Localization for Mobile Robot Teams: A Maximum Likelihood Approach

Andrew Howard, Maja J Matarić and **Gaurav S Sukhatme**

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

Abstract

This paper describes a method for localizing the members of a mobile robot team, using only the robots themselves as landmarks. We assume that robots are equipped with sensors that allow them to measure the relative pose and identity of nearby robots, as well as sensors that allow them to measure changes in their own pose. Using a combination of maximum likelihood estimation and numerical optimization, we can, for each robot, estimate the relative range, bearing and orientation of every other robot in the team. This paper describes the basic formalism and presents experimental results to validate the approach.

Introduction

This papers 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, landmarks, or instrumentation of the environment.

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 this scenario, localization information cannot be obtained using GPS or landmark-based techniques; GPS is generally unavailable or unreliable in urban environments due to multipath effects, while landmark-based techniques require prior models of the environment that are either unavailable, incomplete or inaccurate. For these reasons, we have developed an approach to localization that relies on using the robots *themselves* as landmarks. With this approach, one can obtain good localization information in almost any environment, including those that are undergoing dynamic structural changes. Our only requirement is that the robots are able to maintain at least intermittent line-of-sight contact with other robots in the team.

We make two 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 (sub-

ject 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 of nearby robots and determine their identity. Suitable sensors can be readily 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.

Given these assumptions, the team localization problem can be solved using maximum likelihood estimation. 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.

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.

In this paper, we develop the mathematical machinery required to solve the team localization problem in any number of dimensions, then treat the more specific problem of localization in a plane. We also describe a series of experiments, using both real and simulated robots, aimed at verifying both the accuracy and practicality of this approach.

Related Work

Localization is an extremely well studied area in mobile robotics. The vast majority of this research has concentrated on two problems: localizing a single robot using an a priori map of the environment (Leonard & Durrant-Whyte 1991; Simmons & Koenig 1995; Fox, Burgard, & Thrun 1999), or localizing a single robot whilst simultaneously building a map (Thrun, Fox, & Burgard 1998; Lu & Milios 1997; Yamauchi, Shultz, & Adams 1998; Duckett, Marsland, & Shapiro 2000; Golfarelli, Maio, & Rizzi 1998; Dissanayake *et al.* 2001). Recently, some authors have also considered the related problem of map building with multiple robots (Thrun 2001). All of these authors make use of statistical or probabilistic techniques of one sort or another; the common tools of choice are Kalman filters, maximum likelihood estimation, expectation maximization, and Markovian techniques (using grid or sample-based representations for probability distributions). The team localization problem described in this paper bears many similarities to the simultaneous localization and map building problem, and is amenable to similar mathematical treatments. In this context, the work of Lu and Milios (Lu & Milios 1997) should be noted. These authors describe a method for constructing globally consistent maps by enforcing pairwise geometric relationships between individual range scans; relationships are derived either from odometry, or from the comparison of range scan pairs. MLE is used to determine the set of pose estimates that best accounts this set of relationships. Our mathematical formalism is very similar to that described by these authors, even thought it is directed towards a somewhat different objective; i.e., the localization of mobile robot teams, rather than the construction of globally consistent maps.

Among those who have considered the specific problem of team localization are (Roumeliotis & Bekey 2000) and (Fox *et al.* 2000). Roumeliotis and Bekey present an approach to multi-robot localization in which sensor data from a heterogeneous collection of robots is combined through a single Kalman filter to estimate the pose of each robot in the team. They also show how this centralized Kalman filter can be broken down into N separate Kalman filters (one for each robot) to allow for distributed processing. It should be noted, however, that this method still relies entirely on external landmarks; no attempt is made to sense other robots or to use this information to constrain the pose estimates. In contrast, Fox *et al.* describe an approach to multi-robot localization in which each robot maintains a probability distribution describing its own pose (based on odometry and environment sensing), but is able to refine this distribution through the observation of other robots. This approach extends earlier work on single-robot probabilistic localization techniques (Fox, Burgard, & Thrun 1999). The authors avoid the curse of dimensionality (for N robots, one must maintain a $3N$ dimensional distribution) by factoring the distribution into N separate components (one for each robot). While this step makes the algorithm tractable, it does result in some loss of expressiveness; one cannot, for example, express relationships of the form: "if I am at A then you must be at C, but if I am at B you must be at D". Our approach is able pre-
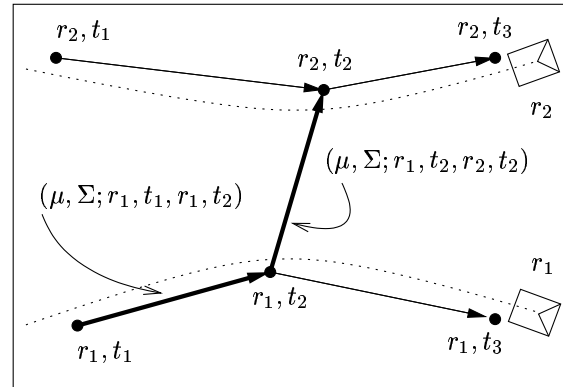

Figure 1: Illustration of the basic formalism. The figure shows two robots, r_1 and r_2, traveling from left to right; at time t_2, robot r_1 observes robot r_2. The nodes represent pose estimates; the arcs represent observations. Two observations are highlighted: a motion observation for robot r_1, and robot observation at time t_2.

serve such relationships by treating the localization problem in its full $3N$ dimensional form. This is feasible only because MLE is a *single* estimate approach; i.e., there is no attempt to maintain a complete probability distribution.

Finally, a number of authors (Kurazume & Hirose 2000; Rekleitis, Dudek, & Milios 1997; Howard & Kitchen 1999) have considered the problem of team localization from a somewhat different perspective. These authors describe *co-operative* approaches to localization, in which team members actively coordinate their activities in order to reduce cumulative odometric errors. The basic method is to keep one subset of the robots stationary, while the other robots are in motion; the stationary robots observe the robots in motion (or vice-versa), thereby providing more accurate pose estimates than can be obtained using odometry alone. While our approach does not require such explicit cooperation on the part of robots, the accuracy of localization can certainly be improved by the adoption of such strategies; we will return to this topic briefly in the final sections of this paper.

Formalism

We formulate the team localization problem as follows. Let h denote the pose estimate for a particular robot at a particular time, and let $H = \{h\}$ be the set of all such estimates. Similarly, let o denote an observation made by some detector, and let $O = \{o\}$ be the set of all such observations. Our aim is to determine the set of estimates H that maximizes the probability of obtaining the set of observations O; i.e., we seek to maximize the conditional probability $P(O \mid H)$. If we assume that observations are statistically independent, we can write this probability as:

$$P(O \mid H) = \prod_{o \in O} P(o \mid H) \qquad (1)$$

where $P(o \mid H)$ is the probability of obtaining the individual measurement o, given the estimates H. Taking the log of

both sides, we can rewrite this equation as:

$$U(O \mid H) = \sum_{o \in O} U(o \mid H) \tag{2}$$

where $U(O \mid H) = -\log P(O \mid H)$ and $U(o \mid H) = -\log P(o \mid H)$. This latter form is somewhat more efficient for numerical optimization. Our aim is now to find the set of estimates H that *minimizes* $U(O \mid H)$. To do this, we must determine the form of the individual observation probabilities $P(o \mid H)$, or their log-probability equivalents $U(o \mid H)$.

We make the following definitions. Let each estimate h be denote by a tuple of the form:

$$h = (\hat{q}; r, t) \tag{3}$$

where $\hat{q}$ is the *absolute pose estimate* for robot r at time t. Note that it is the value of $\hat{q}$ that we are trying to estimate; r and t are constants used for book-keeping purposes only. Let each observation o be denoted by a tuple of the form:

$$o = (\mu, \Sigma; r_a, t_a; r_b, t_b) \tag{4}$$

where μ is the measured pose of robot r_b at time t_b, relative to robot r_a at time t_a; henceforth, we will refer to μ as a *relative pose measurement*. The Σ term is a covariance matrix representing the uncertainty in this measurement. As stated in the Introduction, each robot is equipped with two different detectors: a proprioceptive motion detector, which allows the robot to measure changes in its own pose, and a robot detector, which allows it to measure the identity and relative pose of other robots. Data from the motion detectors are encoded using an observation of the form:

$$o = (\mu, \Sigma; r_a, t_a, r_a, t_b) \tag{5}$$

where μ is the measured change in pose for robot r_a between times t_a and t_b. Data from the robot detectors is encoded using an observation of the form:

$$o = (\mu, \Sigma; r_a, t_a, r_b, t_a) \tag{6}$$

where μ is the measured pose of robot r_b relative to robot r_a, for a measurement taken at time t_a.

If we assume that the measurement uncertainty for all detectors follows a normal distribution, the conditional log-probability $U(o \mid H)$ is given by the quadratic expression:

$$U(o \mid H) = \frac{1}{2}(\mu - \hat{\mu})^T \Sigma (\mu - \hat{\mu}) \tag{7}$$

where μ is the relative pose measurement defined above, and $\hat{\mu}$ is the corresponding *relative pose estimate*; i.e. $\hat{\mu}$ is the estimated pose of robot r_b at time t_b, relative to robot r_a at time t_a. Let $\hat{q}_a$ and $\hat{q}_b$ describe the absolute pose estimates for robot r_a at time t_a, and robot r_b at time t_b, respectively. The relative pose estimate $\hat{\mu}$ is derived from these absolute pose estimates via a simple coordinate transformation Γ:

$$\hat{\mu} = \Gamma(\hat{q}_a, \hat{q}_b) \tag{8}$$

The specific form of Γ depends on the dimensionality of the localization problem (e.g. 2D versus 3D) and on the particular representation chosen for both absolute and relative poses (e.g. Cartesian versus polar coordinates, or cylindrical versus spherical coordinates).

One can visualize this formalism in terms of a directed graph, as shown in Figure 1. We associate each estimate h with a node in the graph, and each observation o with an arc. Each node may have both outgoing arcs, corresponding to observations in which this node was the *observer*, and incoming arcs, corresponding to observations in which this node was the *observee*. Motion observations join nodes representing the same robot at two different points in time, while robot observations join nodes representing two different robots at the same point in time, as indicated in the figure.

Numerical Optimization

Given Equations 2 and 7, together with an appropriate definition for Γ, one can determine the set of poses $\hat{q}$ that minimizes $U(O \mid H)$ using a standard numerical optimization algorithm. The selection of an appropriate algorithm is driven largely by the form of Γ: in general, Γ is non-linear but differentiable, which rules out fast linear algorithms, but permits non-linear gradient-based algorithms (such as steepest descent). The gradient is computed by applying the chain-rule to Equation 2:

$$\frac{\partial}{\partial h} U(O \mid H) = \sum_{o \in O} \frac{\partial \hat{\mu}}{\partial h} \frac{\partial}{\partial \hat{\mu}} U(o \mid H) \tag{9}$$

where the second term (a gradient vector) is computed trivially from Equation 7:

$$\frac{\partial}{\partial \hat{\mu}} U(o \mid H) = -\Sigma(\mu - \hat{\mu}) \tag{10}$$

and the first term (a Jacobian matrix) is computed by differentiating through the Γ function (whatever it may be):

$$\frac{\partial \hat{\mu}}{\partial h} = \begin{cases} \partial \hat{\mu}/\partial \hat{q}_a & \text{if } r = r_a \text{ and } t = t_a \\ \partial \hat{\mu}/\partial \hat{q}_b & \text{if } r = r_b \text{ and } t = t_b \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

This latter equation requires some explanation. Consider once again the directed-graph interpretation of the formalism depicted in Figure 1. Each estimate h is represented by a single node, and each observation o is represented by an arc joining two nodes. The relative pose estimate $\hat{\mu}$ measures the pose of one of these nodes relative to the other. In computing the derivative $\partial \hat{\mu}/\partial h$ for some particular observation, there are three cases to consider: h is the robot making the observation, in which case we compute the derivative with respect to q_a; h is the robot being observed, in which case we compute the derivative with respect to q_b; or else h is neither the observer nor the observee in this particular observation, in which case the derivative must be zero.

We use a conjugate gradient algorithm (Press *et al.* 1999) for optimization. This algorithm is somewhat more complex than a steepest descent algorithm, but has the advantage of being much quicker. In addition, unlike some algorithms, its memory requirements scale linearly, rather than quadratically, with the number of variables being optimized.
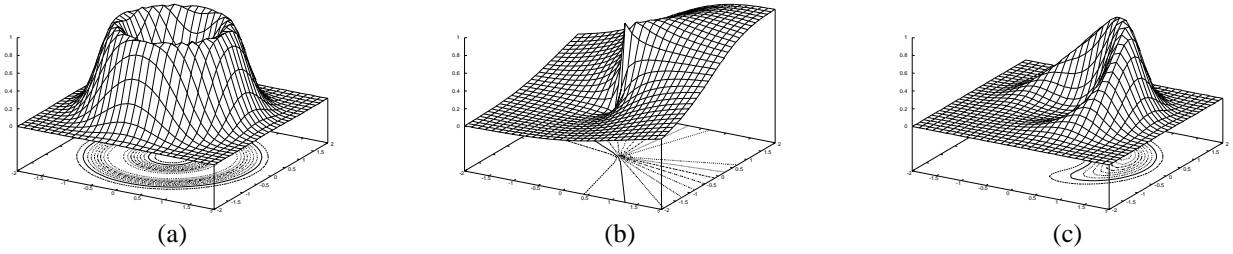
(a)   (b)   (c)

Figure 3: Sample probability distributions for for the planar localization problem. The plots show the probability $P(o \mid H)$ as a function of the estimated absolute position $(\hat{q}_x, \hat{q}_y)$ for the robot being observed. Orientation is not shown. (a) The range is well determined ($\mu_r = 1 \pm 0.1$ m), but the bearing is unknown. (b) The bearing is moderately well determined ($\mu_\phi = 45 \pm 57°$), but the range in unknown. (c) Both range and bearing are well determined ($\mu_r = 1 \pm 0.1$ m, $\mu_\phi = 45 \pm 57°$).

replacements
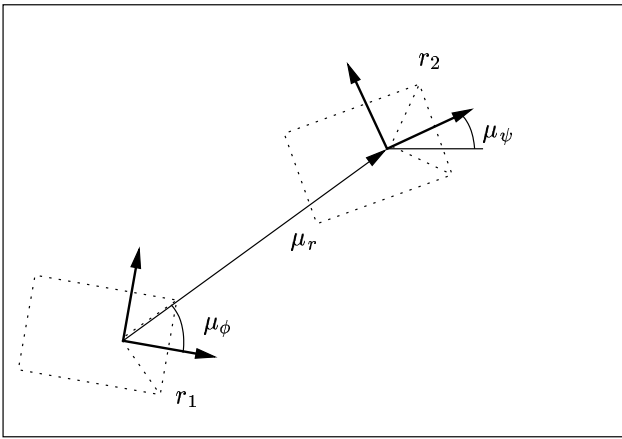
Figure 2: Example of a relative pose measurement for the planar localization problem. Robot r_1 measures the range μ_r, bearing μ_ϕ and orientation μ_ψ of robot r_2.

Localization in a Plane

The formalism described above is quite general, and can be applied to localization problems in two, three, or more dimensions. In order to make the formalism somewhat more concrete, and to lay the theoretical foundations for the experiments described in the next section, we now consider the specific problem of localization in a plane.

We make the following definitions. Let the absolute pose estimate $\hat{q}$ be a 3-vector such that:

$$\hat{q} = [q_x, q_y, q_\theta]^T \tag{12}$$

where q_x and q_y describe the robot's position and q_θ describes its orientation. Let the relative pose measurement μ be a 3-vector such that:

$$\mu = [\mu_r, \mu_\phi, \mu_\psi]^T \tag{13}$$

where μ_r, μ_ϕ and μ_ψ are the range, bearing and orientation of one robot relative to another (or of one robot relative to its earlier pose). Figure 2 illustrates this definition. We choose to express measurements in this particular form, since, for many sensors, the uncertainty in range, bearing and orientation components is effectively uncorrelated. Thus, we can

ignore the off-diagonal terms in the uncertainty matrix Σ, and write:

$$\Sigma = \begin{bmatrix} 1/\sigma_r^2 & 0 & 0 \\ 0 & 1/\sigma_\phi^2 & 0 \\ 0 & 0 & 1/\sigma_\psi^2 \end{bmatrix} \tag{14}$$

Figure 3 shows some of the probability distributions that can be generated using this parameterization. The corresponding coordinate transform function $\Gamma(\hat{q}_a, \hat{q}_b)$ computes the range, bearing and orientation of $\hat{q}_b$ relative to $\hat{q}_a$. Using elementary geometry, we write down the following expression for the relative pose estimate $\hat{\mu}$:

$$
\hat{\mu} = \begin{bmatrix} \hat{\mu}_r \\ \hat{\mu}_\phi \\ \hat{\mu}_\psi \end{bmatrix}
$$

$$
= \begin{bmatrix} \sqrt{(\hat{q}_{bx} - \hat{q}_{ax})^2 + (\hat{q}_{by} - \hat{q}_{ay})^2} \\ \arctan[(\hat{q}_{by} - \hat{q}_{ay})/(\hat{q}_{bx} - \hat{q}_{ax})] - \hat{q}_{a\theta} \\ \hat{q}_{b\theta} - \hat{q}_{a\theta} \end{bmatrix}
$$

for which the corresponding derivatives are:

$$
\frac{\partial \hat{\mu}}{\partial \hat{q}_a} = \begin{bmatrix} \partial\hat{\mu}_r/\partial\hat{q}_{ax} & \partial\hat{\mu}_\phi/\partial\hat{q}_{ax} & \partial\hat{\mu}_\psi/\partial\hat{q}_{ax} \\ \partial\hat{\mu}_r/\partial\hat{q}_{ay} & \partial\hat{\mu}_\phi/\partial\hat{q}_{ay} & \partial\hat{\mu}_\psi/\partial\hat{q}_{ay} \\ \partial\hat{\mu}_r/\partial\hat{q}_{a\theta} & \partial\hat{\mu}_\phi/\partial\hat{q}_{a\theta} & \partial\hat{\mu}_\psi/\partial\hat{q}_{a\theta} \end{bmatrix}
$$

$$
= \begin{bmatrix} -(\hat{q}_{bx} - \hat{q}_{ax})/\hat{\mu}_r & +(\hat{q}_{by} - \hat{q}_{ay})/\hat{\mu}_r^2 & 0 \\ -(\hat{q}_{by} - \hat{q}_{ay})/\hat{\mu}_r & -(\hat{q}_{bx} - \hat{q}_{ax})/\hat{\mu}_r^2 & 0 \\ 0 & -1 & -1 \end{bmatrix}
$$

and

$$
\frac{\partial \hat{\mu}}{\partial \hat{q}_b} = \begin{bmatrix} \partial\hat{\mu}_r/\partial\hat{q}_{bx} & \partial\hat{\mu}_\phi/\partial\hat{q}_{bx} & \partial\hat{\mu}_\psi/\partial\hat{q}_{bx} \\ \partial\hat{\mu}_r/\partial\hat{q}_{by} & \partial\hat{\mu}_\phi/\partial\hat{q}_{by} & \partial\hat{\mu}_\psi/\partial\hat{q}_{by} \\ \partial\hat{\mu}_r/\partial\hat{q}_{b\theta} & \partial\hat{\mu}_\phi/\partial\hat{q}_{b\theta} & \partial\hat{\mu}_\psi/\partial\hat{q}_{b\theta} \end{bmatrix}
$$

$$
= \begin{bmatrix} +(\hat{q}_{bx} - \hat{q}_{ax})/\hat{\mu}_r & -(\hat{q}_{by} - \hat{q}_{ay})/\hat{\mu}_r^2 & 0 \\ +(\hat{q}_{by} - \hat{q}_{ay})/\hat{\mu}_r & +(\hat{q}_{bx} - \hat{q}_{ax})/\hat{\mu}_r^2 & 0 \\ 0 & 0 & +1 \end{bmatrix}
$$

Note that $\partial\hat{\mu}/\partial\hat{q}_a \neq -\partial\hat{\mu}/\partial\hat{q}_b$ as one might naively expect. Note also that the derivatives contain a singularity at $\hat{\mu} = 0$; one must take care to avoid this singularity in the numerical optimization process.

Inserting these definitions into the general formalism described in the previous section, one can solve the planar localization problem in a fairly straight-forward manner.
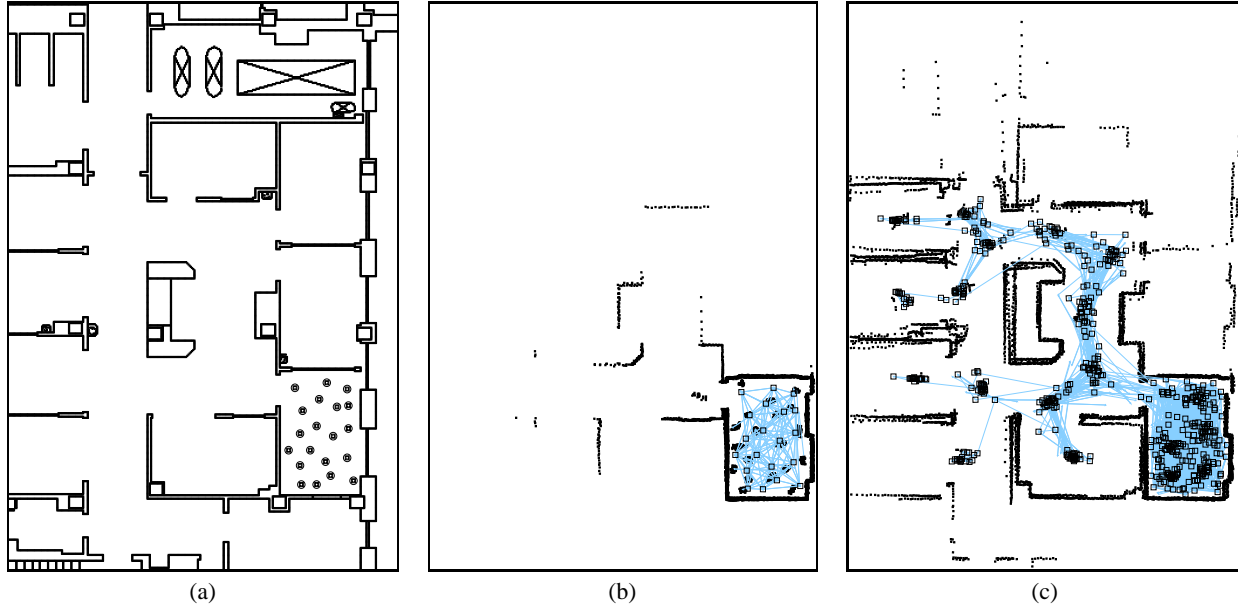
Figure 4: An experiment with simulated robots. (a) Experimental set-up: 20 robots are placed into a single room and allowed to disperse. (b) Combined laser scans at $t = 0$ sec, before any of the robots has moved. Note that this is *not* a stored map: this is live laser data. Pose estimates and observations are also shown, denoted by rectangles and lines respectively. (c) Combined laser scans at $t = 300$ sec, after all of the robots have dispersed.

Experiments

We have conducted a series of experiments aimed at determining both the accuracy and practicality of the approach described in this paper. Two such experiments are presented here. The first experiment was conducted in simulation, using a team of 20 robots performing a deployment task. The second experiment was conducted using a team of 7 real robots performing a simple navigation task. The first of these experiments was chosen to verify that the approach will work for relatively large teams; the second was chosen to verify that the approach will work with real sensor data.

For both experiments, we determine the accuracy of the solution by comparing a subset of relative pose estimates with their corresponding 'true' values. We measure accuracy using *relative* rather than *absolute* pose estimates, since the absolute pose estimates are defined with respect to an arbitrary coordinate system, and hence cannot be meaningfully compared with a 'true' value. We define the *average range error* ϵ_r as follows:

$$\epsilon_r^2 = \frac{1}{N(N-1)} \sum_{h_a \in H'} \sum_{h_b \in H'} (\hat{\mu}_r - \bar{\mu}_r)^2 \qquad (15)$$

where $\bar{\mu}$ is the true relative pose, i.e. the true pose of robot r_b at time t_b, relative to robot r_a at time t_a. The summation is over the subset of H' of H for which we have the true values, and the result is normalized by the number of pose estimates N in H' to generate an average result. One can define similar measures for the bearing error ϵ_ψ and orientation error ϵ_ϕ.

For practical purposes, it is necessary to limit both the number of pose estimates in H and the number of obser-
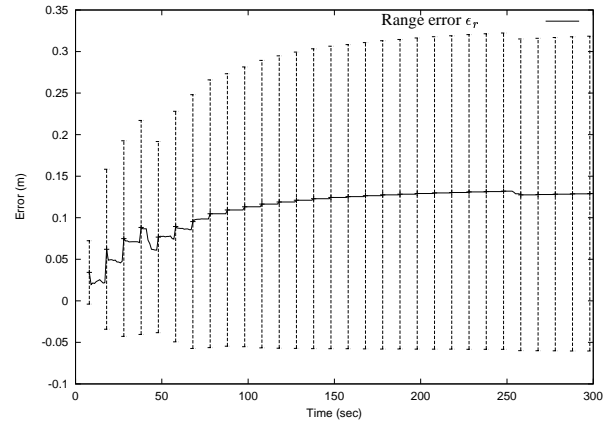

Figure 6: Results for the experiment with simulated robots. Plot of the average range error ϵ_r over the duration of the experiment; the bars denote the standard deviation in ϵ_r.

vations in O. For these experiments, this was done by integrating motion observations over a 10 second period, and discarding most of the robot observations. Thus, in these experiments, H describes the pose of robots at $t = (0, 10, 20, ...)$ sec, and O includes only those observations that occurred at these times.

Experiment with Simulated Robots

The simulated experiment was conducted using the *Stage* (Vaughan 2000) multi-agent simulator. Stage is a high-fidelity simulation capable of accurately mimicking the be-
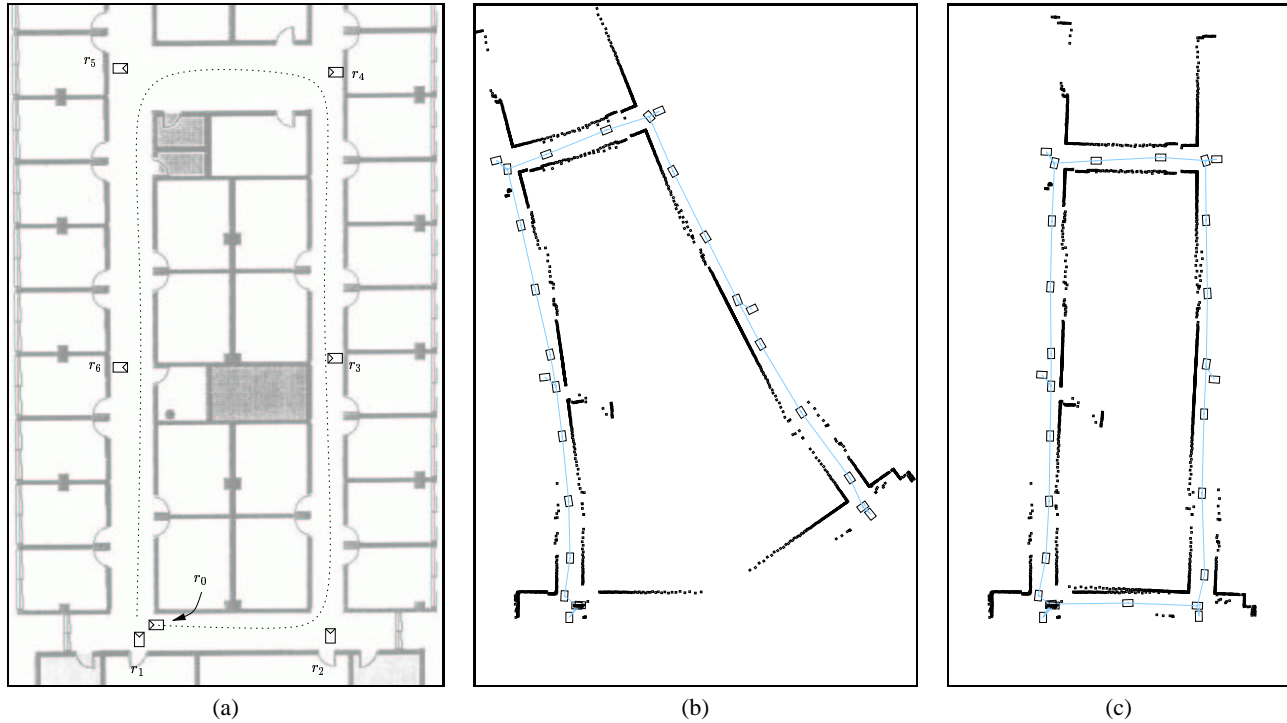
Figure 5: An experiment with real robots. (a) Experimental set-up: six stationary robots (r_1 to r_6) are placed at at strategic locations; the seventh mobile robot (r_0) executes a circuit. (b) Combined laser scans at $t = 200$ sec, after the mobile robot has been seen by all six stationary robots exactly once. Note that this is *not* a stored map: this is live laser data. Pose estimates and observations are also shown, denoted by rectangles and lines respectively. (c) Combined laser scans at $t = 220$ sec, after the mobile robot has been seen by the first stationary robot r_1 for a second time, thus closing the loop.

havior of many real robot sensors and actuators. For this experiment, we used a team of 20 robots, each equipped with odometry, a scanning laser range-finder and a retro-reflective bar-code. The simulated odometry provides motion observations and the laser range-finder/bar-code combination provides robot observations. For added realism, Gaussian noise was added to all measurement (proportional noise of 5% for the range component and a constant $2°$ noise for the bearing and orientation components). All 20 robots were initially positioned in single room in large environment, as shown in Figure 4(a), and subsequently allowed to disperse using the distributed deployment algorithm described in (Howard, Matarić, & Sukhatme 2002).

The quantitative results for this experiment are summarized in Figure 6, which plots both the average range error ϵ_r and its variance as a function of time. Note that the initial error is very low, at around 0.05 ± 0.03 m, but steadily increases over time to around 0.10 ± 0.20 m. This behavior is not surprising, given the way in which the robots disperse. At $t = 0$ sec, all of the robots are crowded into a single room, and many robot observations are generated. One can get a sense for the density of these observations by inspecting Figure 4(b), which shows a plot of both pose estimates and observations, with live laser scan data overlaid. The density of observations is such that the pose of robots is very heavily constrained, and the localization algorithm is able to generate a very accurate set of estimates. As the robots

disperse, however, errors begin to accumulate, and the accuracy of the estimates declines; the variance also increases dramatically. Again, this is not surprising. Any given pair of robots will always be connected by a series of motion and robot observations, and the cumulative uncertainty in these observations is such that the relative pose error for widely separated robots is necessarily greater than that for nearby robots (at least in absolute terms). We note that for most practical applications, knowing the pose of distant robots is much less important than knowing the pose of nearby robots (with whom we are much more likely to interact).

This experiment clearly demonstrates our ability to handle a relatively large robot team and to generate suitably accurate results. At the termination of the experiment, the H and O had 620 and 3058 elements respectively, yet these results were generated in real-time using a standard workstation.

Experiment with Real Robots

The real experiment was conducted using a team of 7 Pioneer 2DX mobile robots, each of which is equipped with a SICK LMS 200 scanning laser range-finder and a retro-reflective bar-code. Motion observations are provided by the robot's on-board odometry; robot observations are provided by the laser range-finder/bar-code combination. For this experiment, 6 of the 7 robots were positioned at fixed locations in the corridors of a building, as shown in Figure 5(a); the
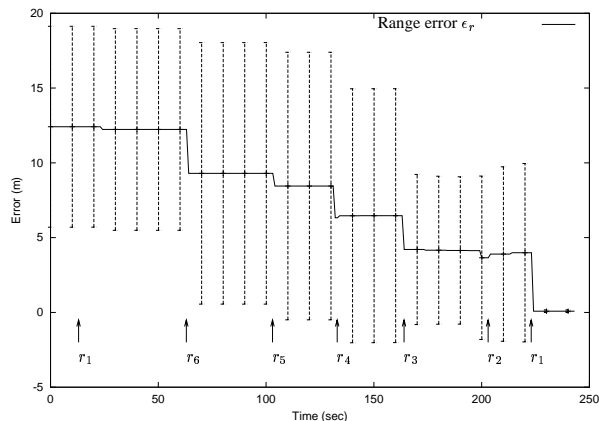
Figure 7: Results for the experiment with real robots. Plot of the average range error ϵ_r over the duration of the experiment; the bars denote the standard deviation in ϵ_r.

remaining robot was then 'joy-sticked' around the circuit, and was thus 'seen' by each of the stationary robots in turn. Note that the stationary robots were positioned outside each other's sensor range, and hence there are no observations that relate the stationary robots directly.

Since these experiments were performed in an un-instrumented environment, ground truth information was obtained by measuring the inter-robot distances between the fixed robots (using a tape-measure). Bearings and orientations were not measured. Our error calculations therefore include only these known values.

The quantitative results for this experiment are shown in Figure 7, which shows a plot of the error ϵ_r as a function of time for this experiment. Also marked on this plot are the times at which each of the stationary robots r_1 to r_6 observed the mobile robot r_0. The most striking feature of this plot is the way in which the error drops immediately after each robot observation. This is to be expected, given that the pose of the stationary robots can only be determined after they have seen the mobile robot at least once. Figure 5(b) shows a plot of the pose estimates, observations, and laser scan data at time $t = 200$ sec. At this point, the mobile robot has been seen by each of the stationary robots exactly once. However, due to the cumulative error in this robot's motion observations, the overall error of the pose estimates remains relatively high. Figure 5(c) shows the same plot at $t = 230$ sec, after the mobile robot has 'closed the loop' by revisiting the first stationary robot. At this point, the error drops dramatically, reaching a final value of 0.08 ± 0.09 m. This figure is quite remarkable when one considers that the loop traversed by the mobile robot is about 80m in length.

While these results serve to verify the accuracy of our approach when applied to real data, they also suggest that, for some applications, and in some environments, good localization can only be achieved through deliberate action on the part of team members (looking for other robots, closing loops, and so on). This raises an interesting set of issues that are, unfortunately, beyond the scope of this paper.

Conclusion and Future Work

The experiments described in the previous section suffice to demonstrate the MLE approach to team localization, and to verify its accuracy in both simulated and real experimental contexts. There are, however, some aspects of this approach that require further experimental investigation. Foremost among these is the impact of local minima, which necessarily plague any non-trivial numerical optimization problem. While the solutions found in the previous section are entirely satisfactory, more experiments need to be done to characterize the overall sensitivity of the approach.

In this paper, we have intentionally omitted many details regarding the practical implementation of the formalism. We note that while it fairly easy to construct a batch-processing algorithm for use off-line, the construction an any-time algorithm suitable for use on-line or 'in-the-loop' is somewhat more involved (and requires some extensions to the formalism). We expect to present details of such an algorithm in the near future.

The mathematical formalism presented in this paper can also be extended and developed 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. If, for example, 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. We are also working on a distributed version of the formalism that will allow a team of robots to collectively localize themselves without the need for any form of centralized computation. We believe that this distributed formalism can be implemented using a constant-time, constant-bandwidth algorithm, and will therefore scale to teams of any size.

References

Dissanayake, M. W. M. G.; Newman, P.; Clark, S.; Durrant-Whyte, H. F.; and Csorba, M. 2001. A solution to the simultaneous localization and map building (slam) problem. *IEEE Transactions on Robotics and Automation* 17(3):229–241.

Duckett, T.; Marsland, S.; and Shapiro, J. 2000. Learning globally consistent maps by relaxation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 4, 3841–3846.

Fox, D.; Burgard, W.; Kruppa, H.; and Thrun, S. 2000. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots, Special Issue on Heterogeneous Multi-Robot Systems* 8(3):325–344.

Fox, D.; Burgard, W.; and Thrun, S. 1999. Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research* 11:391–427.

Golfarelli, M.; Maio, D.; and Rizzi, S. 1998. Elastic correction of dead reckoning errors in map building. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, 905–911.

Howard, A., and Kitchen, L. 1999. Cooperative localisation and mapping. In *International Conference on Field and Service Robotics (FSR99)*, 92–97.

Howard, A.; Matarić, M. J.; and Sukhatme, G. S. 2002. Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem. In *Submitted to the 6th International Conference on Distributed Autonomous Robotic Systems (DARS02)*, submitted.

Kurazume, R., and Hirose, S. 2000. An experimental study of a cooperative positioning system. *Autonomous Robots* 8(1):43–52.

Leonard, J. J., and Durrant-Whyte, H. F. 1991. Mobile robot localization by tracking geometric beacons. *IEEE Transactions on Robotics and Automation* 7(3):376–382.

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

Press, W. H.; Teukolsky, S. A.; Vetterling, W. T.; and Flannery, B. P. 1999. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 2nd edition.

Rekleitis, I. M.; Dudek, G.; and Milios, E. 1997. Multi-robot exploration of an unknown environment: efficiently reducing the odometry error. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, volume 2, 1340–1345.

Roumeliotis, S. I., and Bekey, G. A. 2000. Collective localization: a distributed kalman filter approach. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, 2958–2965.

Simmons, R., and Koenig, S. 1995. Probabilistic navigation in partially observable environments. In *Proceedings of International Joint Conference on Artificial Intelligence*, volume 2, 1080–1087.

Thrun, S.; Fox, D.; and Burgard, W. 1998. A probabilistic approach to concurrent mapping and localisation for mobile robots. *Machine Learning* 31(5):29–55. Joint issue with Autonomous Robots.

Thrun, S. 2001. A probabilistic online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research* 20(5):335–363.

Vaughan, R. T. 2000. Stage: a multiple robot simulator. Technical Report IRIS-00-393, Institute for Robotics and Intelligent Systems, University of Southern California.

Yamauchi, B.; Shultz, A.; and Adams, W. 1998. Mobile robot exploration and map-building with continuous localization. In *Proceedings of the 1998 IEEE/RSJ International Conference on Robotics and Automation*, volume 4, 3175–3720.