A Robust Method of Localization and Mapping Using Only Range

Joseph Djugash, and Sanjiv Singh

The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA
{josephad, ssingh}@ri.cmu.edu

Summary. In this paper we present results in mobile robot localization and simultaneous localization and mapping (SLAM) using range from radio. In previous work we have shown how range readings from radio tags placed in the environment can be used to localize a robot and map tag locations using a standard extended Kalman filter (EKF) that linearizes the probability distribution due to range measurements based on prior estimates. Our experience with this method was that the filter could perform poorly and even diverge in cases of missing data and poor initialization. Here we present a new formulation that gains robustness without sacrificing accuracy. Specifically, our method is shown to have significantly better performance with poor and even no initialization, infrequent measurements, and incorrect data association. We present results from a mobile robot equipped with high accuracy ground truth, operating over several kilometers.

1 Motivation and Problem Statement

Here we focus on the problems of robust localization and SLAM given range data between a mobile robot and static "tags". In the first case, the locations of the tags are known and the robot must accurately localize itself as it measures range to these tags while it moves. In the second case, the location of the tags are not known ahead of time and must be determined in addition to being used for localization. Hence our goal is to achieve a reliable and accurate estimate of pose for a robot operating in an environment containing radio tags to which the robot can measure range. This is a variant of a well known problem of SLAM for a robot using both range and bearing data to both map an environment and localize itself.

Specifically, we would like a method that is robust to poor initialization, noise in measurement and missing data. We use three different datasets with over 2,000 range measurements each from a robot moving in a field of radio tags to compare the results from the proposed method to the classical standard EKF formulation. In each of the datasets, the robot travels over several km

and has highly accurate positioning for groundtruth. We find that our new method is more efficient, accurate and robust than the current standard.

2 Related Work

Most landmark-based localization systems use sensors that measure relative bearing or in some cases both range and bearing to distinct features in the environment. Of particular interest to us are those that use only range to localize the network. For instance, the RADAR system, developed by Bahl and Padmanabhan, utilizes signal strength of packets in the commonly available 802.11b wireless networks for localization of network devices [1]. However, signal strength measurements are often erratic and can be affected by slight changes in the environment. Alternately, the Cricket system uses fixed ultrasound emitters and embedded receivers in the target object to localize the target [2], [3]. Our system uses radio frequencies to measure range– as the robot moves, a transponder periodically sends out a query, and any tags that hear the query respond by sending a reply along with an unique ID number (trivially solving the correspondence problem). The robot can then estimate the distance to each responding tag by determining the time lapsed between sending the query and receiving the response.

Among others, Kantor et al. and Kurth et al. present single filter formulations to combine range measurements with dead reckoning and inertial measurements [4], [5]. These methods formulate the problem as an Extended Kalman Filter (EKF) using linearization in the Euclidean space. Furthermore, these early efforts generally assume that the tag locations were approximately known *a priori.* As can be expected with this formulation, the EKF solution is prone to problems of linearization and multi-modality, two issues particularly common when dealing with range only measurements. In some cases, relatively modest errors in the prior estimate of the tags and noisy measurements can result in an incorrect estimate that will lead the filter to diverge. These issues have been dealt with in the fashion of typical EKF solutions: by assuming that the initial conditions (of the robot and the tag locations) specified are "pretty good" and that the noise in the measurements to be "normal". Such assumptions are limiting and there is a need for a principled and robust method to use range data for localization and SLAM.

One such effort was presented by Leonard et al., in their use of a delayed state filter [6]. This approach does not require any prior information about the locations of the tags. Instead, a batch process on a set of collected measurements within a single EKF update step and allow the collected measurements to identify the linearization point. The main drawback of such an approach is the inability to determine when sufficient measurements have been collected to safely linearize. More specifically, if the batch update chooses the incorrect of the multiple hypotheses due to the lack of sufficient measurements, it is near impossible for the filter to recover from that error. Alternately, Olson et

al. proposed the use of a separate pre-filtering step to approximately locate the landmarks/tags and identify the linearization point [7]. The pre-filtering step involves the use of a hough transform that identifies peaks within the likelihood distribution for a tag, given a collection of measurements. By delaying the initialization of a tag until the hough transform identifies a single peak with a significantly higher likelihood, the method is able to determine, in an online manner, the exact size of the measurement set needed to achieve a good initialization. While this approach seems adequate at first glance, in the presence of large measurement/odometric error and sparse data the peaks within the hough transform become blurred and form a "ridge". This blurring forces the initialization of the tag to be delayed, and with sufficient odometry error it is possible that the distribution never converges to a single peak.

Previously, we have also investigated the use of other filters to address the problem of range-only SLAM [8], [9]. One such method, Monte Carlo localization, or particle filtering, provides a method of representing multimodal distributions for position estimation, with the advantage that the computational requirements can be scaled. However, these sample based methods become intractable (for real-time applications) when modeling the cross-correlation terms sometimes necessary in SLAM. If the particle count in reduced to maintain real-time computation, particle depletion then causes the filter to diverge.

While majority of prior research has focused on formulating the problem in the Euclidean space, we propose the use of a polar parameterization to more accurately represent the nonlinear distributions encountered in the range-only SLAM. This parameterization is similar to one proposed by Funiak et al. [10] for a very different problem– tracking targets from camera networks where the locations and orientations of the cameras are not known. They demonstrate the use of polar parameterization to model the annulus-like distributions that occur while estimating the pose of a camera observing a mobile target even with bearing-only measurements. Civera et al. also address a similar problem in monocular SLAM where they utilize a hybrid estimation scheme to deal with the nonlinear distributions that occur when the camera is either stationary or rotating for significant period of time [11].

Unlike much of the existing methods that linearize the measurements in the Euclidean space, the proposed method operates in the polar space where the linearization is much improved. As such, our method does not require a prior nor does it perform a batch process to identify the linearization point. In addition, the proposed method is also able to model the multi-modal distributions that naturally occur with range measurements.

3 Technical Approach

In the following sections, we present the formulation of the proposed method and compare its results to the standard EKF that has been previously tested and proven. First, we present a quick overview of the standard EKF, which

linearizes the range measurements in the Euclidean space. Then, we detail our proposed method, the relative-over parametrized (ROP) EKF, using the new parametrization tailored to represent range data.

The basic intuition behind the ROP-EKF is that the distributions resulting from range measurement can be more faithfully represented if the linearization of the range measurements is done in polar coordinates instead of Cartesian coordinates. In addition since multiple modes naturally arise in the distributions after multiple measurements, our method includes means for maintaining multiple hypotheses. Lastly, we compare the results of the proposed method against the standard EKF, using real-world experimental data (Section 4) in the area of the size of a football field.

3.1 The Standard EKF

The standard EKF is formulated to estimate the position of the mobile robot given measurements of odometry (from wheel encoders), heading change (from a gyro), and range measurements to stationary tags in the environment. Odometry and gyro measurements are used in the state propagation or the prediction step, while the range measurements are incorporated in the correction step. The method presented here is used as a baseline comparison for our proposed method described in Section 3.2.

Let the robot state (position and orientation) at time k be represented by the state vector $q_k = [x_k^r, y_k^r, \phi_k^r]^T$. The dynamics of the wheeled robot used in the experiments are best described by the following set of nonlinear equations:

$$q_{k+1} = \begin{bmatrix} x_k^r + \triangle D_k \cos(\phi_k^r) \\ y_k^r + \triangle D_k \sin(\phi_k^r) \\ \phi_k^r + \triangle \phi_k \end{bmatrix} + \nu_k = f(\widehat{q}_k, u_k) + \nu_k, \tag{1}$$

where ν_k is a noise vector, $\triangle D_k$ is the odometric distance traveled, and $\triangle \phi_k$ is the orientation change. For every new control input vector, $u(k) = [\triangle D_k, \triangle \phi_k]^T$, that is received, the estimates of robot state and error covariance are propagated using the extended Kalman filter (EKF) prediction equations.

The range measurement received at time k is modeled by:

$$\hat{r}_k^b = \sqrt{(x_k^r - x_k^b)^2 + (y_k^r - y_k^b)^2} + \eta_k \tag{2}$$

where, $\hat{r}_k^b$ is the estimate of the range from the tag to the current state, (x_k^b, y_k^b) is the location of the tag from which the measurement was received and η_k is zero mean Gaussian noise. The measurement is linearized and incorporated into the state and covariance estimates using the EKF update equations.

When performing SLAM, we extend the state vector to include the position estimates of each tag. Thus we get,

$$q_k = \begin{bmatrix} x_k^r, \, y_k^r, \, \phi_k^r, \, x_k^1, \, y_k^1, \, ... \, x_k^N, \, y_k^N \end{bmatrix}^T \tag{3}$$
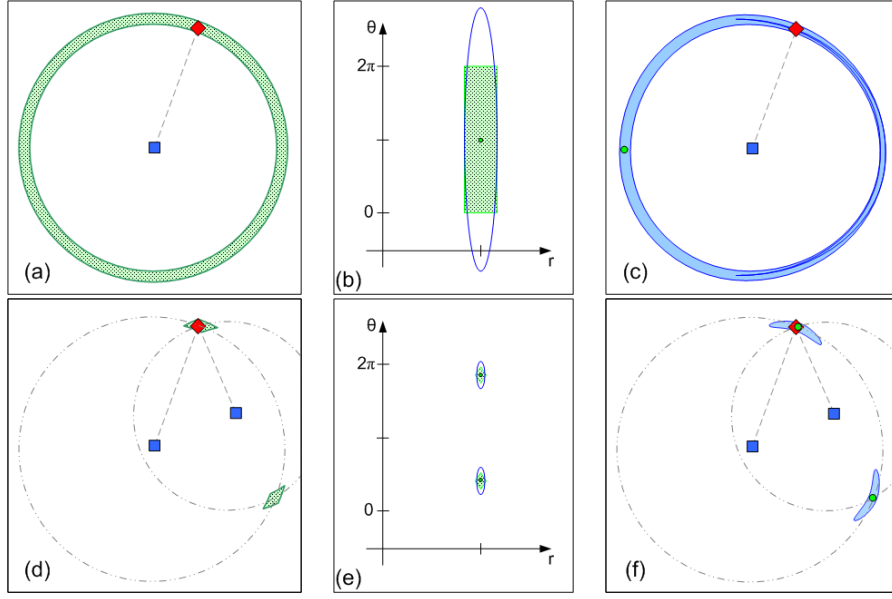
Fig. 1. True and approximated distributions with one range observation (Top Row) and two observations from different locations (Bottom Row). True distributions in Cartesian coordinates (Left Column), true and Gaussian approximated distributions in polar coordinates (Middle Column), and projection of the Gaussian approximated distribution into Cartesian coordinates. In these figures, blue squares represent *observing* tags, whose location is known. Red diamonds represent the true location of the *observed* tag, whose position is being estimated and green circles represent the mean(s) for each mode of the estimated tag. Green shaded regions represent the true uncertainty distribution and blue ellipses represent the estimated uncertainty distribution. The dashed gray lines and circles represent the observed range measurements.

where n is the number of initialized tags at time k. The motion model for the first three states is given by (1). The tags do not move, so the motion model for the remaining states is trivial. The measurement model is the same as that stated in (2) with x_k^b and y_k^b replaced by x_k^i and y_k^i when the measurement is received from the i^{th} tag.

When no prior knowledge of the tag locations exist, it is necessary to perform some type of batch process to initialize tag locations, which are then refined within the filter. We employ an approach similar to [7] that utilizes a two-dimensional probability grid to provide initial estimates of the tag locations. Here a "voting" scheme is used to identify likely locations of the tags, the cells with the greatest number of votes.

3.2 ROP-EKF: Relative-Over Parameterization

While the standard EKF described above operates in the Euclidean space, we formulate our problem in polar coordinates. At each time step, k, the state of tag i is represented by $q_k^i = [c_{x,k}^i, c_{y,k}^i, r_k^i, \theta_k^i]^T$. Each tag's estimate is represented in a polar coordinates, where $(c_{x,k}^i, c_{y,k}^i)$ are the center of the polar coordinate frame and (r_k^i, θ_k^i) are the corresponding range and angle values. The use of this parameterization derives motivation from the polar coordinate system, where annuli, crescents and other ring-like shapes can be easily modeled. In addition to the four variable relative-over parameterization, for the mobile robot, an addition term that represents the current heading of the robot, ϕ_k^r, is maintained. Thus, the complete state vector at time k is represented as:

$$q_k = [q_k^r, \phi_k^r q_k^1, q_k^2, ..., q_k^N]^T.$$

where N is the total number of tags. At each time step, we get some set of motion and range observations, u_k and z_k respectively.

While the motion model remains the same as described in (1), only the $(c_{x,k}^r, c_{y,k}^r)$ terms of the robot's state are updated. The remaining terms, (r_k^r, θ_k^r), are left unchanged by the motion update. The measurement model, however, requires a slight change to equation 2. In order to compute the expected range using the (x_k^i, y_k^i) of the robot/tag, we need to first transform their positions from the polar coordinates into the Cartesian coordinates:

$$x_k^i = c_{x,k}^i + r_k^i \cdot \cos(\theta_k^i).$$
$$y_k^i = c_{y,k}^i + r_k^i \cdot \sin(\theta_k^i). \tag{4}$$

The range measurements can now be incorporated into the state and covariance estimates using the EKF update equations. The measurement matrix for the polar parameterization, used in the EKF update, is given by the Jacobian H:

$$H = \frac{\partial \hat{r}^i}{\partial q^r} = \begin{bmatrix} \frac{(x^r - x^i)}{\hat{z}_{ri}} \\ \frac{(y^r - y^i)}{\hat{z}_{ri}} \\ \cos(\theta^r)\frac{(x^r - x^i)}{\hat{z}_{ri}} + \sin(\theta^r)\frac{(y^r - y^i)}{\hat{z}_{ri}} \\ r^r(\cos(\theta^r)\frac{(y^r - y^i)}{\hat{z}_{ri}} - \sin(\theta^r)\frac{(x^r - x^i)}{\hat{z}_{ri}}) \end{bmatrix} \tag{5}$$

$$\hat{z}_{ri} = \sqrt{(x^r - x^i)^2 + (y^r - y^i)^2}$$

Unlike the standard EKF formulation, the relative-over parameterized EKF does not require a separate batch process to initialize the tags within the filter. Instead, upon the first observation of a particular tag, the annulus of the true distribution is approximated by an elongated Gaussian in polar coordinates ($r\theta$-space). This Gaussian approximation is given an arbitrary mean in θ (within the range $[0, 2\pi)$) with a large variance term, such that the

probability along the θ dimension is near uniform. Figure 1(a-c) reveals the true distributions (green shaded regions) given a single measurement and the corresponding Gaussian ellipse (blue ellipse) depicted in both xy-space and $r\theta$-space.

Next, when a new range observation intersects the annulus at two distinct locations, a multi-modal distribution with two distinct modes (peaks/local maxima in the distribution) is formed. In order to properly deal with multi-modal distributions, we extend our approach to a multi-hypothesis representation using a separate filter/hypothesis for each mode. To elaborate, whenever an annulus is split into separate modes, we simply duplicate the filter and adjust the mean of each filter to represent the two distinct intersection points. Then, by performing a measurement update using the new mean, we are able to appropriately update the covariance terms within the filter. The simple case of splitting a single annulus into two separate modes given a new range observation is depicted in Figure 1(d-f). After two measurements, position estimate of a given tag collapses to two small regions/ellipses within the proposed method. By modeling such multi-modal distributions using the polar parameterization, the proposed method smoothly represents the evolution of the uncertainty distribution. From the time they start as annuli to when they split into multiple modes and till they converge to a single unimodal estimate, the filter accurately represents each tag's distribution.

4 Experiments

We tested our method on range data collected from a mobile robot that traversed long paths in an outdoor field with several stationary radio tags. The objective is for the robot to localize itself using only its own odometry and range measured to radio tags placed in the environment. The initial locations of the RF tags was not known. Range data and odometry, along with high accuracy GPS groundtruth were logged. The log files are then processed by each of the two methods to compare the results.

4.1 Experimental Setup

The system used to perform the experiments was an instrumented autonomous robot with highly accurate (2cm) positioning for ground truth using RTK GPS receivers as well as a fiber optic gyro and wheel encoders. Ground truth position is recorded at 100 Hz.

In our experiments, we use a radio tag system (Pinpoint) [12] to measure range between stationary RF tags (see Figure 2) and a moving transponder. The radio transponder with its 4 antennae are mounted on the four corners of the robot. The transponders send a "chirp", and any tag that receives that signal responds with its unique ID. The range from the transponder to the

tag is estimated based on the time elapsed between the transmitted "chirp" and the received response.

The robot was also equipped with a computer that controlled the tag queries and processes their responses. For each tag response, the system produces a time-stamped distance estimate to the responding tag, along with the unique ID number of that tag and the ID of the antenna that received the response. For experimentation, the RF tags are placed atop traffic cones approximately 46cm above the ground. Figure 2 shows plots from an example dataset of the mean measurements against true ranges using these RF tags and their associated variances. The robot was driven on a flat, grassy area football field. We distributed 7 RF tags throughout the area, and then the robot retraced the path among the tags guided by RTK GPS, traveling several kilometers and observing range measurements to the radio tags.

With this setup, we collected three kinds of data: the groundtruth path of the robot from GPS and inertial sensors, the path from dead reckoning, and the range measurements to the RF tags. The path from dead reckoning is computed by integrating over time incremental measurements of change in the robot's heading from a KVH gyro (with a drift rate of $30 \deg /hr$) and incremental distance traveled measurements from the wheel encoders. Figure 3 shows the three different datasets used to evaluate our proposed method. The datasets were designed to present a variety of robot paths, each with a distinct dead reckoning drift. Dataset #1 presents a path which monotonically increases its heading error due to repeatedly turning in the same direction. In contrast dataset #2 presents a path that minimizes the effect of heading error by balancing the number of left turns with an equal number of right turns. Finally, dataset #3 highlights a much longer trial where the robot was driven in a random manner.
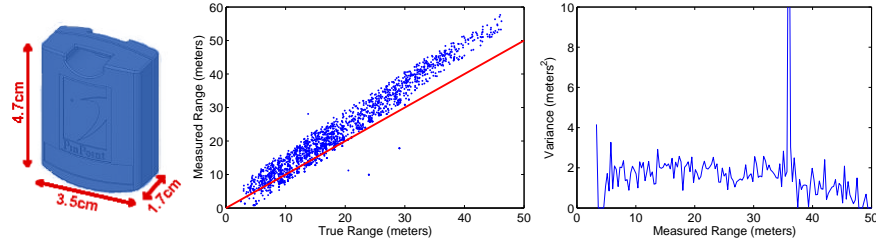

Fig. 2. (Left) Depicts the RF tags from Pinpoint used in our experiments. (Middle) A set of measured ranges plotted against the true measurements. The solid line represents the ideal case when the measured range equals the true range. (Right) the variance associated with various measured ranges. The particular characteristics of the sensor and the non-uniform variance (uncertainty) in the range measurements can be observed.
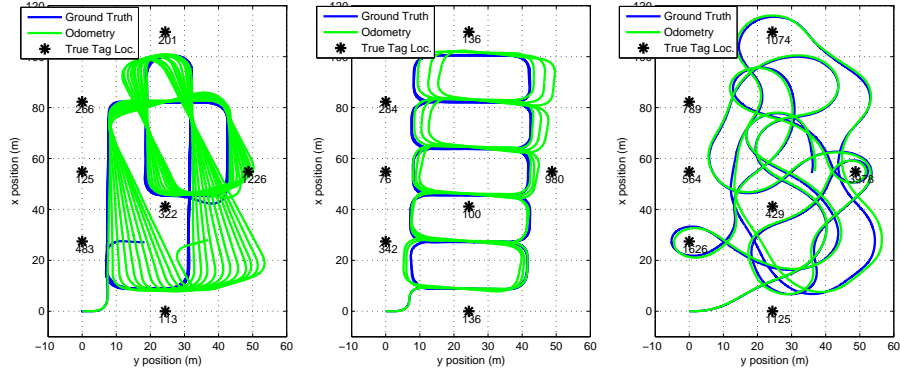
Fig. 3. The robot's true path is shown in blue along with its dead reckoning position in green and the true tag locations for three different datasets. The numbers next to each tag presents the number of measurements received by the robot from the tag. (Left) Dataset #1: the robot traveled 3.7 km receiving 2749 range measurements. (Middle) Dataset #2: the robot traveled 1.36 km receiving 2146 range measurements. (Right) Dataset #3: the robot traveled 6.7 km receiving 10068 range measurements. To reduce the clutter only the initial 1km of the path is shown, however, the full dataset is evaluated and compared numerically.

4.2 Results

Table 1 shows robot path and tag mapping errors from performing localization and SLAM with the two different methods described above. The table reveals that for localization the proposed ROP-EKF performs better. While in SLAM on dataset #1 it produces a higher robot path error, the tag mapping errors are much improved over the standard EKF. Additionally, it should be noted that when comparing the SLAM results it is best to observe errors in the final segment (we report the path errors over the final 10%) of the robot path, instead of the full robot path to illustrate the performance after the tag locations have converged. Note that heuristically we stop updating the location of the tags once their uncertainty falls below a threshold (0.3 m in the experiments described). This has the effect of only updating the position of the robot and the other tags (via the covariances) when range measurements are received. Figure 4 shows the output of the proposed ROP-EKF SLAM on each of the three datasets along with the robot path error over time.

The raw result from SLAM does not produce a globally aligned solution, since no reference coordinate frame is provided apriori to fix the map (location of the tags). Hence a simple affine transform (consisting of only a rotational and translational components) needs to be performed on both the map and path estimates. The transform that is applied, is one that best aligns the estimated map to the real ground truth surveyed map. This produces a reasonable metric by which we can evaluate the accuracy of the SLAM solution.

Table 1. Mean robot path and tag mapping errors (meters) in localization and SLAM with both standard EKF & ROP-EKF.

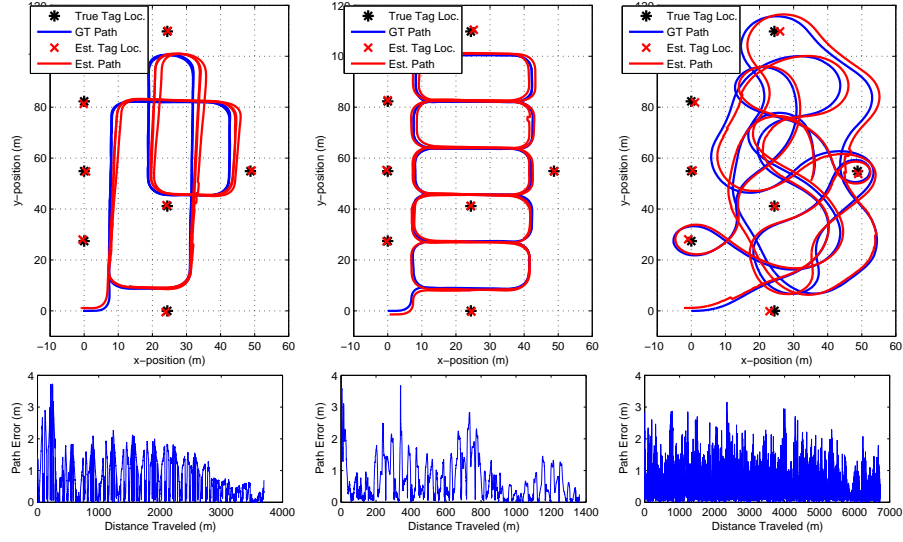| Method	Localization	SLAM		tag map error	
	Path error over entire path	Path error over last 10% of path	Path error over entire path	μ	σ
Dataset #1					
Std. EKF	0.97m	0.43m	0.91m	0.87m	0.41m
ROP-EKF	0.71m	0.36m	0.97m	0.57m	0.24m
Dataset #2					
Std. EKF	0.52m	0.79m	1.03m	0.83m	0.58m
ROP-EKF	0.33m	0.42m	0.78m	0.53m	0.41m
Dataset #3					
Std. EKF	0.63m	0.71m	0.94m	0.98m	0.36m
ROP-EKF	0.37m	0.47m	0.83m	0.62m	0.27m

Fig. 4. (Top) ROP-EKF SLAM results for the three datasets. The plot of the first $1km$ of the robot's path and the final estimated tag locations are shown along with the ground truth. The location of the tags is completely unknown at the start of SLAM. The path and tag estimates shown include an affine transform that re-aligned the local solution into the global coordinate frame. (Bottom) The corresponding robot path errors over time.

Initialization

As mentioned above, the use of standard EKF requires good initialization. In the localization case, this means that the starting estimate of the robot can't be too far from its true location. In the SLAM case, a batch process is necessary to approximately locate the tags before the standard EKF can be used. We find that the need for initialization in both these cases can be relaxed significantly with the use of the ROP-EKF. Figure 5 presents the results of localization with good ($4.25m$ translational and *no* rotational offsets) and bad ($3m$ translational and $20°$ rotational offsets) initialization with the standard EKF along with the un-initialized localization result of ROP-EKF. In both the initialization cases, the standard EKF was initialized with $2m$ standard deviation in (x, y) and $2°$ standard deviation in orientation. As can be observed from the figure, the ROP-EKF is able to accurately estimate the robot path without any prior information about the robot's starting position. The initial mean of the robot's pose (indicated by the green circle in the figure) is simply the mean of the annulus formed by the first measurement from a tag. As additional measurements are observed, this annulus is collapsed into a smaller Gaussian that accurately depicts the robot's position.
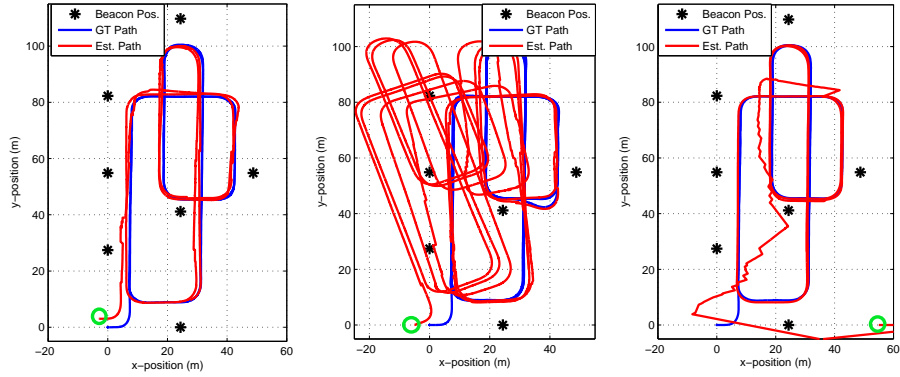

Fig. 5. Robot path estimates and ground truth robot path and tag locations. The green circles indicate the estimate of the robot's pose at the start of the run. In the standard EKF the initial pose was set manually, while in the ROP-EKF it was automatically initialized by the filter upon receiving the first measurement. (Left) Localization result of standard EKF using good initialization of the robot's pose: $4.25m$ translational and *no* rotation offsets. (Middle) Localization result of standard EKF using slightly bad initialization of the robot's pose: $3m$ translational and $20°$ rotational offsets. (Right) Localization result of ROP-EKF without any initialization of the robot's pose.

Bad Correspondence

While our sensor system trivially solves the correspondence problem by sending an unique ID along with the transmitted signal, we show how this method performs even if the data association is significantly degraded by artificially corrupting data association. Figure 6 presents the tag mapping errors of SLAM using both the methods described above for increasing percentages of incorrect data association. We note that even with severe errors in data association, the ROP-EFK maps tags with relatively high accuracy. Figure 7 shows that the path error of the two methods is comparable when the data association is perfect but that the ROP-EKF is significantly better in estimating paths in the presence of significant data association error.
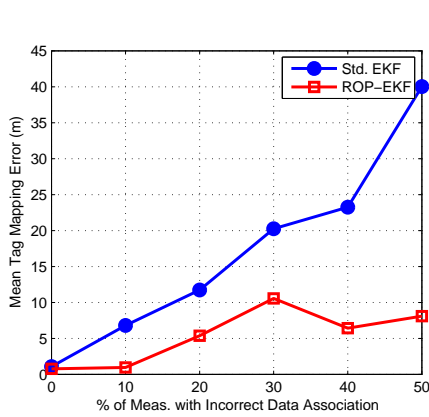

Fig. 6. Mean tag mapping errors of SLAM using both the methods described above for increasing amount of measurements with incorrect data associations.
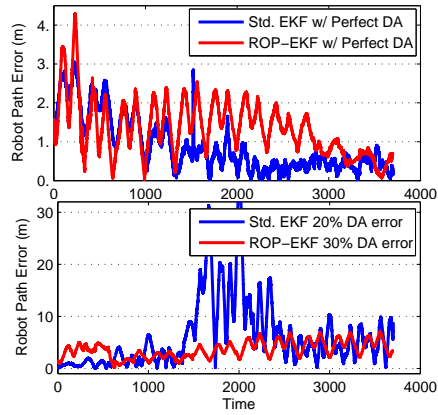

Fig. 7. Mean path error over time for two cases of SLAM (Top) perfect data association (DA) for both methods. (Bottom) 20% incorrect DA with Std. EKF in blue and 30% incorrect DA with ROP-EKF in red.

Sparse Data

In the final evaluation, we examined the performance with a varying amount of input data. We do this by randomly dropping a percentage of the range measurements. Figure 8 shows that the ROP-EKF consistently maps the tags with higher accuracy (50-100%) than the standard EKF. With extremely sparse data ($\leq 30\%$ of the dataset), both methods performs poorly. This is because with few measurements, postions of tags are underconstrained and the robot's estimate of its own position drifts with odometric errors.
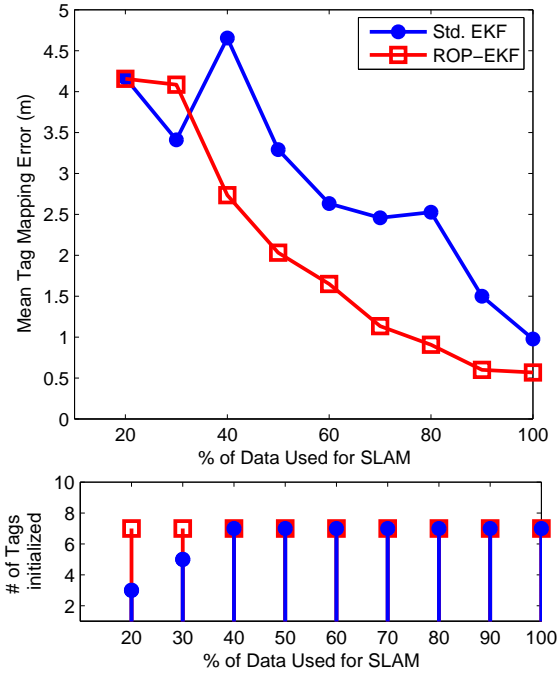
Fig. 8. (Top) Mean tag mapping errors of SLAM using the two methods described above for a varying percentage of data used by the filter. We do this by randomly dropping a percentage of the total 2749 range measurements. (Bottom) Shows the number of tags initialized by each filter given varying amount of data. The standard EKF is not able to initialize all the tags with too little data, thus the mean mapping errors (Top) are computed over the initialized tags only.

5 Conclusions

We have proposed a new method for performing localization and SLAM using range-only data that maintains probability distributions for the locations of the "features" (radio tags) in polar space. Linearization in the polar space is much less problematic than linearization in Euclidean space. The proposed method was compared against the standard Euclidean formulation in the case of poor initialization, sparse data and incorrect data association. The results reveal that the ROP-EKF is much more robust to these difficult test cases due to its improved representation of the state uncertainty. The nature of the parameterization in the ROP-EKF allows the filter to produce reliable localization results (when no initialization information is available) and improved mapping results (without the need for a separate pre-filter/batch processing to initialize the tags). Furthermore, the performance of the ROP-EKF is superior to the standard EKF when range measurements are sparse and there are significant errors in data association.

6 Acknowledgments

References

1. P. Bahl and V. Padmanabhan, "Radar: An in-building RF-based user location and tracking system," in *In Proc. of the IEEE Infocom 2000*, Tel Aviv, Israel, March 2000.
2. N. Priyantha, A. Chakraborty, and H. Balakrishman, "The cricket location support system," in *In Proc. of the 6th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM 2000)*, Boston, MA, August 2000.
3. A. Smith, H. Balakrishnan, M. Goraczko, and N. B. Priyantha, "Tracking Moving Devices with the Cricket Location System," in *2nd International Conference on Mobile Systems, Applications and Services (Mobisys 2004)*, Boston, MA, June 2004.
4. G. Kantor and S. Singh, "Preliminary results in range-only localization and mapping," *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*, vol. 2, 2002.
5. D. Kurth, G. Kantor, and S. Singh, "Experimental results in range-only localization with radio," *Intelligent Robots and Systems, 2003.(IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, vol. 1, 2003.
6. J. Leonard, R. Rikoski, P. Newman, and M. Bosse, "Mapping partially observable features from multiple uncertain vantage points." *International Journal of Robotics Research*, vol. 21, no. 10, pp. 943–975, 2002.
7. E. Olson, J. Leonard, and S. Teller, "Robust range-only beacon localization," in *Proceedings of Autonomous Underwater Vehicles, 2004*, 2004.
8. J. Djugash, S. Singh, and P. I. Corke, "Further results with localization and mapping using range from radio," in *International Conference on Field & Service Robotics*, July 2005.
9. J. Djugash, S. Singh, G. Kantor, and W. Zhang, "Range-only slam for robots operating cooperatively with sensor networks," in *IEEE Int'l Conf. on Robotics and Automation (ICRA '06)*, 2006.
10. S. Funiak, C. E. Guestrin, R. Sukthankar, and M. Paskin, "Distributed localization of networked cameras," in *Fifth International Conference on Information Processing in Sensor Networks (IPSN'06)*, April 2006, pp. 34 – 42.
11. J. Civera, A. Davison, and J. Montiel, "Interacting multiple model monocular slam," in *IEEE International Conference on Robotics and Automation*, May 2008.
12. J. Werb and C. Lanzl, "Designing a positioning system for finding things and people indoors," *Spectrum, IEEE*, vol. 35, no. 9, pp. 71–78, 1998.