

Mutual localization in multi-robot systems using anonymous relative measurements

The International Journal of
Robotics Research
32(11) 1302–1322
© The Author(s) 2013
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/0278364913495425
ijr.sagepub.com



Antonio Franchi¹, Giuseppe Oriolo² and Paolo Stegagno²

Abstract

We propose a decentralized method to perform mutual localization in multi-robot systems using anonymous relative measurements, i.e. measurements that do not include the identity of the measured robot. This is a challenging and practically relevant operating scenario that has received little attention in the literature. Our mutual localization algorithm includes two main components: a probabilistic multiple registration stage, which provides all data associations that are consistent with the relative robot measurements and the current belief, and a dynamic filtering stage, which incorporates odometric data into the estimation process. The design of the proposed method proceeds from a detailed formal analysis of the implications of anonymity on the mutual localization problem. Experimental results on a team of differential-drive robots illustrate the effectiveness of the approach, and in particular its robustness against false positives and negatives that may affect the robot measurement process. We also provide an experimental comparison that shows how the proposed method outperforms more classical approaches that may be designed building on existing techniques. The source code of the proposed method is available within the MLAM ROS stack.

Keywords

Mobile robots, localization, multi-robot localization, unknown data-association, anonymous measurements, relative measurements, particle filters, EKF, RANSAC, range finders

1. Introduction

Mutual localization, defined as the estimation of relative configurations in a group of robots, is a prerequisite for many multi-robot applications such as formation control (Chen et al., 2010; Franchi et al., 2012), cooperative transportation (Fink et al., 2010), coverage/exploration (Howard et al., 2006; Schwager et al., 2009; Durham et al., 2012) and monitoring (Pasqualetti et al., 2012).

External localization tools such as motion capture systems or GPS may sometimes be used to perform mutual localization, at the cost of reducing the system autonomy and flexibility. These solutions are however unfeasible or unreliable in many real-world situations, for example when the robots are operating in unstructured indoor environments or under the surface of the water. A much preferable approach is to preserve the system autonomy by letting the robots perform the localization on their own. In this case, robot-to-robot measurements are collected and broadcast to the group to be used in a cooperative localization scheme. Clearly, this approach requires that each robot is equipped with a sensor that provides some form of relative measurement (e.g. relative distance or bearing) as well

as a communication module for transmitting and receiving information.

Several solutions have been proposed in the literature for different versions of the mutual localization problem. Some authors have considered the localization of static agents, commonly referred to as *network localization* (Aspnes et al., 2006). Shames et al. (2009) address the problem of reconstructing relative positions from distance measurements, whereas the objective of Piovan et al. (2013) is to estimate relative orientations using bearing measurements. Dieudonne et al. (2008) have proven that the network localization problem is NP-hard in the deterministic (no noise) case.

Other works have dealt with moving robots, with the objective of estimating the configurations of the agents in

¹Max Plank Institute for Biological Cybernetics, Tübingen, Germany

²Dipartimento di Ingegneria Informatica, Automatica e Gestionale, Sapienza Università di Roma, Roma, Italy

Corresponding author:

Antonio Franchi, Max Plank Institute for Biological Cybernetics, Spemannstraße 38, 72012 Tübingen, Germany.
Email: antonio.franchi@tuebingen.mpg.de

a common fixed frame. This problem, first addressed by Kurazume et al. (1994), is known as *cooperative localization*, a term introduced by Rekleitis et al. (1998). The use of teammates as landmarks for the localization problem has been proposed by Kurazume and Hirose (2000). Early methods have been mostly proposed in conjunction with navigation and exploration problems (Sanderson, 1998; Grabowski et al., 2000; Rekleitis et al., 2001). Other works have explicitly addressed the localization problem using particle filters (Fox et al., 2000), maximum likelihood estimators (Howard et al., 2002), and Extended Kalman filters (EKF) (Roumeliotis and Bekey, 2002; Martinelli et al., 2005; Mourikis and Roumeliotis, 2006). All of these works show that the agents' ability of sensing each other can improve the localization of the entire system.

A closely related problem, to which we refer as proper *mutual localization*, is to consider different moving frames attached to each robot and then attempt to estimate the relative configurations among the moving frames (Howard et al., 2003). Different estimation strategies have been proposed for this problem, such as the use of hierarchically distributed EKFs (Martinelli, 2007) or closed-form solvers in conjunction with (nonlinear) iterative least squares, both in the 2D (Zhou and Roumeliotis, 2008) and the 3D case (Trawny et al., 2010). An interesting study on the role of the communication in a decentralized localization algorithm has been proposed by Leung et al. (2010).

To the best of the authors' knowledge, in all previous works on network, cooperative, or mutual localization it is assumed that relative measurements include the *identity* of the robots, i.e. that the sensor 'knows' to which robot each measurement refers to. This ubiquitous assumption requires in practice the identification of a distinctive feature for each robot, an ability that is not embedded in simple sensors (e.g. cameras, range finders, RGB-D scanners) and typically calls for sophisticated classification algorithms based on some form of robot tagging.

However, individual identification via tagging has a number of potential drawbacks. First, the process may be prone to failures and false positives whenever the conditions of the surrounding environment are not ideal (e.g. visual tagging may be dramatically affected by slight changes in the lighting conditions). Second, it is difficult to discriminate robots in large groups using features such as shape or color, i.e. most kinds of tagging do not scale well with the number of robots. In any case, assigning an identity via tagging is a form of centralization, and hence decreases the autonomy of the system as well as that of the individual robots.

A more radical approach is to accept the fact that the robots' identities in the relative measurements are unknown, and then attempt to estimate these identities together with the relative configurations by using a more powerful localization strategy. The problem of *mutual localization with anonymous measurements* was first considered by Franchi et al. (2009a), where a two-phase localization system is proposed: first, a geometrical multiple registration algorithm is used to build hypotheses on the

identities of the robot associated with the measurements, and then a multi-hypothesis Extended Kalman filter (MH-EKF) is used to incorporate odometric measurements and reject associations that are inconsistent with the current belief.

In a subsequent work (Franchi et al., 2010a), we have proved that anonymity of relative position measurements causes a combinatorial ambiguity in the geometrical registration when the spatial arrangement of the robots is (or is close to being) rotational symmetric. In this case, the run time of the algorithm proposed by Franchi et al. (2009a) may become a factorial function of the number of agents. In applications where multi-robot systems are required to move in regular formations, this worst-case complexity may indeed materialize, affecting the performance of localization systems based on that algorithm. To overcome this difficulty, we modified the previous localization system by feeding back the current belief about the robot relative configurations into the multiple registration step, so as to narrow its search space and ultimately keep under control the total execution time even in the presence of ambiguities (Franchi et al., 2010b). In addition, the use of particle filters was advocated as a convenient alternative to the MH-EKF.

In this paper we develop to a mature stage the latter approach. In particular, the main contributions of the present work are the following: (1) the introduction of an algorithm which solves a novel, challenging, and practically relevant problem in the field of mutual localization; (2) a theoretical analysis of the implication of anonymity on the solutions of the localization problem; (3) an experimental validation of the proposed method proving its feasibility and effectiveness; and (4) a quantitative comparison of its performance against more classical approaches that may be designed building on existing techniques, such as FastSLAM.

The paper is organized as follows. Section 2 provides a detailed formulation of the problem. The implications of the measurements being anonymous are analyzed in Section 3. Section 4 gives an overview of the proposed approach, whose two main stages are described in detail in Sections 5 and 6. Experimental results are presented in Section 7, while a discussion and a quantitative comparison with a possible alternative approach are given in Section 8. Section 9 describes the source code release and Section 10 concludes the paper.

2. Problem formulation

In this section, we describe the relevant geometry and provide a formulation of the mutual localization problem with anonymous measurements. For the reader's convenience, the main symbols are collected in Table 1. Throughout the section, refer to Figures 1 and 2 for an illustration.

Consider a multi-robot system $\{\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_n\}$ moving in a planar world. The number n of robots is not known and may vary during the operation.




Fig. 1. Relevant frames and configuration variables for two robots $\mathcal{R}_i, \mathcal{R}_j$.

Table 1. Main symbols used in the paper.

\mathcal{R}_i	i th robot
n	number of robots
\mathcal{F}_W	world frame
\mathcal{F}_i	frame attached to \mathcal{R}_i
${}^W q_j$	absolute configuration $({}^W p_j, {}^W \theta_j)$ of \mathcal{R}_j in \mathcal{F}_W
${}^W p_j$	absolute position of \mathcal{R}_j in \mathcal{F}_W
${}^W \theta_j$	absolute orientation of \mathcal{R}_j in \mathcal{F}_W
q_j	relative configuration (p_j, θ_j) of \mathcal{R}_j w.r.t. \mathcal{R}_i (i.e. in \mathcal{F}_i)
p_j	relative position of \mathcal{R}_j w.r.t. \mathcal{R}_i (i.e., in \mathcal{F}_i)
θ_j	relative orientation of \mathcal{R}_j w.r.t. \mathcal{R}_i (i.e., in \mathcal{F}_i)
\oplus, \ominus	additive and subtractive composition of two configurations, or of a configuration and a (set of) position(s)
x^t	value of x at time t
$x^{1:t}$	time history of x from 1 to t
d_i^t	configuration displacement of \mathcal{R}_i between $t-1$ and t
δ_i^t	odometric measure of d_i^t
z	a generic measurement in its space
$p(d_i^t z)$	probabilistic model of the odometer
D	detection region
O_i^t	observation provided by the robot detector of \mathcal{R}_i at t
$o_{i,k}^t$	k th element of O_i^t , representing an anonymous relative position measurement
$p(p_i^t z)$	probabilistic model of the robot detector
C	communication region
N_i^t	neighbors of R_i (i.e. robots communicating with \mathcal{R}_i) at t
$\text{bel}(x)$	belief over variable x
$\overline{\text{bel}}(x)$	belief over variable x before the measurement update
$ \cdot $	cardinality of a set
P	set of points
$l(p)$	label associated to the point p
$\mathcal{L} = (P, \mathfrak{L})$	set of labeled points
M	matching between two sets of labeled points
q_M, \mathcal{L}_M	roto-translation and merged set of labeled points associated to matching M
\hat{q}_j^t	estimate on the relative configuration of R_j computed by the PMR algorithm
n_δ	sample taken from $p(\delta^t \delta)$

As in the works by Howard et al. (2003) and Martinelli (2007), we look at the mutual localization problem from

the subjective perspective of each robot. In the following, we describe the problem from the viewpoint of \mathcal{R}_i , whose objective is to localize every \mathcal{R}_j with respect to itself. Clearly, \mathcal{R}_i is completely generic.

With reference to Figure 1, denote by $\mathcal{F}_W = \{O_W, X_W, Y_W\}$ the world inertial frame and by $\mathcal{F}_j = \{O_j, X_j, Y_j\}$ the frame attached to robot \mathcal{R}_j . The *absolute configuration* of \mathcal{R}_j is ${}^W q_j = ({}^W p_j, {}^W \theta_j)$, with ${}^W p_j$ and ${}^W \theta_j$ respectively the position and orientation of \mathcal{F}_j with respect to (w.r.t.) \mathcal{F}_W ; whereas the *relative configuration* of \mathcal{R}_j w.r.t. \mathcal{R}_i is $q_j = (p_j, \theta_j)$, with p_j and θ_j respectively the position and orientation of \mathcal{F}_j w.r.t. \mathcal{F}_i (see Figure 1). Note that we have $q_i = (\mathbf{0}, 0)$. Both absolute and relative configurations take values in $SE(2)$.

Two robot configurations $q' = (p', \theta')$ and $q'' = (p'', \theta'')$ (absolute or relative) can be composed using the operators $\oplus, \ominus: SE(2) \times SE(2) \rightarrow SE(2)$ introduced by Smith and Cheeseman (1986) and defined as

$$q' \oplus q'' = (p' + R(\theta')p'', \theta' + \theta''), \quad (1)$$

$$q' \ominus q'' = (R(-\theta'')(p' - p''), \theta' - \theta''), \quad (2)$$

where $R(\phi) \in SO(2)$ is the planar rotation matrix associated with angle ϕ . For example, this allows us to express a relative configuration as

$$q_j = {}^W q_j \ominus {}^W q_i. \quad (3)$$

Note that in this context a configuration acts as a roto-translation. This dual interpretation will be used throughout the paper.

The above operators will also be used to compose a configuration with a position; in particular, $q' \oplus p''$ and $p' \ominus q''$ (note the order of the operands) return the position component of the right-hand side of (1) and (2), respectively. Accordingly, a relative position can be written as

$$p_j = {}^W p_j \ominus {}^W q_i.$$

We will also apply \oplus and \ominus to compose a configuration q with a set of positions $P = \{p', p'', \dots\}$, obtaining the new sets of points $q \oplus P$ and $P \ominus q$.

Let us now specify the sensor equipment of the robots. Clearly, we assume a discrete-time structure for the mutual localization system, in view of the use of sensor data with a certain refresh rate. To simplify the notation, time instants are denoted by $1, 2, \dots, t$, but in fact the sampling interval can have any duration. We denote by a superscript t the value of a quantity at time t , and by a superscript $1:t$ the history of its values at times $1, 2, \dots, t$; e.g. we write q_j^t and $q_j^{1:t}$. Refer to Figure 2 for an illustration of the relevant measured quantities.

The generic robot \mathcal{R}_i is equipped with an odometer and a robot detector. The *odometer* returns a measurement δ_i^t of d_i^t , the relative displacement of \mathcal{R}_i between two consecutive time instants:

$$d_i^t = {}^W q_i^t \ominus {}^W q_i^{t-1}.$$




Fig. 2. A system of 3 robots shown at two consecutive time instants. On the generic robot \mathcal{R}_i , the odometer provides a measurement δ_i^t of its own displacement d_i^t , while the robot detector measures the relative positions of other robots falling in its detection region. The robot detector is anonymous and prone to both false positives and negatives. At time $t - 1$, both \mathcal{R}_j and \mathcal{R}_k are detected, and therefore the observation O_i^{t-1} contains two point measurements $o_{i,1}^{t-1}$ and $o_{i,2}^{t-1}$. At time t , \mathcal{R}_j is not seen due to an occlusion, but the detector is deceived by an obstacle that looks as a robot; as a consequence, the observation O_i^t still contains two point measurements $o_{i,1}^t$ and $o_{i,2}^t$, even though there is no relationship between $o_{i,1}^t$ and $o_{i,1}^{t-1}$.

To account for uncertainty, the odometer is represented with a probabilistic model $p(d_i^t|z)$, where z is used, throughout the paper, to indicate a generic measurement in its appropriate space; in this case, $z = \delta_i^t$ is a measured roto-translation. Hence, $p(d_i^t|z)$ represents the probability that the true displacement of \mathcal{R}_i is d_i^t given that the measured displacement is z .

The robot detector of \mathcal{R}_i measures the relative position p_j^t (not the orientation) of every \mathcal{R}_j that falls in the *detection region* D , a planar region attached to \mathcal{R}_i . No assumption is taken on the extension or shape of D . Our experiments will be performed with limited detection regions that are not isotropic, and in particular have a blind zone; therefore, it may happen that \mathcal{R}_i detects \mathcal{R}_j but the reverse is not true. From a probabilistic point of view, the robot detector can be represented by a perception model $p(p_j^t|z)$, where z represents in this case the measured relative position, whose ‘true’ value p_j^t is the position component of (3) at time t .

In addition to the basic sensing uncertainty encoded in the ‘individual’ perception model, we assume that the actual operation of the detector is affected by the following complications, all stemming from the multi-robot nature of our problem (see Figure 2):

1. The detector may detect as robots certain deceiving objects which are similar¹ to robots (*false positives*).
2. The detector may fail to detect robots in D , e.g. due to line-of-sight occlusions (*false negatives*).
3. Most importantly, relative position measurements are *anonymous*, i.e. the detector cannot recognize the identities of the perceived robots.

Summarizing, at each time t the robot detector of \mathcal{R}_i provides a set of measured positions (i.e. points) $O_i^t = \{o_{i,1}^t, \dots, o_{i,m}^t\}$, called *observation* in the following. The number m of points in each observation varies with t ; in addition, each element of O_i^t may represent an actual robot

or a deceiving obstacle in D . Finally, since position measurements are anonymous, the index k in $o_{i,k}^t$ is only associated with the internal ordering of O_i^t , and has no relationship with the identity of the robot it (possibly) represents. For example, if the robot detector is based on a laser range finder it is natural to use an angular ordering.

The last ingredient of our mutual localization system is communication. In particular, it is assumed that at each time t robot \mathcal{R}_i broadcasts a message which can be received by any other robot \mathcal{R}_j falling in its *communication region* C , another planar region rigidly attached to \mathcal{R}_i . The message contains: (1) the integer i , i.e. the identity of \mathcal{R}_i , (2) the measured displacement δ_i^t , (3) the observation O_i^t . In turn, \mathcal{R}_i receives messages from the robots in whose communication region it falls; for compactness, these robots will be called *neighbors* of \mathcal{R}_i at t and collected in the set N_i^t . False negatives may also affect the communication (messages lost between robots that should be able to communicate).

Based on the above description and notation, the problem addressed in this paper may be formulated as follows.

Problem 1 (Mutual localization with anonymous position measurements, *from the viewpoint of \mathcal{R}_i*). For $t = 1, 2, \dots$ and $j \in \cup_{\tau=1}^t N_i^\tau$, compute an estimate of the relative configuration q_j^t of \mathcal{R}_j in the form of a belief function:

$$\text{bel}(q_j^t) := p(q_j^t | \delta_i^{1:t}, O_i^{1:t}, \{\delta_j^\tau, O_j^\tau\}_{\tau=1, \dots, t, j \in N_i^\tau}),$$

that is, the posterior probability for q_j^t given:

1. the measurements of the odometer and of the robot detector of \mathcal{R}_i ;
2. the measurements of the odometers and of the robot detectors of all robots that have been neighbors of \mathcal{R}_i at some time instant.

3. The implications of anonymity

In order to analyze the implications of anonymity in Problem 1 and to justify the proposed solution approach, we consider first its deterministic-static version, often referred in the literature as *network localization* (Aspnes et al., 2006). Depending on whether the position measurements are anonymous or not, we have the two following problems.²

Problem 2 (Network localization with anonymous position measurements, *from the viewpoint of* \mathcal{R}_i). *Modify Problem 1 by assuming that:*

- i) *the robots do not move;*
- ii) *relative position measurements are not affected by uncertainty (no noise, no false positives).*

Compute a set of relative configurations $\{\mathbf{q}_j, j \in N_i\}$ that are consistent with:

1. *the observation of the robot detector of* \mathcal{R}_i ;
2. *the observations of the robot detectors of* $\mathcal{R}_j, \forall j \in N_i$.

Problem 3 (Network localization with non-anonymous position measurements, *from the viewpoint of* \mathcal{R}_i). *Identical to Problem 2 with the following additional assumption:*

- iii) *each observation contains the identity of the detected robots.*

In both problems, localization only concerns robots that are neighbors of \mathcal{R}_i , i.e. robots that communicate their observations to \mathcal{R}_i . Note that Problem 3 is equivalent to the classical formulation of network localization.

Graph theory has been extensively used to analyze network localization problems with planar point (position-only) configurations and non-anonymous relative distance measurements (Aspnes et al., 2006; Jackson, 2007; Shames et al., 2009). In that context, it was shown the problem is uniquely solvable when the framework (i.e. the point formation) possesses a property known as *rigidity*. Proceeding from the work of Jackson (2007), we extend below the concept of rigid framework to our case, i.e. that of planar pose (position and orientation) configurations and relative position measurements, which may be anonymous or not.

A planar *pose-configuration/position-measurements framework* (simply *framework* in the following) is a pair $(\mathcal{G}, \mathcal{W}\mathbf{q})$, where $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a directed graph with vertexes \mathcal{V} and edges \mathcal{E} , and $\mathcal{W}\mathbf{q} = (\mathcal{W}\mathbf{p}, \mathcal{W}\theta)$ is a map from \mathcal{V} to $SE(2)$. The map $\mathcal{W}\mathbf{q}$ associates every vertex $v \in \mathcal{V}$ to an absolute configuration $\mathcal{W}\mathbf{q}(v) = (\mathcal{W}\mathbf{p}(v), \mathcal{W}\theta(v)) \in SE(2)$.

Consider two frameworks $(\mathcal{G}, \mathcal{W}\mathbf{q}_a)$ and $(\mathcal{G}, \mathcal{W}\mathbf{q}_b)$ that share the same graph but map the vertices to different configurations. We say that:

- $(\mathcal{G}, \mathcal{W}\mathbf{q}_a)$ and $(\mathcal{G}, \mathcal{W}\mathbf{q}_b)$ are *equivalent* if

$$\mathcal{W}\mathbf{p}_a(v') \ominus \mathcal{W}\mathbf{q}_a(v) = \mathcal{W}\mathbf{p}_b(v') \ominus \mathcal{W}\mathbf{q}_b(v)$$

for all $(v, v') \in \mathcal{E}$;

- $(\mathcal{G}, \mathcal{W}\mathbf{q}_a)$ and $(\mathcal{G}, \mathcal{W}\mathbf{q}_b)$ are *congruent* if

$$\mathcal{W}\mathbf{p}_a(v') \ominus \mathcal{W}\mathbf{q}_a(v) = \mathcal{W}\mathbf{p}_b(v') \ominus \mathcal{W}\mathbf{q}_b(v)$$

for all $v, v' \in \mathcal{V}$.

Using this terminology, we will say that $(\mathcal{G}, \mathcal{W}\mathbf{q})$ is *rigid* if all of its equivalent frameworks are also congruent frameworks.

Assume now we are given an instance of Problem 2 or 3, and consider the *associated framework* $(\mathcal{G}, \mathcal{W}\mathbf{q})$, defined by:

- $\mathcal{V} = \{v_k \mid k \in \{i\} \cup N_i\}$;
- $(v_k, v_h) \in \mathcal{E}$ if the measurement of the relative position $\mathcal{W}\mathbf{p}(v_h) \ominus \mathcal{W}\mathbf{q}(v_k)$ is available;
- $\mathcal{W}\mathbf{q}(v_k) = \mathcal{W}\mathbf{q}_k$.

The fact that our network localization problems are formulated from the viewpoint of \mathcal{R}_i entails that the associated framework only concerns its neighbors N_i . Note that robots that are detected by \mathcal{R}_i but do not communicate with it are not in N_i ; on the other hand, N_i may include robots that are not detected by \mathcal{R}_i .

Proposition 1. *An instance of Problem 3 is uniquely solvable if and only if its associated framework $(\mathcal{G}, \mathcal{W}\mathbf{q})$ is rigid.*

Proof. If the problem instance is uniquely solvable, the relative configurations of the robots in N_i w.r.t. \mathcal{R}_i are uniquely determined. Now consider a pair of vertices v_k, v_h with $k, h \in \{i\} \cup N_i$. Since we have $\mathcal{W}\mathbf{q}(v_k) \ominus \mathcal{W}\mathbf{q}(v_h) = (\mathcal{W}\mathbf{q}(v_k) \ominus \mathcal{W}\mathbf{q}(v_i)) \ominus (\mathcal{W}\mathbf{q}(v_h) \ominus \mathcal{W}\mathbf{q}(v_i)) = \mathcal{W}\mathbf{q}_k \ominus \mathcal{W}\mathbf{q}_h$, also $\mathcal{W}\mathbf{p}(v_h) \ominus \mathcal{W}\mathbf{q}(v_k)$ is uniquely determined, which in turn implies that any framework which is equivalent to $(\mathcal{G}, \mathcal{W}\mathbf{q})$ must also be congruent to it.

Now suppose that $(\mathcal{G}, \mathcal{W}\mathbf{q})$ is rigid. Then, $\mathcal{W}\mathbf{p}(v_h) \ominus \mathcal{W}\mathbf{q}(v_k)$ is uniquely determined for any pair of vertices v_k, v_h with $k, h \in \{i\} \cup N_i$. Hence, $\mathbf{q}_j, j \in N_i$, can be determined from simple geometry:

$$\mathbf{q}_j = (\mathbf{p}_{ji}, \angle \mathbf{p}_{ij} - \angle \mathbf{p}_{ji} + \pi),$$

where $\mathbf{p}_{ji} = \mathcal{W}\mathbf{p}(v_i) \ominus \mathcal{W}\mathbf{q}(v_j)$, $\mathbf{p}_{ij} = \mathcal{W}\mathbf{p}(v_j) \ominus \mathcal{W}\mathbf{q}(v_i)$, and $\angle \mathbf{p}$ is the bearing angle associated with position vector \mathbf{p} . \square

The following corollary, which descends from the same arguments used in the previous proof, identifies sufficient conditions for the associated framework to be rigid.

Corollary 1. *An instance of Problem 3 is uniquely solvable if the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ of the associated framework admits a directed spanning tree $\mathcal{T} = (\mathcal{V}_T, \mathcal{E}_T)$ such that for any $(v_j, v_k) \in \mathcal{E}_T$ it is both $(v_j, v_k) \in \mathcal{E}$ and $(v_k, v_j) \in \mathcal{E}$. This is certainly true if \mathcal{G} is complete (i.e. all possible measurements are available).*

The importance of rigidity for the unique solvability of Problem 3 is illustrated in Figure 3. In particular, rows 1–3 refer to the case of non-anonymous measurements (Problem 3). In row 1, the graph \mathcal{G} of the associated framework is

	spatial arrangement	measurements	solutions
1. non-anonymous measurements, rigid framework (complete graph)			
2. non-anonymous measurements, rigid framework			
3. non-anonymous measurements, non-rigid framework			
4. anonymous measurements, rigid framework (complete graph)			

Fig. 3. Solvability of Problems 2 and 3. Rows 1–3 show three instances of Problem 3 that share the same spatial arrangement of the robots in terms of positions but, due to the varying orientations, are characterized by different frameworks. In particular, in rows 1 and 2 the framework is rigid, and a unique solution exists in both cases; whereas the framework of row 3 is not rigid, and an infinity of solutions exists. Row 4 refers to Problem 2 and in particular shows the anonymous version of the problem instance in row 1: since the robot position set is rotational symmetric of order 3, unique solvability is lost due to ambiguity.

complete, and therefore the latter is rigid; a unique solution then exists. For the situation in row 2, \mathcal{G} is not complete but the associated framework is still rigid; unique solvability is again guaranteed. The framework associated to the problem instance in row 3 is instead non-rigid, because the connection between \mathcal{R}_i and \mathcal{R}_j is *weak*, and therefore multiple (actually infinite) solutions exist.

Let us now look at Problem 2, where anonymity of the measurements comes into play. To this end, we need a preliminary definition. We say that a set of planar points is *rotational symmetric of order λ* if there exist λ distinct rotations around its centroid that are associated with angles in $(0, 2\pi]$ and map the set into itself. Accordingly, a generic set of points is at least rotational symmetric of order 1 (a 2π rotation will always map the set into itself), whereas actual symmetries are associated to $\lambda \geq 2$.

Proposition 2. *The number of solutions of an instance of Problem 2 is lower bounded by*

$$(\lambda - 1)! (\lambda!)^{\frac{|N_i|+1}{\lambda} - 1} \quad \text{if } \text{centroid}(P) \notin P, \quad (4)$$

$$(\lambda!)^{\frac{|N_i|}{\lambda}} \quad \text{if } \text{centroid}(P) \in P, \quad (5)$$

where λ is the order of rotational symmetry that characterizes the set $P = \{\mathbf{p}_k \mid k \in \{i\} \cup N_i\}$ of the robot cartesian positions.

Proof. Given an instance I_A of Problem 2, denote with I_N its non-anonymous version, and with $(\mathcal{G}, \mathcal{W}, \mathbf{q})$ the framework associated to I_A and I_N (it is the same). Clearly, any solution

of I_N is also a solution of I_A ; denote with σ the number of solutions of I_A . Then consider another instance I'_N of Problem 3 such that the associated framework is $(\mathcal{K}, \mathcal{W}, \mathbf{q})$, where \mathcal{K} is the complete graph built from \mathcal{G} ; let I'_A be the anonymous version of I'_N and σ' the number of solutions of I'_A . Since \mathcal{G} is a subgraph of \mathcal{K} , the input data of I'_A contain all the measurements that are present in the input of I_A ; therefore, it is $\sigma' \leq \sigma$. The proof is concluded using Proposition 4 of (Franchi et al., 2010a), which states that σ' is given by (4)–(5). \square

Note that the above lower bounds on the number of solutions of Problem 2 hold regardless of the rigidity of the graph \mathcal{G} in the associated framework, i.e. of the richness of the available measurements. In particular, both bounds reduce to 1 when $\lambda = 1$, i.e. in the absence of non-trivial rotational symmetries.

Since Problem 3 is a more *informed* version of Problem 2, any solution of the former is a solution of the latter (the reverse is not true). In view of this, we may combine Propositions 1 and 2 in the following statement.

The solution of a given instance of Problem 2 may be not unique for two independent reasons:

1. *Non-rigidity:* The framework associated with the problem is not rigid. In this case, clearly, an infinity of solution exist also for the problem with anonymous measurements.
2. *Ambiguity:* If the set of robot cartesian positions is rotational symmetric of order two or larger, there exists a

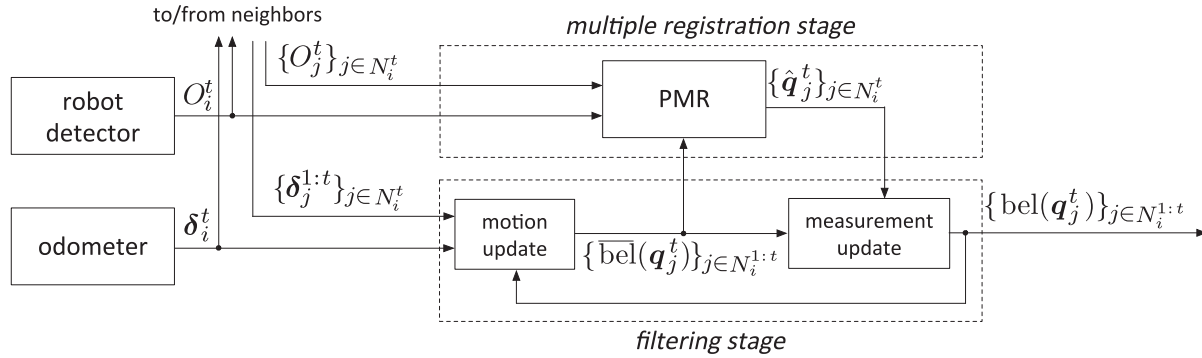


Fig. 4. Scheme of the mutual localization system that runs on the generic robot \mathcal{R}_i . Note the two-stage structure.

finite set of possible solutions whose number grows as a factorial function of the order of symmetry.

Clearly, these situations may occur simultaneously, and their effects cumulate.

We emphasize that ambiguity is a byproduct of the anonymity of the measurements. This is further illustrated in Figure 3, row 4. Even if the robots are in the exact same arrangement of row 1, and the graph \mathcal{G} of the associated framework is still complete, the set of robot cartesian positions is rotational symmetric of order $\lambda = 3$ and therefore, consistently with (4), two distinct solutions arise in the presence of anonymity.

It should be kept in mind that the relevance of the ambiguity phenomenon goes well beyond the deterministic case, in which rotational symmetric situations are a set of zero measure in the space of possible arrangements of the multi-robot system. In fact, when the localization problem is formulated in a stochastic setting to take into account measurement noise, ambiguity arises as soon as when the group is ‘close’ to being in a rotational symmetric arrangement, generating large uncertainties in the probability density of the solution.

Finally, note that rotational symmetry is not the only source of ambiguity for Problem 2. Indeed, if the framework associated with the problem is not complete (be it rigid or not), multiple solutions may exist even for non-symmetric arrangements of the system (Franchi et al., 2010a).

4. Overview of the Proposed Approach

As discussed in the introduction, the main distinctive aspects of Problem 1 with respect to standard versions of multi-robot localization problems found in the literature are the anonymity of the measurements and the presence of false positives/negatives. In principle, the anonymity issue could be addressed by using some form of data association technique. We shall delay the discussion of this classical approach to Section 8, where we will explicitly compare its results to those of our method.

The results of the previous section, and in particular those concerning Problem 2, provide useful guidelines for

designing a solution method for Problem 1, which is the stochastic-dynamic version of the former. In particular:

- To consider all of the possible solutions, which may be multiple in the presence of ambiguity, first enumerate all possible arrangements of the observed robots (including their identities and orientations) that are consistent with the current observations, both endogenous (gathered by \mathcal{R}_i) and hexogeneous (gathered and communicated by its neighbors). Incorporation of other robots’ observations will also allow \mathcal{R}_i to localize robots that it cannot directly detect, either due to the intrinsic limitations (e.g. anisotropy) of its robot detector or to line-of-sight occlusions (false negatives). This geometric step, called *multiple registration*, is also instrumental for rejecting false positives, for which there exists no matching observation.
- Since the number of possible solutions of a multiple registration can be factorial in case of ambiguity, use a probabilistic method such as random sampling consensus (RANSAC) (Fischler and Bolles, 1981) in order to keep the execution time under control.
- Make the multiple registration algorithm aware of the current belief (hence, of past history) to help discriminate among statically ambiguous situations.
- Using observations taken over a time interval may allow to extract a solution even when the multiple registration problem admits an infinite number of solutions due to non-rigidity of the associated framework. This leads to filtering as a way to incorporate successive measurements and odometric data. Since the measurements are uncertain, this is also instrumental in reducing the effect of noise.

This results in a two-phase approach to the problem solution, clearly reflected in the scheme of our mutual localization system shown in Figure 4.

In particular, robot \mathcal{R}_i runs a *probabilistic multiple registration* (PMR) algorithm to compute feasible hypotheses on the identities and relative configurations of the robots belonging to N_i^t , on the basis of the current observations O_i^t , $\{O_j^t\}_{j \in N_i^t}$, and of the current beliefs about $\{q_j^t\}_{j \in N_i^t}$. The

relative configurations thus obtained, together with the odometric measurements δ_i^t and $\{\delta_j^t\}_{j \in N_i^t}$, are used by $|N_i^{1:t}|$ particle filters to update the belief about the configuration of each robot in $N_i^{1:t}$. Note the feedback mechanism that brings the motion-updated belief of the particle filters in input to the multiple registration stage, allowing the latter to prune unreasonable solutions in ambiguous situations and thus reducing its computational load.

The main features of our two-stage approach are:

1. the mutual exclusive structure of the set of measurements is exploited in the registration phase;
2. the increased dimension of the measurements provided to the particle filters (the relative orientation is added by PMR) improves the convergence of the estimation process;
3. the simplicity of the probabilistic model of the robot detector used in the particle filter, which does not have to take into account the identity of the robots;
4. the extension of the perception capabilities of the system members beyond those of the individual robots, without any kind of centralized elaboration.

The multiple registration algorithm and the particle filters are described in detail in the next two sections.

5. Multiple registration stage

At each time instant t , the generic robot \mathcal{R}_i runs the PMR algorithm, that represents the part of the localization system which provides the input to the measurement update in the particle filters (see Figure 4).

In computer vision, *registration* between two images entails computing the roto-translation between the viewpoints from which the images were gathered. In our case, the ‘images’ to be registered are the observations gathered by the neighbors of \mathcal{R}_i , and the different viewpoints are their configurations. Examples of observations with the corresponding frameworks are shown in the ‘measurements’ column of Figure 3.

This section is organized as follows. In Section 5.1 we discuss binary registration and in Section 5.2 we introduce the concept of irreconcilable matchings. Both are prerequisites for describing our multiple registration algorithm PMR in Section 5.3. Examples of applications of PMR are given in Section 5.4, while the algorithm complexity is analyzed in Section 5.5.

5.1. Binary registration

Consider an observation coming from the robot detector of \mathcal{R}_i , i.e. a set of points corresponding to the positions of some neighboring robots (or objects that look like robots) expressed in \mathcal{F}_i (see Figure 2). Since the identities of the detected robots are unknown, let us conventionally assign a zero label to all of the points. Also, we explicitly add to each observation the origin, i.e. the position of \mathcal{R}_i , and assign the i label to it.




Fig. 5. An example of matching. (a) Two sets of labeled points, representing augmented observations by robots \mathcal{R}_i and \mathcal{R}_j ; the sets are assumed to be partially labeled as a result of previous matchings (zero labels not shown). (b) A roto-translation that using Equation (6) results in three pairs (circled) of associated points. (c) The associated pairs shown in the two given sets of labeled points. Since the binary relation defined by these pairs satisfies the conditions of Definition 1, the roto-translation in (b) is actually a matching.

To generalize this concept, define a *set of labeled points* as a pair $\mathcal{L} = (P, l)$ where P is a set of points and $l : P \rightarrow \{0, 1, \dots, n\}$ is an injective labeling map. The label $l(p)$ is the index of the robot to which point p is associated, with $l(p) = 0$ indicating that p is not yet associated with any robot. We may then say that the previous processing transforms an observation into a particular set of labeled points, in which all points have label zero but one.

Consider now two sets of labeled points $\mathcal{L}' = (P', l')$ and $\mathcal{L}'' = (P'', l'')$, representing two ‘views’ of the same scene by different robots. In particular, each view may be either an ‘original’ observation (only one point has a non-zero label), or an ‘augmented’ observation incorporating the identities of some robots as the result of previous computations (multiple points have non-zero labels). Loosely speaking, our objective is to determine the roto-translation between the two robots that produces the best match between the two views.

Denote by q a generic roto-translation acting on \mathcal{L}'' (more precisely, on P''). Under q , a point p' of \mathcal{L}' is associated with a point p'' of \mathcal{L}'' if

$$\|p' - q \oplus p''\| \leq \eta, \quad (6)$$

where $\eta > 0$ is the distance threshold under which the two points are identified as the same. Denote by $B_q \subset P' \times P''$

Algorithm 1: Binary registration.

input : sets of labeled points \mathcal{L}' and \mathcal{L}''
output: multiple matchings, each matching M consisting of a corrected roto-translation \mathbf{q}_M and a merged set of labeled points \mathcal{L}_M

- 1 run a matching algorithm on \mathcal{L}' and \mathcal{L}'' ;
- 2 **foreach** returned roto-translation \mathbf{q} **do**
- 3 correct the roto-translation \mathbf{q} to \mathbf{q}_M ;
- 4 merge \mathcal{L}' and \mathcal{L}'' into a new set of labeled points \mathcal{L}_M ;

the binary relation (i.e. the collection of pairs of associated points) resulting from the application of rule (6).

Definition 1 (Matching between sets of labeled points). *Given two sets of labeled points $\mathcal{L}' = (P', l')$ and $\mathcal{L}'' = (P'', l'')$, a roto-translation \mathbf{q} acting on \mathcal{L}'' is a matching between \mathcal{L}' and \mathcal{L}'' if the corresponding relation $B_{\mathbf{q}} \subset P' \times P''$ is such that:*

1. each point in P' is associated to at most one point in P'' , and vice versa;
2. points that have different non-zero labels are never associated.

The second condition guarantees that two points that have been assigned to different robots are not associated in $B_{\mathbf{q}}$. An example of matching is shown in Figure 5.

The quality of a matching can be measured by the number of points it associates. We assume an algorithm is available that can compute all matchings above a certain quality.

Definition 2 (Matching algorithm). *A matching algorithm is any procedure that, given two sets of labeled points \mathcal{L}' and \mathcal{L}'' , returns all of the matchings \mathbf{q} for which the cardinality of $B_{\mathbf{q}}$ (i.e. the number of pairs of associated points in \mathcal{L}' and \mathcal{L}'') is larger or equal to a specified threshold w .*

A matching algorithm will return each matching as a roto-translation with the corresponding binary relation. The number of returned matchings clearly depends on the cardinality threshold w : a lower threshold will typically produce more matchings. Even when the threshold increases, however, multiple matchings may be returned, e.g., in the presence of rotational symmetry (see Section 3).

With the above definitions, the basic structure of our binary registration algorithm is given in Algorithm 1. A more detailed discussion of each step follows.

Matching (line 1): For this step, we use the technique introduced by Franchi et al. (2009b), which is based on RANSAC (Fischler and Bolles, 1981). Since our matching problem is closely related³ to *point-pattern matching*, several other methods available in the literature can be used directly or with little adaptation, e.g. those by Thrun and Liu (2005) and Cunningham et al. (2012). In any case, the

choice of the matching algorithm is essentially an implementation issue that does not affect the general working principle of the PMR algorithm to be described.

Correction (line 3): This consists in replacing each roto-translation \mathbf{q} with

$$\mathbf{q}_M = \operatorname{argmin}_{\mathbf{q} \in SE(2)} \sum_{(p', p'') \in B_{\mathbf{q}}} \|\mathbf{p}' - \mathbf{q} \oplus \mathbf{p}''\|^2. \quad (7)$$

This means that, independently of how \mathbf{q} was computed, it is corrected by choosing the roto-translation that minimizes the squared error over all of the pairs associated by $B_{\mathbf{q}}$. The minimization problem (7) can be solved by standard methods; in particular, in our implementation we have used the technique proposed by Umeyama (1991).

Merging (line 4): A merged set of labeled points $\mathcal{L}_M = (P_M, l_M)$ is generated as follows:

1. For each pair $(\mathbf{p}' \in P', \mathbf{p}'' \in P'')$ contained in $B_{\mathbf{q}}$, compute the ‘averaged’ point

$$\mathbf{p} = \frac{\kappa}{\kappa + 1} \mathbf{p}' + \frac{1}{\kappa + 1} (\mathbf{q}_M \oplus \mathbf{p}''), \quad (8)$$

and put it in P_M . In this formula, κ is the *generation counter* of \mathbf{p}' ; i.e. $\kappa = 1$ if \mathbf{p}' is an observed point, $\kappa = 2$ if \mathbf{p}' is the result of a previous binary registration between two observed points, and so on. The structure of Equation (8) guarantees that all of the points that have contributed to a certain average, either in the current or in previous binary registrations, are given the same weight.

2. Put in P_M each point $\mathbf{p}' \in P'$ that has not been associated. For each point $\mathbf{p}'' \in P''$ that has not been associated, put in P_M the point $\mathbf{q}_M \oplus \mathbf{p}''$.
3. For averaged points, define the labeling map l_M by letting

$$l_M(\mathbf{p}) = \max\{l(\mathbf{p}'), l(\mathbf{p}'')\}. \quad (9)$$

For the other points in P_M , which come from P' or P'' , the label is simply inherited from l' or l'' .

Note that Equation (9) implies that an averaged point will be unlabeled if it comes from two unlabeled points, and will get the non-zero label if one of the two points is actually labeled. Remember that the two points \mathbf{p}' and \mathbf{p}'' cannot have different non-zero labels due to the definition of matching.

Wrapping up, a binary registration produces multiple pairs of the form $M = (\mathbf{q}_M, \mathcal{L}_M)$, which actually consist of a (corrected) roto-translation \mathbf{q}_M and the corresponding (merged) set of labeled points \mathcal{L}_M . Generalizing the concept introduced in Definition 1, in the following we call matching any pair M that comes out of a binary registration.




Fig. 6. From left to right: three matchings, M_1 , M_2 and M_3 , respectively, coming from different binary registrations. Here M_1 and M_2 are irreconcilable because different points are labeled as \mathcal{R}_j , whereas M_1 and M_3 are irreconcilable because the same point is labeled once as \mathcal{R}_h and once as \mathcal{R}_k .

5.2. Irreconcilable matchings

Our multiple registration algorithm, to be discussed in Section 5.3, builds solutions (matchings) by performing recursive binary registrations on partial matchings. To reduce its running time, it is important to avoid the propagation of partial matchings that will eventually generate to the same solution. To this end, we introduce the following definition.

Definition 3 (Irreconcilability). *Consider two matchings M_1 (between \mathcal{L}' and \mathcal{L}'') and M_2 (between \mathcal{L}' and \mathcal{L}'''), where it may be $\mathcal{L}''' = \mathcal{L}''$, and their merged sets of labeled points \mathcal{L}_{M_1} and \mathcal{L}_{M_2} . M_1 and M_2 are said to be irreconcilable if one of the following holds for at least one pair of points ($\mathbf{p}_1 \in P_{M_1}, \mathbf{p}_2 \in P_{M_2}$):*

1. $\mathbf{p}_1, \mathbf{p}_2$ are the same point (in the sense of (6)) but have different labels;
2. $\mathbf{p}_1, \mathbf{p}_2$ are different points but have the same non-zero label.

Figure 6 illustrates the two sources of irreconcilability.

Irreconcilable matchings will always generate different solutions. In fact, assume we have two irreconcilable matchings M_1 (between \mathcal{L}' and \mathcal{L}'') and M_2 (between \mathcal{L}' and \mathcal{L}'''), and their merged sets of labeled points \mathcal{L}_{M_1} and \mathcal{L}_{M_2} . Consider now another set of labeled points \mathcal{L}''' . Then, all of the matchings produced by binary registration of \mathcal{L}_{M_1} with \mathcal{L}''' will be different from the matchings produced by binary registration of \mathcal{L}_{M_2} with \mathcal{L}''' .

5.3. Probabilistic multiple registration

We may now present PMR, our multiple registration algorithm. The inputs to PMR are the observations available to \mathcal{R}_i at time t , i.e. $O_i^t, \{O_j^t\}_{j \in N_i^t}$, and the motion-updated beliefs on $\{q_j^t\}_{j \in N_i^{1:t}}$ coming from the filtering stage (see Figure 4). The outputs are estimates of the relative configurations $q_j^t, j \in N_i^t$.

A pseudocode description of PMR is shown in Algorithm 2. A detailed step-by-step illustration is given below.

Algorithm 2: Probabilistic Multiple Registration (PMR).

input : observations $O_i^t, \{O_j^t\}_{j \in N_i^t}$, beliefs $\overline{\text{bel}}\{q_j\}_{j \in N_i^{1:t}}$
output: one set of estimated relative configurations
 $Q = \{\hat{q}_j^t\}_{j \in N_i^t}$ for each branch of the algorithm

```

1 generate  $\mathcal{L}_i^t, \{\mathcal{L}_j^t\}_{j \in N_i^t}$  from  $O_i^t, \{O_j^t\}_{j \in N_i^t}$ ;
2  $\mathcal{L} = \mathcal{L}_i^t, Q = \emptyset$ , and  $L = \{\mathcal{L}_j^t\}_{j \in N_i^t}$ ;
3 for  $k = 1$  to  $|N_i^t|$  do
4   perform a binary registration of  $\mathcal{L}$  with each set in  $L$  and
   store the resulting matchings in  $\mathbf{M}$ ;
5   reduce  $\mathbf{M}$  to a maximal subset of irreconcilable
   matchings;
6   remove from  $\mathbf{M}$  all matchings whose fitness is below
 $\gamma \cdot b^*$ , where  $b^*$  is the highest fitness in  $\mathbf{M}$ ;
7   if  $|\mathbf{M}| > 0$  then
8     foreach  $M \in \mathbf{M}$  do
9       create a new branch;
10      remove from  $L$  the set of labeled points that  $M$ 
      matches with  $\mathcal{L}$ ;
11       $\mathcal{L} = \mathcal{L}_M$ ;
12      add  $q_M$  to  $Q$ ;
13   else
14     break

```

Initial sets of labeled points (line 1): As explained at the beginning of Section 5.1, we first generate from the available observations the sets of labeled points $\mathcal{L}_i^t = (P_i^t, l_i^t)$ and $\mathcal{L}_j^t = (P_j^t, l_j^t), j \in N_i^t$ by letting

$$P_i^t = O_i^t \cup \{(0, 0)\}, \quad l_i^t(\mathbf{p}) = \begin{cases} 0 & \text{if } \mathbf{p} \in O_i^t, \\ i & \text{if } \mathbf{p} = (0, 0), \end{cases} \quad (10)$$

$$P_j^t = O_j^t \cup \{(0, 0)\}, \quad l_j^t(\mathbf{p}) = \begin{cases} 0 & \text{if } \mathbf{p} \in O_j^t, \\ j & \text{if } \mathbf{p} = (0, 0). \end{cases} \quad (11)$$

Initialization (line 2): The execution of PMR requires at most $|N_i^t|$ iterations. In each iteration, the state of the algorithm includes the current sets \mathcal{L} of labeled points and Q of relative configurations, plus a collection L of labeled sets that have not been registered in \mathcal{L} . In this initialization step, \mathcal{L} is set to \mathcal{L}_i^t , hence $L = \{\mathcal{L}_j^t\}_{j \in N_i^t}$ (no registration has been performed yet) and Q is set to the empty set.

Binary registrations (line 4): In the k th iteration, the algorithm performs $|N_i^t| - k + 1$ binary registrations between \mathcal{L} and the elements of L . The matchings resulting from these binary registrations are stored in a set \mathbf{M} .

Irreconcilability-based pruning (line 5): \mathbf{M} is reduced by retaining only a maximal subset of irreconcilable matchings.

Belief-based pruning (line 6): A second pruning is performed to eliminate matchings that are inconsistent with the currently available belief. In particular, consider a matching

$M \in \mathbf{M}$ between \mathcal{L} and \mathcal{L}_j^t , in which the roto-translation \mathbf{q}_M represents the estimated relative configuration of \mathcal{R}_j w.r.t. \mathcal{R}_i . Here M is rated using the following fitness function:

$$b_M = \int p(\mathbf{q}_M | \mathbf{q}_j^t) \overline{\text{bel}}(\mathbf{q}_j^t) d\mathbf{q}_j^t, \quad (12)$$

where $p(\mathbf{q}_M | \mathbf{q}_j^t)$ is approximated as a Gaussian distribution whose covariance reflects the uncertainty of the multiple registration process, and $\overline{\text{bel}}(\mathbf{q}_j^t)$ is the motion-updated belief on \mathbf{q}_j^t coming from the filtering stage (see Section 6). All matchings whose fitness is below the threshold $\gamma \cdot b^*$, where $b^* = \max_{M \in \mathbf{M}} b_M$ and $0 < \gamma < 1$, are removed from \mathbf{M} .

Branching (lines 7–12): For each matching M left in \mathbf{M} , a new branch of the algorithm is created, first deleting the corresponding registered set of labeled points from \mathbf{L} , then setting the current set \mathcal{L} of labeled points to \mathcal{L}_M for the next iteration, and finally adding the configuration \mathbf{q}_M to \mathcal{Q} .

Termination (line 14): If $k = |N_i^t|$ and $|\mathbf{M}| = 0$, the current branch has built a solution in which all of the given sets of labeled points have been registered (i.e. \mathbf{L} is empty). However, if the set \mathbf{M} of matchings becomes empty at a certain iteration with $k < |N_i^t|$ (i.e. without \mathbf{L} being empty) the execution of the branch is interrupted and the current set of estimates \mathcal{Q} , which does not contain the configurations of robots whose observations have not been registered, is added to the algorithm outputs (one for each branch).

Note the following:

- Irreconcilability-based pruning (line 5) is aimed at reducing the computational load of PMR by pruning branches that will sooner or later lead to the same solution. No solution is lost by doing this.
- Belief-based pruning (line 6) is particularly useful to discriminate statically ambiguous situations (where the robots are in a rotational symmetric arrangement) by discarding labelings (i.e. solutions) that are inconsistent with the current belief.
- PMR will find and provide partial solutions in situations where a complete solution cannot be found (line 14 with \mathbf{M} empty); in particular, if the framework underlying the given instance of the problem is non-rigid PMR will only return configuration estimates for robots belonging to a rigid component.
- The solution(s) computed by PMR can be further refined by performing a final batch least-squares (BLS) minimization. In fact, each solution entails a particular data association, and one may then consider the problem of adjusting the robot configuration estimates so as to *simultaneously* take into account all measurements that are included in that data association. The solution of this nonlinear least-squares problem, that can be found using numerical methods, would be more accurate than the solution produced by successive pairwise minimizations such as (7).

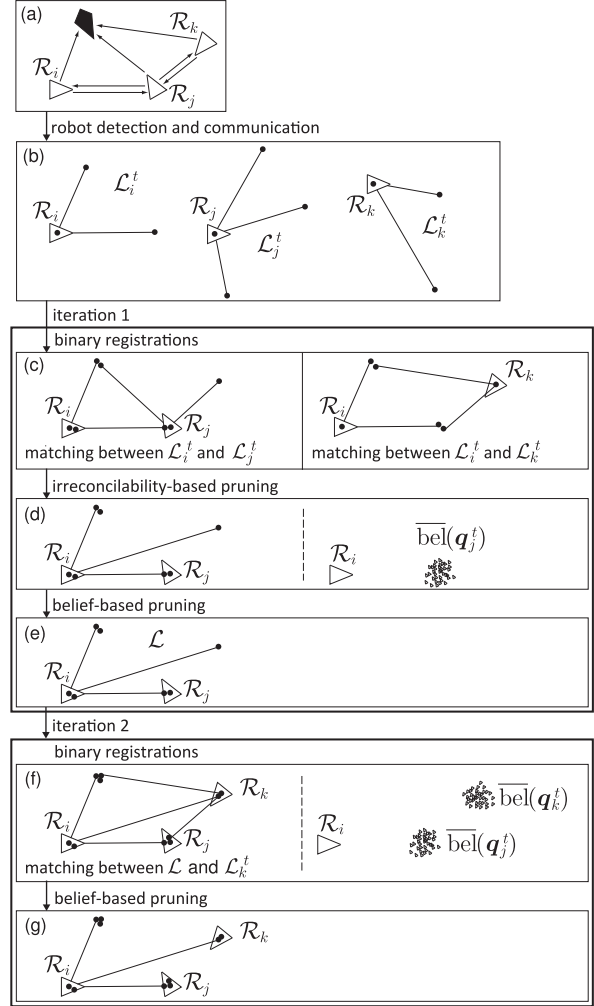


Fig. 7. Execution of PMR in a non-ambiguous situation with three robots and a deceiving obstacle: (a) spatial arrangement and measurements; (b) initial sets of labeled points corresponding to the observations; (c) matchings produced by binary registration of $\mathcal{L}_j^t, \mathcal{L}_k^t$, respectively, with \mathcal{L}_i^t ; (d) matching retained after irreconcilability-based pruning; (e) matching retained after belief-based pruning; (f) matching produced by binary registration of \mathcal{L}_i^t with \mathcal{L} ; (g) final solution.

5.4. Illustrative examples

To clarify how PMR works, we work out two simple examples in the following.

Figure 7 describes the execution of PMR by \mathcal{R}_i in a non-ambiguous situation involving three robots ($\mathcal{R}_i, \mathcal{R}_j, \mathcal{R}_k$) and an obstacle which is erroneously detected as a robot (Figure 7(a)). The objective of the algorithm is to perform a registration of the sets \mathcal{L}_j^t and \mathcal{L}_k^t , obtained from the communicated observations \mathcal{O}_j^t and \mathcal{O}_k^t and stored in \mathbf{L} during the initialization phase, with the set \mathcal{L}_i^t obtained from the observation \mathcal{O}_i^t (Figure 7(b)).

During the first iteration, PMR performs $|N_i^t| = 2$ binary registrations of $\mathcal{L}_j^t, \mathcal{L}_k^t$, respectively, with \mathcal{L}_i^t , obtaining two matchings that are *not* irreconcilable (Figure 7(c)). After

pruning, only one of them (the matching between \mathcal{L}_i^t and \mathcal{L}_j^t) is retained. This matching also survives belief-based pruning (Figure 7(d)). At this point, \mathcal{L}_j^t is deleted from \mathbf{L} and \mathcal{L} is accordingly updated (Figure 7(e)).

In the second iteration, registration with \mathcal{L}_k^t produces a better matching than before (Figure 7(f)), with three associated pairs instead of the previous two. Having passed belief-based pruning (Figure 7(g)), this matching represents the final solution. In particular, the algorithm returns the estimated configurations stored in Q , i.e. \hat{q}_j^t and \hat{q}_k^t .

This example well illustrates the incremental nature of PMR, and in particular how the iterative registration mechanism allows to localize robots that are not directly detected by \mathcal{R}_i . Note also that the deceiving obstacle is automatically ‘rejected’ because no communicated observations fits its location. However, it is useful because it acts as a landmark improving the quality of the matching and, ultimately, the localization accuracy.

To appreciate how PMR handles ambiguous situations, consider instead the example of Figure 8, in which the spatial arrangement of the robots is rotational symmetric (Figure 8(a)). As before, the objective of the algorithm is to perform a registration of the sets \mathcal{L}_j^t and \mathcal{L}_k^t , obtained from the communicated observations O_j^t and O_k^t and stored in \mathbf{L} during the initialization phase, with the set \mathcal{L}_i^t obtained from the observation O_i^t (Figure 8(b)).

During the first iteration, PMR performs $|N_i^t| = 2$ binary registrations of \mathcal{L}_j^t , \mathcal{L}_k^t , respectively, with \mathcal{L}_i^t , obtaining four matchings (Figure 7(c)). A maximal subset of two irreconcilable matchings (Figure 8(d)) is then retained. These are then rated on the basis of their fitness w.r.t. the current belief; in this case, the current belief indicates that only one⁴ of the two matchings can be accepted (Figure 7(e)). At this point, \mathcal{L}_j^t is deleted from \mathbf{L} and \mathcal{L} is accordingly updated.

In the second iteration, registration of \mathcal{L}_k^t with \mathcal{L} produces (as in the previous case) a better matching than with \mathcal{L}_i^t only (Figure 7(f)), with three associated pairs instead of the previous two. After belief-based pruning (Figure 7(g)), this matching survives as the final solution and the outputs of PMR are the estimated configurations stored in Q , i.e. \hat{q}_j^t and \hat{q}_k^t .

5.5. Analysis

The run time of PMR, which accounts for most of the cycle time of our mutual localization system, depends on the number $|N_i^t| \leq n$ of observations that \mathcal{R}_i receives (n is the unknown number of robots).

In normal operation (no ambiguity, or ambiguities that can be resolved by belief-based pruning), PMR expands a single branch, which requires $(|N_i^t| - 1)(|N_i^t| - 2)/2$ binary registrations to produce a solution: in turn, each binary registration requires constant time if performed with the technique described by Franchi et al. (2009b). This leads to a

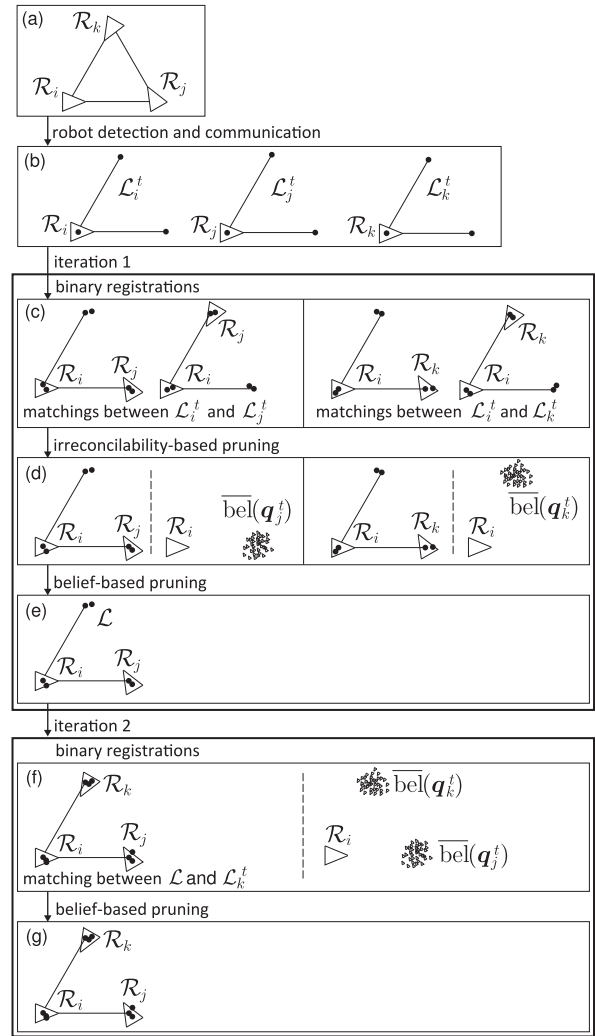


Fig. 8. Execution of PMR in an ambiguous situation (rotational symmetric arrangement) with three robots: (a) spatial arrangement and measurements; (b) initial sets of labeled points corresponding to the observations; (c) matchings produced by binary registration of \mathcal{L}_j^t , \mathcal{L}_k^t , respectively, with \mathcal{L}_i^t ; (d) matchings retained after irreconcilability-based pruning; (e) matching retained after belief-based pruning; (f) matching produced by binary registration of \mathcal{L}_k^t with \mathcal{L} ; (g) final solution.

worst-case complexity $O(n^2)$, while the average-case complexity can be significantly lower if the number of neighbors (robots communicating with \mathcal{R}_i) is smaller than n .

In the presence of ambiguities that cannot be resolved, the algorithm expands a branch for each solution, with the complexity of each branch being that of normal operation. In principle, however, the number of solutions can grow quite large. In particular, it is easy to prove that both the lower bounds of Proposition 2 have a maximum value of $N_i!$ (set $\lambda = N_i + 1$ if the centroid is occupied by a robot, $\lambda = N_i$ otherwise).

One way to keep the complexity polynomial in spite of the presence of ambiguities would be to modify PMR so as to perform at each t a random partial (rather than complete)

exploration of the space of possible solutions. In any case, especially in the early stages of localization, it is advisable to break rotational symmetry whenever possible, e.g., by using an anti-symmetry control law (Franchi et al., 2010a).

6. Filtering stage

In the proposed mutual localization system, the multiple registration stage represented by PMR is followed by a filtering stage that takes into account the odometry measurements to reduce the noise, prune ambiguous solutions based on the current certainty and finally compute a belief on \mathbf{q}_j^t , for $j \in N_i^{1:t}$. The multi-hypothesis nature of the problem suggests the use of particle filters, since they are intrinsically multi-modal.

In principle, \mathcal{R}_i should maintain a single particle filter to compute a joint belief $\text{bel}(\{\mathbf{q}_j^t\}_{j \in N_i^{1:t}})$. However, this is not feasible from a computational point of view, since the dimension of the filter would grow linearly with the number of robots, and the number of particles required to represent its state would grow exponentially. To overcome this problem, \mathcal{R}_i maintains one particle filter for each \mathcal{R}_j , computing a separate belief $\text{bel}(\mathbf{q}_j^t)$, $j \in N_i^{1:t}$, instead of a single joint belief. This choice, already implicit in the formulation of Problem 1, relies on the independence assumption

$$p(\{\mathbf{q}_j^t\}_{j \in N_i^{1:t}}) = \prod_{j \in N_i^{1:t}} p(\mathbf{q}_j^t),$$

which is acceptable in a pure localization scenario. A drawback of this choice is the loss of the explicit mutual exclusion guaranteed by the multiple registration procedure (i.e. the fact that a measurement is associated to a single robot). However, the mutual exclusion implicitly stands, since it is already included in the measurements feeding the filter, which are the outputs of PMR.

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

1. the motion displacements $\delta_j^{(\tau_j+1):t}$ (sent by \mathcal{R}_j), where τ_j is the last time instant before t in which \mathcal{R}_j communicated with \mathcal{R}_i ;
2. the set of all of the relative configuration estimates $\{\hat{\mathbf{q}}_j^t\}$ computed by PMR for \mathcal{R}_j .

The filter equations are very similar to those of Howard et al. (2003). In particular, the update rule that accounts for the motion of \mathcal{R}_i is

$$p(\mathbf{q}_j^t | \delta_i^t) = v_i \int p(\mathbf{d} | \delta_i^t) \text{bel}(\mathbf{q}_j^{t-1}) |_{\mathbf{q}_j^{t-1} = \mathbf{q}_j^t \oplus \mathbf{d}} d\mathbf{d}, \quad (13)$$

where v_i is a normalization factor and $p(\mathbf{d} | \delta_i^t)$ is the model of the odometer. Equation (13) leads to the following update for the single particle of value $\tilde{\mathbf{q}}_j$:

$$\tilde{\mathbf{q}}_j = \tilde{\mathbf{q}}_j \ominus (\delta_i^t \oplus v_i \delta), \quad (14)$$

where n_δ is a sample taken from $p(\mathbf{d} | \delta_i^t)$. Similarly, the update rule that accounts for the motion of \mathcal{R}_j is

$$p(\mathbf{q}_j^t | \delta_j^t) = v_j \int p(\mathbf{d} | \delta_j^t) \text{bel}(\mathbf{q}_j^{t-1}) |_{\mathbf{q}_j^{t-1} = \mathbf{q}_j^t \oplus \mathbf{d}} d\mathbf{d}, \quad (15)$$

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

$$\tilde{\mathbf{q}}_j = \tilde{\mathbf{q}}_j \oplus (\delta_j^t \oplus n_\delta). \quad (16)$$

The consecutive application of Equations (13) and (15) allows the computation of the prediction

$$\overline{\text{bel}}(\mathbf{q}_j^t) = p(\mathbf{q}_j^t | \delta_i^t, \delta_j^t). \quad (17)$$

After the update, the motion of \mathcal{R}_i and \mathcal{R}_j causes a translation of $p(\mathbf{q}_j)$, while the additive noise introduces a blur, i.e. a decrease in sharpness of the probability density by dispersing the particles.

In normal operation, $\tau_j = t - 1$ and \mathcal{R}_i uses δ_j^t as motion measurement for the motion update of the robot \mathcal{R}_j . However, if \mathcal{R}_i and \mathcal{R}_j have not communicated over a non-zero time interval $[\tau_j + 1, \dots, t - 1]$, the motion update of \mathcal{R}_j has not been performed since τ_j . When, at t , communication is resumed, \mathcal{R}_i uses the whole set $\delta_j^{(\tau_j+1):t}$ as motion measurement to perform a cumulative motion update.

As for the measurement update, assume first that the set of relative configuration estimates $\{\hat{\mathbf{q}}_j^t\}$ produced by PMR is composed by a single element. Then, using Bayes law, the measurement update is given by

$$\text{bel}(\mathbf{q}_j^t) = p(\mathbf{q}_j^t | \hat{\mathbf{q}}_j^t) = v'_j p(\hat{\mathbf{q}}_j^t | \mathbf{q}_j^t) \overline{\text{bel}}(\mathbf{q}_j^t), \quad (18)$$

where v'_j is another normalization factor, and the probabilistic model $p(\hat{\mathbf{q}}_j^t | \mathbf{q}_j^t)$ is the same appearing in Equation (12). When the set $\{\hat{\mathbf{q}}_j^t\}$ consists of more than one element, each of them is equally probable, so that using the theorem of total probability we obtain the update equation:

$$\text{bel}(\mathbf{q}_j^t) = p(\mathbf{q}_j^t | \hat{\mathbf{q}}_j^t) = v'_j \sum_{\hat{\mathbf{q}}_j^t \in \{\hat{\mathbf{q}}_j^t\}} p(\hat{\mathbf{q}}_j^t | \mathbf{q}_j^t) \overline{\text{bel}}(\mathbf{q}_j^t), \quad (19)$$

with v'_j different from that in (18).

A number of standard practical techniques have been used to improve the performance of the filter. For example, the initial prior distribution for \mathcal{R}_j is generated using the first configuration estimate(s) $\{\hat{\mathbf{q}}_j\}$ computed by PMR. Moreover, at each step a small percentage of particles are re-initialized using directly the new measurements; for example, this enables the localization system to handle the *kidnapped robot* situation (see the last experiment in the next section).

7. Experimental validation

Before presenting experimental results, we provide a description of our implementation setup.

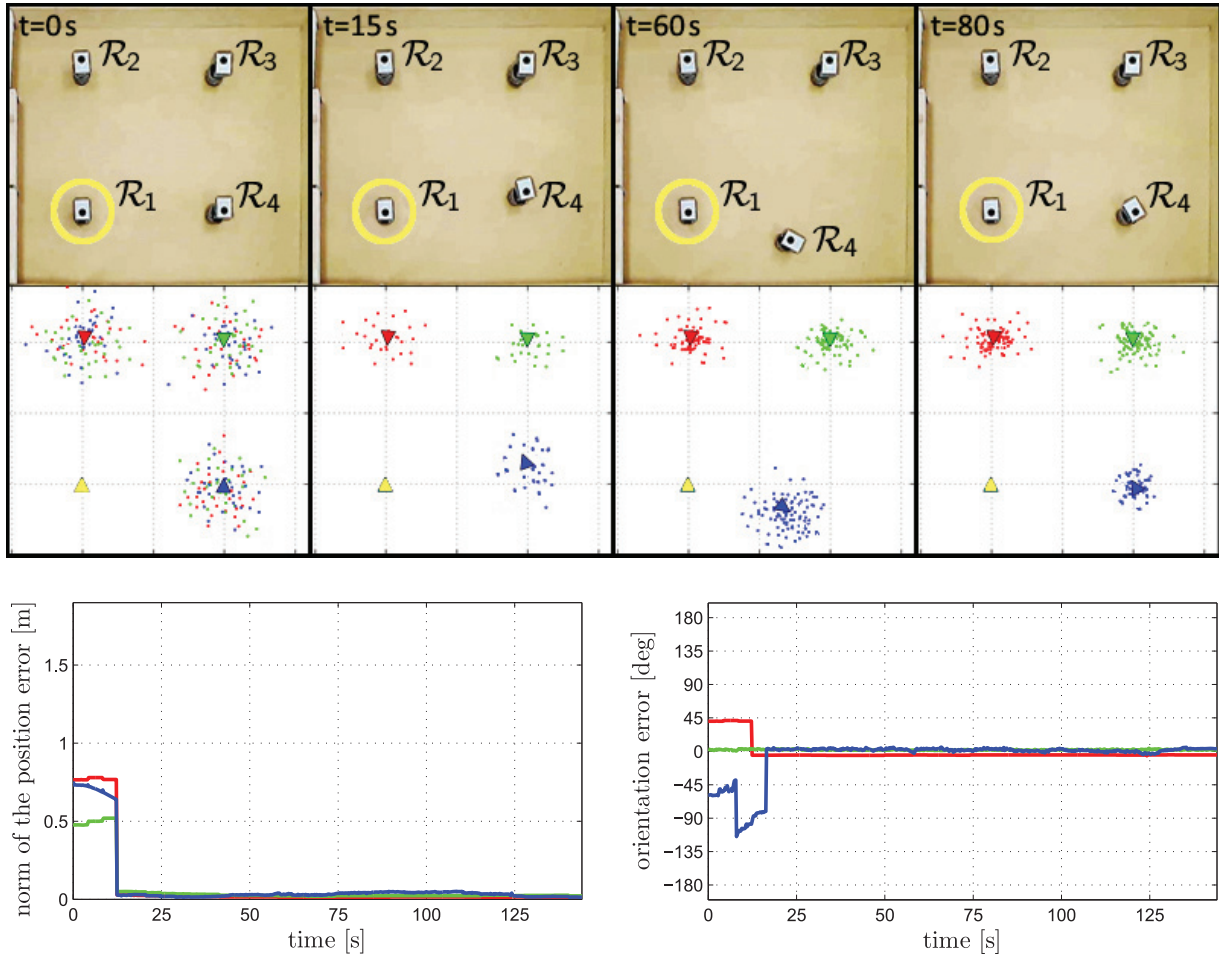


Fig. 9. First experiment. Top: snapshots of the scene and sample distributions computed for \mathcal{R}_2 (red), \mathcal{R}_3 (green) and \mathcal{R}_4 (blue) by the mutual localization system of \mathcal{R}_1 (yellow, circled); the small triangles represent the ground truth. Bottom: norm of the position error and orientation error affecting the configuration estimates of \mathcal{R}_2 , \mathcal{R}_3 and \mathcal{R}_4 .

7.1. Hardware/software setup

The proposed mutual localization algorithm has been tested on a team of four differential-drive Khepera III robots moving in a $3.6 \text{ m} \times 1.9 \text{ m}$ wide arena. In all of the experiments, mutual localization is performed from the perspective of robot \mathcal{R}_1 ($i = 1$). To obtain a ground truth against which to evaluate the accuracy of the proposed localization algorithm, we have used an external three-camera system which is able to track the robots with errors on the position and orientation that are less than 0.01 m and 5° , respectively.

Motion control, sensor data management, and localization have been implemented using MIP, a in-house developed software platform specifically aimed at multi-robot systems.⁵ The source code of the localization algorithm is also available in the Mutual Localization with Anonymous Measurements (MLAM) ROS stack.⁶

Each Khepera III robot has been equipped with a Hokuyo URG-04LX laser sensor that provides bearing-plus-range (i.e. position) measurements at 10 Hz within a 240° field

of view (thus leaving a 120° blind zone behind the robot). The robot detector is then implemented as a memoryless feature extraction algorithm that inspects the laser scan, looking for the indentations caused by a white cardboard box mounted on top of each robot. All cardboard boxes have the same shape and size, making impossible to retrieve the identity of the detected robot on the basis of the profile of the indentation. Each robot measures its own displacement by odometric integration over a 0.1 s time interval. Simultaneous calibration of odometric and sensor parameters was performed using the algorithm by Censi et al. (2013). Communication introduces no significant delay.

The binary registration used in PMR is a RANSAC-based algorithm (Franchi et al., 2009b) with a maximum of 100 trials and a tolerance $\eta = 0.06 \text{ m}$ for point association. Each particle filter is run with 300 particles. The whole mutual localization algorithm by robot \mathcal{R}_1 (PMR + particle filters) runs at 10 Hz on an entry-level GNU-Linux computer, proving that it can be easily executed in real time in a practical applicative scenario.

7.2. Experimental results

Figure 9 summarizes the results of the first experiment. The top part of the figure shows four snapshots of the scene at successive instants of time with the corresponding relative configuration estimates computed by \mathcal{R}_1 in the form of clouds of particles. The orientation of each robot, which is actually included in the estimates, is not shown by the particle. The bottom part of the figure shows the norm of the position error (left) and the orientation error (right), defined respectively as $\|\mathbf{p}_j^t - \bar{\mathbf{p}}_j^t\|$ and $\theta_j^t - \bar{\theta}_j^t$, for $j = 2, 3, 4$, with $\mathbf{q}_j^t = (\mathbf{p}_j^t, \theta_j^t)$ the ground truth and $\bar{\mathbf{q}}_j^t = (\bar{\mathbf{p}}_j^t, \bar{\theta}_j^t)$ the weighted average of the particle distribution for the robot \mathcal{R}_j .

The four robots start in a square spatial arrangement, which is rotational symmetric and hence ambiguous (first snapshot). Only one robot (\mathcal{R}_4) moves over a closed path (second and third snapshots) to conclude the experiment in the starting configuration (fourth snapshot). At the beginning, PMR finds $3! = 6$ solutions (as predicted by Equation (4) with $\lambda = 4$ and $N_i = 3$) which it is not able to prune due to the absence of a priori discriminating knowledge. The result is that the particle filters are initialized at three different configurations for each robot, so that the initial clouds of particles are multi-modal and overlapping. Accordingly, the errors are quite large in the first part of the experiment. When the symmetry is broken ($t \simeq 10$ s) PMR is able to extract a single solution, allowing in turn the filters to disambiguate the estimates; as a consequence, the clouds of particles separate and the errors quickly decrease. In the rest of the experiment the localization system uses new measurements to maintain the accuracy of the estimates. The particle distributions clearly indicate the progressive reduction of uncertainty. When the experiment ends, the spatial arrangement is again rotational symmetric, but PMR is now able to extract a single solution thanks to the use of the current belief.

The objective of the second experiment, shown in Figure 10, is to test the robustness of the system with respect to false positives and negatives of the detection process. To this end, we introduced two obstacles that are seen as robots by the robot detector, and moved all of the robots so as to generate temporary occlusions.

At the beginning of the experiment, robots and obstacles are arranged in a regular lattice (first snapshot) and there is an ambiguity in the configuration of \mathcal{R}_3 and \mathcal{R}_4 . In particular, according to PMR, \mathcal{R}_3 may be placed at the four corners of the rectangle, whereas \mathcal{R}_2 may be located either at its true position or at the specular corner. This particular situation is due to the fact that while \mathcal{R}_4 detects the obstacle in the opposite corner, \mathcal{R}_3 does not due to a momentary failure of the detection process (note that the theoretical analysis of Section 3 does not directly apply due to the presence of false positives). The robots start moving simultaneously (second snapshot), under the action of a decentralized pseudo-random control law with integrated obstacle avoidance. After the symmetry is broken ($t \simeq 8$ s), PMR is once

again able to extract a single solution, the clouds of particles separate, and the errors suddenly decrease. Then, the localization system continues to improve the estimates by the use of the measurements (third snapshot).

Note that \mathcal{R}_1 does not detect all of the other robots through the whole experiment, due to the limited field of view of the range finder and reciprocal occlusions. For example, the second snapshot shows that the line of sight between \mathcal{R}_1 and \mathcal{R}_3 is almost occluded by \mathcal{R}_4 , so that a false negative is generated around the corresponding time instant. However, the estimate is not affected by this temporary lack of measurements, because PMR is still able to place \mathcal{R}_3 by incorporating the observations of other robots that see it. Even when PMR does not produce valid matchings, the particle filters are still updated by incorporating odometric measurements; for example, this happens in the situation of the third snapshot, in which \mathcal{R}_1 is facing the outer wall and all of the other robots are in the blind zone of its detector.

As discussed at the end of Section 5.3, one may perform a batch least-squares optimization with known correspondences to further refine each solution found by PMR before feeding it to the filtering stage. Figure 11 shows the effect of this refinement on a sample solution computed by PMR during the second experiment. The BLS problem was solved by the MATLAB function `lsqnonlin` in 0.204 s. The entity of the correction on the robot configuration estimates is limited: the mean square difference is 0.72 cm for the position and 1.66° for the orientation. Therefore, to avoid the added computational load, we chose not to integrate the BLS refinement step in our localization system.

In the third experiment (Figure 12), the robots start in a non-symmetric spatial arrangement and are driven by the same decentralized pseudo-random control as before. As expected, the particle distributions are well separated and single-modal from the very beginning. At time $t \simeq 20$ s, \mathcal{R}_3 (green) is kidnapped (first snapshot) and released in a different position (second snapshot). The localization system quickly recovers the correct estimate of its relative configurations, as confirmed by the sample distributions in the third snapshot. Note that when the green robot is kidnapped, the belief on its configuration is still updated on the basis of odometric information, even though the underlying framework is momentarily non-rigid due to the lack of the measurements of the green robot.

Extension 1 contains video clips of these experiments. See <http://www.dis.uniroma1.it/~labrob/research/mutLoc.html> for additional material, including other experiments.

8. Discussion and comparison with other possible approaches

This section provides first a short discussion on some general aspects of the proposed localization system and then

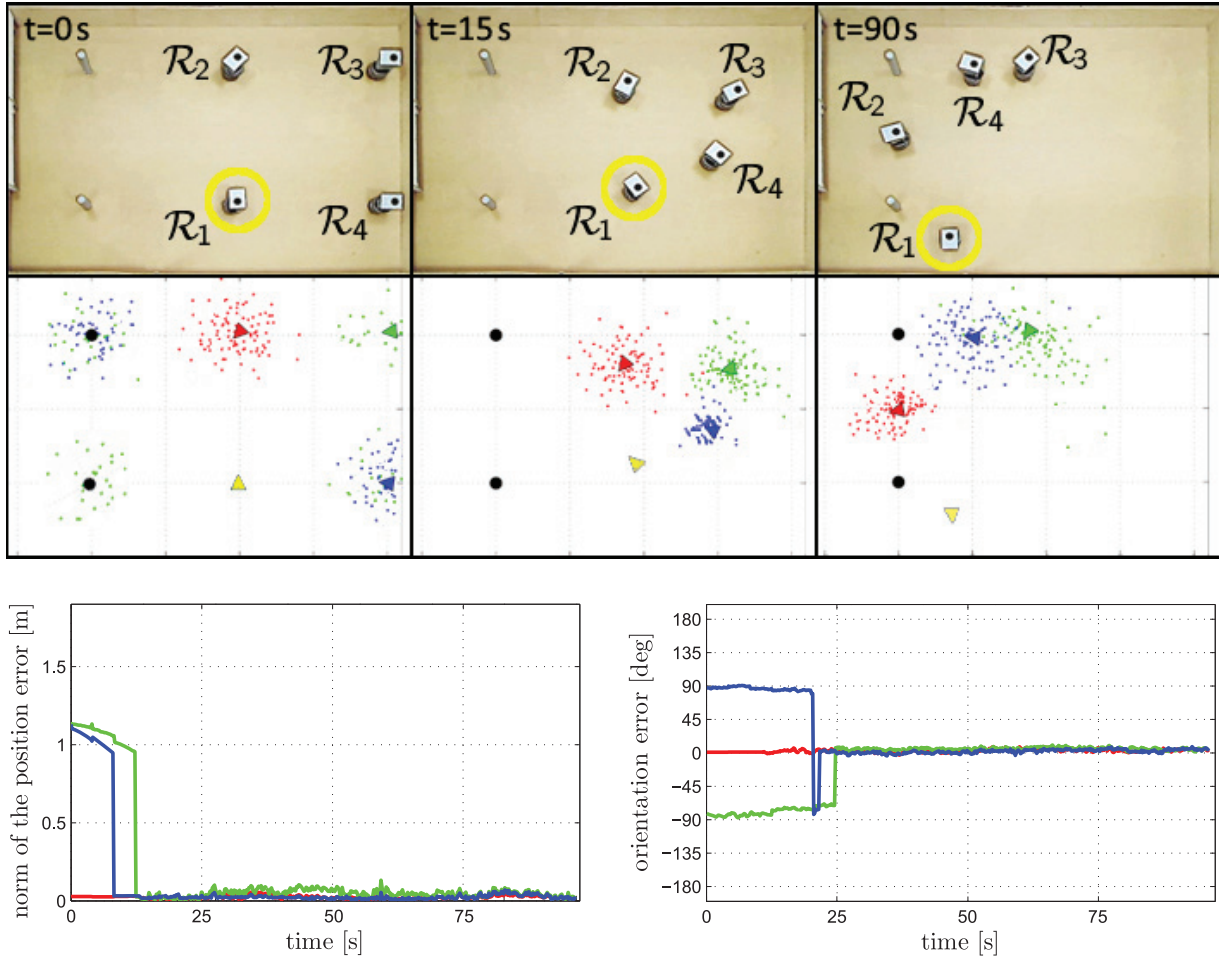


Fig. 10. Second experiment. Top: snapshots of the scene and sample distributions computed for \mathcal{R}_2 (red), \mathcal{R}_3 (green) and \mathcal{R}_4 (blue) by the mutual localization system of \mathcal{R}_1 (yellow, circled); the small triangles represent the ground truth and the black circles represent the obstacles. Bottom: norm of the position error and orientation error affecting the configuration estimates of \mathcal{R}_2 , \mathcal{R}_3 and \mathcal{R}_4 .

a quantitative comparison with a more classical approach developed from existing methods.

8.1. Discussion

The experiments of the previous section highlight the conditions needed in order to have an effective localization.

The presented method requires that every robot sends its measurements to its communication neighbors and that these measurements are sufficiently synchronized. This is not a technical issue, since with the current communication technology it is possible to synchronize two computers in a network with an accuracy of 1 ms or less.

The number n of robots must *not* be known in advance, and may indeed vary during the operation. In fact, whenever a new robot enters in the communication neighborhood of \mathcal{R}_i and the PMR algorithm is able to formulate a hypothesis on its configuration, a new particle filter is initialized for that robot.

Also, rigidity of the framework underlying the problem instance is not necessary for the proposed localization system to work. In fact, when the framework is non-rigid, PMR will still find a partial solution that provides configuration estimates for robots belonging to the its rigid component (see the discussion at the end of Section 5.3). In addition, the filtering stage will exploit odometric measures to propagate current estimates, even for those robots for which PMR is unable to provide an estimate.

Finally, the proposed localization method effectively rejects outliers, since they are either removed by belief-based pruning (if they are inconsistent with the current belief) or they remain unlabeled, hence ‘inactive’, at the end of PMR.

8.2. Comparison with other possible approaches

As mentioned in the introduction, it appears that the problem formulated in this paper (Problem 1) has not been

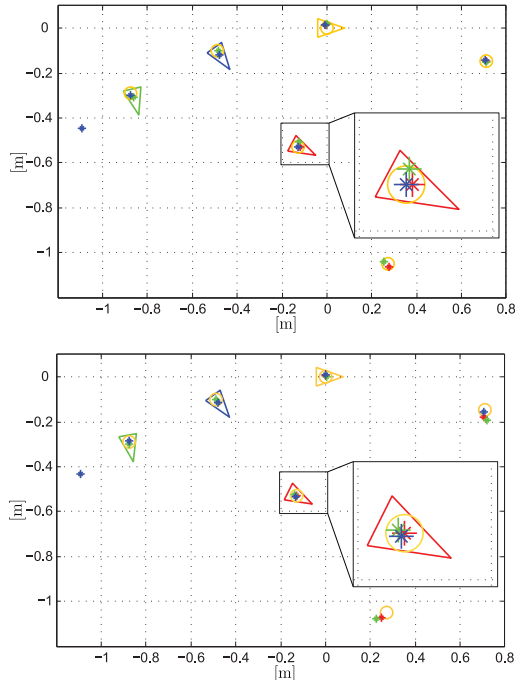


Fig. 11. The effect of the batch least-squares refinement (below) on a sample solution produced by PMR during the second experiment (above). Mutual localization is performed by \mathcal{R}_1 , shown as a yellow triangle, with its observed points marked by yellow circles. The red, green, blue triangles and asterisks represent, respectively, the estimated configurations of \mathcal{R}_2 , \mathcal{R}_3 , \mathcal{R}_4 and their observations. A zoom around \mathcal{R}_2 is provided to highlight the slight refinement.

explicitly considered so far by other researchers. However, some techniques exist in the literature for handling unknown data association in filtering problems.

A relatively standard approach would be to perform data association (i.e. guess the identity of each measured robot) before the filtering phase, using the maximum likelihood criterion (Thrun et al., 2005) or the joint compatibility test (Neira and Tardós, 2001). These approaches are reasonably successful provided that robot encounters are occasional and primarily pairwise, so that the feature sets O_k^i , $k \in \mathcal{N}$, include few or no elements for most of the time. In a densely populated environment, this method, although computationally tractable, may be hazardous, since even only one mismatched data association could jeopardize the quality of the estimate. On the other hand, our objective was to design of a mutual localization method which was effective also in the presence of frequent, not necessarily pairwise encounters among robots (as those occurring in formation control, cooperative exploration, and so on).

A promising idea is to propagate a fixed number of ‘best’ data associations at each step. To do this, one could adopt the strategy of FastSLAM with unknown data association presented by Montemerlo and Thrun (2003), in which feature-based SLAM is performed using a particle filter. The state of each particle consists of the state of the robot

and a bank of EKFs, one for each observed landmark of the map. This approach relies on the conditional independence of the estimates of the landmarks.

Since FastSLAM addresses SLAM with static obstacles, it is not directly applicable to the solution of Problem 1, which requires the localization of moving agents. We have therefore developed a FastSLAM-like method, in which \mathcal{R}_i estimates the relative configurations of the other robots. In practice, each particle represents a different hypothesis on the current state of the system (relative configurations of robots and robot-like obstacles), whose estimate is provided by a bank of EKFs.

When the robot detector of \mathcal{R}_i provides new measurements, data association is performed on a per-particle basis. In each particle, measurements are associated to the robots so as to maximize their likelihood given the current estimates. Then, the filters in each particle are updated, and the particle itself is rated by the likelihood of the measurements. The rating is subsequently used to build the probability of the particles in a drawing, so that particles receiving low ratings will be erased over time. In this way, at any given time, some fraction of the particles may receive plausible (i.e. high-likelihood) but erroneous data associations. If subsequent observations are received that clearly refute these prior assignments, these particles will be assigned low probabilities and removed from the filter. Observations communicated to \mathcal{R}_i by other robots are treated similarly.

Our FastSLAM-like method takes advantage of the main features of the original FastSLAM, i.e. keeping the number of hypotheses upper-bounded, delaying decision making about measurements that can be associated with multiple robots, and allowing the incorporation of mutually exclusive and negative information.

In order to perform a comparison, we have run the FastSLAM-like mutual localization method on the same experiments presented in Section 7. To achieve a rough equivalence to our method from a computational viewpoint, we first implemented the FastSLAM-like method with 100 particles. For example, in the first experiment the maximum value of $|N_i|$ is 3 and this would lead to 300 EKFs, comparable to the 300 particles used by our method. Figure 13 shows the performance of the FastSLAM-like method on the first experiment data. Clearly, localization is not successful, since the orientation estimates are affected by significant errors.

In view of the above results, we increased the number of particles used by the FastSLAM-like method to 1000. Considering that each particle includes a number of EKFs equal to the number of observed features, this brings the computational load to 3000 EKFs in the first experiment and 5000 EKFs in the second experiment, i.e. 10 times the load of our method. As shown by Figure 14, the increase in the number of particles was beneficial, since the performance of the FastSLAM-like method becomes comparable to that of our method (compare with the error plots in Figure 9).

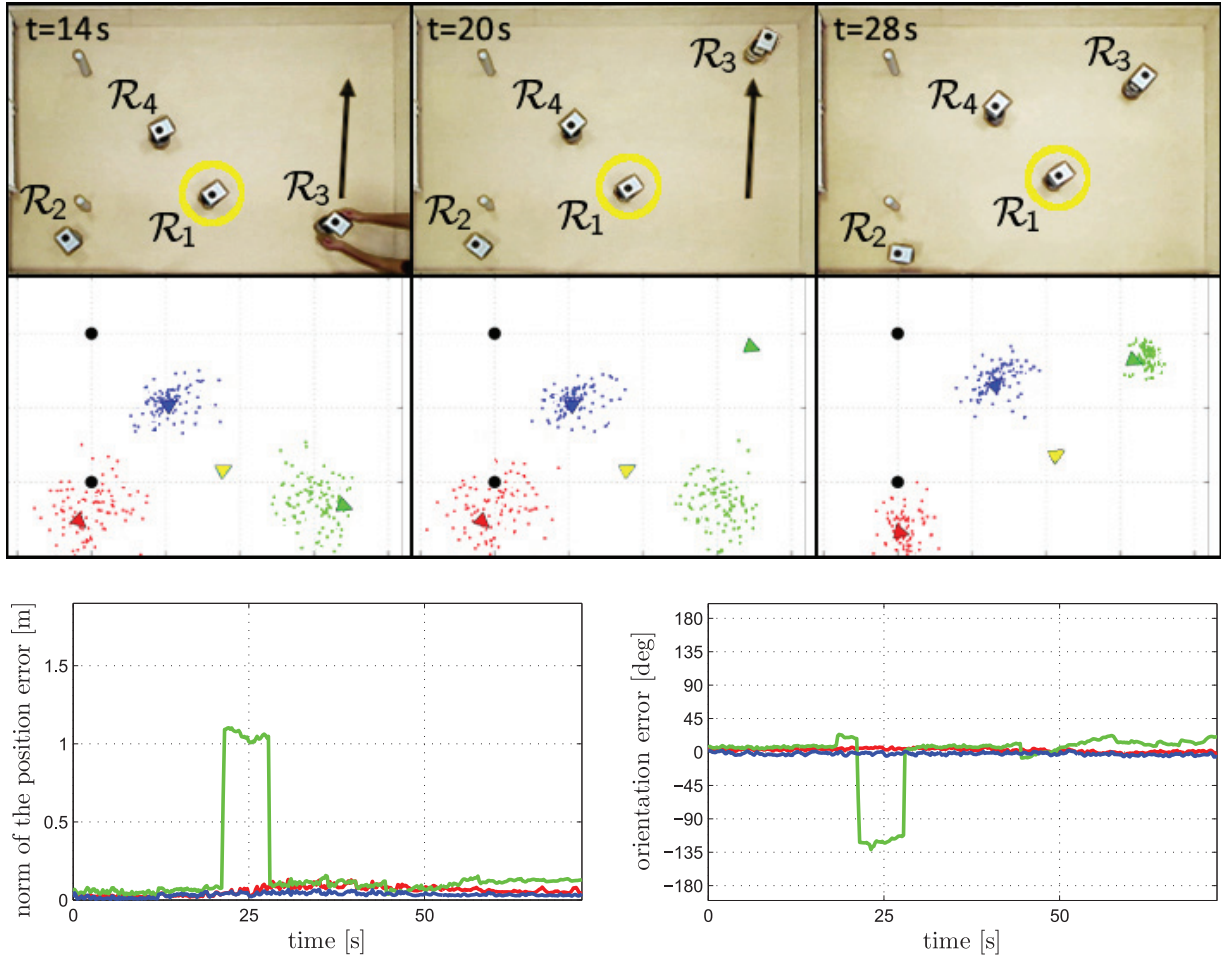


Fig. 12. Third experiment. Top: snapshots of the scene and sample distributions computed for \mathcal{R}_2 (red), \mathcal{R}_3 (green) and \mathcal{R}_4 (blue) by the mutual localization system of \mathcal{R}_1 (yellow, circled); the small triangles represent the ground truth and the black circles represent the obstacles. Bottom: norm of the position error and orientation error affecting the configuration estimates of \mathcal{R}_2 , \mathcal{R}_3 and \mathcal{R}_4 .

Still, the convergence of the FastSLAM-like method once the symmetry is broken is considerably slower.

Results of the FastSLAM-like method (still with 1000 particles) in the second and third experiments are reported in Figures 15 and 16, respectively (to be compared with Figures 10 and 12, respectively). In spite of the much higher computational load, the FastSLAM-like method cannot produce reliable estimates in these two experiments due to their increased difficulty. For example, while our method achieves good performance throughout the third experiment, solving the kidnapping issue in few seconds, the FastSLAM-like method never recovers the actual arrangement of the group, although the correct configuration of some robots is occasionally reconstructed (e.g. in the second experiment, \mathcal{R}_3 after $t = 25$ s).

The superior performance of our method can be explained by the fact that the mutual exclusive structure of the solution (i.e. the fact that the same measurement cannot be associated with different robots) is directly exploited in the multiple registration phase, whose complexity is

quadratic with $|N_i|$ in normal operation. The successive filtering phase ‘forgets’ the mutual exclusion by keeping a bank of independent particle filters. However, as already mentioned, mutual exclusion implicitly stands, since the input data for the filters are the solutions of the multiple registration process. Moreover, the increased dimension of the measurements (augmented with the orientation by PMR) allows the particle filters to converge very quickly.

9. Source code release

A C++ code that implements the mutual localization algorithm described in this paper has been released within the MLAM ROS stack, which is available at <http://www.ros.org/wiki/mlam>. Note that the code is released under the BSD license, therefore it is provided ‘as is’ and we cannot be held responsible for any damage or harm caused by the use of our software.




Fig. 13. First experiment with the FastSLAM-like method using 100 particles. Norm of the position error (top) and orientation error (bottom) affecting the configuration estimates of \mathcal{R}_2 (red), \mathcal{R}_3 (green) and \mathcal{R}_4 (blue).




Fig. 15. Second experiment with the FastSLAM-like method using 1000 particles. Norm of the position error (top) and orientation error (bottom) affecting the configuration estimates of \mathcal{R}_2 (red), \mathcal{R}_3 (green) and \mathcal{R}_4 (blue) by the FastSLAM-like method.




Fig. 14. First experiment with the FastSLAM-like method using 1000 particles. Norm of the position error (top) and orientation error (bottom) affecting the configuration estimates of \mathcal{R}_2 (red), \mathcal{R}_3 (green) and \mathcal{R}_4 (blue) by the FastSLAM-like method.




Fig. 16. Third experiment with the FastSLAM-like method using 1000 particles. Norm of the position error (top) and orientation error (bottom) affecting the configuration estimates of \mathcal{R}_2 (red), \mathcal{R}_3 (green) and \mathcal{R}_4 (blue) by the FastSLAM-like method.

10. Conclusions

We have presented and experimentally validated a novel mutual localization algorithm for the case of anonymous

position measurements. As explained in the paper, this is a challenging and practically relevant operating scenario that has received little attention in the literature.

The method is based on a two-step approach: a probabilistic multiple registration algorithm, which provides all data associations that are consistent with the relative robot measurements and the current belief, as well as the robot orientations; and a dynamic filtering stage, which incorporates odometric data into the estimation process. We designed the proposed method proceeding from a detailed formal analysis of the implications of anonymity on the mutual localization problem.

Experimental results on a team of differentially driven mobile robots have shown the effectiveness of the approach, and in particular its robustness against false positives and negatives that may affect the robot measurement process. We have also provided an experimental comparison that shows how the proposed method outperforms more classical approaches that may be designed building on existing techniques. Finally we have made available the source code of the proposed method within the MLAM ROS stack.

We are currently working on the development of a similar localization approach for groups of 3D agents (i.e. in $SE(3)$) and on its application to swarms of flying robots, possibly equipped with more limited (e.g. bearing-only plus IMU) sensors. In this direction, some preliminary results have been recently presented by Stegagno et al. (2011) and Cognetti et al. (2012).

Notes

1. Here, ‘similar’ must be intended in the sense of the specific attribute used to detect robots, be it shape, color, etc.
2. Since in this section we are considering a static version of the problem we shall omit all of the t superscripts. Note also that, for compactness, symbols introduced and used here but not in the rest of the paper are not included in Table 1.
3. Strictly speaking, the two problems are not equivalent due to the fact that we deal with sets of labeled points.
4. If both matchings had passed belief-based pruning (for example, because the current belief on the configurations of \mathcal{R}_j and \mathcal{R}_k was uniform), the algorithm would have created two branches, ultimately leading to multiple different solutions. This kind of situation occurs in the early stages of the first and second experiments, see Section 7.
5. See <http://www.dis.uniroma1.it/labrob/software/MIP>.
6. See <http://www.ros.org/wiki/mlam>.

Funding

This work was partially supported by the PRIN program of the Italian Ministry of Education, University and Research (MEMO-NET Project).

References

Aspnes J, Eren T, Goldenberg DK, et al. (2006) A theory of network localization. *IEEE Transactions on Mobile Computing* 5(12): 1663–1678.

Censi A, Franchi A, Marchionni L and Oriolo G (2013) Simultaneous maximum-likelihood calibration of odometry and sensor parameters. *IEEE Transactions on Robotics* 29(2): 475–492.

Chen J, Sun D, Yang J and Chen H (2010) Leader–follower formation control of multiple non-holonomic mobile robots incorporating a receding-horizon scheme. *The International Journal of Robotics Research* 29(6): 727–747.

Cognetti M, Stegagno P, Franchi A, Oriolo G and Bühlhoff HH (2012) 3-D mutual localization with anonymous bearing measurements. In: *2012 IEEE International Conference on Robotics and Automation*, St. Paul, MN, pp. 791–798.

Cunningham A, Wurm KM, Burgard W and Dellaert F (2012) Fully distributed scalable smoothing and mapping with robust multi-robot data association. In: *2012 IEEE International Conference on Robotics and Automation*, St. Paul, MN, pp. 1093–1100.

Dieudonne Y, Labbani-Igbida O and Petit F (2008) On the solvability of the localization problem in robot networks. In: *2008 IEEE International Conference on Robotics and Automation*. Pasadena, CA, pp. 480–485.

Durham JW, Franchi A and Bullo F (2012) Distributed pursuit–evasion without global localization via local frontiers. *Autonomous Robots* 32(1): 81–95.

Fink J, Michael N, Kim S and Kumar V (2010) Planning and control for cooperative manipulation and transportation with aerial robots. *The International Journal of Robotics Research* 30(3): 324–334.

Fischler MA and Bolles RC (1981) Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24(6): 381–395.

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

Franchi A, Masone C, Grabe V, Ryll M, Bühlhoff HH and Robuffo Giordano P (2012) Modeling and control of UAV bearing-formations with bilateral high-level steering. *The International Journal of Robotics Research* 31(12): 1504–1525.

Franchi A, Oriolo G and Stegagno P (2009a) Mutual localization in a multi-robot system with anonymous relative position measures. In: *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, St. Louis, MO, pp. 3974–3980.

Franchi A, Oriolo G and Stegagno P (2009b) *Mutual localization in a multi-robot system with anonymous relative position measures*. Technical Report, number 1, 2009, http://ojs.uniroma1.it/index.php/DIS_TechnicalReports/article/view/2797/ Department of Computer and System Sciences Antonio Ruberti.

Franchi A, Oriolo G and Stegagno P (2010a) On the solvability of the mutual localization problem with anonymous position measures. In: *2010 IEEE International Conference on Robotics and Automation*, Anchorage, AK, pp. 3193–3199.

Franchi A, Stegagno P and Oriolo G (2010b) Probabilistic mutual localization in multi-agent systems from anonymous position measures. In: *49th IEEE Conference on Decision and Control*, Atlanta, GA, pp. 6534–6540.

Grabowski R, Navarro-Serment LE, Paredis CJJ and Khosla P (2000) Heterogenous teams of modular robots for mapping and exploration. *Autonomous Robots* 8(3): 293–308.

Howard A, Matarić MJ and Sukhatme GS (2002) Localization for mobile robot teams using maximum likelihood estimation. In: *2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, pp. 434–439.

- Howard A, Mataric MJ and Sukhatme GS (2003) Putting the ‘I’ in ‘team’: An ego-centric approach to cooperative localization. In: *2003 IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, pp. 868–892.
- Howard A, Parker LE and Sukhatme GS (2006) Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment and detection. *The International Journal of Robotics Research* 25(5–6): 431–447.
- Jackson B (2007) Notes on the rigidity of graphs. In: *Levico Conference Notes*. Levico.
- Kurazume R and Hirose S (2000) An experimental study of a cooperative positioning system. *Autonomous Robots* 8(1): 43–52.
- Kurazume R, Nagata S and Hirose S (1994) Cooperative positioning with multiple robots. In: *1994 IEEE International Conference on Robotics and Automation*, San Diego, CA, pp. 1250–1257.
- Leung KYK, Barfoot T and Liu HHT (2010) Decentralized localization of sparsely-communicating robot networks: A centralized-equivalent approach. *IEEE Transactions on Robotics* 26(1): 62–77.
- Martinelli A (2007) Improving the precision on multi robot localization by using a series of filters hierarchically distributed. In: *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Diego, CA, pp. 1053–1058.
- Martinelli A, Pont F and Siegwart R (2005) Multi-robot localization using relative observations. In: *2005 IEEE International Conference on Robotics and Automation*, Barcelona, Spain, pp. 2797–2802.
- Montemerlo M and Thrun S (2003) Simultaneous localization and mapping with unknown data association using FastSLAM. In: *2003 IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, pp. 1985–1991.
- Mourikis AI and Roumeliotis SI (2006) Performance analysis of multirobot cooperative localization. *IEEE Transactions on Robotics* 22(4): 666–681.
- Neira J and Tardós JD (2001) Data association in stochastic mapping using the joint compatibility test. *IEEE Transactions on Robotics* 17(6): 890–897.
- Pasqualetti F, Franchi A and Bullo F (2012) On cooperative patrolling: Optimal trajectories, complexity analysis, and approximation algorithms. *IEEE Transactions on Robotics* 28(3): 592–606.
- Piovan G, Shames I, Fidan B, Bullo F and Anderson BDO (2013) On frame and orientation localization for relative sensing networks. *Automatica* 49(1): 206–213.
- Rekleitis IM, Dudek G and Milios EE (1998) On multiagent exploration. In: *Vision Interface 1998*, Vancouver, Canada, pp. 455–461.
- Rekleitis I, Dudek G and Milios E (2001) Multi-robot collaboration for robust exploration. *Annals of Mathematics and Artificial Intelligence* 31(1): 7–40.
- Roumeliotis SI and Bekey GA (2002) Distributed multirobot localization. *IEEE Transactions on Robotics* 18(5): 781–795.
- Sanderson A (1998) A distributed algorithm for cooperative navigation among multiple mobile robots. *Advanced Robotics* 12(4): 335–349.
- Schwager M, Julian BJ and Rus D (2009) Optimal coverage for multiple hovering robots with downward facing cameras. In: *2009 IEEE International Conference on Robotics and Automation*, Kobe, Japan, pp. 3515–3522.
- Shames I, Fidan B and Anderson BDO (2009) Minimization of the effect of noisy measurements on localization of multi-agent autonomous formations. *Automatica* 45(4): 1058–1065.
- Smith RC and Cheeseman P (1986) On the representation and estimation of spatial uncertainty. *The International Journal of Robotics Research* 5(4): 56–68.
- Stegagno P, Cagnetti M, Franchi A and Oriolo G (2011) Mutual localization using anonymous bearing-only measures. In: *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Francisco, CA, pp. 469–474.
- Thrun S, Burgard W and Fox D (2005) *Probabilistic Robotics*. Cambridge, MA: MIT Press.
- Thrun S and Liu Y (2005) Multi-robot SLAM with sparse extended information filters. In: *11th International Symposium on Robotics Research*, Siena, Italy, pp. 254–266.
- Trawny N, Zhou XS, Zhou K and Roumeliotis SI (2010) Inter-robot transformations in 3D. *IEEE Transactions on Robotics* 26(2): 225–243.
- Umeyama S (1991) Least-square estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13(4): 376–380.
- Zhou XS and Roumeliotis SI (2008) Robot-to-robot relative pose estimation from range measurements. *IEEE Transactions on Robotics* 24(6): 1379–1393.

Appendix: Index to Multimedia Extensions

The multimedia extension page is found at <http://www.ijrr.org>

Table of Multimedia Extensions

Extension	Type	Description
1	Video	Three experiments of mutual localization using anonymous measurements