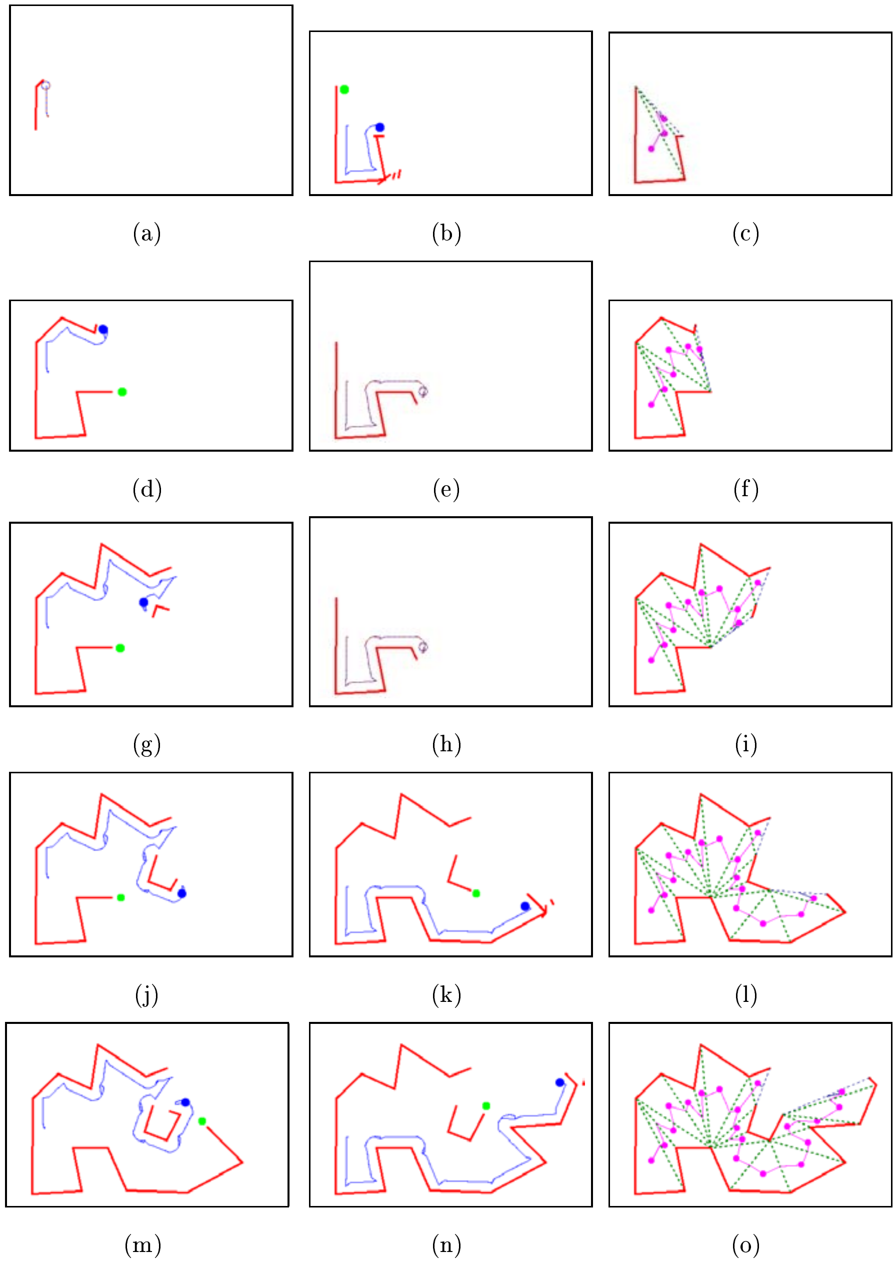

Fig. 3. Exploring an unknown environment with one obstacle: The first column illustrates the trajectory of **Robot 0**. (subfigures a,d,g,j,m). The second column illustrates the trajectory of **Robot 1** (b,e,h,k,n). Finally the third column presents the **map** up to that point (c,f,i,l,o)

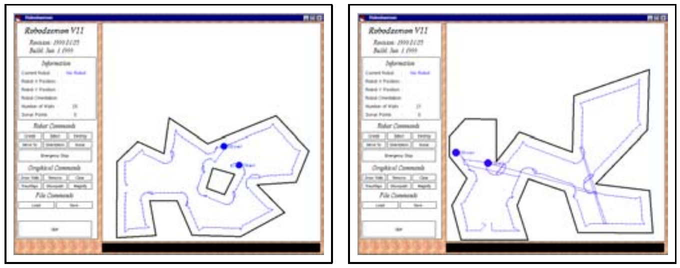

(a) (b)

Fig. 4. The paths of the two robots after the completion of the exploration

5 Discussion

The triangulation of the environment provides a complete decomposition of the free space; no overlapping occurs and no area is left uncovered. The dual graph acts as a guide during the exploration, and after the completion of mapping provides a topological map in addition to the metric map constructed.

The dual graph can be used for efficient path planning after the exploration phase. For any start and finish point, the corresponding triangles could be marked and different strategies could be applied to estimate a path. When the objective is to find the shortest path, a visibility graph could be calculated directly from the triangulation. Moreover, the embedded graph could contain additional information regarding the centroid of each triangle and the estimated odometry errors along this portion of the environment as recorded during the exploration. Thus, the robot path could be used to subsequently compute optimal paths in terms of either robustness or efficiency.

Our exploration/localization approach is especially useful when the internal motion estimates are completely unreliable, as is the case with many legged robots. In that case cooperative localization could provide a robust pose estimation thus allowing the robots to explore or navigate in an unknown environment despite that almost complete absence of accurate odometry.

6 Conclusions

In this paper, we have described an approach to exploring and navigating in *large scale spaces* where positioning and obstacle detection might be difficult using traditional methods. 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. A key requirement is that the environment be small enough so that the robots can see one another from any two points on its boundary that are not occluded from one another (i.e. they are never unable to see one another simply because they are too far away).

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 [11]. 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 [9]. 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 [10]) where they can be used effectively. Doing this efficiently appears feasible but has not been achieved yet.

References

1. Ivan A. Bachelder and Allen M. Waxman. A view-based neurocomputational system for relational map-making and navigation in visual environments. *Robotics & Autonomous 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. J. Borenstein, H R. Everett, L. Feng, and D. Wehe. Mobile robot positioning: sensors and techniques. *Journal of Robotic Systems*, 14(4):231–249, Apr 1997.
4. Gregory Dudek, Michael Jenkin, Evangelos Milios, and David Wilkes. Experiments in sensing and communication for robot convoy navigation. In *Proc.Int. Conf. on Intelligent Robots & Systems (IROS)*, v. 2, p. 268–273, 1995.
5. C. Lin and R. Tummala. Mobile robot navigation using artificial landmarks. *Journal of Robotic Systems*, 14(2):93–106, 1997.
6. F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
7. V. Lumelsky, S. Mukhopadhyay, and K. Sun. Dynamic path planning in sensor-based terrain acquisition. *IEEE Trans. on Robotics & Automation*, 6(4):462–472, 1990.
8. Paul MacKenzie and Gregory Dudek. Precise positioning using model-based maps. In *Proc. of the Int. Conf. on Robotics & Automation*, 1994.
9. Ioannis Rekleitis, Gregory Dudek, and Evangelos Milios. Multi-robot exploration of an unknown environment, efficiently reducing the odometry error. In *Int. Joint Conference in Artificial Intelligence*, v. 2, p. 1340–1345, 1997.
10. Robert Sim and Gregory Dudek. Learning visual landmarks for pose estimation. In *Proc. of Int. Conf. on Robotics & Automation (ICRA)*, May 1999.

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