

Mapping and Localization with RFID Technology

Dirk Hähnel Wolfram Burgard
University of Freiburg
Department of Computer Science
Freiburg, Germany

Dieter Fox
University of Washington
Computer Science and Engineering
Seattle, WA, USA

Ken Fishkin Matthai Philipose
Intel Research Seattle
Seattle, WA, USA




Fig. 1. Typical RFID tags used to label objects. The size of the tag depicted in the center is 11×5 cm.

Abstract—In this paper we analyze whether recent Radio Frequency Identification (RFID) technology can be used to improve the localization of mobile robots and persons in their environment. In particular we study the problem of localizing RFID tags with a mobile platform that is equipped with a pair of RFID antennas. We present a probabilistic measurement model for RFID readers that allow us to accurately localize RFID tags in the environment. We also demonstrate how such maps can be used to localize a robot and persons in their environment. Finally, we present experiments illustrating that the computational requirements for global robot localization can be reduced strongly by fusing RFID information with laser data.

I. INTRODUCTION

Recent advances in the field of radio frequency identification techniques have reached a state that will allow us within the next years to equip virtually every object in an environment with small, cheap Radio Frequency Identification (RFID) tags [6]. Such tags contain circuitry that gain power from radio waves emitted by readers in their vicinity. They use this power to reply their unique identifier to the reader. Figure 1 depicts three different RFID tags that were used to carry out the experiments described in this paper. The detection range of these tags is approximately 6 m.

RFID tags open up a wide variety of applications. For example, an important problem in the health-care sector is the recognition of daily activities a home patient is engaged in. The Guide project [13] uses small RFID readers worn by a person to identify the objects the person touches. The sequence of touched objects is used by a Bayesian reasoning system to estimate the activity of the person and to provide support if needed. Location context can provide important information for the interpretation of RFID readings. For example, touching

the toothpaste has very different meanings depending on whether it happens in the storage room or in the bathroom.

In this paper, we investigate how RFID technology can be enhanced by location information. We use a mobile robot equipped with RFID antennas to determine the locations of RFID tags attached to objects in an indoor environment. Figure 2 (left) depicts the robot built to carry out this research. The robot consists of an off-the-shelf Pioneer 2 robot equipped with a laser range scanner and two RFID antennas. The antennas are mounted on top of the robot and point approximately 45 degrees to the left and to the right with respect to the robot. To use these antennas for estimating the locations of objects, we first learn a sensor model that describes the likelihood of detecting an RFID tag given its location relative to one of the antennas. Since the noise of these sensors is highly non-Gaussian, we represent the measurement likelihood model by a piecewise constant approximation. Then we describe a technique to estimate the locations of RFID tags using a mobile robot equipped with RFID antennas to detect tags. This process uses a map previously learned from laser range data. We then apply Monte Carlo localization [4], [7] to estimate the pose of the robot and even of persons in this environment. Experimental results suggest that it is possible to accurately localize moving objects based on this technology. Further experiments demonstrate that RFID tags greatly reduce the time required for global localization of a mobile robot in its environment. Additionally, this technology can be used to drastically reduce the number of samples required for global localization.

This paper is organized as follows. After discussing related work we will present the sensor model for RFID receivers in Section III. Then we describe how this model can be used in combination with a laser-based FastSLAM [8] approach to effectively determine the locations of RFID tags. In Section V we describe how the resulting beliefs about the locations of the tags can be utilized to determine the position of the robot and of persons in the environment. Finally, we present experimental results illustrating the advantages of RFID tags for robot localization and person tracking.

II. RELATED WORK

In the last years RFID sensors [6] have started to enter the field of mobile robotics. Nowadays RFID readers can detect low-cost passive tags in the range of several meters. These improvements in the detection range of passive tags make this




Fig. 2. Pioneer 2 with Sick Laser Range Finder, RFID reader and two antennas (left). Experimental setup used for learning the likelihood function of tag detections (right).

technology more and more attractive for robotics applications since the information provided by tags can be used to support various tasks like navigation, localization, mapping, and even service applications such as people tracking.

Most of the applications of RFID technology, however, assume that the readers are stationary and only the tags that are attached to objects or persons move. The main focus is to trigger events if a tag is detected by a reader or entering the field of range (for example, to keep track of the contents of storage places [2]). Recently Kantor and Singh used RFID tags for mapping. Their system relies on active beacons which provide distance information based on the time required to receive the response of a tag. Additionally, the positions of the tags have to be known more or less accurately [14], [9]. Tsukiyama [16] also requires given RFID tag positions. Their system assumes perfect measurements and does not include techniques to deal with the uncertainty of the sensor.

The problem considered here is closely related to the simultaneous localization and mapping (SLAM) problem, in which a robot has to generate a map while simultaneously estimating its pose relative to this map. However, due to the limited accuracy of the RFID sensors, SLAM-techniques for range-only [14], [9], bearing-only [3] or range and bearing [5], [11], [15] cannot be applied directly to the data provided by the RFID system. Our algorithm instead uses a variant of FastSLAM [12] to learn the geometric structure of the environment using laser data [8] and then estimates the positions of the tags based on the trajectory computed by the FastSLAM algorithm.

III. LEARNING A PROBABILISTIC SENSOR MODEL FOR THE RFID ANTENNA

To localize an RFID tag in a global reference frame, we estimate the posterior $p(x | z_{1:t}, r_{1:t})$, where x is the position of the tag, $z_{1:t}$ are the observations at time steps $1, \dots, t$, and $r_{1:t}$ are the possibly different locations of the RFID antenna. According to Bayes rule and under the assumption of independence of consecutive measurements given we know the location x of a tag we obtain the following recursive update rule:

$$p(x | z_{1:t}, r_{1:t}) \propto p(z_t | x, r_t) p(x | z_{1:t-1}, r_{1:t-1}) \quad (1)$$

According to this equation, the key term is the quantity $p(z_t | x, r_t)$ which specifies the likelihood of the observation z_t given




Fig. 3. Detection field for the left (upper/green histogram) and right (lower/red histogram) antenna. The middle/blue histogram shows the area where the tag can be seen by both antennas.

the position x of the tag and the location r_t of the antenna. We make the simplifying assumption that this likelihood only depends on the *relative* offset between tag and antenna, that is, it only depends on the difference between x and r_t . The following aspects need to be considered when designing an observation model for RFID tags.

- 1) There are plenty of *false-negative readings*, i.e., situations in which the tag is not detected although it is in the vicinity of the antenna
- 2) Additionally, we obtain *false-positive readings*. In such a case the antenna detects a tag that is not in range specified by the manufacturer. This also includes detection of the RFID tag with the wrong antenna.

There are several reasons for this. For example, the orientation of the tag with respect to the RFID receiver influences the energy absorbed by its own antenna. Depending on this angle, the energy will vary and sometimes not be high enough to power the chip inside the tag. In such a case the tag will simply not respond. Furthermore, the shape and size of the detection range largely depends on the environment. For example, metal typically absorbs the energy sent by the RFID reader and therefore tags attached to metallic objects will be detected only in a short range. But even other, non-metallic objects greatly influence the detectability of tags. For example, if a tag is attached to a concrete wall its detection statistics typically changes drastically. Furthermore, the radio frequency waves emitted by the antenna can be reflected by objects such that the antenna even detects objects outside the specified detection range. Note that the observation model for the RFID antennas should be able to cover this wide range of situations and should not make the robot overly confident in the location of a particular tag or even in its own location during localization.

To determine the observation model for the RFID antennas we generated a statistics by counting frequencies. We proceeded in the following way. We attached an RFID tag to a box and rotated the robot in front of the box. We repeated this for different distances and counted for every point in a discrete




Fig. 4. Simplified sensor model for the left antenna.

grid the frequency of detections of the antenna given the tag was placed at a position covered by this grid cell relative to the robot.

The resulting histogram is shown in Figure 3. This figure contains the detection statistics for both sensors. The histograms were built from 12,822 measurements. As can be seen from the figure, both antennas show quite different behaviors although they were measuring the same RFID tag.

The resulting sensor model used to conservatively approximate the histograms depicted in Figure 3 is shown in Figure 4. This model consists of three components. The major detection range for each antenna consists of an arc with an opening angle of 95 degrees in the direction of the antenna. Additionally, an antenna always detects RFID tags that are close to it even if they are behind the antenna. This is modeled by a circular region around the center of the receiver. The corresponding likelihood for the two detection ranges are also depicted in Figure 4. For locations outside these areas we assume a constant likelihood of 0.5.

IV. MAPPING RFID TAGS

The first application of the sensor model described in the previous section is estimating the location of RFID tags in the environment using a mobile robot. To learn the positions of the tags our system proceeds in two steps. First it learns the geometric structure of the environment using a laser range sensor. Afterwards we estimate the positions of the tags based on the path of the robot.

Since our robot is equipped with a laser range scanner, we apply the FastSLAM algorithm [8] to learn the geometrical structure of the environment. The resulting map used for the experimental results is depicted in Figure 5. Given this map and the maximum likelihood path of the robot computed by the FastSLAM algorithm we can now estimate the locations of the RFID tags. Here we apply the recursive Bayesian filtering scheme given in Equation 1, with $r_{1:t}$ representing the path of the robot.

To represent the belief about the location of an RFID tag we use a set of 1000 randomly chosen positions uniformly distributed in a 25 square meter wide area around the current pose of the robot. This area is independent of the antenna that detected the tag in order to avoid that a detection failure of




Fig. 5. Map of the Intel Research Lab Seattle generated by our FastSLAM routine.

an antenna results in a suboptimal placement of the sampled positions. It is initialized at the first detection of the RFID tag by the robot.

To each of the randomly chosen potential positions we assign a numerical value storing the posterior probability $p(x | z_{1:t}, r_{1:t})$ that this position corresponds to the true position of the tag. Whenever the robot detects a tag, the posterior is updated according to Equation (1) and using the sensor model described in the previous section.

V. LOCALIZATION WITH RFID TAGS

Given the posterior distribution $p(x | z_{1:t}, r_{1:t})$ over potential positions of an RFID tag we are now ready to compute the likelihood of an observation y during localization, given the robot or a person is placed at a location l . According to the law of total probability we obtain

$$p(y | l) = \sum_x p(y | x, l) p(x | z_{1:t}, r_{1:t}). \quad (2)$$

In this equation the term $p(y | x, l)$ corresponds to the relative sensor model described in Section III. The relative offset of the sensor is computed from the global coordinates of the detected RFID tag, x , and the robot position, l . Thus, to determine the likelihood of a tag detection given the robot is at location l , we have to integrate over the posterior probability of the tag's location given the data obtained during the mapping process.

To estimate the pose l of the robot or of persons in the environment, we apply the well-known recursive Bayesian filtering scheme:

$$p(l_t | y_{1:t}, u_{0:t-1}) = \alpha \cdot p(y_t | l_t) \cdot \int_{l'_{t-1}} p(l_t | u_{t-1}, l'_{t-1}) \cdot p(l'_{t-1} | y_{1:t-1}, u_{0:t-2}) d l'_{t-1} \quad (3)$$

Here α is a normalization constant ensuring that $p(l_t | y_{1:t}, u_{0:t-1})$ sums up to one over all l_t . The term $p(l_t | u_{t-1}, l'_{t-1})$ describes the probability that the object is at position l_t given it executed the movement u_{t-1} at position l'_{t-1} . This quantity is computed depending on the object we are tracking. In the case of the robot we compute this




Fig. 6. Evolution of the posterior about the localization of an RFID tag over time. The width of the circles represents the importance weight of the corresponding particle. It is drawn proportional to the ratio between the importance weights of the particular sample and the maximum likelihood sample.

quantity based on the odometry measurements [7]. If we are tracking persons, we simply represent this density by a Gaussian centered around l_t . Furthermore, the quantity $p(y_t | l_t)$ denotes the likelihood of the observation y_t according to our observation model, which is computed using Equation (2). To represent the posterior about the pose of the object being tracked we apply Monte-Carlo localization [4], [7]. In Monte-Carlo localization, the belief of the robot is represented by a set of random samples [1]. Each sample consists of a state vector of the underlying system, which is the pose l of the robot in our case, and a weighing factor ω . The latter is used to store the importance of the corresponding particle. The posterior is represented by the distribution of the samples and their importance factors. The particle filter algorithm used by our system is also known as *sequential importance sampling with resampling* [1]. It updates the belief about the pose of the robot according to the following two alternating steps:

- 1) In the prediction step, we draw for each sample a new sample according to the weight of the sample and according to the model $p(l_t | u_{t-1}, l_{t-1})$ of the robot's dynamics given the movement u_{t-1} executed since the previous update. In the case of localizing a person, this model is simply a Gaussian centered at l_{t-1} .
- 2) In the correction step, the new observation y_t is integrated into the sample set. This is done by bootstrap resampling, where each sample is weighted according to the observation likelihood $p(y_t | l_t)$.

To globally localize the object, we initialize the particle set with a uniform distribution. In the case of RFID sensors, we fortunately can efficiently sample potential locations of the object. We simply place samples only in the potential detection range of the RFID sensor. Such an approach has been applied successfully in the past, for example by Lenser et al. [10].

VI. EXPERIMENTAL RESULTS

Our approach described above has been implemented and tested using a Pioneer 2 robot equipped with a SICK LMS laser range-finder and an Alien Technology's 915 MHz RFID reader with two circularly polarized antennas (see left image of Figure 2). The experiments described here were carried out in the Intel Research Lab, Seattle, WA. Figure 5 shows a two-dimensional occupancy grid map generated with our FastSLAM routine. The size of the environment is 28m by 28m. We installed 100 tags in this environment (see Figure 7). The tags were of the types depicted in Figure 1 and all of them




Fig. 7. RFID tags attached to walls.

were able to communicate with the robot. Most of them were installed along the circular corridor of the environment.

A. Mapping RFID tags

As already mentioned above, we use the trajectory estimated by our FastSLAM routine to determine the posterior about the locations of the tags. When a tag is detected for the first time, we initialize a discrete set of randomly chosen points around the robot and use a uniform distribution to initialize the belief. Whenever a tag is detected, the posterior probability of each sample in that set is multiplied with the likelihood of the observation given the tag is at the position corresponding to that sample. Afterwards we normalize the belief over all samples.

Figure 6 shows a typical example for the evolution of the belief of an RFID tag. The leftmost image shows the initial sample set after the first detection of an RFID tag. The remaining images illustrate how the belief focuses on the true position of the tag as more measurements are obtained. They show the corresponding beliefs after 6, 17, 65, and 200 measurements. Note that the diameter of each circle representing a particle corresponds to its importance weight. As can be seen from the figure, the belief quickly converges to a unimodal distribution. Note that this is not necessarily the case. In principal, our representation can also handle ambiguities in which the location of an RFID tag cannot be determined uniquely, for example, because the robot cannot reach locations which are required to resolve the ambiguity.

Figure 8 depicts the positions of the robot when it detected the tag, for which the beliefs are plotted in Figure 6. Detections of the right antenna are displayed by filled circles and for each




Fig. 8. Places where the robot has detected the RFID tag with the left (unfilled circle) or right antenna (filled circle)

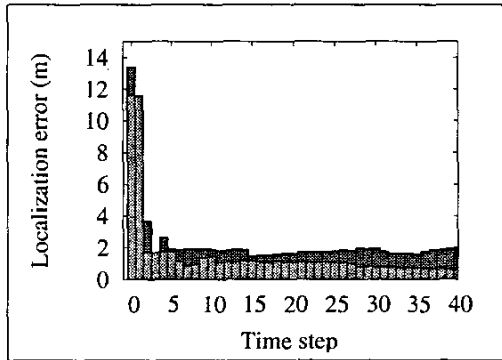


Fig. 9. Error (in m) during global localization with (green or light grey) and without (red or dark grey) odometry using RFID tags only.

detection of the left antenna we draw an unfilled circle. As can be seen from the figure, the measurement noise is quite high and there are several false detections. Nevertheless, our algorithm is able to accurately localize the tag at the wall close to the entrance.

After traveling 791.93m with an average speed of 0.225m/s the robot had processed 50,933 detections of RFID tags. The resulting map of the tags (at their maximum likelihood position) is shown in Figure 11 (left). Thus, our sensor model allows to learn the positions multiple tags in a standard office environment.

B. Localization with RFID Tags

The next set of experiments is designed to illustrate that the RFID map generated in the previous step can be used to localize the robot and even persons equipped with RFID antennas.

In the first experiment we steered the robot through the environment and applied Monte-Carlo localization to globally estimate the position of the vehicle. To simulate the situation in which we localize a person instead of the robot we simply ignored the odometry information and changed the motion model in the Monte Carlo localization procedure. As already mentioned above we used a standard motion model [7] to estimate the pose of the robot. In order to localize and keep track of a person we simply replaced this motion model by a Gaussian distribution centered around the current pose. Note




Fig. 10. Positioning error of the laser based global localization (in m) without (red or dark grey) and with (green or light grey) RFID data.

that this is only a rough approximation of the motions of a person. Better models therefore can be expected to result in more accurate estimates.

Figure 9 shows the localization error during a global localization run using RFID tags only. The two plots show the localization error for global localization without odometry (red/dark grey) and with odometry (green/light grey).

The center image of Figure 11 shows the trajectory for the object being tracked when no odometry information is used. The corresponding ground-truth obtained by laser-based localization is depicted in the right image of the same figure. As can be seen, even with such noisy sensors the estimated trajectory is quite close to the ground truth.

C. Improving Global Localization with RFID Tags

The final experiment is designed to illustrate that the RFID technology can be used to drastically improve the global localization performance even in the case where highly accurate sensors such as laser range finders are used. To analyze this we used a pre-recorded data set to figure out how efficiently the robot can determine its global position in this map. Since the RFID tags are only placed close to the corridor we generated samples only in the corridor of the environment. We compared the time required for global localization using laser data with the time needed when laser and RFID tags were used simultaneously. Figure 10 shows the average localization error for a typical run for both cases. As the figure illustrates, global localization can be achieved much faster when laser and RFID data are combined (green/light grey) compared to a situation in which only laser data is used (red/dark grey).

Additionally, the use of RFID sensors can greatly reduce the number of samples required for global localization. Figure 12 shows the localization error depending on the number of particles for the case in which only laser data is used as well as for the situation in which the laser data is combined with RFID information. It turns out that laser-based global localization is efficient when at least 10,000 particles are used. On the other hand, if we fuse the laser data with the information about the RFID tags, we can globally localize the object with as few as 50 samples.




Fig. 11. Map of Intel Lab with most likely positions of the RFID tags (left), estimated trajectory (without odometry) (center) and the corresponding ground truth (right).




Fig. 12. Localization error (in m) during global localization for different numbers of particles and depending on whether only laser data is used (left image) or whether the combination of laser data and RFID measurements is used (right image).

VII. CONCLUSIONS

In this paper we presented an approach to generate maps of RFID tags with mobile robots. We presented a sensor model that allows us to compute the likelihood of tag detections given the relative position of the tag with respect to the robot. Additionally we described how to compute a posterior about the position of a tag after the trajectory and the map has been generated with a highly accurate FastSLAM algorithm for laser range scans. We furthermore present how the posterior can be used to localize a robot and persons in the environment.

The system has been implemented on a Pioneer 2 robot that was augmented by two RFID antennas. In practical experiments we demonstrated that the system can build accurate maps of RFID tags. We furthermore illustrated that the resulting maps can be used for accurate localization of the robot and moving objects without odometry information. Finally we presented an experiment demonstrating that the combination of a laser-range scanner and RFID technology can greatly reduce the computational demands for the global localization of a moving mobile robot.

ACKNOWLEDGMENTS

This work has partly been supported by the German Science Foundation (DFG) under contract number SFB/TR8-03 and by the EC under contract number IST-2000-29456.

REFERENCES

- [1] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, 2002.
- [2] J. Brusey, M. Harrison, Ch. Floerkemeier, and M. Fletcher. Reasoning about uncertainty in location identification with RFID. In *IJCAI-2003 Workshop on Reasoning with Uncertainty in Robotics*, 2003.
- [3] M. Deans and M. Herbert. Experimental comparison of techniques for localization and mapping using bearing-only sensor. In *Seventh Int. Symp. on Experimental Robotics*, 2000.
- [4] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte Carlo localization for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999.
- [5] G. Dissanayake, H. Durrant-Whyte, and T. Bailey. A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. In *ICRA'2000 Workshop on Mobile Robot Navigation and Mapping*, 2000.
- [6] Klaus Finkenzeller. *RFID Handbook: Radio-Frequency Identification Fundamentals and Applications*. Wiley, New York, 2000.
- [7] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte Carlo localization: Efficient position estimation for mobile robots. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 1999.
- [8] D. Hähnel, W. Burgard, D. Fox, and S. Thrun. An efficient fastslam algorithm for generating maps of large-scale cyclic environments from raw laser range measurements. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.
- [9] George A Kantor and Sanjiv Singh. Preliminary results in range-only localization and mapping. In *Proceedings of the IEEE Conference on Robotics and Automation (ICRA)*, 2002.
- [10] S. Lenser and M. Veloso. Sensor resetting localization for poorly modelled mobile robots. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2000.
- [11] J.J. Leonard and H.J.S. Feder. A computationally efficient method for large-scale concurrent mapping and localization. In *Proc. of the Ninth Int. Symp. on Robotics Research (ISRR)*, 1999.
- [12] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, Edmonton, Canada, 2002. AAAI.
- [13] M. Philipose, K. Fishkin, D. Fox, H. Kautz, D. Patterson, and M. Perkowitz. Guide: Towards understanding daily life via auto-identification and statistical analysis. In *Proc. of the Int. Workshop on Ubiquitous Computing for Pervasive Healthcare Applications (Ub-health)*, 2003.
- [14] Sanjiv Singh, George Kantor, and Dennis Strelow. Recent results in extensions to simultaneous localization and mapping. In *International Symposium on Experimental Robotics*, 2002.
- [15] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning and Autonomous Robots (joint issue)*, 1998.
- [16] T. Tsukiyama. Navigation system for mobile robots using RFID tags. In *Proceedings of the International Conference on Advanced Robotics (ICAR)*, 2003.