

Team Localization: A Maximum Likelihood Approach

Andrew Howard, Maja J Matarić and Gaurav S Sukhatme
 Robotics Research Laboratory, Computer Science Department
 University of Southern California, Los Angeles, California, 90089, U.S.A
 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. 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 assume that robots are equipped with proprioceptive motion sensors (such as odometry or inertial measurement units), together with sensors that will allow them to make occasional measurements of the relative pose and identity of nearby robots (such sensors can be constructed using cameras or scanning laser range-finders). Through a combination of maximum likelihood estimation and numerical optimization, we can consequently infer the relative pose of every robot in the team at any given point in time. This paper describes the basic team localization formalism, its practical implementation, and presents experimental results obtained using real robots in both static and dynamic environments.

Keywords— Mobile robots, cooperative localization, maximum likelihood estimation.

I. 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. Our approach is motivated by the need to localize robots in hostile and 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 occlusions and multi-path effects, while landmark-based techniques require prior models of the environment that are either unavailable, incomplete or inaccurate. In contrast, the method described in this paper can generate good *relative* 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 one another.

We make four basic assumptions. First, we assume that each robot is equipped with a proprioceptive *motion sensor* such that it can measure changes in its own pose (position and orientation). Suitable motion detectors can be

constructed using either odometry, inertial measurement units, or some combination of the two. Second, we assume that each robot is equipped with a *robot sensor* 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). Third, we assume that the identity of robots is always determined correctly (and hence we are not required to solve the data association 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 information back to a central location, where the localization is performed. 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 and numerical optimization. The basic method is as follows. First, we construct a set of estimates $\hat{\mathcal{X}} = \{\hat{\mathbf{x}}\}$ in which each element $\hat{\mathbf{x}}$ represents the estimated pose of a particular robot at a particular time. These pose estimates are defined with respect to some *arbitrary* global coordinate system. Next, we construct a set of observations $\mathcal{M} = \{\mathbf{m}\}$ in which each element \mathbf{m} represents an observation made by a motion sensor, and a set of observations $\mathcal{O} = \{\mathbf{o}\}$ in which each element \mathbf{o} represents an observation made by a robot sensor. Finally, we use numerical optimization to determine the set of estimates $\hat{\mathcal{X}}$ that is most likely to give rise to the combined set of observations $(\mathcal{M}, \mathcal{O})$. Note that this method attempts to determine the pose of every robot at *every point in time*: we are, in effect, ‘unrolling’ the time component and treating this as a static estimation problem. It is therefore necessary, in practice, to bound $\hat{\mathcal{X}}$ such that it includes pose estimates over some finite period of time.

We do not expect robots to use the set of pose estimates $\hat{\mathcal{X}}$ directly; these estimates are defined with respect to an arbitrary global 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 employs this ego-centric viewpoint to coordinate activity. If, on the other hand, some subset of the team chooses to remain stationary, the global coordinate system will become ‘anchored’ in the real world, and the pose estimates in $\hat{\mathcal{X}}$ may be used as global pose

estimates in the standard fashion.

In the remainder of this paper, we describe the basic formalism, its practical implementation, and present experimental results obtained using a small team of real robots in both static and dynamic environments.

II. RELATED WORK

Localization is an extremely well studied topic 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 [1], [2], [3], or localizing a single robot whilst simultaneously building a map of the environment [4], [5], [6], [7], [8], [9]. Recently, some authors have also considered the related problem of map building with multiple robots [10]. All of these authors make use of statistical or probabilistic techniques of one sort or another; the common tools of choice are Kalman filters (KF), maximum likelihood estimation (MLE), expectation maximization (EM), 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 [5] 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 similar to that described by these authors, even though it is directed towards a very different objective, i.e., the localization of mobile robot teams, rather than the construction of globally consistent maps.

Among those who have considered the problem of *co-operative* localization are Roumeliotis and Bekey [11], [12] and Fox *et al.* [13]. Roumeliotis and Bekey [12] 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 on the use of external landmarks. In a similar vein, 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 Markovian localization techniques [3]. 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 results in some loss of expressiveness. One cannot, for example, express relationships of the form: “if




Fig. 1. An illustration of the basic formalism. The figure shows two robots, $r = 1$ and $r = 2$, traveling from left to right and observing each other exactly once. The robots’ activity is encoded in the graph, with nodes representing pose estimates and arcs representing observations. Two observations are highlighted: a motion observation for robot $r = 1$ (between times $t = 1$ and $t = 2$) and a robot observation at time $t = 2$ (between robots $r = 1$ and $r = 2$).

I am at A then you must be at C, but if I am at B you must be at D”.

Finally, a number of authors [14], [15], [16] have described approaches in which team members 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 subset is 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 section of this paper.

III. FORMALISM

A. General Formalism

We formulate the team localization problem as follows. Let $\hat{\mathbf{x}}_t^r$ denote the *absolute pose estimate* for robot r at time t , and let $\hat{\mathcal{X}}$ denote the set of all such estimates. Let $\mathbf{m}_{tt'}^r$ denote an observation made by a motion sensor describing the change in pose of robot r between times t and t' . Let \mathcal{M} denote the set of all such observations. Finally, let $\mathbf{o}_t^{r r'}$ denote an observation made by a robot sensor at time t , in which robot r measures the relative pose of robot r' , and let \mathcal{O} denote the set of all such observations. These definitions are illustrated in Figure 1. Each estimate $\hat{\mathbf{x}}_t^r$ can be thought of as a node in a graph, and each observation $\mathbf{m}_{tt'}^r$ or $\mathbf{o}_t^{r r'}$ can be thought of as an arc between two nodes. Thus, 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. Note also that this is a *directed* graph; an observation $\mathbf{o}_t^{r r'}$ is not equivalent to an observation $\mathbf{o}_t^{r' r}$, for example.

Our aim is to determine the set of pose estimates $\hat{\mathcal{X}}$ that maximizes the probability of obtaining the set of observations $(\mathcal{M}, \mathcal{O})$; i.e., we seek to maximize the conditional probability $P(\mathcal{M}, \mathcal{O} \mid \hat{\mathcal{X}})$. If we assume that observations

are statistically independent, we can write this probability as:

$$P(\mathcal{M}, \mathcal{O} | \hat{\mathcal{X}}) = \prod_{\mathbf{m}_{tt'}^r \in \mathcal{M}} P(\mathbf{m}_{tt'}^r | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r) \times \prod_{\mathbf{o}_t^{rr'} \in \mathcal{O}} P(\mathbf{o}_t^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'}) \quad (1)$$

where $P(\mathbf{m}_{tt'}^r | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r)$ is the probability of obtaining the individual motion observation $\mathbf{m}_{tt'}^r$, given that the estimated initial pose for robot r is $\hat{\mathbf{x}}_t^r$, and the estimated final pose for the same robot is $\hat{\mathbf{x}}_{t'}^r$. Note that we have made the additional (but not unreasonable) assumption that this probability is independent of other pose estimates. In a similar vein, $P(\mathbf{o}_t^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'})$ specifies the probability of obtaining the individual robot observation $\mathbf{o}_t^{rr'}$, given that the estimate pose for the robot r making the observation is $\hat{\mathbf{x}}_t^r$, and the estimated pose for the robot r' being observed is $\hat{\mathbf{x}}_t^{r'}$. Taking the logarithm of Equation 1, we can write:

$$U(\mathcal{M}, \mathcal{O} | \hat{\mathcal{X}}) = \sum_{\mathbf{m}_{tt'}^r \in \mathcal{M}} U(\mathbf{m}_{tt'}^r | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r) + \sum_{\mathbf{o}_t^{rr'} \in \mathcal{O}} U(\mathbf{o}_t^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'}) \quad (2)$$

where

$$\begin{aligned} U(\mathcal{M}, \mathcal{O} | \hat{\mathcal{X}}) &= -\log P(\mathcal{M}, \mathcal{O} | \hat{\mathcal{X}}) \\ U(\mathbf{m}_{tt'}^r | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r) &= -\log P(\mathbf{m}_{tt'}^r | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r) \\ U(\mathbf{o}_t^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'}) &= -\log P(\mathbf{o}_t^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'}) \end{aligned} \quad (3)$$

This latter form is somewhat more convenient for numerical optimization. Our aim now is to find the set of estimates $\hat{\mathcal{X}}$ that *minimizes* $U(\mathcal{M}, \mathcal{O} | \hat{\mathcal{X}})$, for which we need to determine the form of the conditional log-probabilities $U(\mathbf{m}_{tt'}^r | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r)$ and $U(\mathbf{o}_t^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'})$.

If we assume that the uncertainty in motion observations is normally distributed in the coordinate system in which measurements are made, we can describe each motion observation using a tuple of the form:

$$\mathbf{m}_{tt'}^r = (\boldsymbol{\mu}_{tt'}^r, \boldsymbol{\Sigma}_{tt'}^r) \quad (4)$$

where $\boldsymbol{\mu}_{tt'}^r$ is a *relative pose measurement* describing the measured change in pose of robot r between times t and t' , and $\boldsymbol{\Sigma}_{tt'}^r$ is a covariance matrix representing the uncertainty in this measurement. The conditional log-probability for such observations is given by:

$$U(\mathbf{m}_{tt'}^r | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r) = \frac{1}{2}(\boldsymbol{\mu}_{tt'}^r - \hat{\mathbf{y}}_{tt'}^r)^T \boldsymbol{\Sigma}_{tt'}^r (\boldsymbol{\mu}_{tt'}^r - \hat{\mathbf{y}}_{tt'}^r) \quad (5)$$

where $\hat{\mathbf{y}}_{tt'}^r$ is a *relative pose estimate* describing the *estimated* change in pose of robot r between times t and t' . The relative pose estimate is derived from the absolute pose estimates $\hat{\mathbf{x}}_t^r$ and $\hat{\mathbf{x}}_{t'}^r$ via some coordinate transform Γ_m :

$$\hat{\mathbf{y}}_{tt'}^r = \Gamma_m(\hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r) \quad (6)$$

The specific form of Γ_m depends on the coordinate system chosen to represent the absolute pose estimates $\hat{\mathcal{X}}$ and the motion observations \mathcal{M} . We will consider one specific form for Γ_m in Section III-B, where we consider the problem of localization in a plane.

Robot observations are handled in a similar manner to motion observations: each observation is described using a tuple of the form:

$$\mathbf{o}_t^{rr'} = (\boldsymbol{\mu}_t^{rr'}, \boldsymbol{\Sigma}_t^{rr'}) \quad (7)$$

where $\boldsymbol{\mu}_t^{rr'}$ is the relative pose of robot r' , as measured by robot r at time t ; $\boldsymbol{\Sigma}_t^{rr'}$ is the covariance matrix representing the uncertainty in this measurement. The conditional log-probability is given by:

$$U(\mathbf{o}_t^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'}) = \frac{1}{2}(\boldsymbol{\mu}_t^{rr'} - \hat{\mathbf{y}}_t^{rr'})^T \boldsymbol{\Sigma}_t^{rr'} (\boldsymbol{\mu}_t^{rr'} - \hat{\mathbf{y}}_t^{rr'}) \quad (8)$$

where $\hat{\mathbf{y}}_t^{rr'}$ is the estimated pose of robot r' relative to robot r at time t . This relative pose estimate is derived from the absolute pose estimates $\hat{\mathbf{x}}_t^r$ and $\hat{\mathbf{x}}_t^{r'}$ via some coordinate transform Γ_o :

$$\hat{\mathbf{y}}_t^{rr'} = \Gamma_o(\hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'}) \quad (9)$$

As with motion observations, the specific form of Γ_o depends on the coordinate systems being used.

Given appropriate definitions for Γ_m and Γ_o , one can determine the optimal set of pose estimates $\hat{\mathcal{X}}$ using a standard numerical optimization algorithm. The selection of an appropriate algorithm is driven largely by the form of the coordinate transforms, which are, in general, non-linear but differentiable. This rules out fast linear techniques, but allows derivative-based techniques such as steepest descent and conjugate gradient algorithms [17]. The gradient of $U(\mathcal{M}, \mathcal{O} | \hat{\mathcal{X}})$ with respect to $\hat{\mathcal{X}}$ can be computed by differentiating through Equation 2:

$$\begin{aligned} \frac{\partial}{\partial \hat{\mathbf{x}}_t^r} U(\mathcal{M}, \mathcal{O} | \hat{\mathcal{X}}) &= \sum_{\mathbf{m}_{tt'}^r \in \mathcal{M}} \frac{\partial}{\partial \hat{\mathbf{x}}_t^r} U(\mathbf{m}_{tt'}^r | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r) \\ &+ \sum_{\mathbf{m}_{t't}^{r'} \in \mathcal{M}} \frac{\partial}{\partial \hat{\mathbf{x}}_t^r} U(\mathbf{m}_{t't}^{r'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^{r'}) \\ &+ \sum_{\mathbf{o}_t^{rr'} \in \mathcal{O}} \frac{\partial}{\partial \hat{\mathbf{x}}_t^r} U(\mathbf{o}_t^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'}) \\ &+ \sum_{\mathbf{o}_t^{r'r} \in \mathcal{O}} \frac{\partial}{\partial \hat{\mathbf{x}}_t^r} U(\mathbf{o}_t^{r'r} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'}) \end{aligned} \quad (10)$$

The four summations in this equation capture four different cases: for motion observations, the pose estimate $\hat{\mathbf{x}}_t^r$ may correspond to either the initial or the final location of the robot, and for robot observations the pose estimate $\hat{\mathbf{x}}_t^r$ may correspond to either the robot making the observation or the robot being observed. The individual derivatives for the motion observation terms can be evaluated by applying the chain-rule to Equation 5:

$$\frac{\partial}{\partial \hat{\mathbf{x}}_t^r} U(\mathbf{m}_{tt'}^r | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^r) = \frac{\partial \hat{\mathbf{y}}_{tt'}^r}{\partial \hat{\mathbf{x}}_t^r} \boldsymbol{\Sigma}_{tt'}^r (\boldsymbol{\mu}_{tt'}^r - \hat{\mathbf{y}}_{tt'}^r)$$

$$\frac{\partial}{\partial \hat{\mathbf{x}}_t^r} U(\mathbf{m}_{t't}^r | \hat{\mathbf{x}}_{t'}^r, \hat{\mathbf{x}}_t^r) = \frac{\partial \hat{\mathbf{y}}_{t't}^r}{\partial \hat{\mathbf{x}}_t^r} \Sigma_{t't}^r (\boldsymbol{\mu}_{t't}^r - \hat{\mathbf{y}}_{t't}^r) \quad (11)$$

where the derivatives $\partial \hat{\mathbf{y}}_{t't}^r / \partial \hat{\mathbf{x}}_t^r$ and $\partial \hat{\mathbf{y}}_{t't}^r / \partial \hat{\mathbf{x}}_t^r$ must be determined by differentiating Equation 6, given some specific form for Γ_m . The derivatives for the robot observation terms are determined in a similar fashion by applying the chain-rule to Equation 8:

$$\begin{aligned} \frac{\partial}{\partial \hat{\mathbf{x}}_t^r} U(\mathbf{o}_{t't}^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_{t'}^{r'}) &= \frac{\partial \hat{\mathbf{y}}_t^{rr'}}{\partial \hat{\mathbf{x}}_t^r} \Sigma_t^{rr'} (\boldsymbol{\mu}_t^{rr'} - \hat{\mathbf{y}}_t^{rr'}) \\ \frac{\partial}{\partial \hat{\mathbf{x}}_t^r} U(\mathbf{o}_t^{r'r} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^r) &= \frac{\partial \hat{\mathbf{y}}_t^{r'r}}{\partial \hat{\mathbf{x}}_t^r} \Sigma_t^{r'r} (\boldsymbol{\mu}_t^{r'r} - \hat{\mathbf{y}}_t^{r'r}) \end{aligned} \quad (12)$$

where the derivatives $\partial \hat{\mathbf{y}}_t^{rr'} / \partial \hat{\mathbf{x}}_t^r$ and $\partial \hat{\mathbf{y}}_t^{r'r} / \partial \hat{\mathbf{x}}_t^r$ are determined by differentiating Equation 9 for some specific Γ_o .

B. Localization in a Plane

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

Let the absolute pose estimate $\hat{\mathbf{x}}_t^r$ be a 3-vector such that:

$$\hat{\mathbf{x}}_t^r = \begin{bmatrix} \hat{x}_t^r \\ \hat{y}_t^r \\ \hat{\theta}_t^r \end{bmatrix} \quad (13)$$

where $(\hat{x}_t^r, \hat{y}_t^r)$ describes the robot's position (in Cartesian coordinates) and $\hat{\theta}_t^r$ describes its orientation. For motion observations, let the relative pose measurement $\boldsymbol{\mu}_{t't}^r$ be a 3-vector such that:

$$\boldsymbol{\mu}_{t't}^r = \begin{bmatrix} d_{t't}^r \\ \phi_{t't}^r \\ \psi_{t't}^r \end{bmatrix} \quad (14)$$

where $(d_{t't}^r, \phi_{t't}^r)$ describes the measured change in position for robot r (in polar coordinates, i.e. range and bearing) and $\psi_{t't}^r$ describes its change in orientation. The relative pose measurements $\boldsymbol{\mu}_t^{rr'}$ for robot observations are defined in a similar manner:

$$\boldsymbol{\mu}_t^{rr'} = \begin{bmatrix} d_t^{rr'} \\ \phi_t^{rr'} \\ \psi_t^{rr'} \end{bmatrix} \quad (15)$$

where $(d_t^{rr'}, \phi_t^{rr'})$ describes the measured range and bearing of robot r' relative to r , and $\psi_t^{rr'}$ describes the relative orientation of the robots.

We choose to express measurements in terms of polar coordinates since, for many sensors, the uncertainty in range, bearing, and orientation is effectively uncorrelated. Consider, for example, the probability distributions shown in Figure 2; by using polar coordinates, we can accurately model the behavior of sensors that return only range information, sensors that return only bearing information, and

sensors that return both range and bearing information. Thus, for example, polar coordinates are well suited to representing odometric measurements from differential drive robots: Figure 2(c) is a good approximation to the classic 'banana-shaped' distribution that is usually observed with such sensors [3]. Note, however, that if one were to use an IMU as a motion sensor, motion measurements would more naturally be expressed in Cartesian coordinates, and the derivations that follow would have to be suitably modified.

Given the above definitions, we can write down an expression for the relative pose estimates $\hat{\mathbf{y}}_{t't}^r$: the coordinate transform function Γ_m simply transforms from polar to Cartesian coordinates:

$$\begin{aligned} \hat{\mathbf{y}}_{t't}^r &= \begin{bmatrix} \hat{d}_{t't}^r \\ \hat{\phi}_{t't}^r \\ \hat{\psi}_{t't}^r \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{(\hat{x}_{t'}^r - \hat{x}_t^r)^2 + (\hat{y}_{t'}^r - \hat{y}_t^r)^2} \\ \arctan[(\hat{y}_{t'}^r - \hat{y}_t^r) / (\hat{x}_{t'}^r - \hat{x}_t^r)] - \hat{\theta}_t^r \\ \hat{\theta}_{t'}^r - \hat{\theta}_t^r \end{bmatrix} \end{aligned} \quad (16)$$

For optimization, we require the corresponding derivatives, which are given by:

$$\begin{aligned} \frac{\partial \hat{\mathbf{y}}_{t't}^r}{\partial \hat{\mathbf{x}}_t^r} &= \begin{bmatrix} \partial \hat{d}_{t't}^r / \partial \hat{x}_t^r & \partial \hat{\phi}_{t't}^r / \partial \hat{x}_t^r & \partial \hat{\psi}_{t't}^r / \partial \hat{x}_t^r \\ \partial \hat{d}_{t't}^r / \partial \hat{y}_t^r & \partial \hat{\phi}_{t't}^r / \partial \hat{y}_t^r & \partial \hat{\psi}_{t't}^r / \partial \hat{y}_t^r \\ \partial \hat{d}_{t't}^r / \partial \hat{\theta}_t^r & \partial \hat{\phi}_{t't}^r / \partial \hat{\theta}_t^r & \partial \hat{\psi}_{t't}^r / \partial \hat{\theta}_t^r \end{bmatrix} \\ &= \begin{bmatrix} -(\hat{x}_{t'}^r - \hat{x}_t^r) / d_{t't}^r & +(\hat{y}_{t'}^r - \hat{y}_t^r) / (d_{t't}^r)^2 & 0 \\ -(\hat{y}_{t'}^r - \hat{y}_t^r) / d_{t't}^r & -(\hat{x}_{t'}^r - \hat{x}_t^r) / (d_{t't}^r)^2 & 0 \\ 0 & -1 & -1 \end{bmatrix} \\ \frac{\partial \hat{\mathbf{y}}_{t't}^r}{\partial \hat{\mathbf{x}}_t^r} &= \begin{bmatrix} \partial \hat{d}_{t't}^r / \partial \hat{x}_t^r & \partial \hat{\phi}_{t't}^r / \partial \hat{x}_t^r & \partial \hat{\psi}_{t't}^r / \partial \hat{x}_t^r \\ \partial \hat{d}_{t't}^r / \partial \hat{y}_t^r & \partial \hat{\phi}_{t't}^r / \partial \hat{y}_t^r & \partial \hat{\psi}_{t't}^r / \partial \hat{y}_t^r \\ \partial \hat{d}_{t't}^r / \partial \hat{\theta}_t^r & \partial \hat{\phi}_{t't}^r / \partial \hat{\theta}_t^r & \partial \hat{\psi}_{t't}^r / \partial \hat{\theta}_t^r \end{bmatrix} \\ &= \begin{bmatrix} +(\hat{x}_{t'}^r - \hat{x}_t^r) / d_{t't}^r & -(\hat{y}_{t'}^r - \hat{y}_t^r) / (d_{t't}^r)^2 & 0 \\ +(\hat{y}_{t'}^r - \hat{y}_t^r) / d_{t't}^r & +(\hat{x}_{t'}^r - \hat{x}_t^r) / (d_{t't}^r)^2 & 0 \\ 0 & 0 & +1 \end{bmatrix} \end{aligned} \quad (17)$$

There are two features of these derivatives that should be noted. First, $\partial \hat{\mathbf{y}}_{t't}^r / \partial \hat{\mathbf{x}}_t^r \neq -\partial \hat{\mathbf{y}}_{t't}^r / \partial \hat{\mathbf{x}}_t^r$, as one might naively expect. Second, the derivatives contain a singularity at $d_{t't}^r = 0$; one must take care to avoid this singularity during the optimization process.

The relative pose estimates $\hat{\mathbf{y}}_t^{rr'}$ are treated in a similar manner, since Γ_o also transforms from polar to Cartesian coordinates:

$$\begin{aligned} \hat{\mathbf{y}}_t^{rr'} &= \begin{bmatrix} \hat{d}_t^{rr'} \\ \hat{\phi}_t^{rr'} \\ \hat{\psi}_t^{rr'} \end{bmatrix} \\ &= \begin{bmatrix} \sqrt{(\hat{x}_t^{r'} - \hat{x}_t^r)^2 + (\hat{y}_t^{r'} - \hat{y}_t^r)^2} \\ \arctan[(\hat{y}_t^{r'} - \hat{y}_t^r) / (\hat{x}_t^{r'} - \hat{x}_t^r)] - \hat{\theta}_t^r \\ \hat{\theta}_t^{r'} - \hat{\theta}_t^r \end{bmatrix} \end{aligned} \quad (18)$$




Fig. 2. Sample probability distributions for for the planar localization problem. The plots show the probability $P(\mathbf{o}_t^{rr'} | \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'})$ as a function of the estimated position $(\hat{x}_t^{r'}, \hat{y}_t^{r'})$ for the robot being observed. Orientation is not shown. (a) The range is well determined ($d_t^{rr'} = 1 \pm 0.1$ m), but the bearing is unknown. (b) The bearing is moderately well determined ($\phi_t^{rr'} = 45 \pm 57^\circ$), but the range is unknown. (c) Both range and bearing are (moderately) well determined ($d_t^{rr'} = 1 \pm 0.1$ m, $\phi_t^{rr'} = 45 \pm 57^\circ$).




Fig. 3. Plot of the log-probability $U(\mathcal{M}, \mathcal{O} | \hat{\mathcal{X}})$ along a number of lines through $\hat{\mathcal{X}}$ (the lines correspond to successive optimization steps in the conjugate gradient algorithm). Note that the plots are smooth (which aids optimization) and that they contain local minima (which does not). These plots were generated using real data from the experiments described in Section V-A.

The corresponding derivatives are:

$$\frac{\partial \hat{\mathbf{y}}_t^{rr'}}{\partial \hat{\mathbf{x}}_t^r} = \begin{bmatrix} -(\hat{x}_t^{r'} - \hat{x}_t^r)/\hat{d}_t^{rr'} & +(\hat{y}_t^{r'} - \hat{y}_t^r)/(\hat{d}_t^{rr'})^2 & 0 \\ -(\hat{y}_t^{r'} - \hat{y}_t^r)/\hat{d}_t^{rr'} & -(\hat{x}_t^{r'} - \hat{x}_t^r)/(\hat{d}_t^{rr'})^2 & 0 \\ 0 & -1 & -1 \end{bmatrix} \quad (19)$$

$$\frac{\partial \hat{\mathbf{y}}_t^{r'r}}{\partial \hat{\mathbf{x}}_t^r} = \begin{bmatrix} +(\hat{x}_t^{r'} - \hat{x}_t^r)/\hat{d}_t^{r'r} & -(\hat{y}_t^{r'} - \hat{y}_t^r)/(\hat{d}_t^{r'r})^2 & 0 \\ +(\hat{y}_t^{r'} - \hat{y}_t^r)/\hat{d}_t^{r'r} & +(\hat{x}_t^{r'} - \hat{x}_t^r)/(\hat{d}_t^{r'r})^2 & 0 \\ 0 & 0 & +1 \end{bmatrix}$$

Note once again that $\partial \hat{\mathbf{y}}_t^{rr'}/\partial \hat{\mathbf{x}}_t^r \neq -\partial \hat{\mathbf{y}}_t^{r'r}/\partial \hat{\mathbf{x}}_t^r$, and that both derivatives contain a singularity at $\hat{d}_t^{rr'} = 0$.

Inserting these definitions into the general formalism described in the previous sections, one can solve the planar localization problem in a fairly straight-forward manner.

IV. IMPLEMENTATION

A. Optimization: The Conjugate Gradient Algorithm

For optimization, we use the conjugate gradient algorithm described in [17]; this algorithm is somewhat more complex than the standard steepest descent algorithm, but




Fig. 4. An illustration of the extended formalism. The figure shows two robots, $r = 1$ and $r = 2$, traveling from left to right and observing each other exactly once. The robots' activity is encoded in the graph, with nodes representing pose estimates and arcs representing observations. Also shown are the interpolated pose estimates $\hat{\mathbf{z}}_2^1$ and $\hat{\mathbf{z}}_2^2$ for each of the robots at time $t = 2$.

has the advantage of being much faster. In addition, its memory requirements scale linearly with the number of variables being optimized (rather than quadratically, as is the case with some alternatives). This latter feature is important when one considers that the set of pose estimates $\hat{\mathcal{X}}$ over which we are optimizing can be very large.

Figure 3 shows an example of the conjugate gradient algorithm at work on a real data set. Each curve shows the log-probability $U(\mathcal{M}, \mathcal{O} | \hat{\mathcal{X}})$ plotted along a line through the high-dimensional space that is $\hat{\mathcal{X}}$; these lines are computed at successive steps in the algorithm. Note that the plots are smooth (which aids optimization) but that there exist local minima (which does not). For this data set, which contains 20 pose estimates and 50 observations, the algorithm was able to find the global minimum with 61 optimization steps (well under 1 second on a 700 MHz PIII workstation).

Note that we also treat the observations incrementally: as each new observation is added to \mathcal{M} or \mathcal{O} , we optimize over the entire set of pose estimates $\hat{\mathcal{X}}$. This incremental approach helps to ensure that the algorithm converges to the global rather than a local minima, since we have a very good initial 'guess' at the start of each optimization pass. While this approach does not *guarantee* convergence to the global minimum, it appears to suffice in practice.

B. Practical considerations: bounding $\hat{\mathcal{X}}$, \mathcal{M} and \mathcal{O}

The dimensionality of the problem that must be solved by the conjugate gradient algorithm scales linearly with the size of $\hat{\mathcal{X}}$, and the computational cost of each step in this algorithm scales linearly with the size of $\{\mathcal{M} \cup \mathcal{O}\}$. It is therefore necessary, in practice, to bound both the number of estimates in $\hat{\mathcal{X}}$ and the number of observations in $\{\mathcal{M} \cup \mathcal{O}\}$. We use three basic methods for constraining the size of these sets: we discard older estimates and observations, we limit the number of observations relating any pair of estimates, and we limit the rate at which new estimates are generated.

The first two methods are simple and well-defined. First, we consider only those estimates and observations that occur between times t and $t - T$, where t is the current time and T is the duration of a temporal ‘window’. Second, we consider at most m observations for every pair of pose estimates. Thus we can trade off computational cost against localization accuracy by manipulating the duration T and the maximum observations-per-estimate m .

In contrast, the third method, which seeks to limit the rate at which pose estimates are generated, requires some extensions to the formalism described in Section III. Rather than attempting to estimate the pose of every the robot at every point in time, we instead estimate the pose of each robot at only a few discrete points in time; information from the motion sensors is used to ‘fill the gaps’ between these estimates. In effect, we assume that the motion sensors alone can be used to generate good pose estimates, and that these estimates require only occasional corrections.

The formalism is extended as follows. Let $\hat{\mathbf{z}}_{t+}^r$ be the *interpolated pose estimate* for robot r at time t^+ ; this estimate is given by:

$$\hat{\mathbf{z}}_{t+}^r = \Gamma_m^{-1}(\hat{\mathbf{x}}_t^r, \boldsymbol{\mu}_{tt+}^r) \quad (20)$$

where $\hat{\mathbf{x}}_t^r$ is the most recent pose estimate for robot r in $\hat{\mathcal{X}}$ and $\boldsymbol{\mu}_{tt+}^r$ is the measured change in pose between times t and t^+ . The inverse coordinate transform Γ_m^{-1} maps from relative to absolute coordinates. These definitions are illustrated in Figure 4.

The difficulty now lies in the incorporation of information from the robot sensors: as the interval between pose estimates becomes larger, it becomes increasingly unlikely that robot observations will occur at the times represented by these pose estimates. We must therefore make the following modifications to the robot observation model (Equation 8):

$$U(\mathbf{o}_{t+}^{rr'} \mid \hat{\mathbf{x}}_t^r, \hat{\mathbf{x}}_t^{r'}, \boldsymbol{\mu}_{tt+}^r) = \frac{1}{2}(\boldsymbol{\mu}_{t+}^{rr'} - \hat{\mathbf{y}}_{t+}^{rr'})^T \boldsymbol{\Sigma}_t^{rr'} (\boldsymbol{\mu}_{t+}^{rr'} - \hat{\mathbf{y}}_{t+}^{rr'}) \quad (21)$$

where

$$\hat{\mathbf{y}}_{t+}^{rr'} = \Gamma_o(\hat{\mathbf{z}}_{t+}^r, \hat{\mathbf{z}}_{t+}^{r'}) \quad (22)$$

The corresponding derivatives must also be modified to yield:

$$\frac{\partial \hat{\mathbf{y}}_{t+}^{rr'}}{\partial \hat{\mathbf{x}}_t^r} = \frac{\partial \hat{\mathbf{y}}_{t+}^{rr'}}{\partial \hat{\mathbf{z}}_{t+}^r} \frac{\partial \hat{\mathbf{z}}_{t+}^r}{\partial \hat{\mathbf{x}}_t^r} \quad (23)$$

The additional derivative $\partial \hat{\mathbf{z}}_{t+}^r / \partial \hat{\mathbf{x}}_t^r$ can be computed by differentiating through Equation 20.

This extended formalism allows us to approximate the information provided by the motion sensors to an arbitrary degree of fidelity (rather than simply discarding information). It therefore provides us with another means to trade off computational cost against localization accuracy.

V. EXPERIMENTAL VALIDATION

A. Experiment 1

We conducted a series of experiments aimed at validating both the practicality and the accuracy of the team localization algorithm presented in this paper. The first 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 5(a); these totem-poles can be detected from a wide range of angles using the laser range-finders, which can be programmed to return both range and intensity information. This configuration 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, but is subject to a 180° ambiguity. This configuration does *not* allow individual robots to be identified. Given the ambiguity in both orientation and identity, it was necessary, for this experiment, to manually label the data. Ground-truth information for this experiment was provided by an external laser-based tracking system.

The robot team was placed into the environment shown in Figure 5(b) 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 the outer wall able to directly observe one other. The structure of the environment was also 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’. The original structure was later restored, then modified, then restored again.

The qualitative results for this experiment are summarized in Figure 6, which contains a series of ‘snap-shots’ from 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 snap-shot is largely incoherent; at this time, the relative pose of the robots is completely unknown and their absolute poses have been chosen randomly. In the interval $0 < t < 12$ sec, the robots commence wall following; by time $t = 12$ sec, the robots following the outer wall have observed and been observed by 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 hence to correctly align the laser scan data. It should be noted

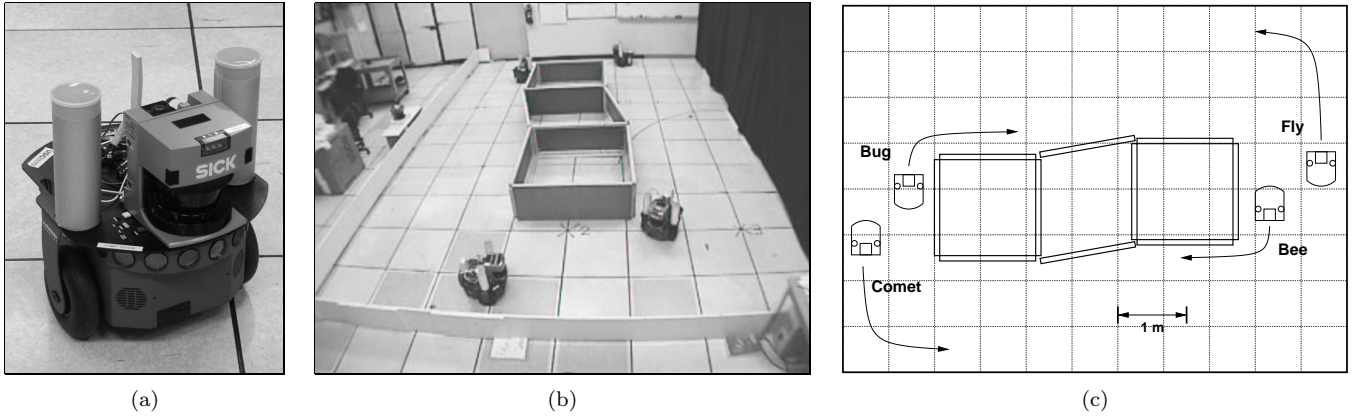


Fig. 5. Setup for Experiment 1. (a) A Pioneer 2DX equipped with a SICK LMS200 scanning laser range-finder and a pair of retro-reflective totem-poles. (b) The arena: the central island is constructed from partitions that can be rearranged during the course of the experiment. (c) Set-up: robots Fly and Comet follow the outer wall, robots Bee and Bug follow the inner wall(s).

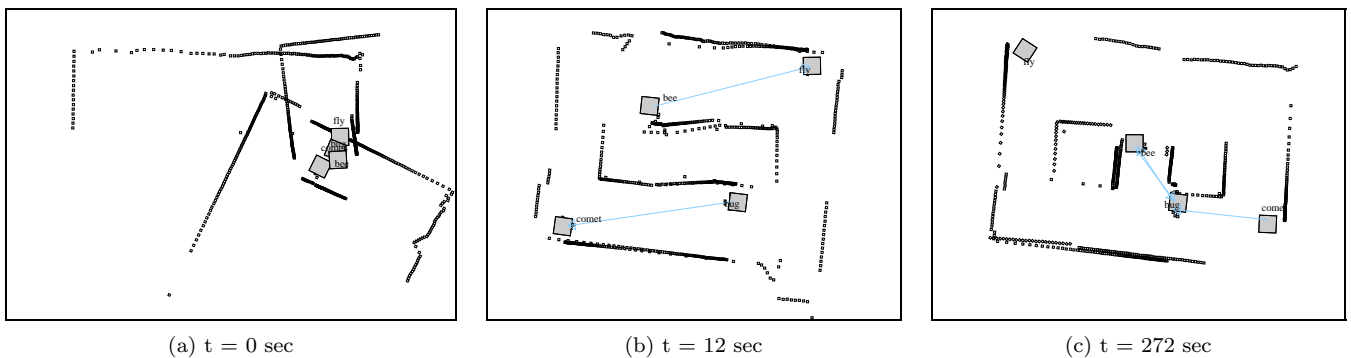


Fig. 6. Snap-shots from Experiment 1. Each sub-figure shows the estimated pose of the robots at a particular point in time, overlaid with the corresponding laser scan data. Arrows denote the observation of one robot by another. Note that these are snap-shots of live data, not cumulative maps of stored data.




Fig. 7. Plots showing the relative pose error as a function of time for Experiment 1. The three plots show the average range, bearing, and orientation errors, respectively.

that the two robots on the outer wall can correctly determine each others' pose, even though they have never observed one other directly. At time $t = 265$ sec, the environment is modified, with the inner wall being restructured to form two separate islands. The localization, however, is unaffected by this change, as shown in the snap-shot at time $t = 272$ sec. The method described in this paper is completely indifferent to such structural changes in the environment.

The accuracy of the localization algorithm was determined by comparing the relative pose estimates $\hat{\mathbf{y}}_t^{rr'}$ against their corresponding true values (as determined by

the ground-truth system). That is, for each time t we compute how accurately each robot has estimated the relative pose of all the other robots. Thus, accuracy is measured by an error term of the form:

$$(\epsilon_t)^2 = \frac{1}{n(n-1)} \sum_r \sum_{r':r' \neq r} (\hat{\mathbf{y}}_t^{rr'} - \mathbf{y}_t^{rr'})^T (\hat{\mathbf{y}}_t^{rr'} - \mathbf{y}_t^{rr'}) \quad (24)$$

where $\mathbf{y}_t^{rr'}$ is the true pose of robot r' relative to robot r at time t . The summation is over all pairs of robots and the result is normalized by the number of pairs $n(n-1)$ to generate an average result. Note that we make no attempt to compare the *absolute* pose estimates $\hat{\mathbf{x}}_t^r$ against the 'true'

values; the arbitrary nature of the global coordinate system renders such comparison meaningless.

The quantitative results for this experiment are summarized in Figure 7, which plots the range, bearing, and orientation components of the total error ϵ_t as a function of time. At time $t = 0$ sec, the relative pose of the robots is completely unknown and the errors are correspondingly high. By time $t = 12$ sec, however, the robots have gathered sufficient information to fully constrain their relative poses, and the errors have fallen to more modest values. Over the duration of the experiment, the range error oscillates around 18 ± 12 cm, while the bearing and orientation errors oscillate around $5.9 \pm 1.9^\circ$ and $6.7 \pm 1.5^\circ$, respectively.

The magnitude the localization errors can be attributed to two factors. First, there is some uncertainty in the relative pose measurements made by the laser-range-finder/retro-reflector combination. These uncertainties are difficult to characterize precisely, but are of the order of ± 2.5 cm and $\pm 2^\circ$. Second, and more significantly, there are uncertainties associated with the temporal synchronization of the laser and odometric measurements. Our low-level implementation is such that the time at which events occur can only be measured to the nearest 100 ms; in this interval, the robot may travel 2 cm and/or rotate through 3° , which will significantly bias the results.

We ascribe the *variation* seen in the error plots to two additional factors. First, we expect that the error will increase during those periods in which the robots cannot see each other and localization is reliant on odometry alone; the odometric accuracy of the robots used in this experiment varies from quite good ($2.5^\circ/\text{lap}$) to quite poor ($30^\circ/\text{lap}$). Second, we expect that errors will decrease during those periods when robots are observing one another. This decrease, however, may be preceded by a ‘spike’ in the error term; the spike is an artifact produced by the optimization algorithm, which may take several cycles (each cycle is 100 ms) to incorporate the new data and determine a new set of pose estimates.

Finally, we note that there is a major spike in the plot at around $t = 300$ sec. This spike corresponds to a collision that occurred between two robots following the first structural change in the environment. As a result of this collision, the robots had to be manually repositioned, leading to gross errors in both robot’s odometry. Nevertheless, as the plot indicates, the algorithm was able to quickly recover.

B. Experiment 2

The second experiment was conducted using a team of 7 Pioneer 2DX mobile robots, each equipped with a SICK LMS 200 scanning laser range-finder and a retro-reflective target. For this experiment, 6 of the 7 robots were positioned at fixed locations in the corridors of a building, as shown in Figure 8(a). The remaining robot was ‘joy-sticked’ around the circuit, and was thus observed by each of the stationary robots in turn. The stationary robots were positioned outside each other’s sensor range, and hence were unable to observe each other directly. Since these ex-

periments were performed in an uninstrumented environment, ground truth information was obtained by measuring the inter-robot distances between the fixed robots. Bearings and orientations were not measured.

Qualitative results for this experiment are shown in Figures 8(b) and (c); each figure is a snap-shot showing the pose estimates, observations, and laser scan data at a particular point in time. In Figure 8b, the mobile robot has been seen by each of the stationary robots exactly once, and the pose of each robot has been approximately determined. The localization error remains relatively high (as evidenced by the warping of the snap-shot) due to cumulative odometric error. In Figure 8(c) the mobile robot has ‘closed the loop’ by revisiting the first stationary robot. At this point, the snap-shot ‘clicks’ into place, forming a highly rectilinear map.

The quantitative results for this experiment are summarized in Figure 9, which shows a plot of the relative range error as a function of time. Marked on this plot are the times at which each of the stationary robots observed the mobile robot. The most striking feature of this plot is the way in which the error drops immediately after each robot observation. This is not unexpected, given that the pose of the stationary robots can only be determined after they have seen the mobile robot (recall that the stationary robots cannot see one another). The final error, after the mobile robot has closed the loop, is 0.08 ± 0.09 m. This result is quite remarkable when one considers that the loop about 80m in length. Note that very low error between times $t = 120$ and $t = 200$ is an artifact: the result of a lucky guess for the initial pose of one of the robots.

While this experiment clearly demonstrates our ability to solve for closed loops, it should be noted that the time interval T over which estimates and observations were maintained (see Section IV-B) was chosen to be larger than the time taken to complete the loop.

VI. CONCLUSION AND FURTHER WORK

The localization method described in this paper allows each member in a mobile robot team to determine the relative range, bearing, and orientation of every other robot in the team. The method does not require the use of GPS, external landmarks, or a priori models of the environment. In addition, this method does not require any of the robots to remain stationary; it is robust to changes in the environment and to poor motion sensing; and robots can use transitive relationships to infer the pose of robots they have never seen. There remain, however, several aspects of both the general method and of our particular implementation that require further 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 implementation, we are yet to fully characterize the scaling properties of the algorithm and the relationship between localization accuracy and factors such as the number of pose estimates in \mathcal{X} .

In closing, we note that the mathematical formalism presented in this paper can be extended in a number of inter-




Fig. 8. Setup and results for Experiment 2. (a) The arena: six stationary robots are placed at strategic locations while the seventh mobile robot (Comet) 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 (Tanis) for a second time, thus closing the loop.




Fig. 9. Plot showing the relative range error as a function of time for Experiment 2; bearing and orientation errors were not measured. In this experiment, the mobile robot (Comet) was seen once by each of the robots Bug, Atom, Milka, Fly, and Bee, and was seen twice by the robot Tanis. The time at which each of these observations was made is indicated on the plot.

esting directions. We can, for example, define a covariance matrix that measures the relative uncertainty in the pose estimates between 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 known, they can undertake actions (such as seeking out other robots) that will reduce this uncertainty.

VII. ACKNOWLEDGMENTS

This work is supported by the DARPA MARS Program grant DABT63-99-1-0015, ONR grant N000140110354, and

ONR DURIP grant N00014-00-1-0638.

REFERENCES

- [1] J. J. Leonard and H. F. Durrant-Whyte, "Mobile robot localization by tracking geometric beacons," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 3, pp. 376–382, 1991.
- [2] R. Simmons and S. Koenig, "Probabilistic navigation in partially observable environments," in *Proceedings of International Joint Conference on Artificial Intelligence*, 1995, vol. 2, pp. 1080–1087.
- [3] D. Fox, W. Burgard, and S. Thrun, "Markov localization for mobile robots in dynamic environments," *Journal of Artificial Intelligence Research*, vol. 11, pp. 391–427, 1999.
- [4] S. Thrun, D. Fox, and W. Burgard, "A probabilistic approach to concurrent mapping and localisation for mobile robots," *Ma-*

- chine Learning*, vol. 31, no. 5, pp. 29–55, 1998, Joint issue with Autonomous Robots.
- [5] F. Lu and E. Miliotis, “Globally consistent range scan alignment for environment mapping,” *Autonomous Robots*, vol. 4, pp. 333–349, 1997.
 - [6] B. Yamauchi, A. Shultz, and W. Adams, “Mobile robot exploration and map-building with continuous localization,” in *Proceedings of the 1998 IEEE/RSJ International Conference on Robotics and Automation*, San Francisco, U.S.A., 1998, vol. 4, pp. 3175–3720.
 - [7] T. Duckett, S. Marsland, and J. Shapiro, “Learning globally consistent maps by relaxation,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco, U.S.A., 2000, vol. 4, pp. 3841–3846.
 - [8] M. Golfarelli, D. Maio, and S. Rizzi, “Elastic correction of dead reckoning errors in map building,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Victoria, Canada, 1998, vol. 2, pp. 905–911.
 - [9] M. W. M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte, and M. Csorba, “A solution to the simultaneous localization and map building (SLAM) problem,” *IEEE Transactions on Robotics and Automation*, vol. 17, no. 3, pp. 229–241, 2001.
 - [10] S. Thrun, “A probabilistic online mapping algorithm for teams of mobile robots,” *International Journal of Robotics Research*, vol. 20, no. 5, pp. 335–363, 2001.
 - [11] S. I. Roumeliotis and G. A. Bekey, “Distributed multi-robot localization,” *IEEE Transactions on Robotics and Automation*, 2002, to appear.
 - [12] S. I. Roumeliotis and G. A. Bekey, “Collective localization: a distributed kalman filter approach,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco, U.S.A., 2000, vol. 3, pp. 2958–2965.
 - [13] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, “A probabilistic approach to collaborative multi-robot localization,” *Autonomous Robots, Special Issue on Heterogeneous Multi-Robot Systems*, vol. 8, no. 3, pp. 325–344, 2000.
 - [14] R. Kurazume and S. Hirose, “An experimental study of a cooperative positioning system,” *Autonomous Robots*, vol. 8, no. 1, pp. 43–52, 2000.
 - [15] I. M. Rekleitis, G. Dudek, and E. Miliotis, “Multi-robot exploration of an unknown environment: efficiently reducing the odometry error,” in *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1997, vol. 2, pp. 1340–1345.
 - [16] A. Howard and L. Kitchen, “Cooperative localisation and mapping,” in *International Conference on Field and Service Robotics (FSR99)*, 1999, pp. 92–97.
 - [17] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, 2nd edition, 1999.