

Probabilistic Mutual Localization in Multi-Agent Systems from Anonymous Position Measures

Antonio Franchi, Giuseppe Oriolo, and Paolo Stegagno

Abstract—Recent research on multi-agent systems has produced a plethora of decentralized controllers that implicitly assume various degrees of agent localization. However, many practical arrangements commonly taken to allow and achieve localization imply some form of centralization, from the use of physical tagging to allow the identification of the single agent to the adoption of global positioning systems based on cameras or GPS. These devices clearly decrease the system autonomy and range of applicability, and should be avoided if possible. Following this guideline, this work addresses the mutual localization problem with anonymous relative position measures, presenting a robust solution based on a probabilistic framework. The proposed localization system exhibits higher accuracy and lower complexity ($O(n^2)$) than our previous method [1]. Moreover, with respect to more conventional solutions that could be conceived on the basis of the current literature, our method is theoretically suitable for tasks requiring frequent, many-to-many encounters among agents (e.g., formation control, cooperative exploration, multiple-view environment sensing). The proposed localization system has been validated by means of an extensive experimental study.

I. INTRODUCTION

Recently, the problem of controlling multi-agent (in particular, multi-robot) systems has attracted increasing attention in view of their wide field of application. Among the most studied topics, one may mention consensus [2], flocking [3], [4], formation control [5], connectivity maintenance [6], [7], distributed estimation [8] and sensing [9], [10]. Most techniques proposed for solving these problems implicitly assume that the system is localized, i.e., that estimates of agent configurations (position, orientation, or both) are available. In real-world applications, external tools such as GPS may be unavailable or unreliable. An alternative strategy aimed at increasing the system autonomy is to perform localization in a cooperative fashion, i.e., broadcasting agent-to-agent measurements (such as relative distance, bearing and orientation) via communication and using this information to jointly estimate the agent configurations. This approach requires that each agent is equipped with an agent detector providing relative measurements, and a communication system to exchange information with other agents.

Some authors have focused on the localization of static agents, commonly referred to as *network localization* [11], [12], [13], [14], [15]. Other works have dealt with moving agents, with the objective of estimating the poses of the

agents in a common fixed frame [16], [17], [18], [19], [20], and have generally shown that the agents' ability of sensing each other can improve the localization of the entire system. In the literature, this kind of multi-agent localization is usually referred as *cooperative localization*. However, agreeing on a common fixed frame already implies a form of centralization. A more decentralized approach to this problem can consider either a moving frame attached to each agent as in [21], or a different fixed frame for each agent as in [1]. We define *relative mutual localization* (RML) as the problem of estimating the relative poses among the moving frames attached to the agents, and *absolute mutual localization* (AML) as the problem of estimating the relative poses among the various fixed frames. RML and AML are equivalent if each agent is localized w.r.t. its own fixed frame.

In most previous works it is assumed that relative measurements include the *identity* of the agents. However, this may be difficult to achieve using simple sensors such as range finders, as it requires the identification of some distinctive feature for each agent, an ability that typically calls for more sophisticated sensors, such as cameras. Moreover, the definition of an identity tagging (such as a chromatic code) amounts to a form of centralization. In [1], we studied AML with *anonymous* relative position measurements, where the agent detector does not provide the identity of the sensed agents. In particular, we proposed a localization system based on two components: (1) MultiReg, a multiple registration algorithm that computes the change of coordinates among the relative frames of the agents using geometrical arguments to invert the measurement map (2) a multi-hypothesis Extended Kalman Filter (MH-EKF) that uses the output of MultiReg and the motion displacements to estimate the change of coordinates among the fixed frames of the agents.

In [22] we have shown that anonymity of the position measures causes a combinatorial ambiguity in the geometrical computation of coordinate changes when the formation is rotational symmetric. In this case, or even when the formation is 'almost' symmetric, the complexity of the MultiReg algorithm becomes exponential with the number of agents. Since in many applications (e.g., entrapment or escorting) multi-agent systems are required to move in regular, possibly rotational symmetric formations, this worst-case complexity may indeed materialize, severely affecting the performance of localization systems based on MultiReg.

In this paper, we address the RML problem with anonymous position measures. Rather than pursuing a relatively straightforward adaptation of our previous method in [1], we significantly modify the localization system in a *probabilistic*

Antonio Franchi is with the Max Plank Institute for Biological Cybernetics, Spemannstraße 38, 72012 Tübingen, Baden-Württemberg, Germany antonio.franchi@tuebingen.mpg.de

Giuseppe Oriolo and Paolo Stegagno are with the Dipartimento di Informatica e Sistemistica, Università di Roma La Sapienza, Roma, Via Ariosto 25, 00185 Italy {oriolo, stegagno}@dis.uniroma1.it




Fig. 1. Mutual localization with anonymous position measures. (a) A group of robots with the associated moving frames (b–e) the feature sets detected by each robot at time t .

sense to overcome the aforementioned difficulties. The key ideas are (1) to use a particle filter to directly estimate the probability function of the agents' relative pose (in other words, to solve RML also in the filter) and (2) to modify MultiReg so that it can use this information as a feedback.

The benefits of our choices are threefold. First, the direct estimation of the relative pose avoids the amplification of the angular error. Second, the particle filter is intrinsically multi-modal and therefore it does not require the use of heuristics for data association. Third, the new framework allows MultiReg to focus on solutions that are most likely according to the current belief, filtering out the effects of rotational symmetries that may arise in the system and avoiding the associated complexity increase.

The paper is organized as follows. Sections II and III present the problem formulation and the localization system architecture, respectively. The novel multiple registration algorithm is introduced in Section IV, while Section V describes the particle filters for estimating the probability density functions of the agent configurations. Section VI reports experimental results, and Section VII outlines a discussion and some future work.

II. PROBLEM FORMULATION

Consider a system of $n \geq 2$ agents (henceforth called *robots*) $\mathcal{R}_1, \dots, \mathcal{R}_n$, with n unknown (hence, it may change during the operation). Denote by $\mathcal{N} = \{1, \dots, n\}$ the robot index set, and let $\mathcal{N}_i = \mathcal{N} \setminus \{i\}$. The robots move in the plane and a moving frame \mathcal{F}_i is rigidly attached to each \mathcal{R}_i (see Fig. 1a). We describe the localization problem from the viewpoint of a generic \mathcal{R}_i , as in [21] and [23]. The superscripts t and $1:t$ denote the value of a variable at the discrete time instant t and all its values at time instants $1, 2, \dots, t$, respectively. The 3-vector describing the position and orientation of \mathcal{F}_j w.r.t. \mathcal{F}_i is the *relative pose* x_j of \mathcal{R}_j , $j \in \mathcal{N}_i$. We use the operators \oplus and \ominus for the composition of poses [24].

Each robot \mathcal{R}_k , $k \in \mathcal{N}$, has a *motion detector* that provides u_k^t , a measure of its displacement between $t-1$ and t . The motion detector is characterized by a probabilistic motion model $p(u'|u)$, where u' and u are, respectively, the 'true' and the measured displacement. In addition, each \mathcal{R}_k is equipped with a *robot detector*, a sensor device that measures the relative position (typically, as bearing and distance) of other robots in \mathcal{F}_k , without the associated identity (see Fig. 1b–e). Robot \mathcal{R}_h , $h \in \mathcal{N}_k$, is detected if it is placed in a *perception set* D_p that is rigidly attached to \mathcal{F}_k . No assumption is taken on the shape D_p . As shown in Fig. 2,




Fig. 2. Robot detection and communication. Triangles are robots, black polygons are occluding objects, the grey region is the perception set and the white region is the communication set.

the robot detector is prone to false positives (it can be deceived by objects that look like robots) and false negatives (robots belonging to D_p which are not detected, e.g., due to line-of-sight occlusions). A probabilistic description of the robot detector is given in the form of a perception model $p(z|x_h \ominus x_k)$, where z is the measured relative position of \mathcal{R}_h in \mathcal{F}_k , and $x_h \ominus x_k$ is the relative pose of \mathcal{F}_h with respect to \mathcal{F}_k . In our method, $p(z|x_h \ominus x_k)$ does not account for false positives/negatives, because the multiple registration algorithm provides robustness to these outliers.

The measures coming from the robot detector will be generically referred to as *features*, as a reminder of the fact that they are anonymous and, in addition, may or may not represent actual robots. We denote by Z_k^t the set of features detected by \mathcal{R}_k at time t (Fig. 1b–e).

Finally, each robot \mathcal{R}_k , $k \in \mathcal{N}$, comes with a *communication module* that can send/receive data to/from any other robot \mathcal{R}_h , $h \in \mathcal{N}_k$, contained in a communication set D_c rigidly attached to \mathcal{F}_k . We will assume that $D_p \subset D_c$, so that if \mathcal{R}_k can detect \mathcal{R}_h it can also communicate with it. Each message sent by \mathcal{R}_k contains: (1) the robot signature (the index k), (2) the current composition of the motion displacements $u_k^{1:t} = u_k^1 \oplus \dots \oplus u_k^t$ incrementally obtained from the elementary measures provided by the motion detector (3) the feature set Z_k^t . False negatives may also affect the communication (robot belonging to D_c that do not receive messages), whereas false positives in the communication may be easily avoided by appropriate message coding. We denote by C_k^t the set of robots which communicate with \mathcal{R}_k at time t (*neighbors* of \mathcal{R}_k) and we let $C_k^{1:t} = \bigcup_{\tau=1}^t C_k^\tau$.

In a probabilistic framework, the RML problem with anonymous position measures requires the generic robot \mathcal{R}_i to compute, at each time instant t , its belief about the relative poses of those robots with which \mathcal{R}_i has communicated, on the basis of the position measures gathered directly by itself or obtained via communication with other robots.

Problem 1: (Probabilistic RML with anonymous position measures) For each $t = 1, 2, \dots$ and $j \in C_i^{1:t}$, compute the belief

$$\text{bel}(x_j^t) := p(x_j^t | u_i^{1:t}, Z_i^{1:t}, \{u_j^\tau, Z_j^\tau\}_{\tau=1, \dots, t, j \in C_i^\tau}),$$

given $u_i^{1:t}$, $Z_i^{1:t}$, u_j^τ and Z_j^τ , $\forall (\tau, j)$ s.t. $\tau = 1, \dots, t$ and $j \in C_i^\tau$.




Fig. 3. Scheme of the mutual localization system that runs on \mathcal{R}_i .

III. PROPOSED APPROACH

The distinctive aspects of Problem 1 with respect to standard versions of the mutual localization problem (e.g., those in [16], [21], [25]) are the anonymity of the measures and the presence of false positives/negatives. A classical approach to deal with these issues would be to guess the identity of the measured robot using a data association technique, such as Maximum Likelihood [26] or Joint Compatibility Test [27], apply an outlier rejection method, and use the labeled measures to feed a recursive estimation algorithm, e.g., a particle filter or an MH-EKF. These approaches are reasonably successful provided that robot encounters are occasional and primarily pairwise, so that the feature sets Z_k^t , $k \in \mathcal{N}$, include few or no elements for most of the time. In fact, data association techniques do not take into account the mutual exclusion constraint of simultaneous measures, e.g., the fact that two different features measured at the same time cannot represent the same robot.

Here, we would like to design a mutual localization method that is effective for tasks requiring frequent, not necessarily pairwise encounters among robots (e.g., formation control, cooperative exploration, multiple-view environment monitoring). With this objective, we adopt the approach outlined in Fig. 3. The generic robot \mathcal{R}_i applies a multiple registration algorithm to compute the most likely relative poses of the robots belonging to C_i^t , on the basis of the sets of features Z_i^t , $\{Z_j^t\}_{j \in C_i^t}$, and the current beliefs about $\{x_j^t\}_{j \in C_i^t}$. The relative poses thus obtained, together with the motion detector measures u_i^t and $\{u_j^{1:t}\}_{j \in C_i^t}$, are used by $|C_i^{1:t}|$ particle filters to update the belief about the pose of each robot in $C_i^{1:t}$. The multiple registration and the particle filters are respectively described in Sections IV and V.

The advantages of our two-stage approach with respect to the more classical possibilities discussed above are:

- 1) the mutual exclusive structure of the set of features is exploited in the registration phase;
- 2) the increased dimension of the measures (the relative angle is also provided) results in a faster convergence of the estimation process;
- 3) the multi-robot system achieves, in a distributed way, a result comparable with the outcome of a single centralized sensor taking simultaneous snapshots of the same scene from different viewpoints. This expands the perception capabilities of the system members beyond those of the individual robots, without asking for a central data processor. For example, \mathcal{R}_i can now

estimate the pose of robots which are occluded by an obstacle or in a blind region of the robot detector, provided that they are seen by other robots.

These advantages come at an acceptable price. In fact, the increased complexity experienced when the initial¹ arrangement of the system is close to being ambiguous (rotational symmetric) would still be unavoidable with the classical approaches (e.g., using a Mahalanobis-distance data association). In any case, as will be shown, the particle filter copes flawlessly with the multiple hypotheses computed in this situation by the registration algorithm for the relative pose of the robots.

IV. PROBABILISTIC MULTIPLE REGISTRATION

At each time instant t , the generic robot \mathcal{R}_i runs *Probabilistic MultiReg* (abbreviated as P-MultiReg), a multiple registration algorithm that is based on MultiReg [1] and represents the part of the localization system which directly interfaces with the particle filters (see Fig. 3). In general, *registration* is the process of computing the relative pose between two or more different viewpoints of the same scene. In our case, P-MultiReg derives an estimate \hat{x}_j^t of each relative pose x_j^t , $j \in C_i^t$, given the sets of features Z_i^t , $\{Z_j^t\}_{j \in C_i^t}$, and the beliefs $\{\text{bel}(x_j^t)\}_{j \in C_i^t}$, which are obtained from the motion model blocks of the particle filters (see Fig. 3). In particular, P-MultiReg uses RANSAC [28] to identify the most likely correspondences between tuple of features belonging to Z_i^t , $\{Z_j^t\}_{j \in C_i^t}$ and a least squares estimation to compute the poses \hat{x}_j^t that best fit those correspondences, i.e., that maximize the likelihood of the measures. A pseudo-code description of P-MultiReg is given in Algorithm 1. Below, we describe in some detail its operation.

The first iteration of P-MultiReg goes as follows. First, the algorithm adds a feature in $(0,0)$, labeled as \mathcal{R}_k , to every Z_k^t , $k \in C_i^t \cup \{i\}$, in order to include the sensing robot in the picture. Since all other features are unlabeled, the feature sets are now partially labeled. Then, P-MultiReg performs $|C_i^t|$ binary registrations² (see Fig. 4) between Z_i^t and every set in $\{Z_j^t\}_{j \in C_i^t}$. The intermediate results obtained from these registrations are: (1) a group of aggregate sets of

¹The situation is completely different if the system comes from a steadily non-ambiguous situation so that the localization method has been able to reduce uncertainty to a minimum, as in the experiment of Fig. 9. In this case, the arising of rotational symmetries does not affect the complexity of the algorithm. A control law that breaks symmetry so as to allow non-ambiguous mutual localization may be used to recover this kind of situation [22].

²Our implementation of binary registration is inspired to the algorithm presented in [29] and can be found in [30].

Algorithm 1: P-MultiReg

input : feature sets Z_i^t , $\{Z_j^t\}_{j \in C_i^t}$, beliefs $\overline{\text{bel}}\{x_j\}_{j \in C_i^t}$
output: estimate relative poses $\{\hat{x}_j^t\}_{j \in C_i^t}$

```

1 for  $i \leftarrow 1$  to  $|C_i^t|$  do
2   perform binary registrations among  $Z_i^t$  and all the
    $Z_j^t, j \in C_i^t$ , not registered yet;
3   select a maximal subset of irreconcilable solutions;
4   select the partial solutions whose metric (1) is above a
   certain threshold;
5   foreach selected solution  $S$  do
6     expand the aggregated set of feature with  $S$ ;
7     tune all the already estimated relative poses taking
     into account new correspondences;
8     create a new branch of the algorithm if more than
     one solution;

```

features (now with two labeled features) obtained from each binary registration (2) a group of (temporary) relative poses, one for each aggregate set. Then, the algorithm prunes the results, discarding those which are either redundant (non-irreconcilable in the terminology of [1], see Fig. 5) or do not fit adequately the corresponding current belief in $\{\overline{\text{bel}}(x_j^t)\}_{j \in C_i^t}$ according to the metric function

$$\int p(\hat{x}_j^t | x_j^t) \overline{\text{bel}}(x_j^t) dx_j^t \quad (1)$$

and a suitable threshold. In the above function, $p(\hat{x}_j^t | x_j^t)$ is the probability to measure \hat{x}_j^t given that the robot is in x_j^t , as computed from the perception model of the robot detector.

Assume for a moment³ that, at the end of this first iteration (binary registrations plus pruning), a single aggregate set of features Z_{ik}^t ($k \in C_i^t$) survives, with the associated temporary estimate of \hat{x}_k^t . The algorithm then performs a second iteration consisting of $|C_i^t| - 1$ binary registrations between Z_{ik}^t and every set in $\{Z_j^t\}_{j \in C_i^t \setminus \{k\}}$, followed by a pruning of their outputs. The intermediate result is an aggregate set of features Z_{ikl}^t ($l \in C_i^t \setminus \{k\}$), now with three labeled features, and two temporary pose estimates \hat{x}_k^t and \hat{x}_l^t . Note that \hat{x}_k^t is updated to account for new correspondences between features of Z_k^t and Z_l^t in the aggregate set Z_{ikl}^t . The algorithm then repeats the basic iteration until no feature sets from $\{Z_j^t\}_{j \in C_i^t}$ are left, producing in the end a pose estimate \hat{x}_j^t for each $j \in C_i^t$ with which registration was successful⁴ and sufficiently consistent with a current belief.

To gain a deeper understanding of how P-MultiReg works, consider now the more complicated situation of Fig. 6. Here, the robots are arranged in a formation close to be rotational symmetric, so that the system configuration is ambiguous from the registration viewpoint (Fig. 6a). The objective of the algorithm is to register the directly perceived feature set Z_i^t with the communicated feature sets $\{Z_j^t, Z_k^t\}$ (Fig. 6b). At the start, P-MultiReg performs $|C_i^t| = 2$ binary registrations between Z_i^t and Z_j^t, Z_k^t , respectively (Fig. 6c), obtaining a




Fig. 4. An example of a binary registration between partially labeled sets of features: (a) initial feature sets Z_1 and Z_2 (b) proposed association (involved features are circled in red) (c) aggregate set Z_{12} and estimated relative pose between \mathcal{R}_1 and \mathcal{R}_2 (the displacement between the two triangles).




Fig. 5. Three non-redundant (irreconcilable) intermediate results (note that the feature sets come from different binary registrations and therefore do not coincide). The first and the second are irreconcilable because \mathcal{R}_2 is assigned to different features, while the first and the third are irreconcilable because the same feature is assigned to different robots (\mathcal{R}_4 and \mathcal{R}_5).

maximum subset with two irreconcilable solutions (Fig. 6d). These are rated on the basis of their fitness w.r.t. the current belief according to the metric function (1); for example, assume that in this case the current belief indicates that only the solution placing \mathcal{R}_j at the rightmost feature can be accepted⁵. The result of the first iteration is therefore the aggregate set of features Z_{ij}^t (Fig. 6e) together with a temporary estimate for the relative pose of \mathcal{R}_j . In the second iteration, the algorithm performs a single binary registration between Z_{ij}^t and the remaining feature set, i.e., Z_k^t (Fig. 6f), and checks the fitness of the result w.r.t. the current belief; if it is above the threshold, the final solution and the associated estimates \hat{x}_j^t and \hat{x}_k^t are produced (Fig. 6g).

The running time of P-MultiReg, which accounts for most of the cycle time of our mutual localization system, depends on the number $|C_i^t| + 1 \leq n$ of feature sets it receives as input. In normal operation (no ambiguity, or ambiguities that can be resolved based on the belief function), P-MultiReg expands only one branch, which executes $(|C_i^t| - 1)(|C_i^t| - 2)/2$ binary registrations to produce a solution: moreover, each binary registration requires constant time. This leads to a worst-case complexity $O(n^2)$, while the average-case complexity can be significantly lower if the number of communicating robots is smaller than n . Clearly, the complexity of P-MultiReg is the same of MultiReg (i.e., exponential in n) in the presence of ambiguities that cannot be resolved.

⁵If both solutions pass the fitness test (for example, because the current belief on the poses of \mathcal{R}_j and \mathcal{R}_j is uniform), the algorithm would expand two branches, leading to two different solutions. In this case, P-MultiReg would behave exactly as MultiReg.

³This is indeed the normal situation when the uncertainty is small.

⁴The registration had at least a solution.

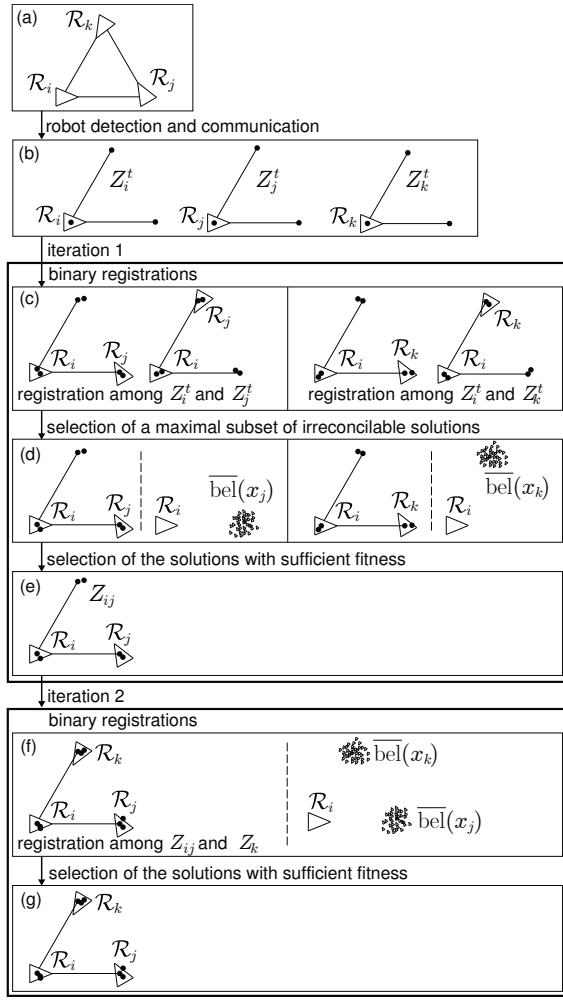


Fig. 6. Execution of P-MultiReg in a simple ambiguous situation: (a) actual configuration (b) initial feature sets with the addition of the labeled features at the origin (c) results of the binary registrations between Z_i and Z_j , Z_k , respectively (d) selection of a maximal subset of irreconcilable solutions (e) selection of the solutions with sufficient fitness w.r.t. the belief (f) result of the binary registration between Z_{ij} and Z_k (g) final result.

V. PARTICLE FILTERS

The generic robot \mathcal{R}_i maintains one particle filter for each \mathcal{R}_j . The use of separate beliefs $p(x_j^t)$, with $j \in C_i^{1:t}$, instead of a single joint belief $\{x_j^t\}_{j \in C_i^{1:t}}$ relies on the independence assumption, i.e., $p(\{x_j^t\}_{j \in C_i^{1:t}}) = \prod_{j \in C_i^{1:t}} p(x_j^t)$. The assumption is true in a pure localization scenario, while in certain situations, e.g., distributed control, it is an acceptable approximation. However, maintaining $p(\{x_j^t\}_{j \in C_i^{1:t}})$ is not feasible from a computational point of view, since it would lead to a distribution whose dimension grows exponentially with the number of robots.

At time t , the j -th filter ($j \in C_i^{1:t}$) receives as inputs the motion displacement u_i of \mathcal{R}_i plus, for each $j \in C_i^t$:

- 1) the motion displacements $u_j^{1:t}$ (sent by \mathcal{R}_j);
- 2) the relative pose estimate \hat{x}_j^t (computed by P-MultiReg).

In particular, \hat{x}_j^t is approximated as a gaussian measurement with a covariance which reflects the uncertainty in the registration steps of the algorithm.

The update rule that accounts for the motion of \mathcal{R}_i is

$$p(x_j|u_i) = N_i \int p(u'|u_i)p(x_j \oplus u')du', \quad (2)$$

where N_i is a normalization factor and $p(u'|u)$ is the motion detector model. Equation (2) leads to the following update for the single particle:

$$x_j = x_j \ominus (u_i \oplus n_u), \quad (3)$$

where n_u is a sample taken by $p(u'|u)$. Similarly, the update rule that accounts for the motion of \mathcal{R}_j is

$$p(x_j|u_j) = N_j \int p(u'|u_j)p(x_j \ominus u')du', \quad (4)$$

where N_j is a normalization factor, and the update equation for the single particle is

$$x_j = x_j \oplus (u_j \oplus n_u). \quad (5)$$

Updates due to the motion of \mathcal{R}_i and \mathcal{R}_j cause a translation of $p(x_j)$, while the additive noise introduces a blur. Using Bayes law, the measurement update is given by

$$p(x_j|\hat{x}_j) = Np(\hat{x}_j|x_j)p(x_j), \quad (6)$$

where N is another normalization factor.

Normally, \mathcal{R}_i uses $u_j^{1:t} \ominus u_j^{1:t-1}$ as motion measurement for the motion update of the robot \mathcal{R}_j . However, when \mathcal{R}_i and \mathcal{R}_j do not communicate for a time interval $[t_a+1, t_b-1]$ (e.g., due to the fact that the robots are far from each other) the motion update of \mathcal{R}_j cannot be performed. When, at t_b , the communication is resumed, \mathcal{R}_i uses $u_j^{1:t_b} \ominus u_j^{1:t_a}$ as motion measurement for the motion update. This explains why the robots send $u_j^{1:t}$ instead of u_j^t .

A number of standard practical techniques have been used to improve the performance of the filter. For example, the initial prior distribution is generated using the first measurements. Moreover, at each step a small percentage of particles are re-initialized using the new measurements; this enables the localization system to deal with the *kidnapped robot* situation (see the next section). We have also reduced the frequency of the measurement update with respect to the motion update to guarantee the independence of subsequent measurements.

VI. EXPERIMENTAL RESULTS

The proposed system for mutual localization with anonymous position measures has been implemented and tested on a team of four Khepera III robots using MIP, a multi-robot oriented software platform⁶ that minimizes the porting effort required to go from simulated to real robots.

Each robot is equipped with a Hokuyo URG-04LX laser sensor that provides range readings at 10 Hz within a 240° field of view (thus leaving a 120° blind zone behind the robot). The robot detector is a simple feature extraction module that inspects the laser scan, looking for the indentations given by the small ‘hat’ mounted atop each robot. Motion

⁶See <http://www.dis.uniroma1.it/labrob/software/MIP>.

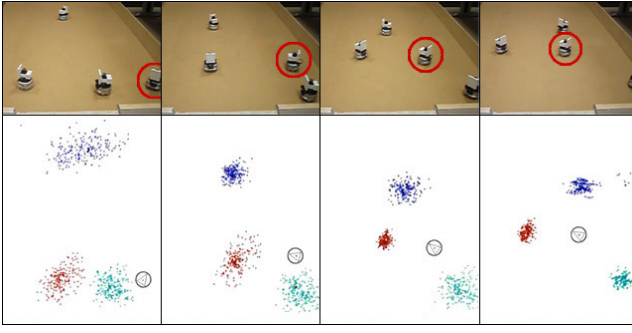


Fig. 7. First experiment. Top: snapshots of the scene; bottom: sample distributions computed by \mathcal{R}_1 (circled).




Fig. 8. Second experiment (kidnapping). Top: snapshots of the scene; bottom: sample distributions computed by \mathcal{R}_1 (circled).

measurements are computed by odometric integration over a 0.1 s time interval to match the robot detector refresh rate. Communication introduces no significant delay. The whole mutual localization system runs at 10 Hz, showing that the associated computations (P-MultiReg and particle filters) can be easily performed in real time.

Results from the first experiment are shown in fig. 7. The robots move in a free arena under the action of a simple pseudo-random control law that includes obstacle avoidance. The sample distributions computed by \mathcal{R}_1 clearly indicate the progressive uncertainty reduction. The orientation of the robots, which is included in the samples, is not shown. In the fourth snapshot, one of the robots (barely visible in the lower right corner) is actually in the blind zone of all the others; still, thanks to P-MultiReg, an estimate for its relative pose is computed and used to update the particle filter.

During the second experiment, shown in Fig. 8, one of the robots is kidnapped (second snapshot) and released in a different position (third snapshot). The localization system quickly recovers the correct formation, as confirmed by the sample distributions in the last frame.

We have also performed extensive mutual localization experiments among obstacles and robot-like objects, concluding that the proposed method exhibits the same kind of robustness seen in [1] with respect to false positives and negatives. While we do not show these results here, we present a different experiment that highlights the specific advantages of the new approach, and in particular the benefit of using P-MultiReg rather than MultiReg for registration.

In the third experiment (Fig. 9) three robots are station-




Fig. 9. Third experiment. Top: snapshots of the scene; Middle: sample distributions computed for \mathcal{R}_2 by \mathcal{R}_1 (circled) using MultiReg; Bottom: sample distributions computed for \mathcal{R}_2 by \mathcal{R}_1 (circled) using P-MultiReg.




Fig. 10. Execution time of P-MultiReg (solid, blue) and MultiReg (dashed, red) during the third experiment.




Fig. 11. Noisy simulation. Solid (blue): distance, bearing and orientation errors affecting the pose estimate generated by the proposed method. Dashed (red): the same errors as generated by the method in [1].

ary while the fourth (\mathcal{R}_2) is driven towards a rotational symmetric (hence, ambiguous) configuration of the system. \mathcal{R}_1 (circled) estimates the pose of \mathcal{R}_2 . When the system gets close to a symmetrical formation (last frame), the mutual localization system that uses MultiReg computes a three-modal sample distribution. As expected, the use of P-MultiReg allows instead to solve the ambiguity.

The execution times of MultiReg and P-MultiReg during the third experiment are shown in Fig. 10. They are comparable until the configuration is far from being symmetrical. When the ambiguity starts to increase, the execution time of MultiReg grows exponentially, while that of P-MultiReg remains unchanged, corroborating our theoretical prediction.

Finally, in the absence of a ground truth system for our experimental setup, we resort to a noisy simulation (including, in particular, a substantial odometric drift) to gather some numerical data on the localization accuracy. Figure 11 reports error data for the pose estimate of a certain robot as generated by the proposed method and by the method in [1]. The former produces a smoother, more accurate position estimate, essentially due to the regularizing effect of feeding back the predicted beliefs to the multiple registration module and to the averaging action over the sample distribution. The larger error on the orientation can be probably reduced with a more careful tuning. The MH-EKF module of the method in [1], which selects the best current pose based on a ranking of the hypotheses, produces instead a more erratic estimate.

See <http://www.dis.uniroma1.it/labrob/research/mutLoc.html> for additional material, including video clips of the experiments.

VII. CONCLUSIONS

We have presented a decentralized and experimentally-validated system for the mutual localization of multi-agent systems with anonymous position measures. We have obtained significant improvements with respect to our previous work on this topic. In particular, the use of a particle filter and the feedback of beliefs in the registration phase have been effective in solving both the ambiguity and complexity issues that affected previous solutions.

An interesting extension of this work would be the development of a similar framework for 3D agents (i.e., in $SE(3)$), for the application to swarms of flying robots, possibly equipped with more limited (e.g., bearing-only) sensors.

REFERENCES

- [1] A. Franchi, G. Oriolo, and P. Stegagno, "Mutual localization in a multi-robot system with anonymous relative position measures," in *2009 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, St. Louis, MO, Oct. 2009, pp. 3974–3980.
- [2] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.
- [3] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Flocking in fixed and switching networks," *IEEE Trans. on Automatic Control*, vol. 52, no. 5, pp. 863–868, 2007.
- [4] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: algorithms and theory," *IEEE Trans. on Automatic Control*, vol. 51, no. 3, pp. 401–420, 2006.
- [5] A. Rahmani, M. Ji, M. Mesbahi, and M. Egerstedt, "Controllability of multi-agent systems from a graph-theoretic perspective," *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 162–186, 2009.
- [6] R. A. Freeman, G. J. Gordon, K. M. Lynch, S. S. Srinivasa, and R. Sankar, "Decentralized estimation and control of graph connectivity in mobile sensor networks," in *2008 American Control Conference*, Seattle, WA, Jun. 2008, pp. 2678–2683.
- [7] T. Gustavi, D. V. Dimarogonas, M. Egerstedt, and X. Hu, "Sufficient conditions for connectivity maintenance and rendezvous in leader-follower networks," *Automatica*, vol. 46, no. 1, pp. 133–139, 2010.
- [8] P. Yang, R. A. Freeman, and K. M. Lynch, "Multi-agent coordination by decentralized estimation and control," *IEEE Trans. on Automatic Control*, vol. 53, no. 11, pp. 2480–2496, 2008.
- [9] S. Martinez, F. Bullo, J. Cortes, and E. Frazzoli, "On synchronous robotic networks - Part II: Time complexity of rendezvous and deployment algorithms," *IEEE Trans. on Automatic Control*, vol. 52, no. 12, pp. 2214–2226, 2007.
- [10] L. C. A. Pimenta, V. Kumar, R. C. Mesquita, and G. A. S. Pereira, "Sensing and coverage for a network of heterogeneous robots," in *47th IEEE Conf. on Decision and Control*, Cancun, Mexico, Dec. 2008, pp. 3947–3952.
- [11] J. Aspnes, T. Eren, D. K. Goldenberg, A. S. Morse, W. Whiteley, Y. R. Yang, B. D. O. Anderson, and P. N. Belhumeur, "A theory of network localization," *IEEE Trans. on Mobile Computing*, vol. 5, no. 12, pp. 1663–1678, 2006.
- [12] T. Eren, W. Whiteley, and P. Belhumeur, "Using angle of arrival (bearing) information in network localization," in *45th IEEE Conf. on Decision and Control*, San Diego, CA, Jan. 2006, pp. 4676–4681.
- [13] G. Piovan, I. Shames, B. Fidan, F. Bullo, and B. Anderson, "On frame and orientation localization for relative sensing networks," in *47th IEEE Conf. on Decision and Control*, Cancun, Mexico, Dec. 2008, pp. 2326–2331.
- [14] J. Fang and A. S. Morse, "Network localization using graph decomposition and rigidity," in *48th IEEE Conf. on Decision and Control*, Dec. 2008, pp. 1091–1096.
- [15] I. Shames, B. Fidan, and B. D. O. Anderson, "Minimization of the effect of noisy measurements on localization of multi-agent autonomous formations," *Automatica*, vol. 45, no. 4, pp. 1058–1065, 2009.
- [16] 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, 2000.
- [17] S. I. Roumeliotis and G. A. Bekey, "Distributed multirobot localization," *IEEE Trans. on Robotics*, vol. 18, no. 5, pp. 781–795, 2002.
- [18] A. Howard, M. J. Mataric, and G. S. Sukhatme, "Localization for mobile robot teams using maximum likelihood estimation," in *2002 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Lausanne, Switzerland, Sep. 2002, pp. 434–439.
- [19] A. Das, J. Spletzer, V. Kumar, and C. Taylor, "Ad hoc networks for localization and control," in *41th IEEE Conf. on Decision and Control*, Dec. 2002, pp. 2978–2983.
- [20] A. I. Mourikis and S. I. Roumeliotis, "Performance analysis of multirobot cooperative localization," *IEEE Trans. on Robotics*, vol. 22, no. 4, pp. 666–681, 2006.
- [21] A. Howard, M. J. Mataric, and G. S. Sukhatme, "Putting the 'I' in 'team': An ego-centric approach to cooperative localization," in *2003 IEEE Int. Conf. on Robotics and Automation*, Taipei, Taiwan, Sep. 2003, pp. 868–892.
- [22] A. Franchi, G. Oriolo, and P. Stegagno, "On the solvability of the mutual localization problem with anonymous position measures," in *2010 IEEE Int. Conf. on Robotics and Automation*, Anchorage, AK, May 2010.
- [23] A. Martinelli, "Improving the precision on multi robot localization by using a series of filters hierarchically distributed," in *2007 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, San Diego, CA, Nov. 2007, pp. 1053–1058.
- [24] R. C. Smith and P. Cheeseman, "On the representation and estimation of spatial uncertainty," *International Journal of Robotics Research*, vol. 5, no. 4, pp. 56–68, 1986.
- [25] X. S. Zhou and S. I. Roumeliotis, "Robot-to-robot relative pose estimation from range measurements," *IEEE Trans. on Robotics*, vol. 24, no. 6, pp. 1379–1393, 2008.
- [26] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [27] J. Neira and J. D. Tardós, "Data association in stochastic mapping using the joint compatibility test," *IEEE Trans. on Robotics*, vol. 17, no. 6, pp. 890–897, 2001.
- [28] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [29] D. Fontanelli, L. Ricciato, and S. Soatto, "A fast ransac-based registration algorithm for accurate localization in unknown environments using lidar measurements," in *2007 IEEE Int. Conf. on Automation Science and Engineering*, 2007, pp. 597–602.
- [30] A. Franchi, G. Oriolo, and P. Stegagno, "Mutual localization of a multi-robot team with anonymous relative position measures," Department of Computer and System Sciences, Tech. Rep. 1, January 2009. [Online]. Available: http://padis2.uniroma1.it:81/ojs/index.php/DIS_TechnicalReports/issue/view/157