Distributed Multi-Robot Localization

Stergios I. Roumeliotis and George A. Bekey

Abstract— **In this paper we present a new approach to the problem of simultaneously localizing a group of mobile robots capable of sensing one another. Each of the robots collects sensor data regarding its own motion and shares this information with the rest of the team during the update cycles. A single estimator, in the form of a Kalman filter, processes the available positioning information from all the members of the team and produces a pose estimate for every one of them. The equations for this centralized estimator can be written in a decentralized form therefore allowing this single Kalman filter to be decomposed into a number of smaller communicating filters. Each of these filters processes sensor data collected by its host robot. Exchange of information between the individual filters is necessary only when two robots detect each other and measure their relative pose. The resulting decentralized estimation schema, which we call *collective localization* constitutes a unique means for fusing measurements collected from a variety of sensors with minimal communication and processing requirements. The distributed localization algorithm is applied to a group of 3 robots and the improvement in localization accuracy is presented. Finally, a comparison to the equivalent decentralized information filter is provided.**

Keywords— **mobile robot localization, multi-robot sensor fusion, distributed Kalman filtering, decentralized estimation.**

I. Introduction

PRECISE localization is one of the main requirements for mobile robot autonomy [1]. Indoors and outdoors, mobile robots need to know their exact position and orientation (pose) in order to perform their required tasks. There have been numerous approaches to the localization problem utilizing different types of sensors [2] and a variety of techniques (e.g. [3], [4], [5], [6], [7]. The key idea behind most of the current localization schemes is to optimally combine measurements from *proprioceptive* sensors that monitor the motion of the vehicle with information collected by *exteroceptive* sensors that provide a representation of the environment and its signals. The various approaches to localization differ primarily in the type of estimator invoked in order to filter out the measurement noise and reduce the uncertainty associated with the interpretation of the sensor signals.

Many robotic applications require that robots work in collaboration [8] in order to perform a certain task [9], [10], [11], [12], [13], [14]. Most existing localization approaches have been developed for the case of a single robot. Even when a group of, say M, robots is considered, the group localization problem is usually resolved by independently solving M pose estimation problems. Each robot estimates its position based on its individual experience (proprioceptive and exteroceptive sensor measurements). Knowledge from the different entities of the team is not combined and each member must rely on its own resources (sensing and processing capabilities). This is a relatively simple approach since it avoids dealing with the complicated problem of fusing information from a large number of independent and interdependent sources. On the other hand, a more coordinated schema for localization has a number of advantages that can compensate for the added complexity.

First, let us consider the case of a homogeneous group of robots. As mentioned before, robotic sensing modalities suffer from uncertainty and noise. When a number of robots equipped with the same sensors detect a particular feature of the environment, such as a door, or measure a characteristic property of the area, such as the local vector of the earth's magnetic field, a number of *independent* measurements originating from the different members of the group is collected. Properly combining all this information will result in a single estimate of increased accuracy and reduced uncertainty. For example, a better estimate of the position and orientation of a landmark can in turn drastically improve the outcome of the robot localization process and thus this group will benefit from this collaboration schema.

The advantages stemming from the exchange of information among the members of a group are more crucial in the case of heterogeneous robotic colonies. When a team of robots is composed of different platforms carrying different proprioceptive and exteroceptive sensors and thus having different capabilities for self-localization, the quality of the localization estimates will vary significantly across the individual members. For example, a robot equipped with a laser rangefinder and expensive INS/GPS modules will outperform another member that must rely on wheel encoders and cheap sonars for localization. Communication and flow of information among the members of the group constitutes a form of *sensor sharing* and can improve the overall positioning accuracy. In fact, as will be evident later, if each robot could sense and communicate with its colleagues at all times then every member of the group would have less uncertainty about its position than the robot with the best localization results (if it were to localize itself without communicating with the rest of the group).

In the following section we outline previous approaches to the group localization problem and state the main differences between these approaches and the presented *collective localization* schema. In Section III we precisely formulate the multi-robot localization problem. Section IV presents the propagation equations for the initially decoupled system and describes the introduction of the covariance cross-correlation terms during the first Kalman filter

S. I. Roumeliotis is with the Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125 USA. E-mail: stergios@robotics.caltech.edu

G. A. Bekey is with the Computer Science Department, University of Southern California, Los Angeles, CA 90089 USA. E-mail: bekey@robotics.usc.edu

update. These terms capture the existing state and sensor data interdependencies. In Section V we present in detail the equations for the (coupled) centralized Kalman estimator that fuses all the sensor information collected by every robot in the group. It is shown that these filter equations can be distributed among the robots while communication is required only during the update cycle of the filter. Section VI contains an observability study for this estimator for different cases of interest, mainly with variations in the type of sensors available and task specifications. Section VII presents the experimental results from the application of the *collective localization* algorithm to the case of three robots moving in the same area. In Section VIII we compare the distributed Kalman filter estimator to its equivalent information filter for this particular application. Finally, Section IX derives the conclusions from this multi-robot localization schema and projects its potential extensions and applications.

II. Previous Approaches

An example of a system that is designed for cooperative localization was first reported in [15]. The authors acknowledge that dead-reckoning is not reliable for long traverses due to the error accumulation and introduce the concept of "portable landmarks". A group of robots is divided into two teams in order to perform cooperative positioning. At each time instant, one team is in motion while the other remains stationary and acts as a landmark. In the next phase the roles of the teams are reversed and this process continues until both teams reach their target. This method can work in unknown environments and the conducted experiments suggest average error of 0.4% for the position estimate and 1° for the orientation over distances of approximately 20m [16]. Improvements over this system and optimum motion strategies are discussed in [17], [18]. A similar realization is presented in [19]. Small size robots equipped with sonars measure their relative distances in order to update their position estimates. At any given time only one robot moves while the rest of the team forms an equilateral triangle of localization beacons. Another variation of this approach is described in [20], [21]. The authors deal with the problem of exploration of an unknown environment using two mobile robots. In order to reduce the odometric error, one robot is equipped with a camera tracking system that allows it to determine its relative position and orientation with respect to a second robot carrying a helix target pattern and acting as a portable landmark. All previous approaches have the following limitations: (a) Only one robot (or team) is allowed to move at any given time, and (b) The two robots (or teams) must maintain visual (or sonar) contact at all times.

A different implementation of a collaborative multi-robot localization schema is presented in [22], [23]. The authors have extended the Monte Carlo localization algorithm [24] to the case of two robots when a map of the area is available to both robots. When these robots detect each other, the combination of their belief functions facilitates their global localization task. The main limitation of this approach is that it can be applied only within known indoor environments. In addition, since information interdependencies are being ignored every time the two robots meet, this method can lead to overly optimistic position estimates. This issue is discussed in detail in [25]. Finally, a Kalman filter based implementation of a cooperative navigation schema is described in [26]. In this case the effect of the orientation uncertainty in both the state propagation and the relative position measurements is ignored resulting in a simplified distributed algorithm.

Although practices like those previously mentioned can be supported within our framework (Sections VI & VII), the advantage of the *collective localization* approach is that it provides a solution to the most general case in which all the robots in the group can move simultaneously while continuous visual contact or a map of the area is not required. In order to treat the group localization problem, we begin from the reasonable assumptions that the robots within the group can communicate with each other (at least 1-to-1 communication) and carry two types of sensors: 1. Proprioceptive sensors that record the self motion of each robot and allow for position tracking, 2. Exteroceptive sensors that monitor the environment for (a) (static) features and identities of the surroundings of the robot to be used in the localization process, and (b) other robots (treated as dynamic features). The goal is to integrate measurements collected by different robots and achieve localization across all the robotic platforms constituting the group.

The key idea for performing *collective localization* is that the group of robots must be viewed as one entity - the "group organism" - with multiple "limbs" (the individual robots in the group; the state of each of them is described by a vector $\vec{\mathbf{x}}_i = [x_i \; y_i \; \phi_i]^T$) and multiple virtual "joints" visualized as connecting each robot with every other member of the team ([27], [28]). The virtual "joints" provide 3 degrees of freedom $(\Delta x, \; \Delta y, \; \Delta\phi)$[1] and thus allow the "limbs" to move in every direction within a plane without any limitations. Considering this perspective, the "group organism" has access to a large number of sensors such as encoders, gyroscopes, cameras etc. In addition, it "spreads" itself across a large area and thus it can collect far more rich and diverse exteroceptive information than a single robot. When one robot detects another member of the team and measures its relative pose, it is equivalent to the "group organism's" joints measuring the relative displacement of these two "limbs". When two robots communicate for information exchange, this can be seen as the "group organism" allowing information to travel back and forth from its "limbs". This information can be fused by a centralized processing unit and provide improved localization results for all the robots in the group. An example of this is shown in Fig. 1. At this point we should note that a realization of a two-member "group organism" would resemble the multiple degree of freedom robot with compliant linkage, which was implemented by J. Borenstein [29], and has been shown to improve localization accuracy [30], [31].

[1] 6 dof in a 3 dimensional space of motion.

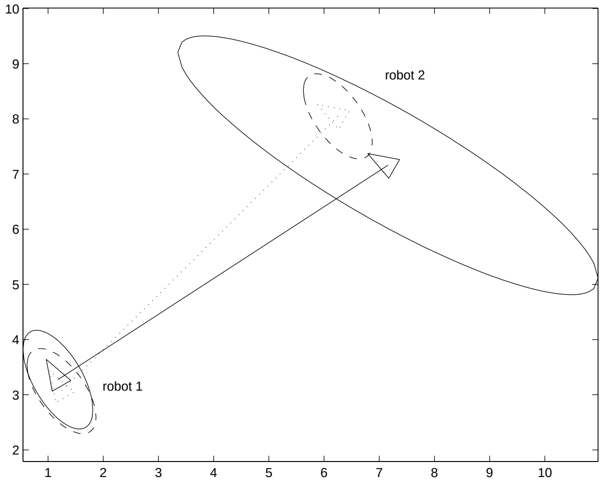

Fig. 1. Here before the two robots detected each other, robot 1 was able to accurately estimate its position while robot 2 was far more uncertain about its own pose estimates. Information flow between the two robots in the form of a relative pose measurement significantly reduces the localization uncertainty region for robot 2. Note that although to a lesser degree, robot 1 also benefits from this exchange. The solid/dashed lines represent the 3σ regions of confidence before/after the relative pose measurement.

The main drawback of addressing the cooperative localization problem as an information combination problem within a single entity ("group organism") is that it requires centralized processing and communication. The solution would be to decentralize the sensor fusion within the group. The *collective localization* approach accomplishes precisely this goal. Since the sensing modalities of the group are distributed, so too should be the processing modules. As will be evident in the following sections, our formulation differs from the aforementioned at its starting point. It is based on a unique characteristic of the multi-robot localization problem, namely that the *state propagation* equations of the centralized system are decoupled. State coupling occurs only when relative pose measurements become available.

III. Problem Statement

Here we specify the mobile robot group localization problem which this paper addresses. First, we state the following assumptions:

1. A group of M independent robots move in an $N-$ *dimensional* space. The motion of each robot is described by its own linear or non-linear equations of motion.

2. Each robot carries proprioceptive and exteroceptive sensing devices in order to propagate and update its own position estimate. These sensors measure the self motion of the robot and monitor the environment for localization features such as landmarks.

3. Each robot also carries exteroceptive sensors that allow it to detect and identify other robots moving in its vicinity and measure their respective displacement (relative position and orientation).

4. All robots are equipped with communication devices that allow exchange of information within the group.

The problem is to determine a principled way to exploit the information exchanged during the interactions among members of a group taking under consideration possible independencies and interdependencies. It is also within our focus to formulate the problem in such a way that it will allow for distributed processing with minimal communication requirements. As we will see in Sections IV and V, communication between the robots is only necessary during the update cycle of collective localization; that is, only when two robots meet.

As previously stated, our starting point is to consider this group of robots as a single centralized system composed of every individual robot moving in the area and capable of sensing and communicating with the rest of the group. In this centralized approach, the motion of the group is described in an $N \times M$-dimensional space and can be estimated by applying Kalman filtering techniques [32], [28]. The goal is to **treat the Kalman filter equations of the centralized system so as to distribute the estimation process among M Kalman filters, each of them operating on a different robot**. Here we will derive the equations for a group of $M = 3$ robots. The same steps describe the derivation for larger groups.

Kalman filter estimation can be divided into two cycles [32]: (a) The *Propagation* cycle where knowledge about the state of the system is propagated to the next (time) step based on the assumptions about the evolution of the system equations, the measured control inputs, and the statistical description of the system noise, and (b) The *Update* cycle where relative pose measurements are processed to update the propagated estimates calculated during the previous cycle.

A. Propagation

First, we describe the propagation equations for the Kalman filter using the velocity measurements from the odometric sensors. The estimated state vector for each member of the team consists of the robot's pose with respect to a fixed reference frame: $\vec{x}_i = [x_i \; y_i \; \phi_i]^T$, $i = 1..3$. For the vehicle's odometry, we use a generic set of equations. However, the method can be easily adapted to a specific vehicle. The continuous time equations for the motion expressed in local coordinates are:

$$\begin{array}{llll} {}^L\dot{x}_i = V_i, & {}^L\dot{y}_i = 0, & \dot{\phi}_i = \omega_i, & i = 1..3 \\ {}^L\dot{\hat{x}}_i = V_{im}, & {}^L\dot{\hat{y}}_i = 0, & \dot{\hat{\phi}}_i = \omega_{im}, & i = 1..3 \end{array} \quad (1)$$

with

$$V_{im} = V_i + w_{V_i} \ , \ \omega_{im} = \omega_i + w_{\omega_i} \ , \quad i = 1..3$$

where V_{im} and ω_{im} are the linear and rotational velocity of robot i as measured by the wheel-encoders' signals, V_i

and ω_i are the real values of these quantities, w_{V_i} and w_{ω_i} are zero-mean white Gaussian processes that represent the noise in the measured signals. Based on Eq. (1), the linearized discrete-time error-state propagation equation for robot i expressed in global coordinates is:

$$\begin{bmatrix} \widetilde{x}_i \\ \widetilde{y}_i \\ \widetilde{\phi}_i \end{bmatrix}_{t_{k+1}} = \begin{bmatrix} 1 & 0 & -V_{im}\delta t \sin\hat{\phi}_i \\ 0 & 1 & V_{im}\delta t \cos\hat{\phi}_i \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \widetilde{x}_i \\ \widetilde{y}_i \\ \widetilde{\phi}_i \end{bmatrix}_{t_k} + \begin{bmatrix} \delta t\cos\hat{\phi}_i & 0 \\ \delta t \sin\hat{\phi}_i & 0 \\ 0 & \delta t \end{bmatrix} \begin{bmatrix} w_{V_i} \\ w_{\omega_i} \end{bmatrix}_{t_k} \quad (2)$$

or

$$\widetilde{\vec{\mathbf{x}}}_{\mathbf{i}}(t_{k+1}) = \Phi_i(t_{k+1}, t_k)\widetilde{\vec{\mathbf{x}}}_{\mathbf{i}}(t_k) + G_i(t_k)\vec{\mathbf{w}}_{\mathbf{i}}(t_k) \quad (3)$$

for $i = 1..3$, where $\delta t = t_{k+1} - t_k$, $\widetilde{\vec{\mathbf{x}}}_{\mathbf{i}}$ is the error-state vector, Φ_i is the system propagation matrix, G_i is the system noise input matrix and $\vec{\mathbf{w}}_{\mathbf{i}}$ is the system noise due to the errors in the linear and rotational velocity measurements of robot i. The system noise covariance is given by:

$$Q_{di}(t_k) = G_i(t_k)E\{\vec{\mathbf{w}}_{\mathbf{i}}(t_k)\vec{\mathbf{w}}_{\mathbf{i}}^{\mathbf{T}}(t_k)\}G_i^T(t_k),\ i = 1..3 \quad (4)$$

Under the assumption that the motion of each robot does not affect, at least not directly, the motion of the other robots, the **error-state propagation** equation for the centralized (augmented) system can be written by combining Eq.s (3) as:

$$\begin{bmatrix} \widetilde{\vec{\mathbf{x}}}_1(t_{k+1}) \\ \widetilde{\vec{\mathbf{x}}}_2(t_{k+1}) \\ \widetilde{\vec{\mathbf{x}}}_3(t_{k+1}) \end{bmatrix} = \begin{bmatrix} \Phi_1 & 0 & 0 \\ 0 & \Phi_2 & 0 \\ 0 & 0 & \Phi_3 \end{bmatrix} \begin{bmatrix} \widetilde{\vec{\mathbf{x}}}_1(t_k) \\ \widetilde{\vec{\mathbf{x}}}_2(t_k) \\ \widetilde{\vec{\mathbf{x}}}_3(t_k) \end{bmatrix} + \begin{bmatrix} G_1 & 0 & 0 \\ 0 & G_2 & 0 \\ 0 & 0 & G_3 \end{bmatrix} \begin{bmatrix} \vec{\mathbf{w}}_1(t_k) \\ \vec{\mathbf{w}}_2(t_k) \\ \vec{\mathbf{w}}_3(t_k) \end{bmatrix}$$

or

$$\widetilde{\mathbf{x}}(t_{k+1}) = \mathbf{\Phi}(t_{k+1}, t_k)\widetilde{\mathbf{x}}(t_k) + G(t_k)\mathbf{w}(t_k) \quad (5)$$

Note that the centralized system matrix $\mathbf{\Phi}$ in the previous equation is diagonal:

$$\Phi(t_{k+1}, t_k) = \begin{bmatrix} \Phi_1(t_{k+1}, t_k) & 0 & 0 \\ 0 & \Phi_2(t_{k+1}, t_k) & 0 \\ 0 & 0 & \Phi_3(t_{k+1}, t_k) \end{bmatrix} \quad (6)$$

Each of the Φ_i matrices describes the motion of robot i. Similarly, the system noise covariance matrix Q_d for the centralized system would be:

$$Q_d(t_k) = \begin{bmatrix} Q_{d1}(t_k) & 0 & 0 \\ 0 & Q_{d2}(t_k) & 0 \\ 0 & 0 & Q_{d3}(t_k) \end{bmatrix} \quad (7)$$

where Q_{di} corresponds to the system noise covariance matrix associated with robot i.

The **covariance propagation** equation for the centralized system is:

$$P(t_{k+1}^-) = \Phi(t_{k+1}, t_k)P(t_k^+)\Phi^T(t_{k+1}, t_k) \ + \ Q_d(t_k) \quad (8)$$

where initially (before any update) we have:

$$P(t_k^+) = \begin{bmatrix} P_{11}(t_k^+) & 0 & 0 \\ 0 & P_{22}(t_k^+) & 0 \\ 0 & 0 & P_{33}(t_k^+) \end{bmatrix} \quad (9)$$

Initially, each robot may know only its own position in global coordinates and the uncertainty related to it. Since there is no *a priori* shared knowledge among the robots, the covariance matrix for the centralized system is diagonal and each of the diagonal elements is the covariance for the state of each of the participating robots. While no update has occurred, i.e. no relative position information has been exchanged, Eq. (8) describes the propagation for the centralized system position uncertainty and by substituting from Eq.s (6), (7), and (9) we have:

$$P(t_{k+1}^-) = \begin{bmatrix} \Phi_1 P_{11}\Phi_1^T + Q_{d1} & 0 & 0 \\ 0 & \Phi_2 P_{22}\Phi_2^T + Q_{d2} & 0 \\ 0 & 0 & \Phi_3 P_{33}\Phi_3^T + Q_{d3} \end{bmatrix}$$

It is obvious from the previous equation that the propagated covariance of the centralized system is also a diagonal matrix as was the initial covariance matrix (Eq. (9)). Therefore the state covariance propagation can easily be decentralized and distributed between the individual robots. Each robot can propagate its own part of the centralized system covariance matrix. This is the corresponding diagonal matrix element P_{ii} that describes the uncertainty associated with the position of robot i:

$$P_{ii}(t_{k+1}^-) = \Phi_i(t_{k+1}, t_k)P_{ii}(t_k^+)\Phi_i^T(t_{k+1}, t_k) + Q_{di}(t_k) \quad (10)$$

for $i = 1..3$. While no relative position information is exchanged between any of the robots in the group (i.e. no relative pose update occurs) the covariance will remain the same until the next propagation step:

$$P(t_{k+1}^+) = P(t_{k+1}^-) \quad (11)$$

Applying Eq.s (8) and (11) repetitively to propagate to the next time steps, the covariance matrix for the centralized system will continue to be diagonal and its computation will be distributed among the 3 robots. All the quantities in Eq.s (1) and (10) are local to each robot i and thus the centralized system state and covariance propagation is fully distributed with all the necessary computations being local.

B. Update

The first time two of the robots meet they measure their relative pose and exchange information regarding their propagated estimates (of state and covariance) in order to update them. For example, when robot 1 detects robot 2, it uses its exteroceptive sensing to measure the relative position and orientation of robot 2 with respect to the frame of reference $\{1\}$ attached to robot 1:

$$\mathbf{z}_{12}(t_{k+1}) \ = \ \begin{bmatrix} {}^1x_2(t_{k+1}) \\ {}^1y_2(t_{k+1}) \\ {}^1\phi_2(t_{k+1}) \end{bmatrix} + \vec{\mathbf{n}}_{12}(t_{k+1})$$

$$= \begin{bmatrix} C^T(\phi_1)\left(\begin{bmatrix} x_2 \\ y_2 \end{bmatrix} - \begin{bmatrix} x_1 \\ y_1 \end{bmatrix}\right) \\ \phi_2 - \phi_1 \end{bmatrix} + \vec{\mathbf{n}}_{12}$$

This measurement will be used to update the pose estimate for the centralized system and the covariance of this estimate as described in the next sections.[2] By linearizing the previous equation, the measurement error can be approximated by the following equation:

$$\widetilde{\mathbf{z}}_{12}(t_{k+1}) = H_{12}(t_{k+1})\widetilde{\mathbf{x}}(t_{k+1}) + \vec{\mathbf{n}}_{12}(t_{k+1}) \qquad (12)$$

where

$$H_{12} = \Gamma_1^T \begin{bmatrix} -\widetilde{H}_{12} & I_{3\times 3} & 0_{3\times 3} \end{bmatrix} \qquad (13)$$

with

$$\widetilde{H}_{12} = \begin{bmatrix} I_{2\times 2} & J(\hat{p}_2 - \hat{p}_1) \\ \vec{0}^T_{2\times 1} & 1 \end{bmatrix} \quad \Gamma_1 = \begin{bmatrix} C(\hat{\phi}_1) & \vec{0}_{2\times 1} \\ \vec{0}^T_{2\times 1} & 1 \end{bmatrix}$$

$$C(\hat{\phi}_1) = \begin{bmatrix} \cos\hat{\phi}_1 & -\sin\hat{\phi}_1 \\ \sin\hat{\phi}_1 & \cos\hat{\phi}_1 \end{bmatrix} \qquad J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$$

$$\hat{p}_i = \left[\hat{x}_i(t^-_{k+1}) \ \hat{y}_i(t^-_{k+1})\right]^T \quad \hat{\phi}_i = \hat{\phi}_i(t^-_{k+1}) \ \ , i = 1, 2$$

I is the identity matrix and $\vec{\mathbf{n}}_{12}(t_{k+1})$ is the noise associated with the relative pose measurement $\mathbf{z}_{12}(t_{k+1})$ between robots 1 and 2. This measurement noise is assumed to be zero-mean white Gaussian with covariance:

$$R_{12}(t_{k+1}) = E\{\vec{\mathbf{n}}_{12}(t_{k+1})\vec{\mathbf{n}}^{\mathbf{T}}_{12}(t_{k+1})\} \qquad (14)$$

IV. Introduction of Cross-correlation Terms

At this point we show how the cross-correlation terms are introduced within the centralized position estimator and how their calculation can be distributed.

By substituting from Eq.s (9) and (13) the **covariance of the residual** is calculated as follows:

$$\begin{aligned} S_{12}(t_{k+1}) &= H_{12}(t_{k+1})P(t^-_{k+1})H^T_{12}(t_{k+1}) + R_{12}(t_{k+1}) \\ &= \Gamma_1^T \widetilde{S}_{12}\Gamma_1 \qquad (15) \end{aligned}$$

with

$$\begin{aligned} \widetilde{S}_{12} &= \widetilde{H}_{12}P_{11}(t^-_{k+1})\widetilde{H}^T_{12} + P_{22}(t^-_{k+1}) + \widetilde{R}_{12} \\ \widetilde{R}_{12} &= \Gamma_1 R_{12}(t_{k+1})\Gamma_1^T \end{aligned}$$

The residual covariance matrix has dimensions $N \times N$, the same as if we were updating the pose estimate of 1 only robot instead of $M = 3$. (In the latter case the dimension of matrix S would be $(N \times M) \times (N \times M)$). As it will be obvious from Eq.s (16) and (18), this results in a *significant* *reduction* of the computations required for both calculating the Kalman filter gain and updating the covariance of the pose estimate.

[2]One of the reasons for initially formulating the group localization problem in a centralized way is that when a robot in the group senses another robot, a *relative* position and orientation measurement is recorded. There is no function that relates the states of only one of the two robots with the relative pose measurement. An observation model for this type of measurement is a function of the position and orientation states of both robots. Motivated by this remark we formulate the group localization problem as a centralized one and then describe how it can be distributed among the robots in the group.

The **Kalman gain** $K(t_{k+1})$ for this update is given by:

$$\begin{aligned} K(t_{k+1}) &= P(t^-_{k+1})H^T_{12}(t_{k+1})S^{-1}_{12}(t_{k+1}) \qquad (16) \\ &= \begin{bmatrix} -P_{11}\widetilde{H}^T_{12}\widetilde{S}^{-1}_{12}\Gamma_1 \\ P_{22}\widetilde{S}^{-1}_{12}\Gamma_1 \\ 0 \end{bmatrix} = \begin{bmatrix} K_1(t_{k+1}) \\ K_2(t_{k+1}) \\ K_3(t_{k+1}) \end{bmatrix} \end{aligned}$$

which suggests that the larger the uncertainty (covariance $P_{ii}(t^-_{k+1})$) of the pose estimate of robot i, $i = 1..2$, the larger the correction coefficient (matrix element $K_i(t_{k+1})$ of the Kalman gain matrix).

The residual of the relative pose measurement is:

$$\mathbf{r}_{12}(t_{k+1}) = \mathbf{z}_{12}(t_{k+1}) - \widehat{\mathbf{z}}_{12}(t_{k+1})$$

where

$$\widehat{\mathbf{z}}_{12}(t_{k+1}) = \Gamma_1^T(\hat{\phi}_1)\begin{bmatrix} -I & I & 0 \end{bmatrix}\widehat{\mathbf{x}}(t^-_{k+1})$$

is the estimated measurement, i.e. the difference between the estimates of the pose $\widehat{\vec{\mathbf{x}}}_1$ of robot 1, and the pose $\widehat{\vec{\mathbf{x}}}_2$ of robot 2 expressed in the coordinate frame $\{1\}$ attached to robot 1.

The **pose estimate update** $\widehat{\mathbf{x}}(t^+_{k+1})$ for the centralized system is given by:

$$\widehat{\mathbf{x}}(t^+_{k+1}) = \widehat{\mathbf{x}}(t^-_{k+1}) + K(t_{k+1})\mathbf{r}_{12}(t_{k+1}) \qquad (17)$$

or

$$\begin{bmatrix} \widehat{\vec{\mathbf{x}}}^+_1 \\ \widehat{\vec{\mathbf{x}}}^+_2 \\ \widehat{\vec{\mathbf{x}}}^+_3 \end{bmatrix} = \begin{bmatrix} \left(I - P_{11}\widetilde{H}^T_1\widetilde{S}^{-1}_{12}\right)\widehat{\vec{\mathbf{x}}}^-_1 + P_{11}\widetilde{H}^T_1\widetilde{S}^{-1}_{12}\left(\widehat{\vec{\mathbf{x}}}^-_2 - \Gamma_1\mathbf{z}_{12}\right) \\ \left(I - P_{22}\widetilde{S}^{-1}_{12}\right)\widehat{\vec{\mathbf{x}}}^-_2 + P_{22}\widetilde{S}^{-1}_{12}\left(\widehat{\vec{\mathbf{x}}}^-_1 + \Gamma_1\mathbf{z}_{12}\right) \\ \widehat{\vec{\mathbf{x}}}^-_3 \end{bmatrix}$$

Finally, the **covariance update** $P(t^+_{k+1})$ for the centralized system is calculated as:

$$P(t^+_{k+1}) = P(t^-_{k+1}) - P(t^-_{k+1})H^T_{12}(t_{k+1})S^{-1}_{12}(t_{k+1})H_{12}(t_{k+1})P(t^-_{k+1}) \qquad (18)$$

Substituting for $P(t^-_{k+1})$ from Eq. (9) and for $H_{12}(t_{k+1})$ from Eq. (13) we have:

$$P^+ = \begin{bmatrix} P^-_{11} - P^-_{11}\widetilde{H}^T_{12}\widetilde{S}^{-1}_{12}\widetilde{H}_{12}P^-_{11} & P^-_{11}\widetilde{H}^T_{12}\widetilde{S}^{-1}_{12}P^-_{22} & 0 \\ P^-_{22}\widetilde{S}^{-1}_{12}\widetilde{H}_{12}P^-_{11} & P^-_{22} - P^-_{22}\widetilde{S}^{-1}_{12}P^-_{22} & 0 \\ 0 & 0 & P^-_{33} \end{bmatrix} \qquad (19)$$

By inspection of Eq. (19) we derive the following conclusions:

1. The covariances of robots 1 and 2 are the only ones that change:

$$\begin{aligned} P_{11}(t^+_{k+1}) &= P_{11}(t^-_{k+1}) - P_{11}(t^-_{k+1})\widetilde{H}^T_{12}\widetilde{S}^{-1}_{12}\widetilde{H}_{12}P_{11}(t^-_{k+1}) \\ P_{22}(t^+_{k+1}) &= P_{22}(t^-_{k+1}) - P_{22}(t^-_{k+1})\widetilde{S}^{-1}_{12}P_{22}(t^-_{k+1}) \end{aligned}$$

Notice that matrix $S_{12}(t_{k+1})$, as calculated in Eq. (15), depends only on quantities local to robots 1 and 2. Thus this update can be performed locally at robots 1 and 2 which actively participate in the update. It is not necessary for the rest of the team to know about this update. Therefore the communication is limited to robots 1 and 2 only. They must exchange their state estimates $\widehat{\vec{\mathbf{x}}}_1(t_{k+1}^-)$ and $\widehat{\vec{\mathbf{x}}}_2(t_{k+1}^-)$ and their corresponding covariances $P_{11}(t_{k+1}^-)$ and $P_{22}(t_{k+1}^-)$ in order for them to calculate the covariance of the residual, $S_{12}(t_{k+1})$, required for the update.

2. The covariances of the rest of the robots (in this case robot 3) remain the same:

$$P_{33}(t_{k+1}^+) = P_{33}(t_{k+1}^-) \tag{20}$$

Thus, no computations are required for robot 3 and no information from the exchange among robots 1 and 2 needs to be communicated to any of the rest of the group.

3. Cross-correlation (information coupling) terms appear, changing the form of the overall (centralized system) covariance matrix. The introduced submatrices are:

$$\begin{aligned} P_{12}(t_{k+1}^+) &= P_{11}(t_{k+1}^-)\widetilde{H}_{12}^T(t_{k+1})\widetilde{S}_{12}^{-1}(t_{k+1})P_{22}(t_{k+1}^-) \\ &= P_{21}^T(t_{k+1}^+) \end{aligned} \tag{21}$$

These new elements represent the *shared knowledge* in the robot group and need to be included in the calculations during the next propagation and update steps.

V. Collective Localization After the First Update

In this section we present the propagation and update cycles of the Kalman filter estimator for the centralized system after the first update. Since cross-correlation elements have been introduced in the covariance matrix of the state estimate, this matrix is now written as:

$$P(t_k^+) = \begin{bmatrix} P_{11}(t_k^+) & P_{12}(t_k^+) & P_{13}(t_k^+) \\ P_{21}(t_k^+) & P_{22}(t_k^+) & P_{23}(t_k^+) \\ P_{31}(t_k^+) & P_{32}(t_k^+) & P_{33}(t_k^+) \end{bmatrix} \tag{22}$$

A. Propagation

Each robot i continues to move independently of the others and its motion is described again by Eq.s (1). Since the measured quantities are local to robot i, the state propagation equations can be distributed among the robots. The same is not true for the covariance of the state estimates. In the previous section, we derived the equations for the propagation of the initial, fully decoupled system. Here we will examine how Eq. (8) is modified in order to include the cross-correlation terms introduced after a few updates of the system. The **error-state propagation** is provided by Eq (5):

$$\widetilde{\mathbf{x}}(t_{k+1}) = \Phi(t_{k+1}, t_k)\widetilde{\mathbf{x}}(t_k) + G(t_k)\mathbf{w}(t_k) \tag{23}$$

Substituting from Eq. (22) the **covariance propagation** is given by:

$$P(t_{k+1}^-) = \Phi(t_{k+1}, t_k)P(t_k^+)\Phi^T(t_{k+1}, t_k) + Q_d(t_k) \tag{24}$$

$$= \begin{bmatrix} \Phi_1 P_{11}^+ \Phi_1^T + Q_{d1} & \Phi_1 P_{12}^+ \Phi_2^T & \Phi_1 P_{13}^+ \Phi_3^T \\ \Phi_2 P_{21}^+ \Phi_1^T & \Phi_2 P_{22}^+ \Phi_2^T + Q_{d2} & \Phi_2 P_{23}^+ \Phi_3^T \\ \Phi_3 P_{31}^+ \Phi_1^T & \Phi_3 P_{32}^+ \Phi_2^T & \Phi_3 P_{33}^+ \Phi_3^T + Q_{d3} \end{bmatrix}$$

Eq. (24) is repeated at every propagation step and it can be distributed among the robots.

(a) The calculation of each of the propagated diagonal submatrix elements of the overall covariance matrix requires processing of odometric measurements only from the corresponding robot:

$$P_{ii}(t_{k+1}^-) = \Phi_i(t_{k+1}, t_k)P_{ii}(t_k^+)\Phi_i^T(t_{k+1}, t_k) + Q_{di}(t_k) \tag{25}$$

for $i = 1..3$.

(b) The computation of the propagated cross-correlation submatrices can be distributed between the robots after appropriate factorization of these terms. For example, by employing singular value decomposition the cross-correlation matrix $P_{ij}(t_k^+)$ can be written as:

$$P_{ij}(t_k^+) = U_{ij}W_{ij}V_{ij}^T = (U_{ij}W_i)(V_{ij}W_j)^T = \mathcal{P}_i(t_k^+)\mathcal{P}_j^T(t_k^+)$$

where W_{ij} is the diagonal matrix of the eigenvalues of $P_{ij}(t_k^+)$ and W_i, W_j are diagonal matrices such that:

$$W_i W_j^T = W_i W_j = W_{ij} \tag{26}$$

From Eq. (24), after m steps of propagation without any intermediate update, the cross-correlation term $P_{ij}(t_{k+m}^-)$ is given by:

$$\begin{aligned} P_{ij}(t_{k+m}^-) &= \Phi_i(t_{k+m}, t_{k+m-1})..\Phi_i(t_{k+1}, t_k)P_{ij}(t_k^+) \times \\ &\quad \Phi_j^T(t_{k+1}, t_k)..\Phi_j^T(t_{k+m}, t_{k+m-1}) \\ &= \Phi_i(t_{k+m}, t_{k+m-1})..\Phi_i(t_{k+1}, t_k)\mathcal{P}_i(t_k^+) \times \\ &\quad \left(\Phi_j(t_{k+m}, t_{k+m-1})..\Phi_j(t_{k+1}, t_k)\mathcal{P}_j(t_k^+)\right)^T \\ &= \mathcal{P}_i(t_{k+m}^-)\mathcal{P}_j^T(t_{k+m}^-) \end{aligned} \tag{27}$$

Each of the matrices $\mathcal{P}_i(t_{k+m}^-)$ and $\mathcal{P}_j(t_{k+m}^-)$ can be computed by processing odometric measurements local to robot i and robot j correspondingly during the propagation phase.

This process repeats until the next update. Note that all the matrices in the previous equations have dimensions $N \times N$ compared to the centralized estimator the covariance of which has dimensions $(N \times M) \times (N \times M)$, where N is the number of estimated states for each robot and M is the number of robots in the group.

This result is very important since the state and covariance propagations described by Eq.s (1), (25) and (27) allow for a **fully distributed estimation algorithm during the propagation cycle**. The computational gain is particularly significant if we consider that most of the time the robots propagate their pose and covariance estimates based on their own perception while updates are usually rare and they take place only when two robots meet.

B. Update

If now, for example, robots 2 and 3 meet and measure their relative position and orientation, then the **residual covariance matrix update** will be:

$$\begin{aligned} S_{23}(t_{k+1}) &= H_{23}(t_{k+1})P(t_{k+1}^-)H_{23}^T(t_{k+1}) + R_{23}(t_{k+1}) \\ &= \Gamma_2^T(\hat{\phi}_2)\widetilde{S}_{23}\Gamma_2(\hat{\phi}_2) \end{aligned}$$

with

$$\widetilde{S}_{23} = \widetilde{H}_{23}P_{22}\widetilde{H}_{23}^T - P_{32}\widetilde{H}_{23}^T - \widetilde{H}_{23}P_{23} + P_{33} + \widetilde{R}_{23}$$
$$\widetilde{R}_{23} = \Gamma_2(\hat{\phi}_2)R_{23}(t_{k+1})\Gamma_2^T(\hat{\phi}_2)$$

$R_{23}(t_{k+1})$ is the measurement noise covariance matrix associated with the relative pose measurement between robots 2 and 3; it is defined similarly to Eq. (14). As during the first update, here also the dimension of the residual covariance matrix is only $N \times N$ (as if we were updating the pose estimate of one only robot) and not $(N \times M) \times (N \times M)$ where $M = 3$ is the number of robots in the group. For large teams of robots a smaller dimension of S significantly reduces the processing required for computing its inverse. In order to calculate matrix $S_{23}(t_{k+1})$, only the covariances of the two meeting robots are needed along with their cross-correlation terms. These terms are exchanged when the two robots detect each other.

The **Kalman gain** $K(t_{k+1})$ for this update is given by:

$$K(t_{k+1}) = P(t_{k+1}^-)H_{23}^T(t_{k+1})S_{23}^{-1}(t_{k+1}) \quad (28)$$
$$= \begin{bmatrix} \left(P_{13} - P_{12}\widetilde{H}_{23}^T\right)\widetilde{S}_{23}^{-1}\Gamma_2 \\ -\left(P_{22}\widetilde{H}_{23}^T - P_{23}\right)\widetilde{S}_{23}^{-1}\Gamma_2 \\ \left(P_{33} - P_{32}\widetilde{H}_{23}^T\right)\widetilde{S}_{23}^{-1}\Gamma_2 \end{bmatrix}$$
$$= \begin{bmatrix} K_1(t_{k+1}) \\ K_2(t_{k+1}) \\ K_3(t_{k+1}) \end{bmatrix} = \widetilde{K}(t_{k+1})\Gamma_2$$

The correction coefficients (matrix elements $K_i(t_{k+1})$, $i = 2, 3$, of the Kalman gain matrix) in the previous equation are smaller compared to the corresponding correction coefficients of Eq. (16) calculated during the first update. Here the correction coefficients are reduced by the cross-correlation terms $P_{23}(t_{k+1}^-)$ and $P_{32}(t_{k+1}^-)$ respectively. This can be explained by examining what the information contained in these cross-correlation matrices is. As previously described, the cross-correlation terms represent the information common to the two meeting robots acquired during a previous direct (e.g. robot 2 met robot 3) or indirect (e.g. robot 1 met robot 2 and then robot 2 met robot 3) exchange of information. The more knowledge these two robots (here 2 and 3) already share, the less gain can be obtained from this update session. In addition, by observing that $K_1(t_{k+1}) = (P_{13}(t_{k+1}^-) - P_{12}(t_{k+1}^-)\widetilde{H}_{23}^T(t_{k+1}))\widetilde{S}_{23}^{-1}(t_{k+1})\Gamma_2$ we should infer that robot 1 is affected by this update to the extent that the information shared between robots 1 and 2 differs from the information shared between robots 1 and 3.

In the extreme case that the two robots share the same knowledge, the corresponding submatrices $P_{ij}(t_{k+1}^-)$, whose differences appear in Eq. (28) when calculating the Kalman gains, will be equal (proofs for these are given in [25]) and the Kalman gains for all 3 robots will become zero.[3] This is intuitively correct: if every robot shares the same positioning information, any exchange of information will have no effect on the current position estimates.

The **pose estimate update** $\widehat{\mathbf{x}}(t_{k+1}^+)$ is computed as follows:

$$\widehat{\mathbf{x}}(t_{k+1}^+) = \widehat{\mathbf{x}}(t_{k+1}^-) + K(t_{k+1})\mathbf{r}_{k+1} \quad (29)$$
$$= \widehat{\mathbf{x}}(t_{k+1}^-) + \widetilde{K}(t_{k+1}) \times \left[\Gamma_2\mathbf{z}(t_{k+1}) - \left(\widehat{\mathbf{x}}_3(t_{k+1}^-) - \widehat{\mathbf{x}}_2(t_{k+1}^-)\right)\right]$$

The **covariance update** is calculated as before by applying Eq. (18). The centralized system covariance matrix calculation can be divided into $3(3+1)/2 = 6$, $N \times N$ matrix calculations and distributed among the robots of the group.[4]

In conclusion, during any consecutive update, after the first one, the computations that take place are:

- Robot 2 or 3:

$$\widetilde{S}_{23}(t_{k+1}) = \widetilde{H}_{23}P_{22}(t_{k+1}^-)\widetilde{H}_{23}^T - \widetilde{H}_{23}P_{23}(t_{k+1}^-) - P_{32}(t_{k+1}^-)\widetilde{H}_{23}^T + P_{33}(t_{k+1}^-) + \widetilde{R}_{23}$$

$$P_{23}(t_{k+1}^+) = P_{23}(t_{k+1}^-) - \left[P_{23}(t_{k+1}^-) - P_{22}(t_{k+1}^-)\widetilde{H}_{23}^T\right] \times\widetilde{S}_{23}^{-1}(t_{k+1}^-)\left[P_{33}(t_{k+1}^-) - \widetilde{H}_{23}P_{23}(t_{k+1}^-)\right]$$

- Robot 1 or 2:

$$P_{12}(t_{k+1}^+) = P_{12}(t_{k+1}^-) - \left[P_{13}(t_{k+1}^-) - P_{12}(t_{k+1}^-)\widetilde{H}_{23}^T\right] \times\widetilde{S}_{23}^{-1}(t_{k+1})\left[P_{32}(t_{k+1}^-) - \widetilde{H}_{23}P_{22}(t_{k+1}^-)\right]$$

- Robot 1 or 3:

$$P_{13}(t_{k+1}^+) = P_{13}(t_{k+1}^-) - \left[P_{13}(t_{k+1}^-) - P_{12}(t_{k+1}^-)\widetilde{H}_{23}^T\right] \times\widetilde{S}_{23}^{-1}(t_{k+1})\left[P_{33}(t_{k+1}^-) - \widetilde{H}_{23}P_{23}(t_{k+1}^-)\right]$$

- Robot 1:

$$\widetilde{K}_1(t_{k+1}) = \left[P_{13}(t_{k+1}^-) - P_{12}(t_{k+1}^-)\widetilde{H}_{23}^T\right]\widetilde{S}_{23}^{-1}(t_{k+1})$$

$$\widehat{\mathbf{x}}_1(t_{k+1}^+) = \widehat{\mathbf{x}}_1(t_{k+1}^-) + \widetilde{K}_1\left[\Gamma_2\mathbf{z}(t_{k+1}) - \left(\widehat{\mathbf{x}}_3(t_{k+1}^-) - \widehat{\mathbf{x}}_2(t_{k+1}^-)\right)\right]$$

$$P_{11}(t_{k+1}^+) = P_{11}(t_{k+1}^-) - \left[P_{13}(t_{k+1}^-) - P_{12}(t_{k+1}^-)\widetilde{H}_{23}^T\right] \times\widetilde{S}_{23}^{-1}(t_{k+1})\left[P_{31}(t_{k+1}^-) - \widetilde{H}_{23}P_{21}(t_{k+1}^-)\right]$$

[3] A realistic scenario that approximates this is the case where the two robots are moving slowly or standing still while continuously measuring their relative position and orientation. Since there is almost no new "influx" of pose uncertainty, the repeated measurements trigger updates that only suppress the uncertainty associated with the relative pose measurement itself and not with the robots' pose estimates.

[4] In general $M(M+1)/2$ matrix equations distributed among M robots, thus $(M+1)/2$ matrix calculations per robot.

- Robot 2:

$$\widetilde{K}_2(t_{k+1}) = \left[P_{23}(t^-_{k+1}) - P_{22}(t^-_{k+1})\widetilde{H}^T_{23}\right] \widetilde{S}^{-1}_{23}(t_{k+1})$$

$$\widehat{\vec{\mathbf{x}}}_2(t^+_{k+1}) = \widehat{\vec{\mathbf{x}}}_2(t^-_{k+1}) + \widetilde{K}_2 \left[\Gamma_2 \mathbf{z}(t_{k+1}) - \left(\widehat{\vec{\mathbf{x}}}_3(t^-_{k+1}) - \widehat{\vec{\mathbf{x}}}_2(t^-_{k+1})\right)\right]$$

$$P_{22}(t^+_{k+1}) = P_{22}(t^-_{k+1}) - \left[P_{23}(t^-_{k+1}) - P_{22}(t^-_{k+1})\widetilde{H}^T_{23}\right]$$
$$\times \widetilde{S}^{-1}_{23}(t_{k+1}) \left[P_{32}(t^-_{k+1}) - \widetilde{H}_{23} P_{22}(t^-_{k+1})\right]$$

- Robot 3:

$$\widetilde{K}_3(t_{k+1}) = \left[P_{33}(t^-_{k+1}) - P_{32}(t^-_{k+1})\widetilde{H}^T_{23}\right] \widetilde{S}^{-1}_{23}(t_{k+1})$$

$$\widehat{\vec{\mathbf{x}}}_3(t^+_{k+1}) = \widehat{\vec{\mathbf{x}}}_3(t^-_{k+1}) + \widetilde{K}_3 \left[\Gamma_2 \mathbf{z}(t_{k+1}) - \left(\widehat{\vec{\mathbf{x}}}_3(t^-_{k+1}) - \widehat{\vec{\mathbf{x}}}_2(t^-_{k+1})\right)\right]$$

$$P_{33}(t^+_{k+1}) = P_{33}(t^-_{k+1}) - \left[P_{33}(t^-_{k+1}) - P_{32}(t^-_{k+1})\widetilde{H}^T_{23}\right]$$
$$\times \widetilde{S}^{-1}_{23}(t_{k+1}) \left[P_{33}(t^-_{k+1}) - \widetilde{H}_{23} P_{23}(t^-_{k+1})\right]$$

VI. Observability Study

Here we examine the centralized system observability matrix for a set of different cases of interest in order to evaluate the convergence properties of the corresponding Kalman filter estimators.

A. None of the robots has absolute positioning capabilities

The observability matrix of the system is:

$$M_{DTI} = \left[\begin{array}{ccc} H^T & \Phi^T H^T & (\Phi^T)^2 H^T \end{array}\right] \quad (30)$$

For small errors in orientation the system matrix for each robot can be approximated by $\Phi_i(t_{k+1}, t_k) \simeq I$, where I is the 3×3 identity matrix.[5] The centralized system propagation matrix is now written as:

$$\Phi = \left[\begin{array}{ccc} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{array}\right] \quad (31)$$

Similarly, for small errors in orientation the observation submatrix $\widetilde{H}_{ij}$ that corresponds to a relative pose measurement between robots i and j can be approximated by $\widetilde{H}_{ij}(t_{k+1}) \simeq I$.

If each of the three robots sees another robot at some time instant, the combined measurement matrix would be:

$$H = \left[\begin{array}{ccc} \Gamma_1^T & 0 & 0 \\ 0 & \Gamma_2^T & 0 \\ 0 & 0 & \Gamma_3^T \end{array}\right] \left[\begin{array}{ccc} -I & I & 0 \\ 0 & -I & I \\ I & 0 & -I \end{array}\right] \quad (32)$$

[5]When the orientation of the robot is accurately known (e.g. outdoors measured with a compass, or indoors computed using laser scan matching techniques [4]) the error-state propagation equation can be approximated by: $\widetilde{\vec{\mathbf{x}}}_i(t_{k+1}) = I\, \widetilde{\vec{\mathbf{x}}}_i(t_k) + G_i(t_k)\vec{\mathbf{w}}_i(t_k)$, $i = 1..3$.

In this case, the rank of matrix M_{DTI} is $6 < 9$ and thus the system is not observable. This is expected since none of the robots has access to absolute positioning information. By measuring their relative poses, the three robots improve their **position tracking** accuracy but they are not able to bound the overall uncertainty. The results reported in [31] are examples of this case.

B. At least one of the robots has absolute positioning capabilities

The main difference between this and the previous case is evident in the form of matrix H. If robot 1, for example, has absolute positioning capabilities then matrix H will be:

$$H = \left[\begin{array}{cccc} I & 0 & 0 & 0 \\ 0 & \Gamma_1^T & 0 & 0 \\ 0 & 0 & \Gamma_2^T & 0 \\ 0 & 0 & 0 & \Gamma_3^T \end{array}\right] \left[\begin{array}{ccc} I & 0 & 0 \\ -I & I & 0 \\ 0 & -I & I \\ I & 0 & -I \end{array}\right] \quad (33)$$

Here, the rank of M_{DTI} is 9 and thus the system is observable when at least one of the robots has access to absolute positioning information (e.g. by using GPS or a map of the environment). The map-based collaborative localization problem presented in [22], [23] is an example of such case.

C. At least one of the robots remains stationary

If at any time instant at least one of the robots in the group remains stationary and acts as a landmark for the rest of the team, the position estimate for this robot and the uncertainty regarding this estimate will be constant. This is equivalent to the idle robot directly measuring its own position and orientation. Therefore this case falls into the previous category and the system is considered locally observable. Examples of this scenario are the implementations found in [15] through [21].

VII. Experimental Results

The *collective localization* algorithm was implemented and tested for the case of three Pioneer II mobile robots shown in Fig. 2. The experiments were conducted in a lab environment with an overhead camera recording the ground truth (i.e. tracking the absolute pose) for each of the three robots. The accuracy of the tracking system (described in [33]) is $5cm$ for the position and $5°$ for the orientation. Every Pioneer II is equipped with wheel-encoders on the two front wheels that measure the translational and rotational velocity of the vehicle. Here, instead of implementing a system that directly measures the relative position and orientation between two robots such as the one presented in [34], we use the information from the tracking system to produce these measurements. Each relative pose measurement is the difference between the poses of two robots, expressed in local coordinates, with the addition of noise. This way we were able to select the accuracy level of the sensor data and control their frequency. For the experiments reported here the relative position and orientation measurements had accuracy of $300mm$ and $34°$ respectively

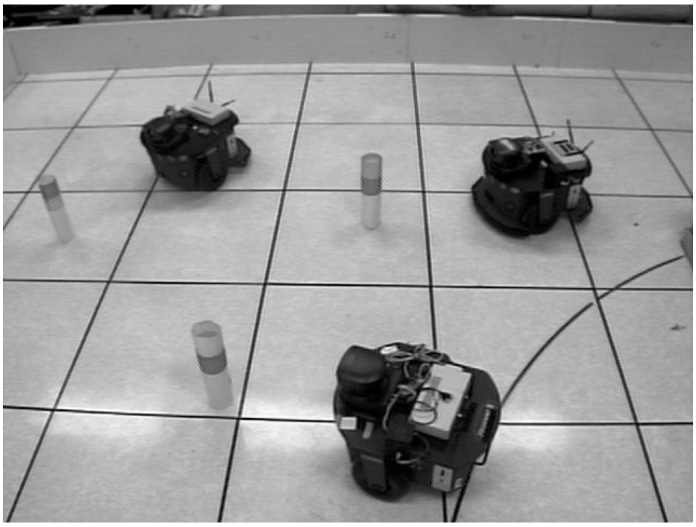

Fig. 2. The Pioneer II mobile robots used for testing the *collective localization* algorithm. Wheel-encoders are the only sensors measuring the self motion of each of the robots.

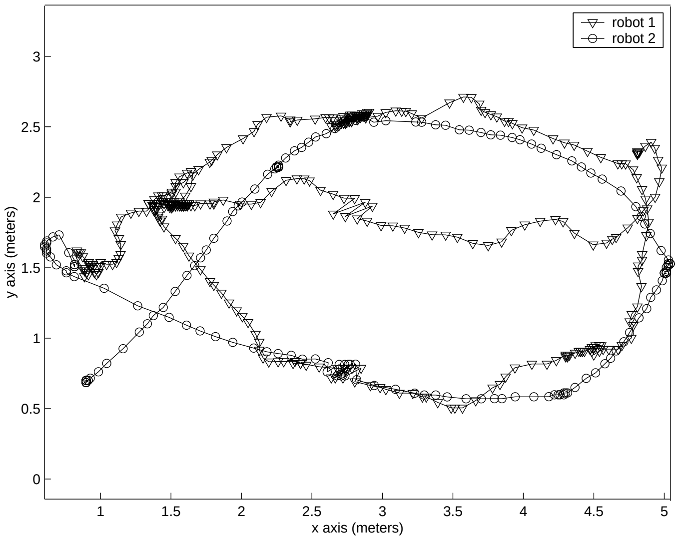

Fig. 3. The trajectories of two of the robots moving within the same area as recorded by the overhead camera tracking system.

(3σ values).

The three robots start from three different locations and move within the same area. An example of the trajectories followed by two of these robots is shown in Fig. 3. We examine the following three different scenarios.

A. None of the robots has absolute positioning information

In this case none of the robots has access to absolute positioning information and therefore the Kalman filter estimator will eventually diverge. The observability study reported in Section VI-A corresponds to this situation.

A.1 Dead-reckoning

First we examine the case where there is no exchange of information between the three robots. Each of them *dead-reckons* its pose estimates independently from the others by appropriately integrating its translational and rotational velocity measured by the wheel-encoders. Fig. 4 presents the recorded error along the x coordinate for robot 2 and the 3σ region of confidence for this error.[6] While in motion, the uncertainty for the position of this robot will grow continuously without bound. The recorded error at the end of this trial was $425mm$ ($\sim 3.4\%$) over traveled distance of approximately $12.5m$ while the maximum expected error[7] was almost $930mm$ ($\sim 7.4\%$) for the same distance. Fig. 7 shows the dead-reckoned (DR) trajectory of robot 2 for this trial.

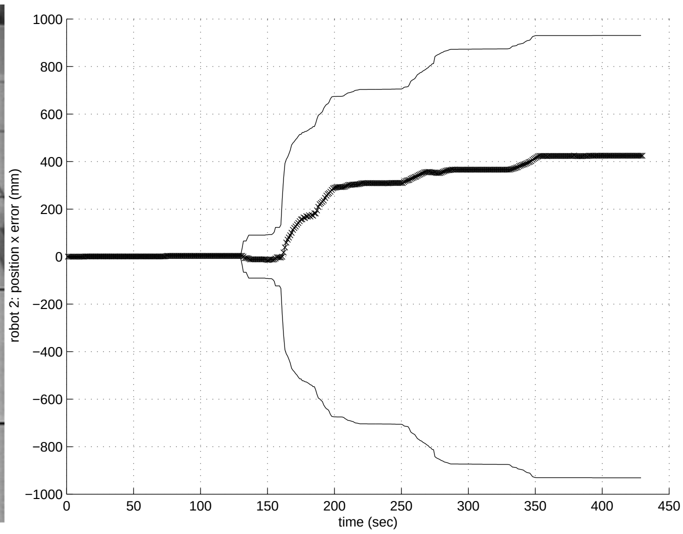

Fig. 4. Position x error for robot 2 when no relative pose measurements are available. The robot integrates the rotational and translational velocity measured by the encoders in order to estimate its position (dead-reckoning). The two bounding lines determine the 3σ region of confidence for the position x error and they are calculated based on the covariance of the position x estimate. The flat-line portions of the region of confidence (constant position uncertainty) correspond to time intervals that the robot was moving very slowly.

A.2 Continuous relative pose measurements

The alternative to the previous case is when the three robots, while moving, continuously exchange relative pose information and perform *collective localization.* As it is shown in Fig. 5, when each robot keeps track of the other members in the team, the rate that the position uncertainty grows is considerably lower compared to when there was no communication between them. The recorded position x error for robot 2 remains less than $39mm$ ($\sim 0.31\%$) over the same distance as before. The maximum expected error for this trial was $175mm$ ($\sim 1.4\%$). The significant reduction in localization error is due to the fact that three robots estimate their positions based not only on their own dead-reckoned estimates but also on the information collected

[6]The 3σ region of confidence for the pose error of robot i is calculated based on the values of the diagonal elements of corresponding covariance submatrix P_{ii}. For example, at time t_k the error along the x direction for robot 2 would be bounded by $+/-3\sqrt{P_{22\ (1,1)}(t_k)}$.

[7]This is the maximum value that the expected error can take with probability 99.7%.

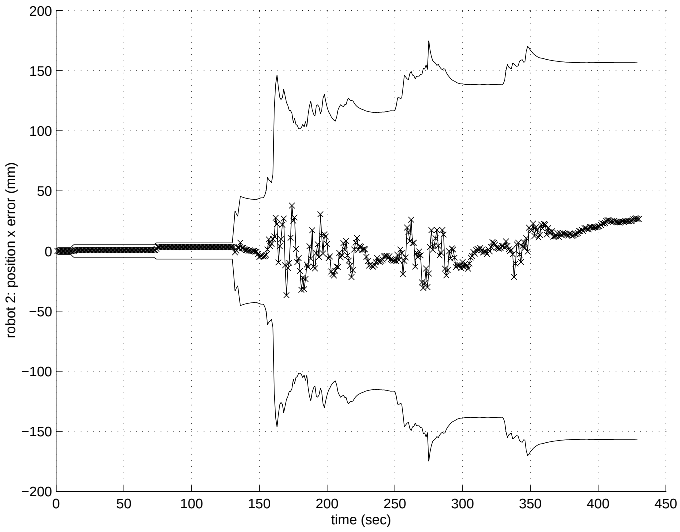

Fig. 5. Position x error for robot 2 when the robots continuously measure relative position and orientation. The position uncertainty is almost one order of magnitude smaller compared to the case of simple dead-reckoning. Note the difference in scale between this figure and Fig.s 4 and 6. The two bounding lines determine the 3σ region of confidence for the position x error and they are calculated based on the covariance of the position x estimate. The flat-line portions of the region of confidence (constant position uncertainty) correspond to time intervals that the robot was moving very slowly.

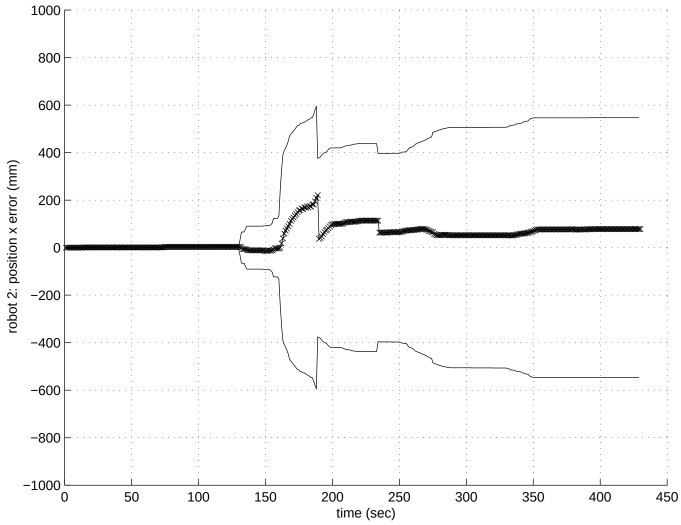

Fig. 6. Position x error for robot 2 when the group collects relative position and orientation measurements intermittently (during the time instants $t = 189sec$ and $t = 234sec$). The rest of the time, each of the robots dead-reckons its position. Note the sharp decrease in uncertainty for $t = 189sec$ and $t = 234sec$. The two bounding lines determine the 3σ region of confidence for the position x error and they are calculated based on the covariance of the position x estimate. The flat-line portions of the region of confidence (constant position uncertainty) correspond to time intervals that the robot was moving very slowly.

from the rest of the team regarding each robot's relative motion with respect to every other member in the group. The resulting trajectory as estimated by the collective localization (CL) algorithm is shown in Fig. 7.

A.3 Intermittent relative pose measurements

Positioning errors are reduced even when the three robots exchange positioning information only intermittently. That is when they detect each other only at certain time instants. For example, as it is shown in Fig. 6, the three robots measure their relative position and orientation only twice ($t = 189sec$ and $t = 234sec$). Still, this is enough to reduce the rate of increase in the positioning uncertainty. The corrections provided are proportional to the uncertainty accumulated in their previous estimates. This can be seen in Fig. 8 where the dead-reckoned (DR) and the estimated by collective localization (CL) trajectories of robot 2 are depicted along with the actual one recorded by the overhead tracking system. Initially, between the locations S and A, DR and CL produce the same position estimates. When at A ($t = 189sec$), robot 2 measures its relative pose with respect to the other two (also moving) robots and CL uses the extra information from robots 1 and 3 to produce better pose estimates for every robot in the group. The correction for the position estimate of robot 2 is shown in Fig. 8. A second correction is provided by CL at location B ($t = 234sec$). It is clear that CL outperforms DR throughout the portion of the trajectory between the first update at A and the end location E.[8] Nevertheless, in this case that relative measurements are available only intermittently, the trajectory estimated by CL is less accurate than the one computed during the previous experiment when relative poses were measured continuously (Fig. 7). During this last experiment (Fig. 6) the position error along the x direction for robot 2 was approximately $224mm$ ($\sim 1.8\%$). The maximum expected error was less than $600mm$ ($\sim 4.8\%$), significantly smaller compared to the corresponding error when the robots were localizing independently (Fig. 4).

In the previous experiments robot 2 covered less distance ($12.5m$) than either robot 1 ($21m$) or 3 ($26m$). Therefore it had the least uncertainty about its location compared to the other two when each of them had to localize independently. By direct comparison of the CL results shown in Fig.s 5 and 6 to the DR results in Fig. 4 it is evident that even the robot with the best localization results (here robot 2) experiences lower positioning uncertainty when it communicates with the rest of the team (CL) compared to when it only relies on its own sensors for localization (DR).

B. One of the robots has absolute positioning information

In this case, only robot 2 is provided with absolute measurements[9] of its position. In order to better demonstrate

[8] In [35] it was shown that smoothing of pose estimates right after updates can significantly improve the localization results. Here, if we post-process the CL pose estimates within a smoother right after the correction at A, this will decrease the errors for the portion of the trajectory between S and A which robot 2 had originally dead-reckoned. Finally, smoothing of the estimates between A and B will further refine this already corrected portion of the trajectory of robot 2.

[9] These are synthetic measurements. The pose of robot 2 is tracked by the overhead camera and the recorded consecutive pose measurements, corrupted with additive noise, are then provided to robot 2.

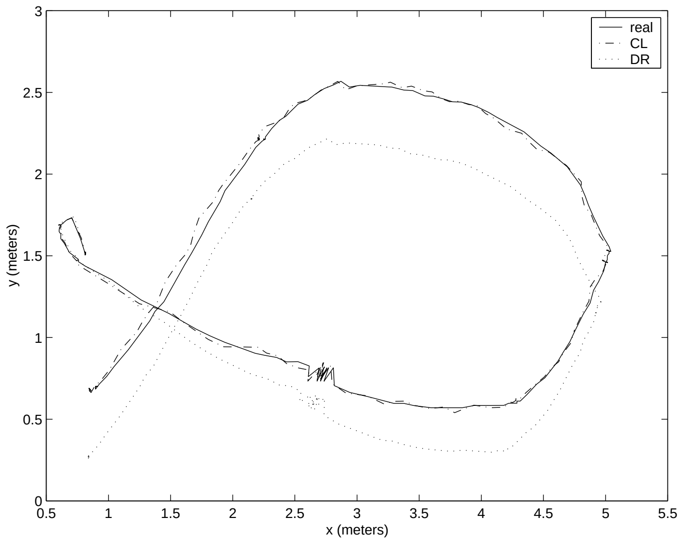

Fig. 7. The trajectory of robot 2 when the three robots measure their relative poses continuously. The solid line is the real trajectory as recorded by the overhead camera, the dashed-dotted line is the estimated by the collective localization (CL) algorithm, and the dotted line is the dead-reckoned (DR) one. Robot 2 zig-zags around $(x, y) = (2.7, 0.8)$ as it maneuvers to avoid collision with robot 1.

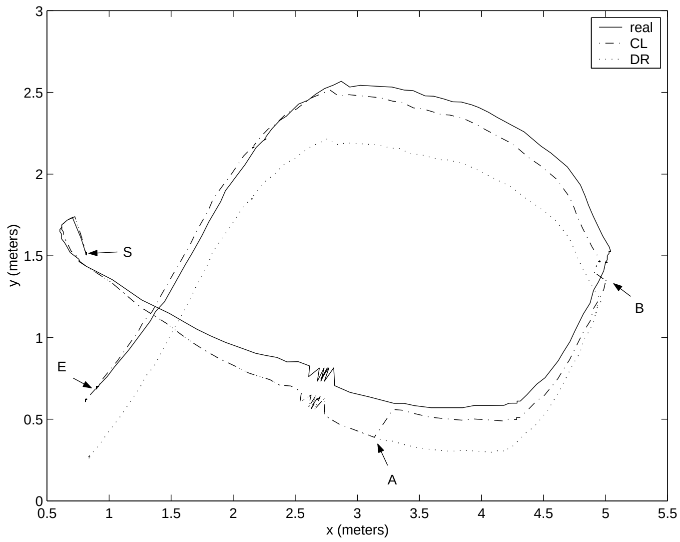

Fig. 8. The trajectory of robot 2 when the three robots measure their relative poses intermittently. The solid line is the real trajectory as recorded by the overhead camera, the dashed-dotted line is the estimated by the collective localization (CL) algorithm, and the dotted line is the dead-reckoned (DR) one. Relative pose measurements are available only twice at locations A and B.

the ability of the CL algorithm to process and distribute relative and absolute positioning information across this group of robots we tested the following two scenarios.

B.1 Continuous communication between the robots

As discussed in Section VI-B, the Kalman filter estimator for this case *converges* locally and therefore we expect that the positioning error will be bounded for *all* three robots. This is confirmed by the results depicted in Fig. 9. At the beginning of this experiment none of the robots knew its initial location. Only robot 2 receives absolute positioning information (from the overhead tracking system) and at the same time it measures its relative pose with respect to robots 1 and 3. This exchange of information among the three robots provides robots 1 and 3 with indirect access to the absolute localization data available only to robot 2. Since robot 2 knows its location in global coordinates so will robots 1 and 3 by measuring their relative displacement from robot 2. This constitutes a form of *sensor sharing.* CL allows all robots to have access to sensor data collected by any other robot in the group.

As shown in Fig. 9, both the recorded position error and the maximum expected error along the x coordinate is bounded and does not exceed $20mm$ ($\sim 0.16 - 0.08\%$) for all three robots. Fig. 10 depicts the CL estimated trajectory of robot 1 of this trial. Finally, we should mention that similar results would be acquired if instead of robot 2 directly measuring its absolute position, one or more robots in the group were performing map-based localization.[10]

The accuracy of these measurements was $30mm$ for the position and $17°$ for the orientation (3σ values). This is equivalent to having an indoor GPS available.

[10]In our previous work [6], [28] we have shown that a robot with unknown initial position is capable of globally localizing within a previously mapped area by maintaining a number of hypotheses about its probable location. A *single* Kalman filter is necessary to track and update these hypotheses. When sufficient information about the robot's surroundings is collected, these hypotheses are reduced to one: the hypothesis that corresponds to the actual position of the robot.

B.2 No communication between the robots

Here robot 2 continues to receive absolute positioning information but there is no communication among the members of the group. Since no relative pose measurements are available, robots 1 and 3 simply dead-reckon their pose estimates (the initial location for all three robots is known in this case). Fig. 11 presents the x position errors for all three robots. Even though the position errors for robot 2 are still bounded (as expected since it receives absolute positioning information), they reach higher values compared to the previous case (Fig. 9) when it was allowed to communicate with the rest of the team. The main degradation in performance comes for robots 1 and 3. Since they have no access to absolute positioning information the maximum expected errors for their pose estimates will continuously grow. The recorded errors along the x direction were $767mm$ ($\sim 3.7\%$) and $-891mm$ ($\sim 3.4\%$) for robots 1 and 3 respectively. Finally, the maximum expected errors were $1722mm$ ($\sim 8.2\%$) for robot 1 and $2053mm$ ($\sim 7.9\%$) for robot 3.

C. One of the robots remains stationary

In this case one of the robots (robot 2) remains stationary while the other two move in the same area and measure their relative positions and orientations with respect to the standing one. Even though no absolute positioning information is available to any of the robots in the group, the system is observable. This was explained in Section VI-

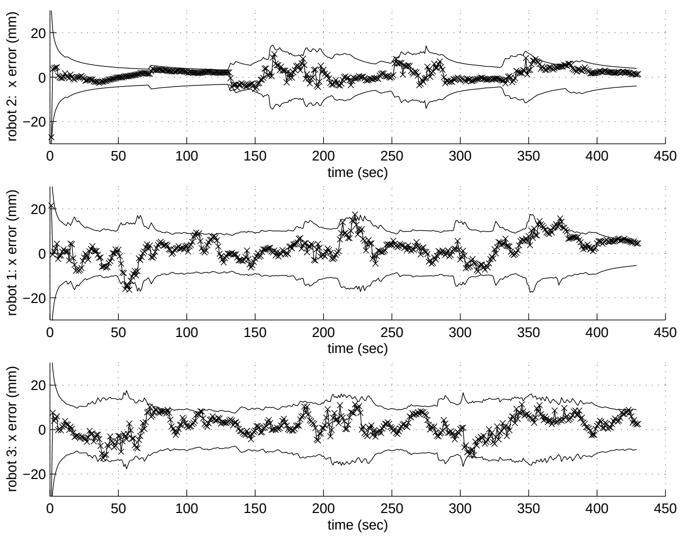

Fig. 9. Position x error for robots 2, 1, and 3 when robot 2 receives absolute positioning information and continuously measures relative position and orientation with respect to robots 1 and 3. The bounding lines around the positioning errors determine the 3σ regions of confidence for the position x error and they are calculated based on the covariance of the position x estimates. For all three robots the position error is bounded.

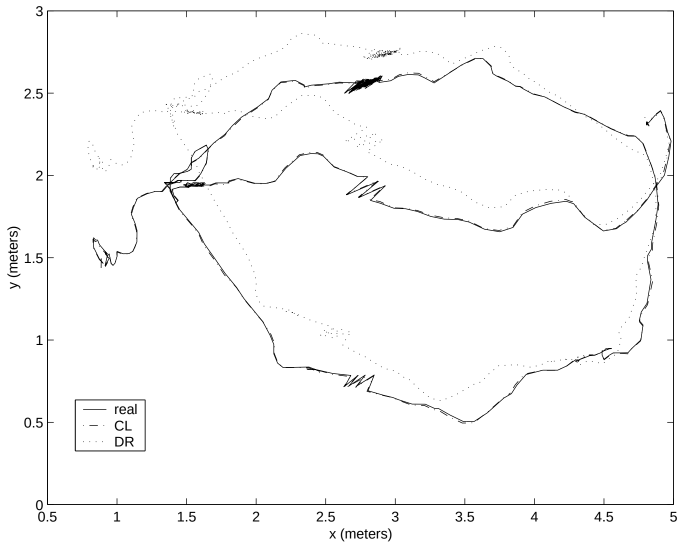

Fig. 10. The trajectory of robot 1 when the three robots measure their relative poses continuously and robot 2 has access to absolute positioning information. The solid line is the real trajectory as recorded by the overhead camera, the dashed-dotted line is the estimated by the collective localization (CL) algorithm, and the dotted line is the dead-reckoned (DR) one. Robot 1 zig-zags at three different locations as it maneuvers to avoid collision with the other two robots in the group.

C. Robot 2 acts as a landmark for robots 1 and 3 which perform CL and thus correct their dead-reckoned pose estimates by measuring their relative displacement with respect to robot 2. The position uncertainty for robot 2 is constant and therefore the position errors for both robots 1 and 3 will be bounded. This is verified by the recorded errors depicted in Fig. 12. Note also that since the location of robots 1 and 3 was initially unknown, the maximum expected error at the beginning was large but it decreases rapidly when the first relative pose measurements (with respect to robot 2) are recorded. Through the length of the experiment, the recorded errors and the maximum expected errors along the x direction for robots 1 and 3 were less than $141mm$ ($\sim 0.67 - 0.54\%$).

Finally, we should note that the results presented for this last case would be the same if a landmark was being detected instead of robot 2. A stationary feature that appears on a map of the environment would have the same effect on the positioning errors of a group of robots that measure their relative poses with respect to this landmark. It is our belief that an extension of the *collective localization* algorithm would be a suitable solution to the problem of simultaneous localization and mapping.

VIII. Discussion

At this point it is worth mentioning that a variant of the Kalman filter with decentralized *local loops*, one for each sensor, was presented by H. W. Sorenson in [36] and later revisited in its inverse, information filter, form in [37] for sequential processing of incoming sensor data for the case of a *single robot*. These modifications of the Kalman filter are particularly useful when dealing with asynchronous measurements originating from a variety of sensing modalities (an application of this can be found in [38]). Our focus here is distributed state estimation across a group of robots rather than sequential sensor processing on a single robot. Nevertheless, the latter can be easily incorporated in the presented distributed localization schema.

The information filter in particular has certain advantages compared to the Kalman filter for specific estimation applications [39] and since they are equivalent they can potentially be applied interchangeably. For the case of distributed multi-robot localization, the Kalman filter performs significantly better mainly due to the reduced number of computations required. The single matrix inversion necessary is that of the residual covariance matrix $S(t_{k+1})$ of dimensions $N \times N$ and this occurs only when a relative pose measurement is available (update cycle). The information filter requires large matrix inversions during every propagation step. More specifically the information matrix $\mathcal{I}(t_{k+1}^-) = P^{-1}(t_{k+1}^-)$ propagation equation is:

$$\mathcal{I}(t_{k+1}^-) = \Theta(t_{k+1}) - \Theta(t_{k+1})G_d(t_k)\Xi^{-1}(t_{k+1})G_d^T(t_k)\Theta(t_{k+1})$$

where

$$\Xi(t_{k+1}) = G_d^T(t_k)\Theta(t_{k+1})G_d(t_k) + Q_d^{-1}(t_k)$$

and

$$\Theta(t_{k+1}) = \Phi^T(t_k, t_{k+1})\mathcal{I}(t_k^+)\Phi(t_k, t_{k+1})$$

For a group of M robots, the matrix $\Xi(t_{k+1})$ of dimensions $(M \times L) \times (M \times L)$, where L is the dimension of matrix $Q_d(t_k)$, has to be inverted during each propagation step

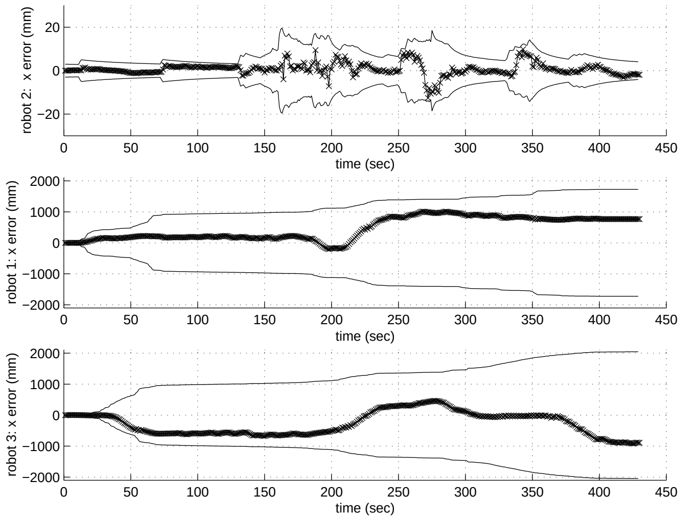

Fig. 11. Position x error for robots 2, 1, and 3 when robot 2 receives absolute positioning information and there is *no communication* with robots 1 or 3. The bounding lines around the positioning errors determine the 3σ regions of confidence for the position x error and they are calculated based on the covariance of the position x estimates. The position error is bounded only for robot 2 while for robots 1 and 3 it grows continuously while these robots move.

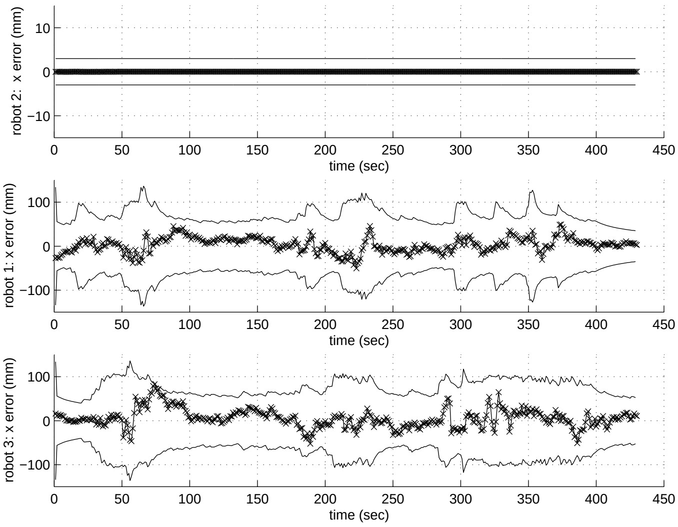

Fig. 12. Position x error for robots 2, 1, and 3 when robot 2 is standing still while robots 1 and 3 continuously measure their relative position and orientation with respect to robots 2. The bounding lines around the positioning errors determine the 3σ regions of confidence for the position x error and they are calculated based on the covariance of the position x estimates. Since robot 2 is not moving, its position uncertainty is constant. For the other two robots that use robot 2 as a landmark, the position error is bounded.

and for a large group of robots this becomes computationally inefficient.

In addition, the information filter produces estimates of the quantity:

$$\hat{y}(t_{k+1}^{+}) = \mathcal{I}(t_{k+1}^{+})\hat{x}(t_{k+1}^{+}) \qquad (34)$$

instead of $\hat{x}(t_{k+1}^{+})$ and therefore the estimated by the filter information matrix $\mathcal{I}(t_{k+1}^{+}) = P^{-1}(t_{k+1}^{+})$ (of dimensions $(M \times N) \times (M \times N)$) must also be inverted at every step in order to solve Eq. (34) for the estimates $\hat{x}(t_{k+1}^{+})$ of the poses of all the robots in the group. Finally, as it is described in [25], there are situations where the covariance matrix (and therefore the information matrix) may lose rank and become singular. In these cases the required matrix inversion for retrieving $\hat{x}(t_{k+1}^{+})$ is not be feasible.

IX. Conclusions & Extensions

In this paper a new approach to the multi-robot localization problem was presented. Initially, a centralized Kalman filter estimator was introduced that optimally combines information collected by a group of robots and produces positioning estimates of uniform accuracy for every member of the team. We showed that the equations of this single estimator can be distributed into M *reduced dimension* filters, one for each robot.

Every filter processes sensor measurements gathered by its host robot and exchanges information with the rest of the estimators *only* when relative pose measurements become available. Consequently, this distributed form of the filter has reduced computational and communication requirements (proportional to the number of relative pose measurements) and it can be realized for the case of large teams of robots. In addition, depending on the processing capabilities of the individual robots, the collective localization architecture allows for adaptive distribution of the required computations. For example, the reduced dimension filter of a robot with limited processing power can run on a second robot that receives the sensor data gathered by the first one. This feature of distributed estimation is particularly important for heterogeneous groups of robots that handle tasks with various demands on their resources.

The reduced order estimators process sensor data local to their host robot when propagating estimates of position and uncertainty. This brings flexibility to the localization process since the implementation of each of these decentralized filters can vary across different platforms. For example, the motion of different robots can be described by various models (planar or 3-dimensional) depending on their particular capabilities, mission, morphology of local area, etc. Similarly there is no uniformity constraint on the type of sensors employed. Sensing modalities of different precision and noise profile can be incorporated and therefore the collective localization approach is applicable to both homogeneous and heterogeneous teams of robots.

The exteroceptive sensors that allow for relative positioning introduce localization coupling within the group. Every time two robots meet and measure their relative poses, they exchange information not only for themselves but also for the rest of the team (cross-correlation terms). As shown in Section VII, robots that carry precise sensors and exhibit better localization results transfer positioning knowledge to poorly equipped robots. This constitutes a form of sensor sharing and can potentially reduce the cost of constructing teams of robots. In addition, it increases the robustness of the group towards sensor failures. Robots can rely on their teammates for positioning when they experience problems with their own sensors. The relative update sessions cause information diffusion across the group and in effect form a more homogeneous, in terms of localization accuracy, robotic colony. This is evident in the experimental results presented previously and it is also explained by the equality of the covariance matrices after a precise relative pose update [25].

To the best of our knowledge this is the first multi-robot localization approach to explicitly address the problem of sensor data interdependencies that appear when robots exchange information regarding their pose estimates. Different forms of cooperative localization investigated in the past were described as special cases of this unifying framework conditioned on availability of sensors and motion strategies employed. In Section VI we examined the convergence properties of the various realizations of this distributed estimator and we then validated these theoretical results through corresponding experiments employing a group of three mobile robots.

The current implementation of the distributed Kalman filter estimator relies on relative position and orientation measurements in order to propagate positioning information across the group. With proper modification of the relative observation model, the collective localization approach can exploit a variety of exteroceptive sensors and use measurements of: (i) relative position only (e.g. when relative orientation cannot be specified), (ii) relative distance [19] (e.g. when the robots are far apart from each other but they are able to measure distances between them based on time of flight, or interferometry on the exchanged signals), and (iii) relative bearing [40] (e.g. when they can only measure the direction in which a signal emitted by another robot is received).

A generalization of this method can be applied to the problem of multi-robot concurrent localization and mapping. When a group of robots is assigned the task of mapping an area of interest, collaboration among them is desired. This will save the group time and will increase the accuracy of the map while resolving potential inconsistencies. As aforementioned, instead of some of the robots standing still, features of the environment (landmarks) can be used to improve localization results. The poses of these landmarks can be estimated simultaneously with the poses of the robots in the team within the same distributed framework. When different robots detect each other, they will be able to combine information acquired during their individual mapping experiences. Some of this information will be collected from the same source (correlated) and some will be from different sources (independent) in the environment. It is our belief that the collective localization methodology can be extended to describe the fusion of information collected by the different members of the group in order to produce a combined, uniformly accurate map of the area.

Finally, as part of our future work, we plan to investigate different communication architectures, within large groups of robots, when broadcasting is not possible [41]. We also intend to address the problem of temporary communication loss, e.g. when one of the robots is beyond reach from the rest of the group for some time, and examine how a delayed information exchange with the other team members can be realized.

References

[1] I. J. Cox, "Blanche-an experiment in guidance and navigation of an autonomous robot vehicle," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 2, pp. 193–204, April 1991.

[2] H. R. Everett, *Sensors for Mobile Robots*, AK Peters, 1995.

[3] J. Borenstein and L. Feng, "Measurement and correction of systematic odometry errors in mobile robots," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 6, pp. 869–880, Dec. 1996.

[4] F. Lu and E. Milios, "Robot pose estimation in unknown environments by matching 2d range scans," *Journal of Intelligent and Robotic Systems: Theory and Applications*, vol. 18, no. 3, pp. 249–275, March 1997.

[5] C.F. Olson and L.H. Matthies, "Maximum likelihood rover localization by matching range maps.," in *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, Leuven, Belgium, 16-20 May 1998, pp. 272–277.

[6] S.I. Roumeliotis and G.A. Bekey, "Bayesian estimation and kalman filtering: A unified framework for mobile robot localization," in *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Fransisco, CA, April 24-28 2000, pp. 2985–2992.

[7] S. Thrun, "Bayesian landmark learning for mobile robot localization," *Machine Learning*, vol. 33, no. 1, pp. 41–76, Oct. 1998.

[8] L. E. Parker, "Current state of the art in distributed autonomous mobile robotics," in *Proceedings of the 5th International Symposium on Distributed Autonomous Robotic Systems (DARS 2000)*, Knoxville, TN, Oct. 4-6 2000, pp. 3–12.

[9] L.E. Parker, "Alliance: An architecture for fault tolerant multirobot cooperation," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 2, pp. 220–240, April 1998.

[10] C. Ferrari, E. Pagello, J. Ota, and T. Arai, "Multirobot motion coordination in space and time," *Robotics and Autonomous Systems*, vol. 25, no. 3-4, pp. 219–229, Nov. 1998.

[11] M.S. Fontan and M.J. Mataric, "Territorial multi-robot task division," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 5, pp. 815–822, Oct. 1998.

[12] P. Pirjanian and M. Mataric, "Multi-robot target acquisition using multiple objective behavior coordination," in *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Fransisco, CA, April 24-28 2000, pp. 2696–2702.

[13] Y. Arai, T. Fujii, H. Asama, H. Kaetsu, and I. Endo, "Realization of autonomous navigation in multirobot environment," in *Proceedings of the 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Victoria, B.C., Canada, Oct. 13-17 1998, pp. 1999–2004.

[14] W. Williamson, J. Min, J. L. Speyer, and J. Farrell, "A comparison of state space, range space, and carrier phase differential GPS/INS relative navigation," in *Proceedings of the American Control Conference*, Chicago, IL, June 2000, pp. 2932–2938.

[15] R. Kurazume, S. Nagata, and S. Hirose, "Cooperative positioning with multiple robots," in *Proceedings of the 1994 IEEE International Conference in Robotics and Automation*, Los Alamitos, CA, 8-13 May 1994, vol. 2, pp. 1250–1257.

[16] R. Kurazume, S. Hirose, S. Nagata, and N. Sashida, "Study on cooperative positioning system (basic principle and measurement experiment)," in *Proceedings of the 1996 IEEE International Conference in Robotics and Automation*, Minneapolis, MN, April 22-28 1996, vol. 2, pp. 1421–1426.

[17] R. Kurazume and S. Hirose, "Study on cooperative positioning system: optimum moving strategies for CPS-III.," in *Proceedings of the 1998 IEEE International Conference in Robotics and Automation*, Leuven, Belgium, 16-20 May 1998, vol. 4, pp. 2896–2903.

[18] R. Kurazume and S. Hirose, "An experimental study of a cooperative positioning system," *Autonomous Robots*, vol. 8, no. 1, pp. 43–52, January 2000.

[19] R. Grabowski, L. E. Navarro-Serment, C. J.J. Paredis, and P. K. Khosla, "Heterogeneous teams of modular robots for mapping

and exploration," *Autonomous Robots*, vol. 8, no. 3, pp. 293–308, June 2000, Special Issue on Heterogeneous Multirobot Systems.

[20] I.M. Rekleitis, G. Dudek, and E.E. Milios, "Multi-robot exploration of an unknown environment, efficiently reducing the odometry error," in *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*, M.E. Pollack, Ed., Nagoya, Japan, 23-29 Aug. 1997, vol. 2, pp. 1340–1345.

[21] I.M. Rekleitis, G. Dudek, and E.E. Milios, "On multiagent exploration," in *Proceedings of Vision Interface*, Vancouver, Canada, June 1998, pp. 455–461.

[22] D. Fox, W. Burgard, H. Kruppa, and S. Thrun., "Collaborative multi-robot localization," in *In Proc. of the 23rd Annual German Conference on Artificial Intelligence (KI)*, Bonn, Germany, 13-15 Sept. 1999, pp. 255–266.

[23] D. Fox, W. Burgard, H. Kruppa, and S. Thrun, "A probabilistic approach to collaborative multi-robot localization," *Autonomous Robots*, vol. 8, no. 3, pp. 325–344, June 2000, Special Issue on Heterogeneous Multirobot Systems.

[24] S. Thrun, D. Fox, and W. Burgard, "Monte carlo localization with mixture proposal distribution," in *Proc. of the AAAI National Conference on Artificial Intelligence*, Austin, TX, 2000.

[25] S. I. Roumeliotis, "Distributed Multi-Robot Localization," Tech. Rep., California Institute of Technology, 2002, http://robotics.caltech.edu/~stergios/tech_reports/tr_collective.pdf.

[26] A.C. Sanderson, "A distributed algorithm for cooperative navigation among multiple mobile robots," *Advanced Robotics*, vol. 12, no. 4, pp. 335–349, 1998.

[27] S.I. Roumeliotis and G.A. Bekey, "Collective localization: A distributed kalman filter approach to localization of groups of mobile robots," in *Proceedings of the IEEE International Conference on Robotics and Automation*, San Fransisco, CA, April 24-28 2000, pp. 2958–2965.

[28] Stergios I. Roumeliotis, *Robust Mobile Robot Localization: From single-robot uncertainties to multi-robot interdependencies*, Ph.D. thesis, Electrical Engineering Department, University of Southern California, Los Angeles, California, May 2000.

[29] J. Borenstein, "Control and kinematic design of multi-degree-of freedom mobile robots with compliant linkage," *IEEE Transactions on Robotics and Automation*, vol. 11, no. 1, pp. 21–35, Feb. 1995.

[30] J. Borenstein, "Internal correction of dead-reckoning errors with a dual-drive compliant linkage mobile robot," *Journal of Robotic Systems*, vol. 12, no. 4, pp. 257–273, April 1995.

[31] J. Borenstein, "Experimental results from internal odometry error correction with the omnimate mobile robot," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 6, pp. 963–969, Dec. 1998.

[32] P. S. Maybeck, *Stochastic Models, Estimation and Control*, vol. 141-1 of *Mathematics in Science and Engineering*, Academic Press, 1979.

[33] P. Pirjanian, "Multi-objective multi-robot target acquisition," http://robotics.usc.edu/~paolo/presentations/icra2000.ppt.

[34] S. Premvuti and Jing Wang, "Relative position localizing system for multiple autonomous mobile robots in distributed robotic system: system design and simulation," *Robotics and Autonomous Systems*, vol. 18, no. 3, pp. 319–326, Aug. 1996.

[35] S. I. Roumeliotis, G. S. Sukhatme, and G. A. Bekey, "Smoother-based 3-d attitude estimation for mobile robot localization," in *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, Detroit, MI, May 10-15 1999, vol. 3, pp. 1979–1986.

[36] H. W. Sorenson, *Advances in Control Systems*, vol. 3, chapter Kalman Filtering Techniques, Academic Press, 1966.

[37] M. Bozorg, E.M. Nebot, and H.F. Durrant-Whyte, "Decentralized architecture for asynchronous sensors," *Autonomous Robots*, vol. 6, no. 2, pp. 147–164, April 1999.

[38] S. I. Roumeliotis, G. S. Sukhatme, and G. A. Bekey, "Sensor fault detection and identification in a mobile robot," in *Proceedings of the 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Victoria, BC, Canada, 13-17 Oct. 1998, vol. 3, pp. 1383–1388.

[39] A.G.O. Mutambara and M.S.Y. Al-Haik, "State and information space estimation: A comparison," in *Proceedings of the American Control Conference*, Albuquerque, New Mexico, June 1997, pp. 2374–2375.

[40] D. W. Gage, "Minimum-resource distributed navigation and mapping," in *SPIE Mobile Robots XV*, Boston, MA, Nov. 2000, vol. 4195-12.

[41] S. I. Roumeliotis and M. J. Mataric, "Small-world networks of mobile robots," in *Proc. of the Seventeenth National Conference on Artificial Intelligence*, Austin, TX, July 30 - Aug. 3 2000, pp. 25–26.

Stergios I. Roumeliotis received his Diploma in Electrical Engineering from the National Technical University of Athens, Greece in 1995, and his M.S. and Ph.D. degrees in Electrical Engineering from the University of Southern California, Los Angeles, CA in 1996 and 2000, respectively. He is currently a Postdoctoral Fellow at the Division of Engineering and Applied Science at the California Institute of Technology, Pasadena, CA. His research has focused on sensing and estimation for distributed autonomous systems, mobile robot localization, mapping and navigation, fault detection and identification, and human-robot interaction.

George A. Bekey is Gordon Marshall Professor of Computer Science, Electrical Engineering and Biomedical Engineering at the University of Southern California. He received the B.S. degree in Electrical Engineering from the University of California, Berkeley, in 1950 and the M.S. and Ph.D. degrees in Engineering from UCLA, in 1952 and 1962, respectively. His research interests are focused on intelligent robotic systems and their applications. He served as Chairman of the USC Electrical Engineering Department (1978-82), as Chairman of the Computer Science Department (1984-89) and as Associate Dean for Research of the USC School of Engineering (1996-99). He was the Founder of the Robotics Research Laboratory at USC.

In addition to his teaching and university research, Dr. Bekey's professional experience includes four years as a Research Engineer with the Department of Engineering at UCLA, where he worked primarily with computers; three years as a Computer Applications Engineer at Beckman Instruments and four years as Head of the Simulation Section and Senior Staff Engineer at TRW Systems in Los Angeles. At TRW he worked on simulation and control of space vehicles.

Dr. Bekey has published over 200 papers in the areas of simulation, biomedical engineering, control systems and robotics. He is co-author (with W.J. Karplus) of the text Hybrid Computation (Wiley, 1968) and co-editor of several books, the most recent being Neural Networks and Robotics (with K. Goldberg, Kluwer, 1993) and Robot Colonies (with R. Arkin, Kluwer, 1997). He is Editor-in-Chief of the international journal Autonomous Robots and Founding Editor of the IEEE Transactions on Robotics and Automation. He is a Member of the US National Academy of Engineering and a Fellow of several societies, including the Institute of Electrical and Electronics Engineers (IEEE), the American Association for the Advancement of Science (AAAS), the American Institute for Medical and Biological Engineering (AIMBE) and the American Association for Artificial Intelligence (AAAI). From 1996 to 1998 he served as President of the IEEE Robotics and Automation Society (RAS). In 2001 he received the Engelberger Award from the Robotics Institute of America for contributions to education in robotics. In 2002 he received the "Pioneer in Robotics and Automation" award from the IEEE RAS.