Extended Abstract for the 8th International Symposium on Experimental Robotics (ISER '02)
Sant'Angelo d'Ischia, Italy, July 2002

Localization for Mobile Robot Teams: A Distributed MLE Approach

Andrew Howard and Maja J Matarić

Robotics Research Laboratories, Department of Computer Science, University of Southern California, Los Angeles, CA 90089-0781
ahoward@usc.edu, mataric@usc.edu

1 Introduction

This paper describes a method for localizing the members of a mobile robot team, using only the robots themselves as landmarks. That is, we describe a method whereby each robot can determine the relative range, bearing and orientation of every other robot in the team, without the use of GPS, external landmarks, or instrumentation of the environment. We also describe a distributed implementation of this method that has the potential to scale to large teams, and to be robust to the failure or destruction of individual robots.

Our approach is motivated by the need to localize robots in hostile and sometimes dynamic environments. Consider, for example, a search-and-rescue scenario in which a team of robots must deploy into a damaged structure, search for survivors, and guide rescuers to those survivors. In such environments, localization information cannot be obtained using GPS or landmark-based techniques: GPS is generally unavailable or unreliable due to signal obstructions or multi-path effects, while landmark-based techniques require prior models of the environment that are either unavailable, incomplete or inaccurate. The environment may also be undergoing dynamic structural changes that render such models invalid. In contrast, by using the robot themselves as landmarks, the method described in this paper can generate good localization information in almost any environment, including those that are undergoing structural changes. Our only requirement is that the robots are able to maintain at least intermittent line-of-sight contact with one-another. The distributed nature of the implementation also offers the possibility that this method may be robust to the failure individual of robots.

We make three basic assumptions. First, we assume that each robot is equipped with a proprioceptive *motion detector* such that it can measure changes in its own pose (subject to some degree of uncertainty). Suitable motion detectors can be constructed using either odometry or inertial measurement units. Second, we assume that each robot is equipped with a *robot detector* such that it can measure the relative pose and identity of nearby robots. Suitable sensors can be constructed using either vision (in combination with color-coded markers) or scanning laser range-finders (in combination with retro-reflective bar-codes). We further assume that the identity of robots is always determined correctly (which eliminates what would otherwise be a combinatoric labeling problem) but that there is some uncertainty in the relative pose measurements. Finally, we assume that each robot is equipped with some form of transceiver that can be used to broadcast messages to every other robot in the team. Standard IEEE 802.11b wireless network adapters can be used for this purpose.

Given these assumptions, the team localization problem can be solved using a combination of maximum likelihood estimation (MLE) and numerical optimization. The basic method is as follows. First, we construct a set of estimates $H = \{h\}$ in which each element h represents a pose estimate for a particular robot at a particular time. These pose estimates are defined with respect to some *arbitrary* global coordinate system. Second, we construct a set of observations $O = \{o\}$ in which each element o represents a relative pose measurement made by either a motion or robot detector. For motion detectors, each observation o represents the measured change in pose of a single robot; for robot detectors, each observation o represents the measured pose of one robot relative to another. Finally, we use numerical optimization to determine the set of estimates H that is most likely to give rise to the set of observations O.

Note that, in general, we do not expect robots to use the set of pose estimates H directly; these estimates are defined with respect to an arbitrary coordinate system whose relationship with the external environment is undefined. Instead, each robot uses these estimates to compute the pose of every other robot *relative to itself*, and uses this ego-centric viewpoint to coordinate activity. We note, however, that some subset of the team may choose to remain stationary, thereby ‘anchoring’ the global coordinate system in the real world. In this case, the pose estimates in H may be used as global coordinates in the standard fashion.

The localization method described above can be implemented in an entirely distributed manner. Basically, each robot is given responsibility for maintaining and optimizing its own pose estimates, while broadcast communication is used to ensure consistency between the pose estimates generated by different robots. In effect, this algorithm partitions the set H into N non-intersecting subsets (one for each robot), which are then optimized in parallel. The final result is equivalent to that obtained using a single centralized optimization algorithm.

In the full paper, we will describe both the MLE formulation and its distributed implementation, and discuss related work. In this extended abstract, we restrict ourselves to presenting the results of a controlled experiment conducted with a team of four mobile robots, and noting that details of the formalism can be found in [1].

2 Experiments

This section presents the results of a controlled experiment aimed at determining the accuracy of the distributed team localization algorithm. The experiment was conducted using a team of four Pioneer 2DX mobile robots equipped with SICK LMS200 scanning laser range-finders. Each robot was also equipped with a pair of retro-reflective ‘totem-poles’ as shown in Figure 1(a). These totem-poles can be detected from a wide range of angles using the SICK lasers (which can be

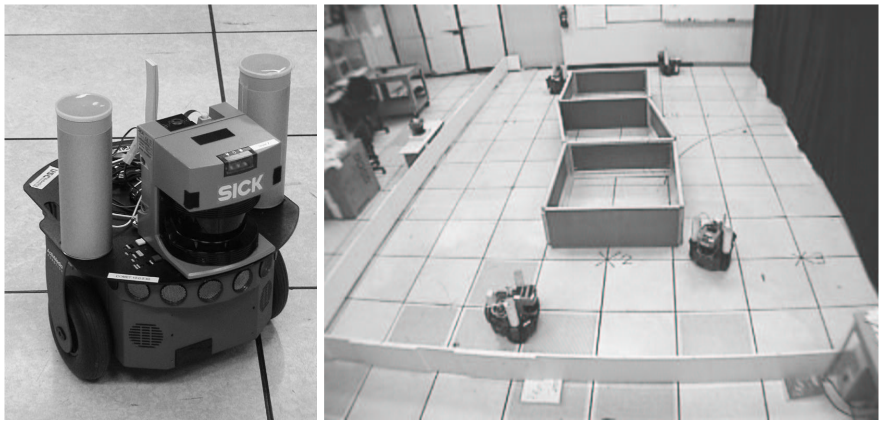

(a) (b)

Fig. 1. (a) A Pioneer 2DX equipped with a SICK LMS200 scanning laser range-finder and a pair of retro-reflective totem-poles. (b) The arena for the experiment; the central island is constructed from partitions that can be re-arranged during the course of the experiment.

programmed to return intensity information in addition to range measurments). This arrangement allows each robot to detect the presence of other robots and to determine both their range (to within a few centimeters) and bearing (to within a few degrees). Orientation can also be determined to within a few degrees, but is subject to a $180°$ ambiguity. This arrangement does *not* allow individual robots to be identified. Given the ambiguity in both orientation and identity, it was necessary, for this experiment, to group robot observations into 'tracks', which were then manually labelled.

The team was placed into the environment shown in Figure 1(c) and each robot executed a simple wall following algorithm. Two robots followed the inner wall, and two followed the outer wall. The robots were arranged such that at no time were the two robots on outer wall able to directly sense each other. Ground-truth information was provided by an external laser-based tracking system; the paths generated by this system are shown in Figure 2(a). Total duration of the experiment was 16 minutes.

The structure of the environment was modified a number of times during the course of the experiment. At time $t = 265$ sec, for example, the inner wall was modified to form two separate 'islands', with one robot circumnavigating each. The original structure was later restored, then broken, the restored again.

The accuracy of the algorithm was determined by comparing the robot's *relative pose estimates* with their corresponding true values (as determined by the ground-truth system). Thus, we define the *average range error* ϵ_r to be:

$$\epsilon_r^2(t) = \frac{1}{N(N-1)} \sum_{r_a} \sum_{r_b} (\hat{\mu}_r - \bar{\mu}_r)^2 \tag{1}$$

where μ is the estimated pose of robot r_b relative to robot r_a at time t, and $\bar{\mu}$ is the true pose of robot r_b relative to robot r_a at the same time. The summation is over all

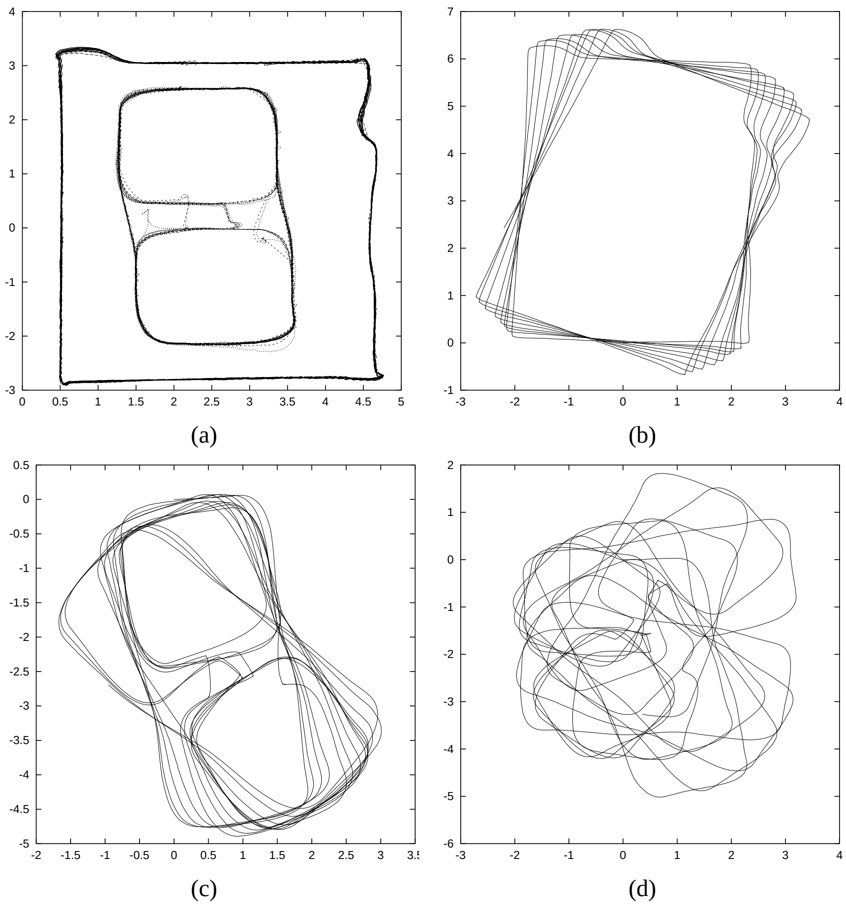

Fig. 2. (a) True paths for the four robots. Two robots follow the outer wall, two follow the inner wall. The inner wall is changed during the expierment, giving rise to the different paths seen in the figure. (b) Odometric path for robot ‘fly’, which follows the outer wall. Note the relatively slow drift. (c) Odometric path for robot ‘bee’, which follows the inner wall. (d) Odometric path for robot ‘bug’, which also follows the inner wall. Note the rapid drift: this robot has extremely poor odometry.

pairs of robots and the result is normalized by the number of robots N to generate an average result. One can define similar measures for the bearing error ϵ_ψ and orientation error ϵ_ϕ. Collectively, these error terms measure the average accuracy with which robots are able to determine each others relative pose. Note that we make no attempt to compare the *absolute* pose estimates $\{h\}$ against some ‘true’ value; these estimates are defined with respect to an arbitrary coordinate system which renders such comparison meaningless.

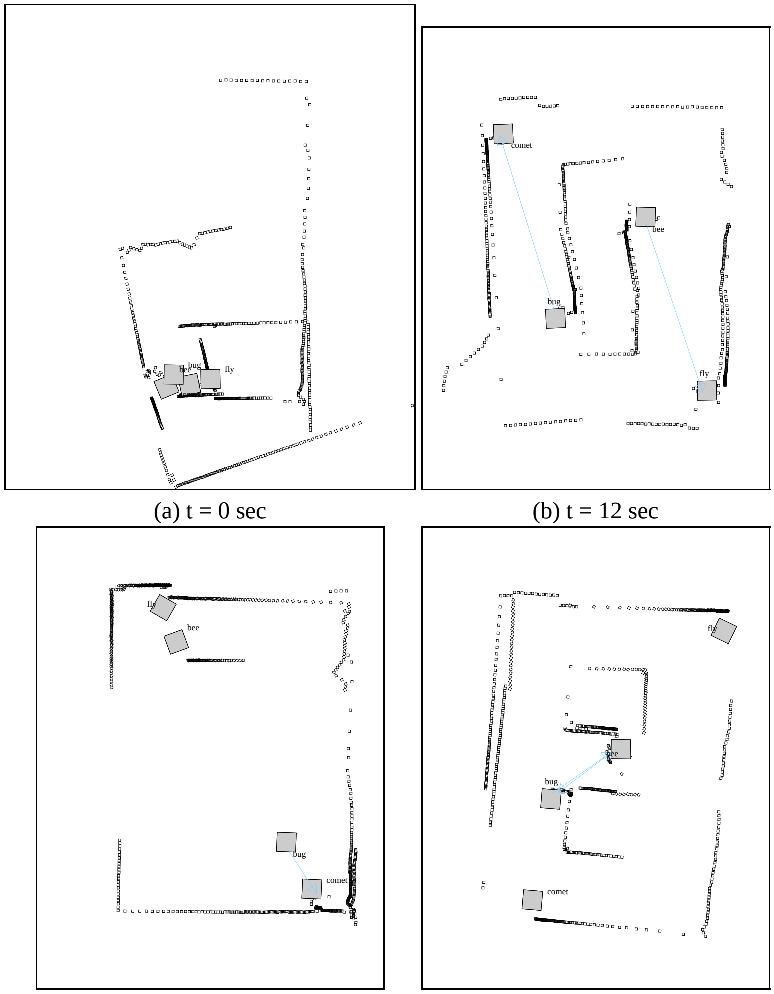

(a) t = 0 sec (b) t = 12 sec

(c) t = 61 sec (d) t = 272 sec

Fig. 3. Experimental snap-shots. Each sub-figure shows the estimated pose of the robots at a particular point in time, overlaid with the corresponding laser scan data. Arrows denote the observation of one robot by another. Note that these are snap-shots of live data; they are *not* cummulative maps of stored data.

2.1 Results

The qualitative results for this experiment are summarized in Figure 3, which contains a series of 'snap-shots' of the experiment. Each snap-shot shows the estimated pose of the robots at a particular point in time, overlaid with the corresponding laser scan data. Note that these are snap-shots of *live* data, not cumulative maps of *stored* data. At time $t = 0$, the relative pose of the robots is completely unknown, the snap-

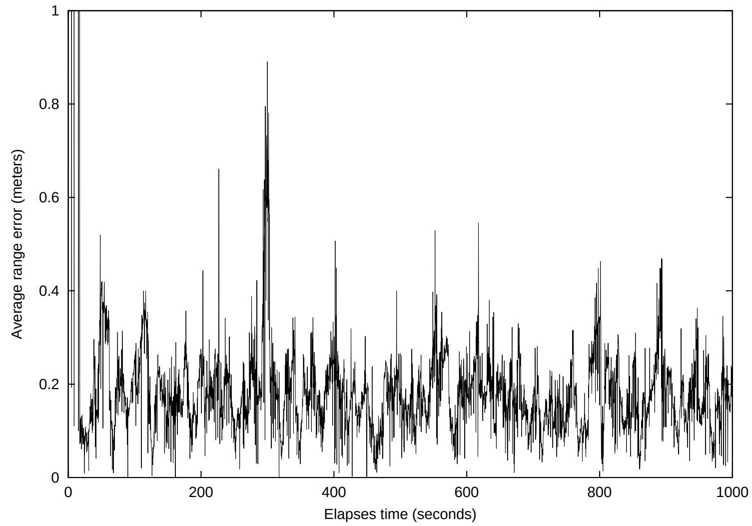

Fig. 4. Plot of the average range error ϵ_r as a function of time.

shot at this time is therefore incoherent; the pose of the robots is largely random, and the laser scan data is completely mis-aligned. In the interval $0 < t < 12$ sec, the robots commence wall following. The robots Fly and Comet follow the outer wall, whilst Bee and Bug follow the inner wall. By time $t = 12$ sec, both of the robots following the outer wall have observered both of the robots following the inner wall. As the snap-shot from this time indicates, there is now sufficient information to fully constrain the relative poses of the robots, and to correctly align the laser scan data. It should be that the two robots on the outer wall can correctly determine each others pose, even though they have never seen each other. At time $t = 265$ sec, the environment is modified, with the inner wall being re-structured to form two separate islands. The two robots following the inner wall now follow different paths, but the localization is un-affected, as shown in the snap-shot at time $t = 272$ sec. The algorithm described in this paper is completely indifferent to such structural changes in the environment.

The quantitative results for this experiment are summarized in Figure 4, which plots the average range error ϵ_r for the team. At time $t = 0$ sec, when the relative pose of the robots is completely unknown, this error is high. However, as the robots make observations, this error quickly falls. Overall, the eror tends to oscillate in the range 0.20 ± 0.15 m.

The variation seen in this plot can be ascribed to a number of factors. First, we expect that the error will rise during those periods in which the robots cannot see each other and localization is reliant on odometry alone. The odometry for this set of robots is, in fact, quite variable. Figure 2 shows the odometric paths generated by three of the four robots. The quality of the odometry ranges from very good (on robot Fly) to very bad (on robot Bug). It is our suspicion that most of the variation seen in the error plot can be ascribed to the poor odometry on latter robot. Second, we expect that the error will fall during those periods when robots are observing one another. This fall, however, may be proceeded by a 'spike' in the error term; this is spike is an artifact of the distributed optimization algorithm, which may take several communication cycles to generate self-consistent results. Finally, we note that there is an artifact in the plot at around time $t = 300$ sec. This artifact corresponds to a collision that occurred between robots Bee and Bug. As a result of this collision, the robots had to be manually re-positioned, leading to gross errors in both robots odometry. Nevertheless, the system quickly recovered.

3 Conclusion and Further Work

The experiment described in the previous section demonstrates several key capabilities of the team localization method described in this paper: this method does not require external landmarks, nor does it require that any of the robots remain stationary; it is robust to changes in the environment and to poor motion sensing; and robots can infer the pose of robots they have never seen. While the accuracy of the localization is not high, it is certainly good enough to facilitate many forms of coorperative behavior. It also should be noted that these are *preliminary* results, and that we expect to see improvements in accuracy as we 'tweak' the algorithm.

There remain many aspects of both the general method and of our distributed implementation of this method that require further experimental analysis. With regards to the method, we have not yet analyzed the impact of local minima (which necessarily plague any non-trivial numerical optimization problem). With regards to the distributed implementation, we are yet to measure how the algorithm scales with team size, although we suspect that both computation and bandwidth requirements will scale linearly.

Finally, we note that the mathematical formalism presented in this paper can be extended in a number of interesting directions. We can, for example, define a covariance matrix that measures the relative uncertainty in the pose estimates for pairs of robots. This matrix can then be used as a *signal* to actively control the behavior of robots. Thus, for example, if two robots need to cooperate, but their relative pose is not well know, they can undertake actions (such seeking out other robots) that will reduce this uncertainty.

References

1. A. Howard, M. J. Matarić, and G. S. Sukhatme. Localization for mobile robot teams: A maximum likelihood approach. Technical Report IRIS-01-407, Institute for Robotics and Intelligent Systems Technical Report, University of Southern California, 2001.