# Team Localization: A Maximum Likelihood Approach 

Andrew Howard, Maja J Matarić and Gaurav S Sukhatme<br>Robotics Research Laboratory, Computer Science Department, University of Southern California<br>ahoward@usc.edu, mataric@usc.edu, gaurav@usc.edu


#### Abstract

This paper describes a method for localizing the members of a mobile robot team by 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 sensor (such as odometry or inertial measurement units) and some form of sensor 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). By employing a combination of maximum likelihood estimation and numerical optimization, we can subsequently infer the relative pose of every robot in the team at any given point in time.

This paper describes the basic formalism, its practical implementation, and presents experimental results obtained using both real and simulated robots in both static and dynamic environments.


## 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. Suitable motion detectors can be constructed using either odometry or inertial measurement units. 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 retroreflective 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 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 $X=\{\hat{\mathbf{x}}\}$ 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 $M=\{\hat{\mathbf{m}}\}$ in which each element $\hat{\mathbf{m}}$ represents an observation made by a motion sensor; and a set of observations $O=\{o\}$ in which each element $o$ represents an observation made by a robot sensor. Finally, we use numerical optimization to determine the set of estimates $X$ that is most likely to give rise to the combined set of observations $(M, O)$.

Note that this method attempts to determine the pose of robots 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 $X$ such that it includes pose estimates over some finite period of time. Note also that we do not expect robots to use the set of pose estimates $X$ directly; these estimates are defined with respect to an arbitrary global coordinate system whose relationship with the external environment is undefined. Instead, each robot employs these estimates to compute the pose of every other robot relative to itself, and uses this ego-centric viewpoint to coordinate activity. On the other hand, if 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 $X$ may be used as global pose estimates in the standard fashion.

In the remainder of this paper, we describe the basic for-
malism, its practical implementation, and present experimental results obtained using small teams 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, however, 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 [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 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 problem of cooperative localization are Roumeliotis and Bekey [11] and Fox et al. [12]. Roumeliotis and Bekey [11] 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 problem of dimensionality (for $N$ robots, one must maintain a $3 N$ dimensional distribution) by factoring the distribution into $N$ separate components (one for each robot). While this step makes the algorithm tractable, it also results in some loss of expressiveness. The algorithm also relies on the use of external landmarks.

Finally, a number of authors [13], [14], [15] have described approaches in which team members actively coor-


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 $r=2$ ).
dinate 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 sections 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 $X$ denote the set of all such estimates. Let $\hat{\mathbf{m}}_{t t^{\prime}}^{r}$ denote an observation made by a motion sensor describing the change in pose of robot $r$ between times $t$ and $t^{\prime}$. Let $M$ denote the set of all such observations. Finally, let $\hat{\mathbf{o}}_{t}^{r r^{\prime}}$ denote an observation made by a robot sensor at time $t$, in which robot $r$ measures the relative pose of robot $r^{\prime}$, and let $O$ 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 $\hat{\mathbf{m}}_{t t^{\prime}}^{r}$ or $\hat{\mathbf{o}}_{t}^{r r^{\prime}}$ can be thought of as a link 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 $\hat{\mathbf{o}}_{t}^{r r^{\prime}}$ is not equivalent to an observation $\hat{\mathbf{o}}_{t}^{r^{\prime} r}$, for example.

Our aim is to determine the set of pose estimates $X$ that maximizes the probability of obtaining the set of observations $(M, O)$; i.e., we seek to maximize the conditional probability $P(M, O \mid X)$. If we assume that observations are statistically independent, we can write this probability
as:

$$
\begin{align*}
P(M, O \mid X) & =\prod_{\hat{\mathbf{m}}_{t t^{\prime}}^{r} \in M} P\left(\hat{\mathbf{m}}_{t t^{\prime}}^{r} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right) \\
& \times \prod_{\hat{\mathbf{o}}_{t}^{r r^{\prime}} \in O} P\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right) \tag{1}
\end{align*}
$$

where $P\left(\hat{\mathbf{m}}_{t t^{\prime}}^{r} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right)$ is the probability of obtaining the individual motion observation $\hat{\mathbf{m}}_{t t^{\prime}}^{r}$, 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 $\hat{\mathbf{x}}_{t^{\prime}}^{r}$. Note that we have made the additional (but not unreasonable) assumption that this probability is independent of other pose estimates. In a similar vein, $P\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right)$ specifies the probability of obtaining the individual robot observation $\hat{\mathbf{o}}_{t}^{r r^{\prime}}$, 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 $r^{\prime}$ being observed is $\hat{\mathbf{x}}_{t}^{r^{\prime}}$. Taking the logarithm of Equation 1, we can write:

$$
\begin{align*}
U(M, O \mid X) & =\sum_{\hat{\mathbf{m}}_{t t^{\prime}}^{r} \in M} U\left(\hat{\mathbf{m}}_{t t^{\prime}}^{r} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right) \\
& +\sum_{\hat{\mathbf{o}}_{t}^{r r^{\prime}} \in O} U\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right) \tag{2}
\end{align*}
$$

where

$$
\begin{align*}
U(M, O \mid X) & =-\log P(M, O \mid X) \\
U\left(\hat{\mathbf{m}}_{t t^{\prime}}^{r} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right) & =-\log P\left(\hat{\mathbf{m}}_{t t^{\prime}}^{r} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right) \\
U\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right) & =-\log P\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right) \tag{3}
\end{align*}
$$

This latter form is somewhat more convenient for numerical optimization. Our aim is now to find the set of estimates $X$ that minimizes $U(M, O \mid X)$, for which we will need to determine the form of the conditional log-probabilities $U\left(\hat{\mathbf{m}}_{t t^{\prime}}^{r} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right)$ and $U\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right)$.

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

$$
\begin{equation*}
\hat{\mathbf{m}}_{t t^{\prime}}^{r}=\left(\boldsymbol{\mu}_{t t^{\prime}}^{r}, \boldsymbol{\Sigma}_{t t^{\prime}}^{r}\right) \tag{4}
\end{equation*}
$$

where $\boldsymbol{\mu}_{t t^{\prime}}^{r}$ is a relative pose measurement and $\boldsymbol{\Sigma}_{t t^{\prime}}^{r}$ is a covariance matrix representing the uncertainty in this measurement. The conditional log-probability for such observations is given by:

$$
\begin{equation*}
U\left(\hat{\mathbf{m}}_{t, t^{\prime}}^{r} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right)=\frac{1}{2}\left(\boldsymbol{\mu}_{t t^{\prime}}^{r}-\hat{\mathbf{y}}_{t t^{\prime}}^{r}\right)^{T} \boldsymbol{\Sigma}_{t t^{\prime}}^{r}\left(\boldsymbol{\mu}_{t t^{\prime}}^{r}-\hat{\mathbf{y}}_{t t^{\prime}}^{r}\right) \tag{5}
\end{equation*}
$$

where $\hat{\mathbf{y}}_{t t^{\prime}}^{r}$ is a relative pose estimate describing the estimated change in pose of robot $r$ between times $t$ and $t^{\prime}$. The relative pose estimate is derived from the absolute pose estimates $\hat{\mathbf{x}}_{t}^{r}$ and $\hat{\mathbf{x}}_{t^{\prime}}^{r}$ via some coordinate transform $\Gamma_{m}$ :

$$
\begin{equation*}
\hat{\mathbf{y}}_{t t^{\prime}}^{r}=\Gamma_{m}\left(\hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right) \tag{6}
\end{equation*}
$$

The specific form of $\Gamma_{m}$ depends on the coordinate system chosen to represent the absolute pose estimates $X$ and the motion observations $\hat{M}$. We will consider one specific form
for $\Gamma_{m}$ 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:

$$
\begin{equation*}
\hat{\mathbf{o}}_{t}^{r r^{\prime}}=\left(\boldsymbol{\mu}_{t}^{r r^{\prime}}, \boldsymbol{\Sigma}_{t}^{r r^{\prime}}\right) \tag{7}
\end{equation*}
$$

where $\boldsymbol{\mu}_{t}^{r r^{\prime}}$ is the relative pose of robot $r^{\prime}$, as measured by robot $r$ at time $t ; \boldsymbol{\Sigma}_{t}^{r r^{\prime}}$ is the covariance matrix representing the uncertainty in this measurement. The conditional logprobability is given by:

$$
\begin{equation*}
U\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right)=\frac{1}{2}\left(\boldsymbol{\mu}_{t}^{r r^{\prime}}-\hat{\mathbf{y}}_{t}^{r r^{\prime}}\right)^{T} \boldsymbol{\Sigma}_{t}^{r r^{\prime}}\left(\boldsymbol{\mu}_{t}^{r r^{\prime}}-\hat{\mathbf{y}}_{t}^{r r^{\prime}}\right) \tag{8}
\end{equation*}
$$

where $\hat{\mathbf{y}}_{t}^{r r^{\prime}}$ is the estimated pose of robot $r^{\prime}$ relative to robot $r$ at time $t$. The relative pose estimate is derived from the absolute pose estimates $\hat{\mathbf{x}}_{t}^{r}$ and $\hat{\mathbf{x}}_{t}^{r^{\prime}}$ via some coordinate transform $\Gamma_{o}$ :

$$
\begin{equation*}
\hat{\mathbf{y}}_{t}^{r r^{\prime}}=\Gamma_{o}\left(\hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right) \tag{9}
\end{equation*}
$$

As with the motion observations, the specific form of $\Gamma_{o}$ depends on the coordinate systems being used.

Given appropriate definitions for $\Gamma_{m}$ and $\Gamma_{o}$, one can determine the optimal set of pose estimates $X$ using a standard numerical optimization algorithm. The selection of an appropriate algorithm is driven largely by the form of these coordinate transforms, which are, in general, non-linear by differentiable. This rules out fast linear algorithms, but allows gradient-based techniques such as steepest descent and conjugate gradient algorithms [16]. Such algorithms require that we compute the gradient of $U(M, O \mid X)$ with respect to $X$; this can be done by differentiating through Equation 2:

$$
\begin{align*}
\frac{\partial}{\partial \hat{\mathbf{x}}_{t}^{r}} U(M, O \mid X) & =\sum_{\hat{\mathbf{m}}_{t t^{\prime}}^{r} \in M} \frac{\partial}{\partial \hat{\mathbf{x}}_{t}^{r}} U\left(\hat{\mathbf{m}}_{t t^{\prime}}^{r} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right) \\
& +\sum_{\hat{\mathbf{m}}_{t^{\prime} t}^{r} \in M} \frac{\partial}{\partial \hat{\mathbf{x}}_{t}^{r}} U\left(\hat{\mathbf{m}}_{t^{\prime} t}^{r} \mid \hat{\mathbf{x}}_{t^{\prime}}^{r}, \hat{\mathbf{x}}_{t}^{r}\right) \\
& +\sum_{\hat{\mathbf{o}}_{t}^{r r^{\prime}} \in O} \frac{\partial}{\partial \hat{\mathbf{x}}_{t}^{r}} U\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right) \\
& +\sum_{\hat{\mathbf{o}}_{t}^{r^{\prime} r} \in O} \frac{\partial}{\partial \hat{\mathbf{x}}_{t}^{r}} U\left(\hat{\mathbf{o}}_{t}^{r^{\prime} r} \mid \hat{\mathbf{x}}_{t}^{r^{\prime}}, \hat{\mathbf{x}}_{t}^{r}\right) \tag{10}
\end{align*}
$$

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. Individual derivatives for the motion observation terms can be evaluated by applying the chain-rule to Equation 5:

$$
\begin{align*}
\frac{\partial}{\partial \hat{\mathbf{x}}_{t}^{r}} U\left(\hat{\mathbf{m}}_{t t^{\prime}}^{r} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t^{\prime}}^{r}\right) & =\frac{\partial \hat{\mathbf{y}}_{t t^{\prime}}^{r}}{\partial \hat{\mathbf{x}}_{t}^{r}} \boldsymbol{\Sigma}_{t t^{\prime}}^{r}\left(\boldsymbol{\mu}_{t t^{\prime}}^{r}-\hat{\mathbf{y}}_{t t^{\prime}}^{r}\right) \\
\frac{\partial}{\partial \hat{\mathbf{x}}_{t}^{r}} U\left(\hat{\mathbf{m}}_{t^{\prime} t}^{r} \mid \hat{\mathbf{x}}_{t^{\prime}}^{r}, \hat{\mathbf{x}}_{t}^{r}\right) & =\frac{\partial \hat{\mathbf{y}}_{t^{\prime} t}^{r}}{\partial \hat{\mathbf{x}}_{t^{\prime}}^{r}} \boldsymbol{\Sigma}_{t^{\prime} t}^{r}\left(\boldsymbol{\mu}_{t^{\prime} t}^{r}-\hat{\mathbf{y}}_{t^{\prime} t}^{r}\right) \tag{11}
\end{align*}
$$

where the derivatives $\partial \hat{\mathbf{y}}_{t t^{\prime}}^{r} / \partial \hat{\mathbf{x}}_{t}^{r}$ and $\partial \hat{\mathbf{y}}_{t^{\prime} t}^{r} / \partial \hat{\mathbf{x}}_{t}^{r}$ can be determined by differentiating Equation 6, given some specific form for $\Gamma_{m}$, The derivatives for the robot observation terms can be determined in a similar fashion by applying the chain-rule to Equation 8:

$$
\begin{align*}
\frac{\partial}{\partial \hat{\mathbf{x}}_{t}^{r}} U\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right) & =\frac{\partial \hat{\mathbf{y}}_{t}^{r r^{\prime}}}{\partial \hat{\mathbf{x}}_{t}^{r}} \boldsymbol{\Sigma}_{t}^{r r^{\prime}}\left(\boldsymbol{\mu}_{t}^{r r^{\prime}}-\hat{\mathbf{y}}_{t}^{r r^{\prime}}\right) \\
\frac{\partial}{\partial \hat{\mathbf{x}}_{t}^{r}} U\left(\hat{\mathbf{o}}_{t}^{\boldsymbol{\prime}^{\prime} r} \mid \hat{\mathbf{x}}_{t}^{r^{\prime}}, \hat{\mathbf{x}}_{t}^{r}\right) & =\frac{\partial \hat{\mathbf{y}}_{t}^{r^{\prime} r}}{\partial \hat{\mathbf{x}}_{t}^{r}} \boldsymbol{\Sigma}_{t}^{r^{\prime} r}\left(\boldsymbol{\mu}_{t}^{r^{\prime} r}-\hat{\mathbf{y}}_{t}^{r^{\prime} r}\right) \tag{12}
\end{align*}
$$

where the derivatives $\partial \hat{\mathbf{y}}_{t}^{r r^{\prime}} / \partial \hat{\mathbf{x}}_{t}^{r}$ and $\partial \hat{\mathbf{y}}_{t}^{r^{\prime} r} / \partial \hat{\mathbf{x}}_{t}^{r}$ are determined by differentiating Equation 9 for some specific $\Gamma_{o}$.

## 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:

$$
\hat{\mathbf{x}}_{t}^{r}=\left[\begin{array}{l}
x_{t x}^{r}  \tag{13}\\
x_{t y}^{r} \\
x_{t \theta}^{r}
\end{array}\right]
$$

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

$$
\boldsymbol{\mu}_{t t^{\prime}}^{r}=\left[\begin{array}{l}
\mu_{t t^{\prime} d}^{r}  \tag{14}\\
\mu_{t t^{\prime} \phi}^{r} \\
\mu_{t t^{\prime} \psi}^{r}
\end{array}\right]
$$

where $\left(\mu_{t t^{\prime} d}^{r}, \mu_{t t^{\prime} \phi}^{r}\right)$ describes the measured change in position for robot $r$ (in polar coordinates; i.e. range and bearing) and $\mu_{t t^{\prime} \psi}^{r}$ describes its change in orientation. The relative pose measurements $\boldsymbol{\mu}_{t}^{r r^{\prime}}$ for robot observations are defined in a similar manner:

$$
\boldsymbol{\mu}_{t}^{r r^{\prime}}=\left[\begin{array}{l}
\mu_{t d}^{r r^{\prime}}  \tag{15}\\
\mu_{t \phi}^{r r^{\prime}} \\
\mu_{t \psi}^{r r^{\prime}}
\end{array}\right]
$$

where $\left(\mu_{t d}^{r r^{\prime}}, \mu_{t \phi}^{r r^{\prime}}\right)$ describes the measured range and bearing of robot $r^{\prime}$ relative to $r$, and $\mu_{t \psi}^{r r^{\prime}}$ describes their relative orientation.

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 is a good approximation to the classic 'banana-shaped' distribution that is usually observed with such sensors [3]. Note, however, that if one was 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 these definitions, we can write down an expression for the relative pose estimates $\hat{\mathbf{y}}_{t t^{\prime}}^{r}$ : the coordinate transform function $\Gamma_{m}$ simply transforms from polar to Cartesian coordinates:

$$
\hat{\mathbf{y}}_{t t^{\prime}}^{r}=\left[\begin{array}{c}
y_{t t^{\prime} d}^{r}  \tag{16}\\
y_{t t^{\prime} \phi}^{r} \\
y_{t t^{\prime} \psi}^{r}
\end{array}\right]=\left[\begin{array}{c}
\sqrt{\left(\Delta_{t t^{\prime} x}^{r}\right)^{2}+\left(\Delta_{t t^{\prime} y}^{r}\right)^{2}} \\
\arctan \left(\Delta_{t t^{\prime} y}^{r} / \Delta_{t t^{\prime} x}^{r}\right)-x_{t \theta}^{r} \\
\Delta_{t t^{\prime} \theta}^{r}
\end{array}\right]
$$

where $\hat{\boldsymbol{\Delta}}_{t t^{\prime}}^{r}=\hat{\mathbf{x}}_{t^{\prime}}^{r}-\hat{\mathbf{x}}_{t}^{r}$. For optimization, we require the corresponding derivatives, which are given by:

$$
\begin{align*}
& \frac{\partial \hat{\mathbf{y}}_{t t^{\prime}}^{r}}{\partial \hat{\mathbf{x}}_{t}^{r}}= {\left[\begin{array}{lll}
\partial y_{t t^{\prime} d}^{r} / \partial x_{t x}^{r} & \partial y_{t t^{\prime} \phi}^{r} / \partial x_{t x}^{r} & \partial y_{t t^{\prime} \psi}^{r} / \partial x_{t x}^{r} \\
\partial y_{t t^{\prime} d}^{r} / \partial x_{t y}^{r} & \partial y_{t t^{\prime} \phi}^{r} / \partial x_{t y}^{r} & \partial y_{t t^{\prime} \psi}^{r} / \partial x_{t y}^{r} \\
\partial y_{t t^{\prime} d}^{r} / \partial x_{t \theta}^{r} & \partial y_{t t^{\prime} \phi}^{r} / \partial x_{t \theta}^{r} & \partial y_{t t^{\prime} \psi}^{r} / \partial x_{t \theta}^{r}
\end{array}\right] } \\
&= {\left[\begin{array}{ccc}
-\Delta_{t t^{\prime} x}^{r} / y_{t t^{\prime} d}^{r} & +\Delta_{t t^{\prime} y}^{r} /\left(y_{t t^{\prime} d}^{r}\right)^{2} & 0 \\
-\Delta_{t t^{\prime} y}^{r} / y_{t t^{\prime} d}^{r} & -\Delta_{t t^{\prime} x}^{r} /\left(y_{t t^{\prime} d}^{r}\right)^{2} & 0 \\
0 & -1 & -1
\end{array}\right] } \\
& \frac{\partial \hat{\mathbf{y}}_{t^{\prime} t}^{r}}{\partial \hat{\mathbf{x}}_{t}^{r}}=\left[\begin{array}{ccc}
+\Delta_{t t^{\prime} x}^{r} / y_{t t^{\prime} d}^{r} & -\Delta_{t t^{\prime} y}^{r} /\left(y_{t t^{\prime} d}^{r}\right)^{2} & 0 \\
+\Delta_{t t^{\prime} y}^{r} / y_{t t^{\prime} d}^{r} & +\Delta_{t t^{\prime} x}^{r} /\left(y_{t t^{\prime} d}^{r}\right)^{2} & 0 \\
0 & 0 & +1
\end{array}\right] \tag{17}
\end{align*}
$$

There are two features of these derivatives that should be noted. First, $\partial \hat{\mathbf{y}}_{t t^{\prime}}^{r} / \partial \hat{\mathbf{x}}_{t}^{r} \neq-\partial \hat{\mathbf{y}}_{t^{\prime} t}^{r} / \partial \hat{\mathbf{x}}_{t}^{r}$ as one might naively expect. Second, the derivatives contain a singularity at $y_{t t^{\prime} d}^{r}=0$; one must take care to avoid this singularity during the optimization process.

The relative pose estimates $\hat{\mathbf{y}}_{t}^{r r^{\prime}}$ are treated in a similar manner, since $\Gamma_{o}$ also transforms from polar to Cartesian coordinates:

$$
\hat{\mathbf{y}}_{t}^{r r^{\prime}}=\left[\begin{array}{c}
y_{t d}^{r r^{\prime}}  \tag{18}\\
y_{t \phi}^{r r^{\prime}} \\
y_{t \psi}^{r r^{\prime}}
\end{array}\right]=\left[\begin{array}{c}
\sqrt{\left(\Delta_{t x}^{r r^{\prime}}\right)^{2}+\left(\Delta_{t y}^{r r^{\prime}}\right)^{2}} \\
\arctan \left(\Delta_{t y}^{r r^{\prime}} / \Delta_{t x}^{r r^{\prime}}\right)-x_{t \theta}^{r} \\
\Delta_{t \theta}^{r r^{\prime}}
\end{array}\right]
$$

where $\hat{\boldsymbol{\Delta}}_{t}^{r r^{\prime}}=\hat{\mathbf{x}}_{t}^{r^{\prime}}-\hat{\mathbf{x}}_{t}^{r}$. The corresponding derivatives are:

$$
\begin{align*}
\frac{\partial \hat{\mathbf{y}}_{t}^{r r^{\prime}}}{\partial \hat{\mathbf{x}}_{t}^{r}} & =\left[\begin{array}{ccc}
-\Delta_{t x}^{r r^{\prime}} / y_{t d}^{r r^{\prime}} & +\Delta_{t y}^{r r^{\prime}} /\left(y_{t d}^{r r^{\prime}}\right)^{2} & 0 \\
-\Delta_{t y}^{r r^{\prime}} / y_{t d}^{r r^{\prime}} & -\Delta_{t x}^{r r^{\prime}} /\left(y_{t d}^{r r^{\prime}}\right)^{2} & 0 \\
0 & -1 & -1
\end{array}\right] \\
\frac{\partial \hat{\mathbf{y}}_{t}^{r^{\prime} r}}{\partial \hat{\mathbf{x}}_{t}^{r}} & =\left[\begin{array}{ccc}
+\Delta_{t x}^{r r^{\prime}} / y_{t d}^{r r^{\prime}} & -\Delta_{t y}^{r r^{\prime}} /\left(y_{t d}^{r r^{\prime}}\right)^{2} & 0 \\
+\Delta_{t y}^{r r^{\prime}} / y_{t d}^{r r r^{\prime}} & +\Delta_{t x}^{r r^{\prime}} /\left(y_{t d}^{r r^{\prime}}\right)^{2} & 0 \\
0 & 0 & +1
\end{array}\right] \tag{19}
\end{align*}
$$

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

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.


Fig. 2. Sample probability distributions for for the planar localization problem. The plots show the probability $P\left(\hat{\mathbf{o}}_{t}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}\right)$ as a function of the estimated position $\left(\hat{\mathbf{x}}_{t x}^{r^{\prime}}, \hat{\mathbf{x}}_{t y}^{r^{\prime}}\right)$ for the robot being observed. Orientation is not shown. (a) The range is well determined $\left(\mu_{t d}^{r r^{\prime}}=1 \pm 0.1 \mathrm{~m}\right)$, but the bearing is unknown. (b) The bearing is moderately well determined $\left(\mu_{t \phi}^{r r^{\prime}}=45 \pm 57^{\circ}\right)$, but the range in unknown. (c) Both range and bearing are (moderately) well determined ( $\mu_{t d}^{r r^{\prime}}=1 \pm 0.1 \mathrm{~m}, \mu_{t \phi}^{r r^{\prime}}=45 \pm 57^{\circ}$ ).


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

## IV. Implementation

## A. Optimization: The Conjugate Gradient Algorithm

For optimization, we use the conjugate gradient algorithm described in [16]; this algorithm is somewhat more complex than the standard steepest descent algorithm, but 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 $X$ 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 $U(M, O \mid X)$ plotted along a line through the high-dimensional space that is $X$; 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


Fig. 4. An illustration of the extended formalism. The figure shows two robots, $r=11$ 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 $\hat{p}_{1}$ and $\hat{p}_{2}$ for each of the robots at time $t=2$.
workstation).
We also adopt a incremental approach to optimization: as each new observation is added to $M$ or $O$, we optimize over the entire set of pose estimates $X$. While it is somewhat inefficient, this incremental approach helps to ensure that the algorithm converges to a global rather than a local minimum (we have a very good initial 'guess' at the start of each optimization pass). While this approach does guarantee convergence to the global minimum, it appears to suffice in practice.

## B. Practical considerations: bounding $X, M$ and $O$

The dimensionality of the problem that must be solved by the conjugate gradient algorithm scales linearly with the size of $X$; the computational cost of each step in this algorithm scales linearly with the size of $\{M \cup O\}$. It is therefore necessary, in practice, to bound both the number of estimates in $X$ and the number of observations in $\{M \cup$ $O\}$. 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: 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 the temporal 'window'; and 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 number of observations $m$.

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 the robots at every point in time, we instead estimate the pose at only a few discrete points in time, and use information from the motion sensors to 'fill the gaps' between these estimates. In effect, we make two additional assumptions about the motion sensors: that they produce observations at a very high rate, and that these observations are relatively accurate. Thus, the motion sensors alone can be used to generate good pose estimates that require only occasional corrections. Let $\hat{\mathbf{z}}_{t+}^{r}$ be the interpolated pose estimate for robot $r$ at time $t^{+}$; this estimate is given by:

$$
\begin{equation*}
\hat{\mathbf{z}}_{t^{+}}^{r}=\Gamma_{m}^{-1}\left(\hat{\mathbf{x}}_{t}^{r}, \boldsymbol{\mu}_{t t^{+}}^{r}\right) \tag{20}
\end{equation*}
$$

where $\hat{\mathbf{x}}_{t}^{r}$ is the most recent pose estimate for robot $r$ in $X$ and $\boldsymbol{\mu}_{t t^{+}}^{r}$ is the measured change in pose between times $t$ and $t^{+}$. The inverse coordinate transform $\Gamma_{m}^{-1}$ maps from relative to absolute coordinates. 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 Equations 8:

$$
\begin{equation*}
U\left(\hat{\mathbf{o}}_{t^{+}}^{r r^{\prime}} \mid \hat{\mathbf{x}}_{t}^{r}, \hat{\mathbf{x}}_{t}^{r^{\prime}}, \boldsymbol{\mu}_{t t^{+}}^{r}\right)=\frac{1}{2}\left(\boldsymbol{\mu}_{t}^{r r^{\prime}}-\hat{\mathbf{y}}_{t^{+}}^{r r}\right)^{T} \boldsymbol{\Sigma}_{t}^{r r^{\prime}}\left(\boldsymbol{\mu}_{t}^{r r^{\prime}}-\hat{\mathbf{y}}_{t^{+}}^{r r}\right) \tag{21}
\end{equation*}
$$

where

$$
\begin{equation*}
\hat{\mathbf{y}}_{t^{+}}^{r r^{\prime}}=\Gamma_{o}\left(\hat{\mathbf{z}}_{t^{+}}^{r}, \hat{\mathbf{z}}_{t^{+}}^{r^{\prime}}\right) \tag{22}
\end{equation*}
$$

The corresponding derivatives must also be modified to yield:

$$
\begin{equation*}
\frac{\partial \hat{\mathbf{y}}_{t^{+}}^{r r^{\prime}}}{\partial \hat{\mathbf{x}}_{t}^{r}} \longrightarrow \frac{\partial \hat{\mathbf{y}}_{t^{+}}^{r r^{\prime}}}{\partial \hat{\mathbf{x}}_{t}^{r}}=\frac{\partial \hat{\mathbf{y}}_{t^{+}}^{r r^{\prime}}}{\partial \hat{\mathbf{z}}_{t^{+}}^{r}} \frac{\partial \hat{\mathbf{z}}_{t^{+}}^{r}}{\partial \hat{\mathbf{x}}_{t}^{r}} \tag{23}
\end{equation*}
$$

Tthe additional derivative $\partial \hat{\mathbf{z}}_{t^{+}}^{r} / \partial \hat{\mathbf{x}}_{t}^{r}$ can be computed by differentiating through Equation 20, once the form of $\Gamma_{m}^{-1}$ is known. These definitions are illustrated in Figure 4.

This extended formalism allows us 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 tradeoff computational cost against localization accuracy.

## V. Experimental Validation

## A. Experiment 1

We have conducted a series of experiments aimed at validating both the practicality and the accuracy of the team localization algorithm presented in this paper. Our 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 SICK lasers, which can be programmed to return intensity information in addition to the normal range information. 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 $180^{\circ}$ 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 manually label the data. Ground-truth information 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; importantly, the robots were arranged such that at no time were the two robots on the outer wall able to directly sense one other. The structure of the environment was modified a number of times during the course of the experiment; at time $t=265 \mathrm{sec}$, for example, the inner wall was modified to form two separate 'islands', with one robot circumnavigating each island. The original structure was later restored, then modified, the 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 $0<t<12 \mathrm{sec}$, the robots commence wall following; by time $t=12 \mathrm{sec}$, the robots following the outer wall have observered 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 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 $t=265 \mathrm{sec}$, the environment is modified, with the inner wall being re-structured to form two separate islands. The localization, however, is un-affected by this change, as shown in the snap-shot at time $t=272 \mathrm{sec}$. Thus, a key feature of the method described in this paper is that it is largely indifferent to such structural changes in the environment.

The accuracy of the localization algorithm is determined by comparing the relative pose estimates $\hat{\mathbf{y}}_{t}^{r r^{\prime}}$ against their corresponding true values (as determined by the groundtruth system). That is, at each time $t$ we measure how accurately each robot has estimated the relative pose of all


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 re-arranged 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 cummulative 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.
the other robots. Accuracy is measured by an error term of the form:

$$
\begin{equation*}
\left(\epsilon_{t}\right)^{2}=\frac{1}{n(n-1)} \sum_{r} \sum_{r^{\prime}}\left(\hat{\mathbf{y}}_{t}^{r r^{\prime}}-\mathbf{y}_{t}^{r r^{\prime}}\right)^{T}\left(\hat{\mathbf{y}}_{t}^{r r^{\prime}}-\mathbf{y}_{t}^{r r^{\prime}}\right) \tag{24}
\end{equation*}
$$

where $\mathbf{y}_{t}^{r r^{\prime}}$ is the true pose of robot $r^{\prime}$ relative to robot $r$ at time $t$. The summation is over all pairs of robots and the result is normalized by the number of robots $n$ to generate an average result. Note that we make no attempt to compare the absolute pose estimates $\hat{\mathbf{x}}_{t}^{r}$ against some 'true' value; 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 $\epsilon_{t}$ as a function of time. At time $t=0 \mathrm{sec}$, the relative pose of the robots is completely unknown, and the errors are correspondingly high. By time $t=12 \mathrm{sec}$, 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 $5.5 \pm 5.2 \mathrm{~cm}$, while the bearing and orientation errors oscillate around $1.7 \pm 0.7^{\circ}$ and $1.9 \pm 0.6^{\circ}$ respectively. The magnitude these errors can be attributed to two key factors. First, there is some uncertainty in the relative pose measurements made by the laser-range-finder/retroreflector combination. These uncertainties are difficult to characterize precisely, but are of the order of $\pm 2.5 \mathrm{~cm}$ and


Fig. 9. Plot showing the relative range error as a function of time for Experiment 2; bearing and orientation errors were not measured.
$\pm 2^{\circ}$. 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 time, the robot may travel 2 cm and/or rotate through $3^{\circ}$, which will significantly bias the results.

We ascribe the variation seen in the error plots two different 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 odometric accuracy of the robots used in this experiment varies from quite good to quite poor: drift rates for orientation, for example, vary from $2.5^{\circ}$ per lap to $30^{\circ}$ per lap. Second, we expect that errors 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 produced by the optimization algorithm, which may take several cycles (each cycle is 100 ms ) to incoporate the new data and relax to a new set of pose estimates.

Finally, we note that there is a major spike in the plot at around $t=300 \mathrm{sec}$. 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 re-positioned, leading to gross errors in both robot's odometry. Nevertheless, as the plot indicates, the localization algorithm was able to quickly recover.

## B. Experiment 2

The second experiment was conducted using a team of 7 Pioneer 2DX mobile robots, each of which was 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 then 'joy-sticked' around the circuit, and was thus observed 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 tapemeasure. Bearings and orientations were not measured.

The quantitative results for this experiment are shown in Figure 9, which shows a plot of the relative range error $\epsilon_{t r}$ as a function of time. Also 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 at least once (recall that these robots cannot see one-another). 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.
Qualitative results for this experiment are shown in Figures 8(b) and (c); each figure is a snap-shot showing pose estimates, observations, and the laser scan data for a particular point in time. In Figure 8 b , the mobile robot has been seen by each of the stationary robots exactly once. As a result of the cumulative error in this robot's motion observations, however, the overall error in the pose estimates remains relatively high. In Figure 8(c) 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 $0.08 \pm 0.09 \mathrm{~m}$. This value is quite remarkable when one considers that the loop traversed by the mobile robot is about 80 m in length.

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

## VI. Conclusion and Further Work

The experiments described in the previous section demonstrate several key capabilities of the team localization method: 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 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 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 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 estimates in $X$.

In closing, 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 between pairs of robots. This matrix can then be


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 $t=200 \mathrm{sec}$, 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 $t=220 \mathrm{sec}$, after the mobile robot has been seen by the first stationary robot (Tanis) for a second time, thus closing the loop.
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.

## 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] Reid Simmons and Sven 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] Sebastian Thrun, Dieter Fox, and Wolfram Burgard, "A probabilistic approach to concurrent mapping and localisation for mobile robots," Machine Learning, vol. 31, no. 5, pp. 29-55, 1998, Joint issue with Autonomous Robots.
[5] F Lu and E Milios, "Globally consistent range scan alignment for environment mapping," Autonomous Robots, vol. 4, pp. 333349, 1997.
[6] Brian Yamauchi, Alan Shultz, and William 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. DurrantWhyte, 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, "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.
[12] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localization," $A u$ tonomous Robots, Special Issue on Heterogeneous Multi-Robot Systems, vol. 8, no. 3, pp. 325-344, 2000.
[13] Ryo Kurazume and Shigeo Hirose, "An experimental study of a cooperative positioning system," Autonomous Robots, vol. 8, no. 1, pp. 43-52, 2000.
[14] I. M. Rekleitis, G. Dudek, and E. Milios, "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.
[15] Andrew Howard and Les Kitchen, "Cooperative localisation and mapping," in International Conference on Field and Service Robotics (FSR99), 1999, pp. 92-97.
[16] 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.

