DRAFT VERSION - To appear in IEEE/RSJ IROS, September 22-26, 2008, Nice, France

Efficient Probabilistic Range-Only SLAM

Jose-Luis Blanco, Juan-Antonio Fernández-Madrigal, and Javier González

***Abstract*— This work addresses Range-Only SLAM (RO-SLAM) as the Bayesian inference problem of sequentially tracking a vehicle while estimating the location of a set of beacons without any prior information. The only assumptions are the availability of odometry and a range sensor able of identifying the different beacons. We propose exploiting the conditional independence between each beacon distribution within a Rao-Blackwellized Particle Filter (RBPF) for maintaining independent Sum of Gaussians (SOGs) for each map element. It is shown then that a proper probabilistic observation model can be derived for online operation with no need for delayed initializations unlike other approaches. We provide a rigorous statistical comparison of this proposal with previous work of the authors where a Monte-Carlo approximation was employed instead for the conditional densities. As verified experimentally, this new proposal represents a significant improvement in accuracy, computation time, and robustness against outliers.**

I. Introduction

The problem of Simultaneous Localization and Mapping (SLAM) has been a topic with a very intense research in the last years. Virtually all of the reported approaches rely on probabilistic Bayesian filtering [13], since in any realistic scenario we only have noisy measurements and imperfect actions.

Regarding the kind of sensors employed for SLAM, a lot of attention has been payed to precise laser scanners, stereo cameras (both providing bearing and range estimates) and monocular cameras (providing bearing-only information). For a survey of works in the field please refer to [2], [21]. However, relatively few works have addressed the problem of building maps with sensors providing range-only (RO) data, in despite of their important applications such as submarine autonomous vehicles [18] or ground vehicles in industrial environments [9].

There are two fundamental characteristics that render RO-SLAM specially challenging: the existence of outliers due to the sensor nature (typically sonar or radio pulses), and more importantly the high ambiguity of the measurements. To illustrate this later issue, consider the example in Fig. 1, where a robot measures the distance to the indicated beacon from three different positions along a straight path. For each position it is shown the ring-shape area of the estimated position of the beacon that is consistent with the measurement (concretely, the figure represents the probabilistic inverse sensor model, but further details are irrelevant at this point). Therefore, two problems with these sensors are: (i) the large portion of the environment where a beacon could be, given just one observation, and (ii) the very likely possibility of multiple plausible hypotheses (as shown in the figure). These issues become more severe when building 3D maps, a requisite when beacons are placed at different heights. Thus, RO-SLAM may become a problem even more challenging than bearing-only SLAM ([6], [16]) where the landmark estimates typically converge to a single mode. On the other hand, an advantage of RO-SLAM is the non-existence of the data association problem, since the usage of active beacons allows most sensors to establish unique correspondences between sensed ranges and particular beacons.

Regarding previous proposals for RO-SLAM, in [20] it is reported a geometric method for adding new beacons to a map using delayed initialization, but a partially known map is required at the beginning. Range-only localization is also addressed in [12] and [15] under the classic EKF-based implementation of SLAM. The authors propose an approximation of the sensor model inspired by the circular-shaped distributions obtained for range sensors. They also address map building but assuming a prior knowledge about the beacon locations. Sub-sea RO-SLAM is demonstrated in [18] with good results, even with the lack of a reliable ego-motion estimation (e.g. from odometry). The main difference with the present work is the usage of a least-square error minimization procedure instead of a probabilistic filter where several robot path hypotheses are considered simultaneously. The work in [19] achieves RO-SLAM through a different strategy: firstly, an initial estimation of the position of each

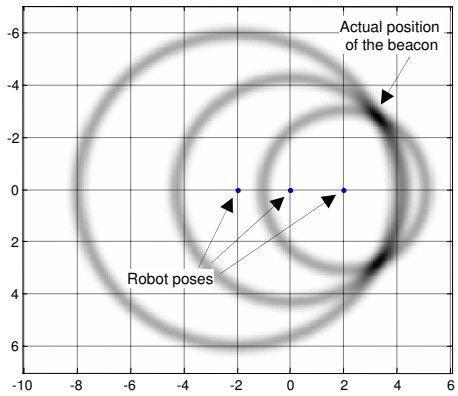

Fig. 1. One peculiarity of Range-Only SLAM is that map estimations may converge to multi-modal densities. In this example, the symmetry in the observations made by a robot over a straight path leads to two regions with a high probability of containing the sensed beacon.

This work was supported by the Spanish Government under research contract DPI2005-01391 and the FPU grand program.

The authors are with the Department of System Engineering and Automation, Málaga University, Málaga, 29071, Spain {jlblanco,jafma,jgonzalez}@ctima.uma.es

beacon is computed using a voting scheme over a 2D grid (a interesting contribution of that work is a preliminary robust filtering of outliers using a graph cut approach). Once the initial estimation converges, a standard EKF deals with the SLAM problem. A similar scheme is adopted in [7], where the authors also explore the possibility of inter-beacon range measurements to improve map building.

In one of our previous works [4] we introduced a probabilistic formulation of RO-SLAM based on a Rao-Blackwellized Particle Filter (RBPF), a solution with the advantage of decoupling the conditional distributions of each beacon in the map for each path particle. This entails freedom in the design of each of these distributions in such a way that at some given instant several already well-localized beacons may coexist in the map with more recently added beacons that have higher uncertainty. One advantage of our solution is that a robot can employ all the available information to perform localization without waiting until the beacon distributions converge as in other proposals. Additionally, new beacons can be inserted at any time in the filter, which is unfeasible under EKF-based solutions to SLAM.

The contributions of this paper over previous works are: (i) a new inverse sensor model for initializing map distributions as weighted Sums of Gaussians (SOGs), (ii) the explanation of how to update those Gaussians and their weights using a multi-hypothesis EKF, and (iii) the derivation of the corresponding observation model required for the RBPF. The present approach has the advantage of a reduced computation burden due to the limited number of Gaussians required for a proper representation, while still providing an accurate approximation of the strongly non-Gaussian and frequently multi-modal distributions found in RO-SLAM.

We also provide statistical analysis from simulations demonstrating that our new approach can cope with highly noisy sensors and reduces in one order of magnitude the average errors of the MC approximation, at a fraction of its computation time.

II. Problem statement

Following the standard notation in the SLAM literature [22], we denote the robot path as the sequence of poses in time $x^t = \{x_1, x_2, ..., x_t\}$, while robot actions (odometry) and observations (range measurements) for each time step t will be represented by u_t and z_t, respectively.

The RBPF approach to estimating the joint SLAM posterior $p(x^t, m|u^t, z^t)$ consists on approximating the marginal distribution of the robot path x^t using importance sampling, then computing the map as a set of conditional distributions given each path hypothesis. That is, if we denote the i'th particle as $x^{t,[i]}$ and its corresponding importance weight as $\omega_t^{[i]}$, we are assuming the following Monte-Carlo approximation of the path:

$$p(x^t|u^t, z^t) \approx \sum_{i=1}^{M} \omega_t^{[i]} \delta_{x^{t,[i]}}(x^t) \tag{1}$$

The motivation for choosing a RBPF approach is that it allows factoring the distribution of the map associated to each particle as:

$$p\left(m|x^t, z^t, u^t\right) = \prod_l p\left(m_l|x^t, z_l^t\right) \tag{2}$$

being m_l the different individual beacon positions in the map m. The factorization in (2) follows from the conditional independence of each beacon in the map given the robot path, hence for each of them we can employ the kind of representation most convenient at each time step without affecting either the given path or other beacons. Section III is devoted to the computation and update of these map densities within the main RBPF.

For completeness we summarize next the sequential algorithm to be executed at each time step. Assume that the set of particles for the previous time step $x^{t-1,[i]}$ are approximately distributed according to the real posterior. In the case of the first time step all the particles can be arbitrarily initialized to the origin. At each iteration, new particles are drawn using the robot motion model (in our case, derived from odometry readings), that is, $x_t^{[i]} \sim p(x_t|x_{t-1}^{[i]}, u_t)$. Next, the importance weights are updated as:

$$\omega_t^{[i]} \propto \omega_{t-1}^{[i]} p\left(z_t|x^{t,[i]}, z^{t-1}\right) \tag{3}$$

with the probabilistic observation model $p\left(z_t|x^{t,[i]}, z^{t-1}\right)$ being derived in section III-D. If necessary, the particles may be resampled to preserve their diversity. This is typically performed whenever the effective sample size falls below a given threshold [17]. After updating the estimate of the robot path, the corresponding conditional distributions of the map must be also updated to account for the new range readings, as discussed next.

III. Derivation of the Solution

In a RBPF, each robot path hypothesis is associated with the corresponding conditional map distribution $p(m_l|x^{t,[i]}, z^t)$ for each beacon l. Note that in the following we drop the l subscript for clarity, since subsequent derivations apply equally and independently to any number of beacons in the map.

Incorporating new information (the new robot pose x_t and the observation z_t) into the map belief m of one beacon is carried out applying the Bayes rule, as follows:

$$\begin{aligned} & \underbrace{p(m|x^t, z^t)}_{\text{Posterior}} \overset{\text{Bayes}}{\propto} \\ & p(m|x^{t-1}, z^{t-1}) p(x_t, z_t|m, x^{t-1}, z^{t-1}) \\ = \; & p(m|x^{t-1}, z^{t-1}) \underbrace{p(x_t|m, x^{t-1}, z^{t-1})}_{\text{Constant}} p(z_t|m, x^t, z^{t-1}) \\ \propto \; & \underbrace{p(m|x^{t-1}, z^{t-1})}_{\text{Prior}} \underbrace{p(z_t|m, x^t, z^{t-1})}_{\text{Inverse sensor model}} \end{aligned} \tag{4}$$

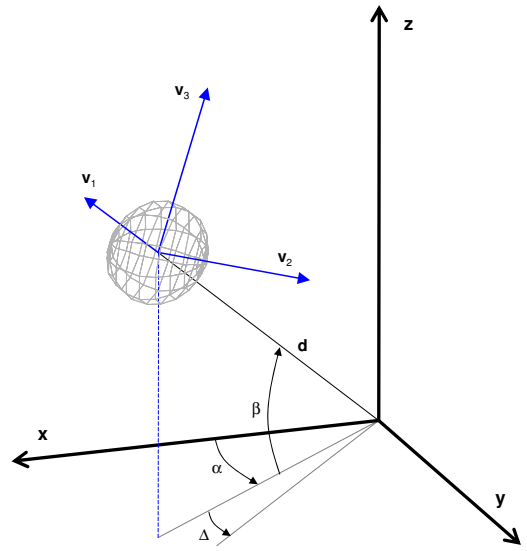

Fig. 2. The variables involved in the generation of the Gaussian modes within each SOG. Azimuth and elevation angles from the sensor position (the coordinate origin in the figure) are represented by α and β, respectively. The covariance matrix is computed by mapping the uncertainties in the radial ($\mathbf{v}_1$) and tangent ($\mathbf{v}_2$, $\mathbf{v}_3$) directions using the appropriate transformation matrix. Refer to section III-A for further details.

where the definition of conditional probability has been applied in the third line, and the marked term is constant over all the RBPF particles since neither the last action u_t nor the last observation z_t appear. Put in words, the posterior belief is obtained by multiplying the previous belief by the inverse sensor model, which determines the likelihood of finding the beacon at any position m given the observed range value z_t and assuming x_t as some fixed hypothesis for the robot pose. Next subsections discuss how to implement the above general equation.

A. Inserting a new beacon

When a beacon is firstly observed at some given instant of time t (not necessarily the first time step), the prior belief in (4) is undefined. Thus, if there is no a priori information at all about the spatial disposition of beacons it is plausible to assume a uniform prior. There is however one interesting exception with important practical consequences. In the case of a ground vehicle building a 3D map, if we know in advance that all the beacons have been placed at a height above (or below) the robot (the typical situation in an industrial environment), the prior becomes a uniform distribution over half of the space (and zero in the complementary part). This is important since, as shown in the experimental results, a vehicle moving on a flat scenario can build a 3D map only up to a symmetry with respect to the robot plane: it cannot be disambiguated whether a given beacon is above or below the robot.

Once defined the prior belief as a uniform, it follows from (4) that the initial map distribution becomes simply the inverse sensor model $p(z_t|m, x^t, z^{t-1})$ (or its evaluation over half of the 3D space in the special case commented above). Assuming a range sensor model with additive Gaussian noise of variance σ_s^2, the probabilistic model is:

$$p(z_t|m, x^t, z^{t-1}) \propto \exp\left(-\frac{1}{2}\frac{|x_t - m|^2}{\sigma_s^2}\right) \tag{5}$$

The evaluation of this model gives the typical ring shapes as those in Fig. 1. The problem is that this density cannot be filtered iteratively in any analytical form, thus in RO-SLAM we must rely on approximations. In this work we propose to adopt a Sum of Gaussians (SOG) approximation, such as:

$$p(z_t|m, x^t, z^{t-1}) \approx \sum_{k=1}^{N} \upsilon_t^k \mathcal{N}(m; \hat{m}_t^k, \Sigma_t^k) \tag{6}$$

being υ_t^k the weight of each Gaussian mode, and $\hat{m}_t^k$ and Σ_t^k its mean and covariance matrix, respectively.

With the aim of approximating (5) for a range measurement $z_t = r$, we have to generate a number of Gaussians equally spaced over a sphere centered at the sensor position and with radius r, following the procedure described next. Each of the SOG modes is placed at a direction stated by a discrete set of azimuth ($-\pi < \alpha \leq \pi$) and elevation ($-\pi/2 < \beta \leq \pi/2$) angles – refer to Fig. 2 for a schematic representation. Let Δ denote the angular increments between consecutive Gaussians along either α or β. Since the model approximation will become poorer for larger distances between Gaussians, let d_m represent the maximum distance allowed between adjacent Gaussians in the regular grid over the sphere.

The angular increment Δ can be computed as $\Delta = 2\pi/B$, being B the integer number of modes along any great circle of the sphere, determined as:

$$B = 2\left\lceil \frac{\pi r}{d_m} \right\rceil \tag{7}$$

where the 2 is out of the ceiling operator to force an even number of modes and assure symmetry.

At this point we have computed the discrete sequences of angles α_i and β_j for $i = 1, .., B$ and $j = 1.., \frac{B}{2}$ that define a regular grid over the sphere. Thus, the mean of each SOG mode is given simply by:

$$\hat{m}_t^{ij} = \begin{pmatrix} x_0 + r\cos\alpha_i \cos\beta_j \\ y_0 + r\sin\alpha_i \cos\beta_j \\ z_0 + r\sin\beta_j \end{pmatrix} \tag{8}$$

Here $(x_0\ y_0\ z_0)^T$ stands for the absolute coordinates of the robot range sensor, which are known since we are assuming a robot pose hypothesis $x_t^{[i]}$. To compute the covariance Σ_t^{ij}, let define three unit orthogonal vectors with origin the mean of the Gaussian. For convenience, the first vector $\mathbf{v}_1$ will be always pointing radially, hence the others ($\mathbf{v}_2$ and $\mathbf{v}_3$) are tangential to the sphere, as illustrated in Fig. 2. Such a vectorial base is well-defined for any sphere of non-zero radius. Note that the uncertainty in the radial direction (the "thickness" of the sphere) is determined by the noise σ_s in the sensor model (5). Furthermore, since we desire radial

symmetry by design, the uncertainty σ_t in both tangential directions should be equal and proportional to the separation between Gaussians, that is, $\sigma_t = r\Delta K$, with K being a proportionality factor. The covariance can be now computed as:

$$\Sigma_t^{ij} = \begin{pmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \mathbf{v}_3 \end{pmatrix} \begin{pmatrix} \sigma_s^2 & 0 & 0 \\ 0 & \sigma_t^2 & 0 \\ 0 & 0 & \sigma_t^2 \end{pmatrix} \begin{pmatrix} \mathbf{v}_1^T \\ \mathbf{v}_2^T \\ \mathbf{v}_3^T \end{pmatrix} \quad (9)$$

Finally, the weights υ_t^{ij} of all the newly generated Gaussians are equal since all of them are equally probable. It is worth noting that the above process can be applied to 2D maps by fixing β to 0 and discarding one of the tangent vectors.

The constant K deserves further comment, since it determines the quality of the approximation to the real inverse sensor model. To illustrate graphically the effects of this constant, please refer to Fig. 4(a)–(b) where a SOG is generated for $K = 0.5$ and $K = 0.3$. Simple visual inspection reveals that the second case leads to a worse approximation of the actual "ring" density. We have computed the Kullback-Leibler divergence (KLD) [14] as a well grounded measure of similarity between the approximate and the actual densities for different ranges r from 1 to 10 meters, and for different values of K. The results, in Fig. 4(c), reveal that good approximations can be obtained, but unfortunately the optimal K varies with the sensed range r. However, values of K near 0.4 always lead to reduced KLD values, hence that is the value employed for the rest of experiments in this work.

B. Updating the SOG

Once the beacon has been inserted in the map, subsequent range measurements z_t update the belief through (4), which in this particular case of the prior density defined as a SOG can be efficiently implemented as a multi-hypothesis EKF. Basically, the mean and covariance are updated using the standard EKF formulas [11] (omitted here for brevity), while the weights are updated by evaluating the actual reading z_t into the Gaussian of the observation predicted by each SOG mode [1]:

$$\upsilon_t^k \propto \upsilon_{t-1}^k \mathcal{N}\left(z_t; h_t^k, {\sigma_t^k}^2\right) \quad (10)$$

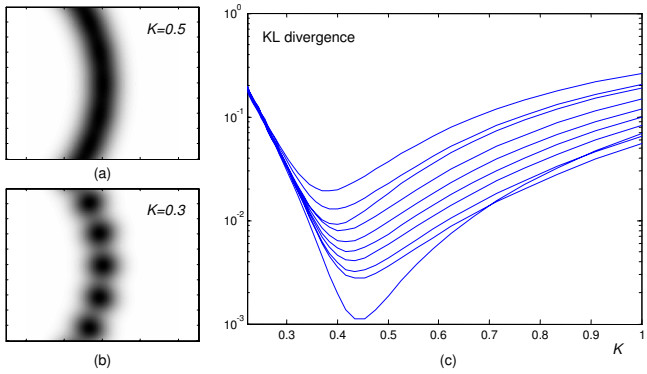

Fig. 4. (a)–(b) Inverse sensor model for a sensed range of $r = 1$ meter computed from our SOG approximation and two different values of the constant K, which determines the standard deviation of Gaussian nodes in tangential directions. (c) The Kullback-Leibler divergence (KLD) of our SOG approximation as a function of K and for r between 1 and 10.

where the mean $h_t^k = h(x_t^{[i]}, \hat{m}_t^k)$ is computed using the sensor model $h(\cdot)$, and the variance ${\sigma_t^k}^2$ includes the sensor noise σ_s and the projection of the beacon uncertainty:

$${\sigma_t^k}^2 = H\Sigma_t^k H^T + \sigma_s^2 \quad (11)$$

being H the Jacobian of the sensor model. Note the proportionality in (10), due to the fact that all the weights are normalized to sum the unity at each iteration.

An important insight in this approach is that, eventually, most of the Gaussians in a SOG will have negligible weights as they become inconsistent with the complete sequence of observations. Since the contribution of these modes to the overall density will be negligible as well, we can simply remove them from the SOG, hence our approach automatically adapts its computational burden to the real uncertainty present at each instant of time, as will be shown later on with real experiments (this is analogous to approaches that scale the number of samples in particle filters [10]).

C. An illustrative example

To illustrate and clarify all the ideas discussed up to this point, consider the example depicted in Fig. 3, where a well-localized robot estimates the position of just one beacon (i.e. there is only mapping, no localization). At the initial instant t_1, the SOG is created as the robot observes the beacon for the first time, following the procedure described in section

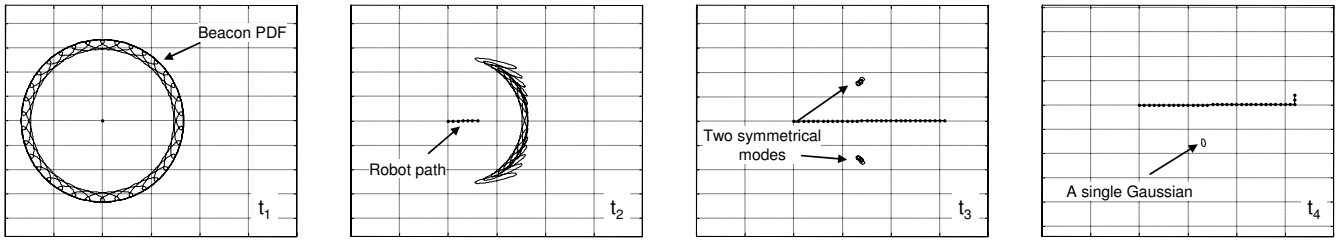

Fig. 3. An example of how our approach iteratively estimates the location of one beacon. It is shown how the Gaussian modes with lowest weight in the SOG are discarded with time, and how symmetrical results are obtained for straight paths. In t_4 it can be seen how this symmetry quickly disappears when the robot deviates from the straight path.

III-A (in this case reduced to 2D for clarity in the representation). Next, as the robot moves along a straight path, the mean and covariances of all the Gaussians are modified by subsequent observations, as can be appreciated at instant t_2. At this point many modes already have negligible weights, hence they have been removed from the filter. Later on, at time t_3, the beacon has converged to two symmetrical modes with respect to the robot path. However, the symmetry is broken when the robot moves apart from its previous straight path, as it can be seen in t_4, where the beacon has been successfully localized without ambiguities.

D. The observation model

Taking into account our approximation of each map distribution as a SOG (6), we can expand the observation model required in (3) to update the weights of the RBPF, as follows:

$$\begin{aligned} & p\left(z_t|x^{t,[i]}, z^{t-1}\right) \\ = & \int p\left(z_t|x_t^{[i]}, m\right) p\left(m|x^{t-1,[i]}, z^{t-1}\right) dm \\ = & \sum_{k=1}^{N} v_{t-1}^k \mathcal{N}\left(z_t; h_t^k, {\sigma_t^k}^2\right) \end{aligned} \tag{12}$$

where each normal distribution represents the predicted observation for one mode within the SOG. The parameters of these Gaussians have been already described in (10)–(11).

IV. Experiments

To validate our approach we have performed extensive simulations in order to (i) characterize its performance against different levels of input noise, and (ii) to compare it with our previous work [4]. Additionally, we present the construction of a 3D map from data gathered by a real robot. A video illustrating the following results and the source code are available online in [3] [1].

[1]Since this work is unpublished, please use the direct link http://babel.isa.uma.es/mrpt/index.php/Paper:RO-SLAM_with_SOG

A. Performance characterization

We have carried out three series of simulations in order to characterize statistically the accuracy in the maps generated by our approach. The first experiment characterizes the average localization error for beacons in the final map as a function of the sensor noise σ_s. The results are represented in Fig. 6(a) through the mean errors and their corresponding 67% confidence interval. To obtain statistically significant results we have executed our approach 50 times for each parameter value, each time with 20 randomly placed beacons. Average errors as a function of odometry noise and the number of particles in the RBPF have been computed similarly and the results are plotted in Fig. 6(b)–(c). Errors in odometry are represented as the ratio between the standard deviation of the Gaussian noise and the actual increment between consecutive robot poses.

The interpretation of these statistical results is that, as expected, lower levels of noise or more particles lead to lower average errors. Our method can cope with heavily corrupted range observations (e.g. σ_s up to $1m$ in a $20\times20m$ map) exhibiting map errors approximately proportional to the sensor noise (see Fig. 6(a)). However, smaller errors in odometry – above 5% in our experimental setup, see Fig. 6(b) – lead to a faster raise in the map errors. The reason is the usage of the robot motion model (based on odometry) as the proposal density in the RBPF. More optimal choices [5], [8] would give increased robustness at the cost of a larger computational burden.

B. Comparison to Monte-Carlo approximation

Next we compare our method, based on a SOG representation of the map conditional densities, to a previous work [4] where a MC approximation was employed instead. In order to provide a fair comparison between both methods, we have analyzed the average errors in the maps for similar computational burdens and the average execution times for similar map errors. The results are summarized in Fig. 5(a) and (b), respectively. For similar computation times, the MC approach obtained an average error of 0.28m while for our new proposal this error reduces to 0.03m. Furthermore,

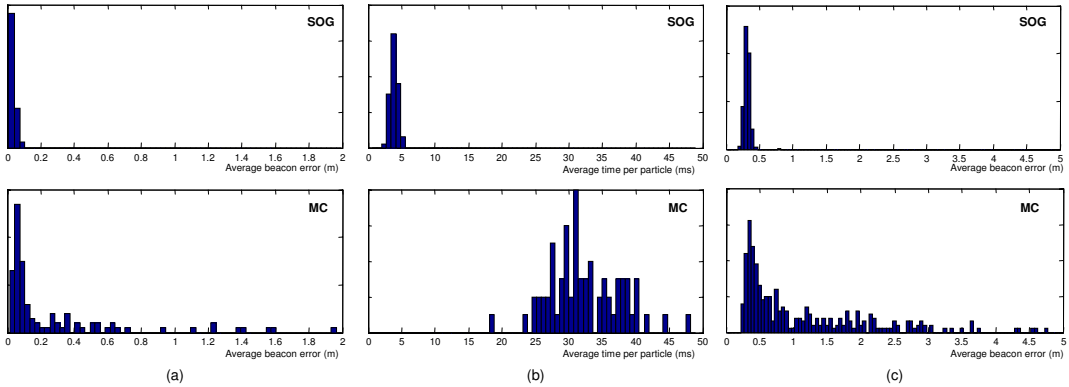

Fig. 5. A comparison of the present approach (SOG) and the method in [4] (MC). (a) Average beacon error (m) for similar computation time in both approaches. (b) Average computation time per particle (ms) for similar final beacon errors. (c) Beacon errors for a simulated sensor heavily corrupted with noise and outliers.

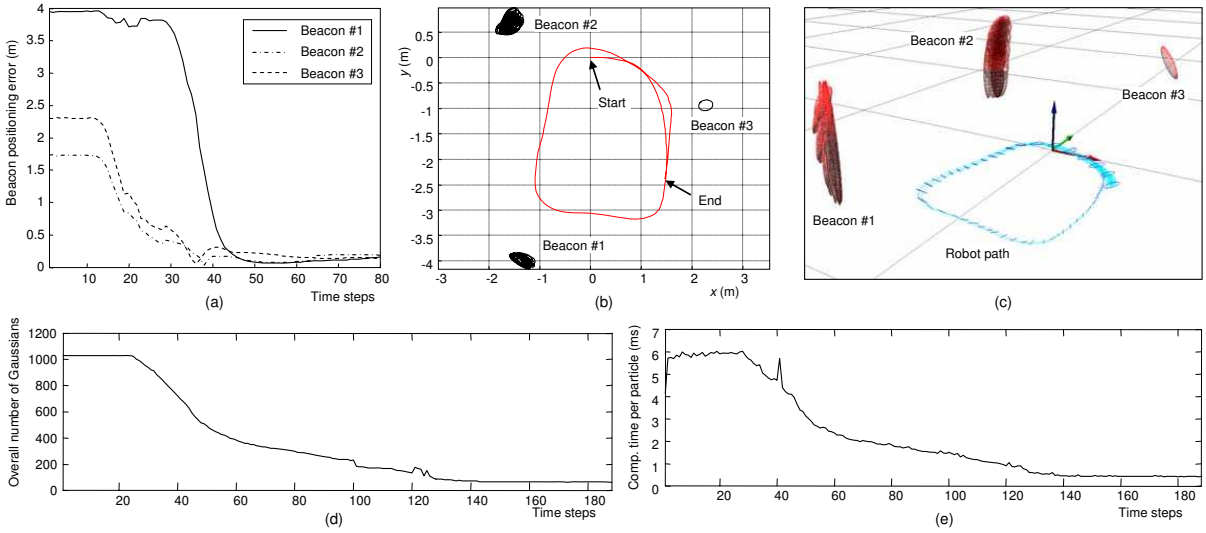

Fig. 7. Results from data gathered by a real robot equipped with a UWB transceiver. (a) Errors in each beacon estimate with respect to the ground truth. (b)–(c) A bird-view and a 3D representation of the final state of the filter, respectively. The map shown is that one associated to the particle with the highest probability. (d)–(e) Overall number of Gaussians in the most likely map and computation time, respectively, for each iteration of the algorithm.

in the second situation (similar map errors) the average execution times are 3.9 and 32.4 ms per particle for the SOG and MC solutions, respectively. Therefore, the new proposal outperforms [4] by one order of magnitude in both efficiency and accuracy.

We have also analyzed the tolerance of both methods to the presence of outliers in the range data. In our approach, unusually large errors in the sensed range are specially problematic if they occur the first time a beacon is observed, since we rely on this first range to initialize the map distribution. With this idea in mind, we have setup an experiment simulating a highly noisy sensor, $\sigma_s = 1m$, whose ratio of outliers is 30% the first time a beacon is observed and 5% for the rest of observations. Outliers have been simulated by adding to the ranges a uniform noise between 1 and 10 meters. The average map error histograms for 50 repetitions are shown in Fig. 5(c). In this case, a mean error of 1.12m is obtained for the MC approximation, while the present approach achieves a significant reduction up to 0.32m. The better results of the SOG representation arise from the better adaptability of Gaussian modes to newer observations (i.e. they can be "displaced"), even recovering from a heavily corrupted initial distribution.

C. A 3D map from a real dataset

Finally, we have applied our approach to a sequence of UWB range measurements [9] within an indoor environment. The experimental setup consists of a Pioneer mobile robot with a UWB transceiver onboard, while other three UWB devices act as static radio beacons. The position of the three beacons has been measured manually to provide the ground truth required to evaluate the results. In this case we have employed the a priori knowledge that beacons are above the robot to limit the map prior distribution to one half of the 3D space, as discussed in section III-A.

TABLE I
SUMMARY OF ERRORS FOR THE 3D MAP BUILT FROM UWB DATA

Beacon coordinate		Ground truth (m)	Estimate (m)	Error (m)
#1	x	0	-0.059	0.059
	y	0	-0.278	0.278
	z	0.912	1.17	0.257
#2	x	-0.320	-0.392	0.072
	y	4.332	4.419	0.087
	z	1.374	1.214	0.159
#3	x	3.403	3.513	0.109
	y	2.802	2.740	0.062
	z	2.175	2.108	0.067

After completing one loop along the room, the estimates for all three beacons have converged to unimodal distributions, as shown in Fig. 7(b)–(c). In despite of the noisy sensor ($\sigma_s = 0.10m$), it can be observed in Fig. 7(a) how the errors in the location of each beacon quickly vanish as the robot moves just a few meters. The better estimation of the beacon #3 (refer to Fig. 7(c)) is a consequence of its higher height relative to the robot, which makes range observations more discriminant. The errors between the final beacon estimates (the means) and the ground truth are summarized in Table I.

Finally, observe in Fig. 7(d)–(e) how the computational burden of the algorithm decreases with time as the map estimate becomes more precise and fewer Gaussians are required to represent the maps densities. Unlike in our previous approach [4], a reduced number of Gaussians will not lead to degeneracy of the filter as in a MC approximation. At the point of maximum complexity, just after initializing three beacons, the computation time is of 6ms per particle, thus our method is efficient enough to perform online on a real robot.

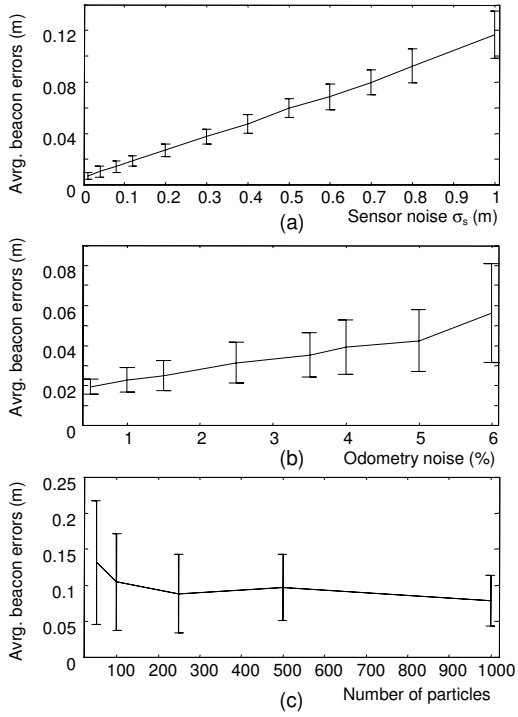

Fig. 6. Statistical results from simulations sweeping different parameters: (a) the sensor noise (σ_s), (b) odometry errors, and (c) the number of particles in the RBPF. Results present average errors for beacon locations among 67% confidence intervals.

V. Discussion

In this work we have presented a novel representation for the map conditional densities associated to the particles of a RBPF, which is based on a SOG. This approach has revealed well-suited to the particular problems that arise in RO-SLAM, being significantly more efficient and accurate than other alternatives based on MC approximations. We have verified experimentally the robustness of the method against readings heavily corrupted with noise and outliers, as well as its ability to build 3D maps from a real dataset gathered using UWB radio transceivers. Future research will address the problem of the dependency on odometry as well as analyzing better alternatives to the proposal density in the RBPF.

References

[1] D. Alspach and H. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," *IEEE Transactions on Automatic Control*, vol. 17, no. 4, pp. 439–448, 1972.

[2] T. Bailey and H. Durrant-Whyte, "Simultaneous localisation and mapping (SLAM): Part II-State of the art," *Robotics and Automation Magazine*, vol. 13, pp. 108–117, 2006.

[3] J. Blanco, "The Mobile Robot Programming Toolkit (MRPT) website," 2008. [Online]. Available: "http://babel.isa.uma.es/mrpt/"

[4] J. Blanco, J. Gonzalez, and J. Fernández-Madrigal, "A Pure Probabilistic Approach to Range-Only SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2008.

[5] ——, "An Optimal Filtering Algorithm for Non-Parametric Observation Models in Robot Localization," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2008.

[6] A. Davison, I. Reid, N. Molton, and O. Stasse, "MonoSLAM: Real-Time Single Camera SLAM," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 6, June 2007.

[7] J. Djugash, S. Singh, G. Kantor, and W. Zhang, "Range-only SLAM for robots operating cooperatively with sensor networks," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2006, pp. 2078–2084.

[8] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.

[9] J. Fernández-Madrigal, E. Cruz-Martin, J. Gonzalez, C. Galindo, and J. Blanco, "Application of UWB and GPS Technologies for Vehicle Localization in Combined Indoor-Outdoor Environments," *International Symposium on Signal Processing and Its Applications (ISSPA)*, 2007.

[10] D. Fox, "Adapting the Sample Size in Particle Filters Through KLD-Sampling." *International Journal of Robotics Research*, vol. 22, no. 12, pp. 985–1003, 2003.

[11] S. Julier and J. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," *Int. Symp. Aerospace/Defense Sensing, Simul. and Controls*, vol. 3, 1997.

[12] 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.

[13] S. Kay, *Fundamentals of statistical signal processing: estimation theory*. Prentice-Hall, Inc. Upper Saddle River, NJ, USA, 1993.

[14] S. Kullback and R. Leibler, "On Information and Sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, 1951.

[15] 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.

[16] N. Kwok and G. Dissanayake, "An efficient multiple hypothesis filter for bearing-only SLAM," vol. 1, 2004.

[17] J. Liu, "Metropolized independent sampling with comparisons to rejection sampling and importance sampling," *Statistics and Computing*, vol. 6, no. 2, pp. 113–119, 1996.

[18] P. Newman and J. Leonard, "Pure range-only sub-sea SLAM," *Robotics and Automation, 2003. Proceedings. ICRA'03. IEEE International Conference on*, vol. 2, 2003.

[19] E. Olson, J. Leonard, and S. Teller, "Robust range-only beacon localization," *Autonomous Underwater Vehicles, 2004 IEEE/OES*, pp. 66–75, 2004.

[20] S. Singh, G. Kantor, and D. Strelow, "Recent results in extensions to simultaneous localization and mapping," *International Symposium on Experimental Robotics*, 2002.

[21] S. Thrun, *Robotic Mapping: A Survey*. School of Computer Science, Carnegie Mellon University, 2002.

[22] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. The MIT Press, September 2005.