

Putting the ‘I’ in ‘Team’: an Ego-Centric Approach to Cooperative Localization

Andrew Howard, Maja J Matarić and Gaurav S. Sukhatme
Robotics Research Laboratory, Computer Science Department
University of Southern California
Los Angeles, California 90089-0781, USA
ahoward@usc.edu gaurav@usc.edu mataric@usc.edu

Abstract—This paper describes a cooperative method for relative localization of mobile robot teams; that is, it describes a method whereby every robot in the team can estimate the pose of every other robot, relative to itself. This method does not require the use of GPS, landmarks, or maps of any kind; instead, robots make direct measurements of the relative pose of nearby robots, and broadcast this information to the team as a whole. Each robot processes this information independently to generate an *ego-centric* estimate for the pose of other robots. Our method uses a Bayesian formalism with a particle filter implementation, and is, as a consequence, very robust. It is also completely distributed, yet requires relatively little communication between robots. This paper describes the basic ego-centric formalism, sketches the implementation, and presents experimental results obtained using a team of four mobile robots.

I. INTRODUCTION

Localization is a key capability for mobile robot teams no less than it is for individual mobile robots. For teams, however, we must distinguish between two different forms of localization: *global* or *absolute* localization, in which each of the robots is localized with respect to some external coordinate system, and *relative* localization, in which each of the robots is localized with respect to the other robots in the team. For many team oriented tasks, it is this latter class of localization that is the more important. Robots executing a formation behavior, for example, need to know their pose with respect to other robots in the formation; they do not need to know their global pose. In a similar vein, robots wishing to share sensor data with one another require knowledge of their relative pose if they are to make sense of that data. With good relative localization, sensor data from all of the robots in a team can be combined into common, coherent, spatially extended representation.

Naturally, given a set of absolute pose estimates for the robots in a team, one can always derive a set of relative pose estimates. Global localization, using either GPS or model-based methods, is a well studied problem [9], [13], [4], and will not be treated here. Instead, we propose an approach in which the robots make relative pose estimates through direct observation of other robots; the robots make no attempt to determine their absolute pose.

Since this approach does not require external landmarks or environment models, it is highly suited for use in environments that are either unknown, non-static, or both.

We make the following assumptions. First, we assume that each robot in the team is equipped with a *motion sensor* with which it can measure changes in its own pose; odometry or inertial measurement units are typically used for this purpose. Second, we assume that each robot is equipped with a *robot sensor* that allows it to measure the relative pose and identity of nearby robots; robot sensors can be readily constructed using cameras and/or scanning laser range finders in combination with coded fiducials placed on the robots. Finally, we assume that robots are able to communicate using some broadcast medium (such as a wireless network). For simplicity, we assume that communication is both reliable and complete, but note that this is not essential (the approach described in this paper can be extended to handle the unreliable communication case).

The basic method is as follows. Consider a single robot taken from a team of n robots; this robot attempts to estimate the pose of every other robot in the team *relative to itself*. To do this, it maintains a set of $n - 1$ independent probability distributions, each of which describes the pose of one robot in an ego-centric coordinate system. These distributions are updated in response to observations obtained from both the robot’s own sensors, and from the sensors of other robots (robots are required to broadcast their observations to the team). There are five basic types of observations that must be considered: from the motion sensors we may observe either that *this* robot has moved, or that some *other* robot has moved; from the robot sensors we may learn that this robot has observed another robot, that another robot has observed this robot, or that two other robots have observed each other. The update rules used for the first four types of observations can be derived using a fairly simple Bayesian analysis; they are similar to the update rules used in Monte-Carlo localization [4].

The update rules for the last type of observation, however, are somewhat more complicated. This type of observation defines a relationship between two probability distributions, and one must pay careful attention to the

dependencies that exist between these two distributions. Thus, for example, we must avoid the situation in which distribution A is used to update distribution B, which is then used to update distribution A, and so on. In Section III-D, we introduce the concept of a *dependency tree* as an approximate, but efficient, method for preventing such circular reasoning.

The ego-centric approach described in this paper has a number of attractive features. It is highly decentralized (each robot maintains its own set of $n - 1$ distributions), yet requires very little communication bandwidth (robots need only transmit their observations). It is also robust to communication failure: pose estimates simply get more uncertain as observations are lost. Our particular implementation, which makes use of both particle filters and mixture sampling [16], is also extremely robust. It does not require prior knowledge of the robot poses, and can recover from tracking failures. That is, it is capable of both self-initializing and self-correcting.

The chief drawback of the approach is its computational cost: particle filters are relatively expensive, and each robot in a team of n robots must maintain $n - 1$ separate filters. Our empirical results, however, are promising: we can easily run the system in real-time on a team of four mobile robots equipped with Pentium-class computers.

II. RELATED WORK

Localization is an extremely well studied problem in mobile robotics. The vast majority of this research, however, has concentrated on just two problems: localizing a single robot using an a priori map of the environment [9], [13], [4], or localizing a single robot while simultaneously building a map [15], [10], [17], [2], [1]. Recently, some authors have also considered the related problem of map building with multiple robots [14]. All of these authors make use of statistical or probabilistic techniques; the common tools of choice are Kalman filters, particle filters, maximum likelihood estimation and expectation maximization.

Among those who have considered the specific problem of cooperative localization are [12] and [3]. Roumeliotis and Bekey [12] present an approach to multi-robot localization in which sensor data from a heterogeneous collection of robots are combined through a single Kalman filter to estimate the pose of each robot in the team. They also show how this centralized Kalman filter can be broken down into n separate Kalman filters (one for each robot) to allow for distributed processing. Similarly, Fox *et al.* [3] describe an approach to multi-robot localization in which each robot maintains a probability distribution describing its own pose (based on odometry and environment sensing), but is able to refine this distribution through the observation of other robots. This approach extends earlier work on single-robot probabilistic localization techniques

[4]. Note that both of these authors are principally concerned with determining the *absolute* pose of robots (with or without the use of external landmarks).

Finally, a number of authors [8], [11], [7] have considered the problem of cooperative localization from a somewhat different perspective. These authors describe *active* approaches to localization, in which team members carefully coordinate their activities in order to reduce cumulative odometric errors. The basic method is to keep one subset of the robots stationary, while the other robots are in motion; the stationary robots observe the robots in motion (or vice-versa), thereby providing more accurate pose estimates than can be obtained using odometry alone. While our approach does not require such explicit cooperation on the part of robots, the accuracy of localization can certainly be improved by the adoption of such strategies.

III. FORMALISM

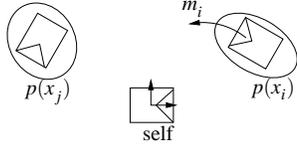
Consider a single robot drawn from a team of n robots; we will refer to this robot as *self*. Let x_i denote the current pose of some other robot i relative to self. Our aim is to compute the probability distribution $p(x_i)$ for each such robot in the team, based on observations made both by self and by other robots.

As previously indicated, we assume that each robot is equipped with two sensors: a *robot sensor* that allows it to measure the relative pose and identity of nearby robots, and a *motion sensor* that allows it to measure changes in its own pose. Together, these two sensors give rise to five distinct types of observations. For the motion sensor there are observations made by self, and observations made by another robot. For the robot sensor there are observations in which self observes another robot; observations in which another robot observes self; and observations in which another robot observes a third robot. We refer to this last class of observations as *transitive*, since there is an intermediary between self and the robot being observed. For each kind of observation, we require a different update rule for the probability distributions $p(x_i)$. Since transitive observations in fact require *two* update rules (to be used in different situations) there are a total of six update rules. These rules are summarized in Figure 1 and are described in detail below.

A. On Coordinate Transforms

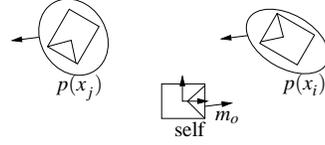
The formalism described in this paper relies heavily on the use of coordinate transformations, for which we define a pair of coordinate transform operators \ominus and \oplus . Let x_i and x_j denote the pose of robots i and j in some coordinate system, and let z_{ij} denote the pose of robot i relative to robot j (i.e., the pose of robot i in the coordinate system centered on robot j). The operators \oplus and \ominus are defined such that:

$$z_{ij} = x_i \ominus x_j \quad \text{and} \quad x_i = z_{ij} \oplus x_j \quad (1)$$



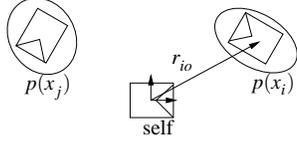
Motion observation m_i : robot i observes motion (indirect).

$$p(x_i | m_i) = N_i \int p(m'_i | m_i) p(x_i \ominus m'_i) dm'_i$$



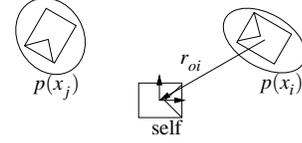
Motion observation m_o : self observes motion (direct).

$$\forall i \ p(x_i | m_o) = N_o \int p(m'_o | m_o) p(x_i \oplus m'_o) dm'_o$$



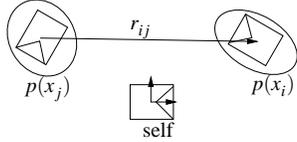
Robot observation r_{io} : self observes robot i (direct).

$$p(x_i | r_{io}) = N_{io} p(r_{io} | x_i \ominus x_o) p(x_i)$$



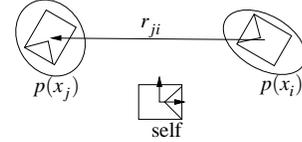
Robot observation r_{oi} : robot i observes self (indirect).

$$p(x_i | r_{oi}) = N_{oi} p(r_{oi} | x_o \ominus x_i) p(x_i)$$



Robot observation: r_{ij} robot j observes robot i (indirect).

$$p(x_i | r_{ij}) = N_{ij} \int p(r_{ij} | x_i \ominus x_j) p(x_i) p(x_j) dx_j$$



Robot observation: r_{ji} robot i observes robot j (indirect).

$$p(x_i | r_{ji}) = N_{ji} \int p(r_{ji} | x_j \ominus x_i) p(x_i) p(x_j) dx_j$$

Fig. 1. Summary of the distribution update rules. There are 6 rules corresponding to 5 different types of observations (transitive robot observations have two update rules). The ellipses indicate the probability distribution associated with other robots.

These operators adhere to the normal rules of association:

$$(z_{ij} \oplus x_j) \ominus x_j = z_{ij} \oplus (x_j \ominus x_j) \quad (2)$$

but do not necessarily commute:

$$z_{ij} \oplus x_j \neq x_j \oplus z_{ij}. \quad (3)$$

If one ignores robot orientation and considers position only, these operators reduce to standard vector addition and subtraction. In the discussion that follows, these operators can be interpreted in this manner with minimal loss of generality.

B. Motion Observations

Let m_i denote a motion observation made by robot i . This observation affects the distribution $p(x_i)$, which must be both translated (to account for the robot's motion) and 'blurred' (to account for the uncertainty in that motion). Mathematically, one can show that the update rule for such an observation is given by the convolution:

$$p(x_i | m_i) = N_i \int p(m'_i | m_i) p(x_i \ominus m'_i) dm'_i \quad (4)$$

where the integral is over all possible changes m'_i in the robot's pose, and N_i is a normalization constant. The second term in the integral is the robot's pose distribution prior to motion; i.e., probability that the robot was at pose $x_i \ominus m'_i$. The first term in the integral is the *motion sensor model*: $p(m'_i | m_i)$ is the probability that the *true* change in robot pose is m'_i , given that the *measured* change is m_i . For a perfect motion sensor (i.e., one for which there is no uncertainty in m_i) the sensor model reduces to a delta function, and the integral collapses to yield $p(x_i | m_i) = p(x_i \ominus m_i)$. That is, the distribution $p(x_i)$ is translated through a distance m_i .

Let m_o denote a motion observation made by self. Since self is, by definition, at the origin of the coordinate system, this observation is used to update the pose distributions $p(x_i)$ for the *other* robots. Thus, all distributions must be both translated (to account for self's motion) and 'blurred' (to account for the uncertainty in that motion). Thus the update rule for such observations is given by the convolution:

$$\forall i \ p(x_i | m_o) = N_o \int p(m'_o | m_o) p(x_i \oplus m'_o) dm'_o \quad (5)$$

where the integral is over all possible changes m_o in the self's pose and N_o is the normalization constant. The second term in the integral describes the pose distribution for robot i prior to motion by self, while the first term is the motion sensor model described above. In the case of perfect motion sensors, the integral collapses to $p(x_i | m_o) = p(x_i \oplus m_o)$; that is, all of the probability distributions are translated through a distance $0 \ominus m_o$.

Note that we have made the simplifying assumption that the motion sensor model depends *only* on the measured change in pose, and not on the robot's state. We also assume (without loss of generality) that the robots are homogeneous, and hence the same sensor model can be used for motion observations from all robots.

C. Robot Observations (Non-Transitive)

Let r_{io} denote an observation made by self of some robot i ; r_{io} describes the measured pose of robot i relative to self. This observation can be used to update the distribution $p(x_i)$; the update rule is derived through the application of Bayes' Law:

$$p(x_i | r_{io}) = N_{io} p(r_{io} | x_i \ominus x_o) p(x_i) \quad (6)$$

where N_{io} is a normalization constant, and $p(x_i)$ is the prior distribution. The conditional term is the *robot sensor model*, which describes the probability of obtaining the measurement r_{io} , given that the true pose of robot i relative to self is $x_i \ominus x_o$. Note that x_o denotes the pose of self, which is equal to zero by definition; we use this notation purely for the sake of symmetry.

Let r_{oi} denote an observation made by robot i of self; i.e., r_{oi} describes the measured pose of self relative to robot i . Once again, we use this observation to update the distribution $p(x_i)$; the update rule is:

$$p(x_i | r_{oi}) = N_{oi} p(r_{oi} | x_o \ominus x_i) p(x_i) \quad (7)$$

This rule makes use of the same robot sensor model used in Equation 7, but applies it in the reverse sense; i.e., to the pose of self relative to i instead of the pose of i relative to self. We assume that the robot sensor model depends only on this relative pose, and that the model is the same for all robots.

D. Robot Observations (Transitive)

Transitive robot observations, i.e., observations in which some robot i observes another robot j (with neither robot being self), must be treated with some care. Consider, for example, a scenario in which robot i observes robot j , which then observes robot i . If we were to use the first observation to update the distribution $p(x_j)$, and subsequently use the second observation to update $p(x_i)$, we would be guilty of double-counting: the first observation would effectively be used to update $p(x_i)$ not once, but twice. This kind of circular reasoning can easily lead to

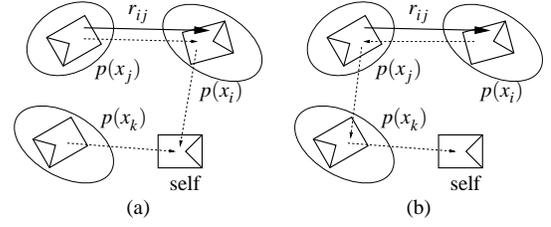


Fig. 2. Two possible dependency trees; the arrows indicate child-parent relationships. (a) Robot j observes robot i , and since i is an ancestor of j we update $p(x_j)$. (b) Robot j observes robot i , and since j is an ancestor of i , we update $p(x_i)$.

over-convergence, with the pose estimates for a pair of robots i and j quickly converging to some precise but inaccurate value.

In order to avoid such circular reasoning, we maintain a *dependency tree* for all $n - 1$ distributions. In this tree, each distribution has exactly one *parent* distribution, and zero or more *child* distributions (see, for example, Figure 2). A distribution cannot be used to update its ancestors, but may be used to update its descendants; furthermore, whenever a distribution is used in this manner, it becomes the new parent of the distribution being updated. Thus the topology of the tree may change as new observation are added.

Given some transitive observation r_{ij} , the dependency tree is used to decide which of the two distributions $p(x_i)$ or $p(x_j)$ will be updated. If i is an ancestor of j , we update j ; if j is an ancestor of i , we update i . In the case that neither distribution is an ancestor of the other, we have a free choice as to which distribution should be updated. In this case latter case, we compute the moments of the two distributions and update the one with the greatest variance.

Let r_{ij} be an observation made by robot j of robot i ; i.e., the observation r_{ij} describes the measured pose of i relative to j . Let $p(x_i)$ be the distribution we have chosen to update (according to the rules described above). The update rule is given by the convolution:

$$p(x_i | r_{ij}) = N_{ij} \int p(r_{ij} | x_i \ominus x_j) p(x_i) p(x_j) dx_j \quad (8)$$

where the integral is over all possible poses for robot j . The first term in the integral is the robot sensor model defined in the previous section.

Let r_{ji} be an observation made by robot i of robot j ; i.e., the observation r_{ji} describes the measured pose of j relative to i . Once again, let $p(x_i)$ be the distribution we chosen to update. The update rule is given by the convolution:

$$p(x_i | r_{ji}) = N_{ji} \int p(r_{ji} | x_j \ominus x_i) p(x_i) p(x_j) dx_j \quad (9)$$

where the integral is over all possible poses for robot j .

E. Discussion

While the dependency tree prevents the most obvious circular updates, it does have some limitations. The dependency tree assumes that distributions are dependent *only* on the distribution that was last used to update them, and are therefore independent of all other distributions. As a consequence, it is still possible to construct situations in which circular updates occur. While these situations tend to be somewhat contrived, the dependency graph must nevertheless be treated as an approximate solution whose effectiveness must be determined empirically.

There does exist a more complete solution to this problem: each robot can maintain a set of pair-wise probability distributions $p(x_i \ominus x_j)$ for all (i, j) pairs, and project these distributions when necessary to generate $p(x_i \ominus x_o)$ and $p(x_j \ominus x_o)$. This solution is, of course, much more expensive, since each robot must maintain $O(n^2)$ pair-wise distributions. Alternatively, each robot can maintain only $O(n)$ distributions, but must periodically broadcast these distributions to the team as a whole (resulting in $O(n^2)$ broadcasts). Nevertheless, we are currently investigating this solution and comparing it with the dependency tree approach.

IV. IMPLEMENTATION

The formalism described in the previous section has only three inputs: the motion and robot sensor models (which must be known a priori), and a stream of observations. Observations may come either from a robot’s own sensors, or from the sensors of other robots. To ensure that all robots have access to all sensor observations, every robot on the team must broadcast its observations. The bandwidth required for this is small (of the order of a few hundred bytes/sec/robot) and scales linearly with the number of robots. In practice, we use UDP broadcast for transmitting observations; with appropriate time-stamping of the observations, one can easily construct a system that is robust to occasional packet loss.

The most difficult aspect of the implementation is, of course, maintaining the probability distributions $p(x_i)$. Since these distributions may be both multi-model and non-Gaussian, we model them using *particle filters*. Consider once again an individual robot ‘self’. Self maintains a set of $n - 1$ independent particle filters, each of which represents the relative pose distribution $p(x_i)$ for another robot. Each filter, in turn, contains a relatively large set of *samples*, each of which has a *pose* and a *weight*. Roughly speaking, each sample pose represents one possible pose for the robot, while each sample weight reflects the probability that this is the true pose. For each observation, one or more of the particle filters must be updated: motion observations modify the sample pose, while robot observations modify the sample weight. The

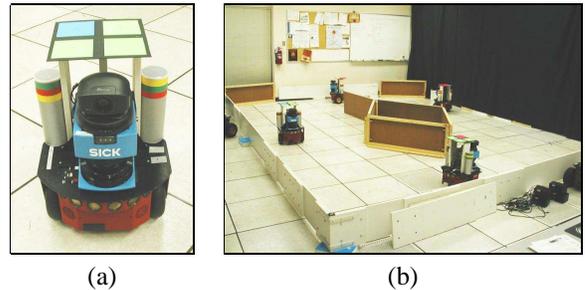


Fig. 4. (a) One of the robots used in the experiment. The robot is equipped with a scanning laser range-finder, a camera, a pair of retro-reflective/color-coded fiducials, and a color-coded hat (for the overhead tracking system). (b) The experimental environment.

full details of particle filter implementation are relatively complex, and will not be described here.

V. EXPERIMENTS

We have conducted a preliminary experiment aimed at determining the accuracy of the ego-centric approach and the robustness of our particular implementation. With regard to accuracy, we wish to determine the quality of both the pose estimates and the uncertainty in those estimates. With regard to robustness, we wish to ensure that the system will both self-initialize and self-correct.

A. Method

The experiment was performed using four Pioneer 2DX mobile robots running the *Player* [5] robot device server (Figure 4(a)). Each robot was equipped with a SICK scanning laser range-finder and a Sony pan-tilt-zoom camera; the laser and camera are used together to detect coded fiducials placed on the robots. These fiducials have both retro-reflective and colored patches: the retro-reflective patches are detected by the laser (which determines the fiducial’s range and bearing), while the colored patches are detected by the camera (which determines the fiducial’s identity). With appropriately placed fiducials, the range, bearing, orientation and identity of each robot can be uniquely determined.

Ground-truth data was generated by the *Mezzanine* tracking package [6], which uses multiple overhead cameras to track color-code fiducials placed on top of the robots. The tracking system is accurate to about ± 10 cm and $\pm 2^\circ$. Note that this system was used solely to generate comparative ground-truth data; no information from the tracking system was provided to the robots.

The environment for the experiment was a pen measuring approximately 8 by 6 meters, containing various obstacles (Figure 4(b)). The robots executed a simple wall-following behavior in this pen, performing 6 complete circuits without pausing. Note that since the algorithm described in this paper does not make use of environmental features, the structure of the environment is irrelevant

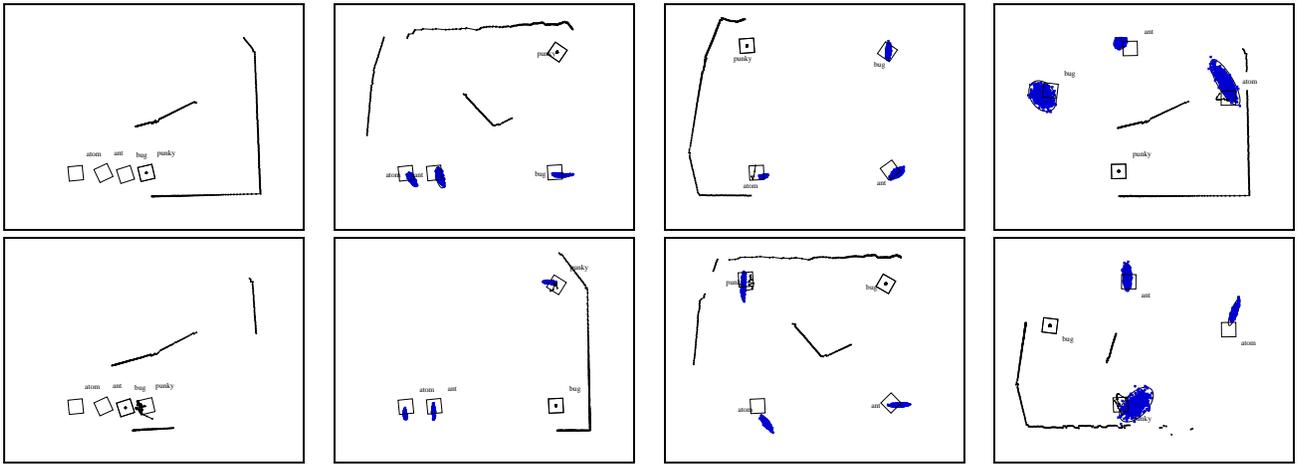


Fig. 3. Snap-shots showing the pose estimates generated by two of the four robots; each row contains a time-series of the estimates generated by one robot. The ‘clouds’ show the particle filter estimates, the ellipses show the 3σ interval, and the squares denote the true robot poses (as determined by the overhead tracking system). These figures have been rendered in the ground-truth coordinate system to simplify comparison.

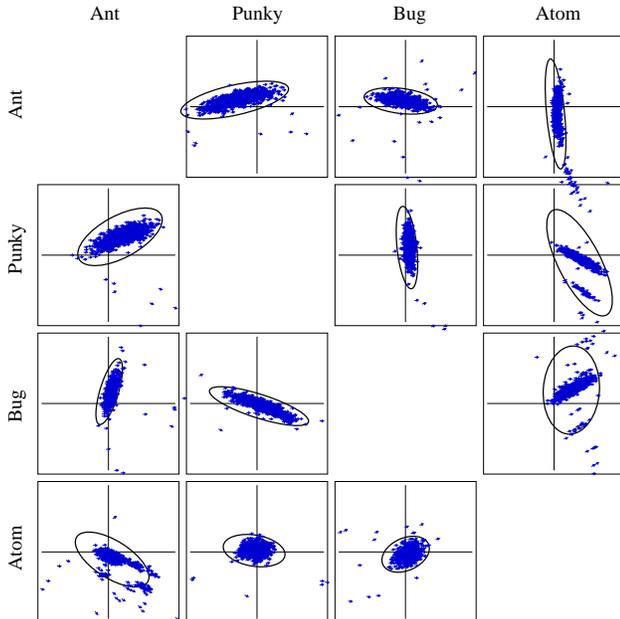


Fig. 5. A snap-shot taken at time $t = 50$ s summarizing the difference between the pose estimates and the true pose. Each row contains the estimates generated *by* one of the robots; each column contains the pose estimates generated *for* one of the robots. The particle clouds show the particle filter estimates relative to the true pose; the ellipses show the 3σ interval. Each plot is 2 m on a side.

except insofar as it determines the possible line-of-sight relationships between robots. Thus, for example, we expect localization to be less accurate in very cluttered environments, where robots will have fewer opportunities to observe one another. The particular environment chosen for this experiment permitted only intermittent robot observations.

B. Results

Figure 3 shows a set of ‘snap-shots’ taken during the experiment. Each row in this figure shows a time-series of probability distributions generated by one of the robots. The top row, for example, shows the probability distributions generated by the robot Punky for each of the robots Bug, Ant and Atom; i.e., these distributions correspond to Punky’s estimates for the relative pose of each of the other robots. The true pose of each of the robots is indicated by the square, and the distributions have been plotted in the ground-truth coordinate system in order to simplify comparison.

There are a number of features in this figure that should be noted. First, all of the pose distributions have been correctly initialized, despite the fact that no a priori pose information was provided; the distributions in these figures are generated entirely by observations made during the course of the experiment. Second, every robot is able to estimate the pose of every other robot, *including those robots that have never been observed directly*. Inspecting the top row, for example, it is apparent that the lead robot (Punky) has accurately determined the pose of the other three robots; this in spite of the fact that Punky does not make any observations itself until the final frame of this sequence. Finally, the distributions maintained by different robots are not necessarily identical; this is to be expected given that the only information shared between the robots is the set of raw observations.

Figure 5 shows a somewhat different view of the data. This is a snap-shot taken at time $t = 50$ s, with each row showing the set of estimates generated *by* one of the robots; and each column showing the set of estimates generated *for* one of the robots. Furthermore, each cell shows the *difference* between the relative pose estimate

and the true relative pose. Thus, for example, the top-right cell shows the difference between Ant's estimate of Atom's relative pose and the true relative pose of these two robots. The ellipses in Figure 5 show the 3σ intervals for each of the distributions; if our sensors models are accurate, and if the distributions are Gaussian (which they are not), this ellipse should capture the origin 99.9% of the time. The actual figure for this experiment is 98%,

VI. CONCLUSION AND FURTHER WORK

The cooperative, ego-centric localization method described in this paper has a number of attractive features: it is fully distributed, requires relatively little communication, and scales linearly with the number of robots in the team. Furthermore, the experimental results presented in Section V, while preliminary, are extremely promising.

We are currently in the process of extending this research, both experimentally and theoretically. From an experimental perspective, we are exploring cooperative localization problems involving range-only or bearing-only sensors. In principle, the approach described in this paper should work just as well for these sensors as it does for the range-bearing-and-orientation sensors described in Section V. This hypothesis, however, has yet to be validated. From a theoretical perspective, we are investigating more complete solutions to the problem of ego-centric localization, such as the one described in Section III-E. The relative merits of these (more expensive) solutions remain to be determined.

ACKNOWLEDGMENTS

This work is sponsored in part by DARPA grants DABT63-99-1-0015 and 5-39509-A (via UPenn) under the Mobile Autonomous Robot Software (MARS) program.

VII. REFERENCES

- [1] M. W. M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte, and M. Csorba. A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Transactions on Robotics and Automation*, 17(3):229–241, 2001.
- [2] T. Duckett, S. Marsland, and J. Shapiro. Learning globally consistent maps by relaxation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 4, pages 3841–3846, San Francisco, U.S.A., 2000.
- [3] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots, Special Issue on Heterogeneous Multi-Robot Systems*, 8(3):325–344, 2000.
- [4] D. Fox, W. Burgard, and S. Thrun. Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research*, 11:391–427, 1999.
- [5] B. P. Gerkey, R. T. Vaughan, K. Støy, A. Howard, G. S. Sukhatme, and M. J. Matarić. Most valuable player: A robot device server for distributed control. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS01)*, pages 1226–1231, Wailea, Hawaii, Oct. 2001.
- [6] A. Howard. Mezzanine user manual; version 0.00. Technical Report IRIS-01-416, Institute for Robotics and Intelligent Systems Technical Report, University of Southern California, 2002.
- [7] A. Howard and L. Kitchen. Cooperative localisation and mapping. In *International Conference on Field and Service Robotics (FSR99)*, pages 92–97, 1999.
- [8] R. Kurazume and S. Hirose. An experimental study of a cooperative positioning system. *Autonomous Robots*, 8(1):43–52, 2000.
- [9] J. J. Leonard and H. F. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Transactions on Robotics and Automation*, 7(3):376–382, 1991.
- [10] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
- [11] I. M. Rekleitis, G. Dudek, and E. Milios. Multi-robot exploration of an unknown environment: efficiently reducing the odometry error. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, volume 2, pages 1340–1345, 1997.
- [12] S. I. Roumeliotis and G. A. Bekey. Collective localization: a distributed Kalman filter approach. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, pages 2958–2965, San Francisco, U.S.A., 2000.
- [13] R. Simmons and S. Koenig. Probabilistic navigation in partially observable environments. In *Proceedings of International Joint Conference on Artificial Intelligence*, volume 2, pages 1080–1087, 1995.
- [14] S. Thrun. A probabilistic online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 20(5):335–363, 2001.
- [15] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localisation for mobile robots. *Machine Learning*, 31(5):29–55, 1998. Joint issue with *Autonomous Robots*.
- [16] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust Monte Carlo localization for mobile robots. *Artificial Intelligence Journal*, 128(1–2):99–141, 2001.
- [17] B. Yamauchi, A. Shultz, and W. Adams. Mobile robot exploration and map-building with continuous localization. In *Proceedings of the 1998 IEEE/RSJ International Conference on Robotics and Automation*, volume 4, pages 3175–3720, San Francisco, U.S.A., 1998.