Improving Grid-based SLAM with Rao-Blackwellized Particle Filters by Adaptive Proposals and Selective Resampling

Giorgio Grisetti[†‡]

[†] *Dipartimento Informatica e Sistemistica*
Universitá "La Sapienza"
I-00198 Rome, Italy

Cyrill Stachniss[‡] Wolfram Burgard[‡]

[‡] *University of Freiburg*
Department of Computer Science
D-79110 Freiburg, Germany

***Abstract*— Recently Rao-Blackwellized particle filters have been introduced as effective means to solve the simultaneous localization and mapping (SLAM) problem. This approach uses a particle filter in which each particle carries an individual map of the environment. Accordingly, a key question is how to reduce the number of particles. In this paper we present adaptive techniques to reduce the number of particles in a Rao-Blackwellized particle filter for learning grid maps. We propose an approach to compute an accurate proposal distribution taking into account not only the movement of the robot but also the most recent observation. This drastically decrease the uncertainty about the robot's pose in the prediction step of the filter. Furthermore, we present an approach to selectively carry out re-sampling operations which seriously reduces the problem of particle depletion. Experimental results carried out with mobile robots in large-scale indoor as well as in outdoor environments illustrate the advantages of our methods over previous approaches.**

I. Introduction

Building maps is one of the fundamental task of mobile robots. In the past a lot of researchers focused on this problem. In literature the mobile robot mapping problem is often referred to as the *simultaneous localization and mapping (SLAM) problem* [2, 4, 5, 7, 8, 12, 14, 16, 18]. It is considered to be complex problem, because for localization a robot needs a consistent map and for acquiring the map the robot requires a good estimate of its location. This mutual dependency among the pose and the map estimates makes the SLAM problem hard, and requires to search for a solution in a high-dimensional space.

Recently Murphy, Doucet and colleagues [16, 4] introduced Rao-Blackwellized particle filters as an effective means to solve the simultaneous localization and mapping (SLAM) problem. The main problem of the Rao-Blackwellized approaches is their complexity, measured in terms of the number of particles required to build an accurate map. Therefore reducing this quantity is one of the major challenges of this family of algorithms. Additionally, the resampling step can eliminate the correct particle. This effect is also known as the particle depletion problem [19].

In this work we present two approaches to drastically increase the performance of Rao-Blackwellized particle filters applied to SLAM with grid maps:

- A proposal distribution that considers the accuracy of the robot's sensors and allows to draw particles in an highly accurate manner, as well as
- an adaptive resampling technique, which maintains a reasonable variety of particles and this way reduces the risk of particle depletion.

The proposal distribution is computed by evaluating the likelihood around a particle-dependent most likely pose obtained by a scan registration procedure. In this way the last reading is taken into account while generating the new particle, allowing to estimate the evolution of the system according to a more informed (and thus more accurate) model than the one obtained using only the last odometry reading as in [5] and [8]. The use of this refined model has two effects. The map is more accurate, since the current laser scan is incorporated into the map related to each particle, *after* considering its effect on the robot pose. The estimation error accumulated over time is significantly lower and less particles are required to represent the posterior. The second approach, the adaptive resampling strategy allows to perform a resampling step only when needed, keeping a reasonable particle diversity. This results in a significant reduction of the particle depletion problem.

Our approach has been validated by a set of systematic experiments in large scale indoor and outdoor environments In all experiments our approach generated highly accurate metric maps. Additionally the number of the required particles is one order of magnitude lower than with the previous approaches.

This paper is organized as follows. After a discussion of the related works in the following section, we explain how a Rao-Blackwellized filter can be used to solve the SLAM problem. Section IV describes our improvements. Experiments carried out on real robots as well as in simulation are discussed in Section V.

II. Related Work

The mapping algorithms proposed so far can be roughly classified according to the map representation and the underlying estimation technique. One popular map representation is the occupancy grid [15]. Whereas such grid-based approaches are computationally expensive and typically require

a huge amount of memory, they are able to represent arbitrary features and provide detailed representations. Feature based representations such as topological maps or landmark-based maps are attractive because of their compactness. They, however, rely on predefined feature extractors, which assumes that some structures of the environments are known in advance.

The estimation algorithms can be roughly classified according to their underlying basic principle. The most popular approaches are Extended Kalman filters (EKFs), maximum likelihood techniques and Rao Blackwellized particle filters. The effectiveness of the EKF approaches comes from the fact that they estimate a fully correlated posterior over landmark maps and robot poses. Their weakness lies in the strong assumptions that have to be made on both the robot motion model and the sensor noise. Moreover, the landmarks are assumed to be uniquely identifiable. If these assumptions are violated the filter is likely to diverge [6].

A popular maximum likelihood algorithm computes the most likely map given the history of sensor readings by constructing a network of relations that represents the spatial constraints among the robot poses. Gutmann *et al.* [7] proposed an effective way for constructing such a network, and for detecting loop closures, while running an incremental maximum likelihood algorithm. When a loop closure is detected a global optimization on the relation network is performed. Recently Hähnel *et al.* [9], proposed an approach which is able to track several map hypotheses using an association tree. However the necessary expansions of this tree can prevent the approach from being feasible for real-time operation.

In a work by Murphy [16], Rao-Blackwellized particle filters (RBPF) have been introduced as an effective means to solve the SLAM problem. Each particle in a RBPF represents a possible robot trajectory and a map. The framework has been subsequently extended by Montemerlo et al. [14, 13] for approaching the SLAM problem with landmark maps. To learn accurate grid maps RBPF have been used by Eliazar and Parr [5] and Hähnel *et al.* [8]. Whereas the first work describes an efficient map representation, the second presents an improved motion model that reduces the number of required particles.

The work described in this paper is an improvement of the algorithm proposed by Hähnel *et al.* [8]. Instead of using a fixed proposal distribution our algorithm computes on the fly an improved proposal distribution on a per particle base. This allows to directly use most of the information obtained from the sensors while generating the particles. The computation of the proposal distribution is similar to that of the FastSLAM-2 algorithm [12], which, however, relies on predefined landmarks.

The advantage of our approach is twofold. First, our algorithm draws the particles in a more effective way. Second, the highly accurate proposal distribution allows to utilize the number of effective particles as a robust indicator to decide whether or not a resampling has to be carried out. This further reduces the particle depletion problem.

III. Rao-Blackwellized Mapping

The key idea of the Rao-Blackwellized particle filter for SLAM is to estimate a posterior $p(x_{1:t} \mid z_{1:t}, u_{0:t})$ about potential trajectories $x_{1:t}$ of the robot given its observations $z_{1:t}$ and its odometry measurements $u_{0:t}$ and to use this posterior to compute a posterior over maps and trajectories:

$$p(x_{1:t}, m \mid z_{1:t}, u_{0:t}) = p(m \mid x_{1:t}, z_{1:t}) p(x_{1:t} \mid z_{1:t}, u_{0:t}). \tag{1}$$

This can be done efficiently, since the posterior over maps $p(m \mid x_{1:t}, z_{1:t})$ can be computed analytically [15] given the knowledge of $x_{1:t}$ and $z_{1:t}$.

To estimate the posterior $p(x_{1:t} \mid z_{1:t}, u_{0:t})$ over the potential trajectories Rao-Blackwellized mapping uses a particle filter in which an individual map is associated to every sample. Each map is built given the observations $z_{1:t}$ and the trajectory $x_{1:t}$ represented by the corresponding particle. The trajectory of the robot evolves according to the robot motion and for this reason the proposal distribution is chosen to be equivalent to the probabilistic odometry motion model.

One of the most common particle filtering algorithms is the Sampling Importance Resampling (SIR) filter. A Rao-Blackwellized SIR filter for mapping incrementally processes the observations and the odometry readings as they are available. This is done by updating a set of samples representing the posterior about the map and the trajectory of the vehicle. This is done by performing the following four steps:

1) *Sampling*: The next generation of particles $\{x_t^{(i)}\}$ is obtained from the current generation $\{x_{t-1}^{(i)}\}$, by sampling from a proposal distribution $\pi(x_t \mid z_{1:t}, u_{0:t})$.
2) *Importance Weighting*: An individual importance weight $w^{(i)}$ is assigned to each particle, according to

$$w^{(i)} = \frac{p(x_t^{(i)} \mid z_{1:t}, u_{0:t})}{\pi(x_t^{(i)} \mid z_{1:t}, u_{0:t})}. \tag{2}$$

   The weights $w^{(i)}$ account for the fact that the proposal distribution π in general is not equal to the true distribution of successor states.
3) *Resampling*: Particles with a low importance weight w are typically replaced by samples with a high weight. This step is necessary since only a finite number of particles are used to approximate a continuous distribution. Furthermore, resampling allows to apply a particle filter in situations in which the true distribution differs from the proposal.
4) *Map Estimating*: for each pose sample $x_t^{(i)}$, the corresponding map estimate $m_t^{(i)}$ is computed based on the trajectory and the history of observations according to $p(m_t^{(i)} \mid x_{1:t}^{(i)}, z_{1:t})$.

In the literature of particle filters several methods for computing better proposal distributions and for reducing the

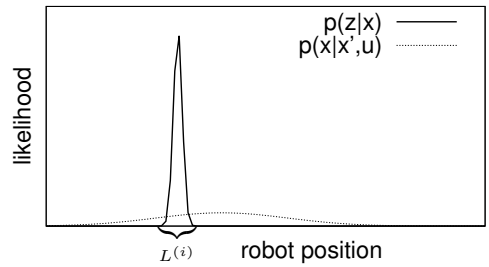

Fig. 1. The two components of the motion model. Within the interval $L^{(i)}$ the product of both functions is dominated by the observation likelihood. Accordingly the model of the odometry error can safely be approximated by a constant value.

particle depletion problem have been described [3, 17]. Our approach applies two concepts also described by Doucet [3]: The computation of the optimal proposal distribution and an adaptive resampling technique. The first of this two concepts has been successfully applied in FastSLAM-2 [12] for landmark-based mapping, while to the best of our knowledge, the adaptive resampling has never been investigated in the context of mapping with Rao-Blackwellized particle filters.

IV. RBPF with Improved Proposals and Adaptive Resampling

Whereas the generic algorithm specifies a framework for Rao-Blackwellized mapping, it leaves open how the proposal distribution is computed and when the resampling should be carried out. Throughout the remainder of this paper, we describe techniques to compute accurate proposal distributions, and to adaptively determine when to resample. Both approaches lead to a highly efficient Rao-Blackwellized mapping technique, which requires one order of magnitude less particles than the ones required by previous grid-based approaches of that type.

A. Computing the Improved Proposal Distribution

As described in Section III, one needs to draw samples from a proposal distribution π in the prediction step. This distribution should approximate the true distribution $p(x_t \mid z_{1:t}, u_{0:t})$ and can be chosen arbitrarily.

According to Doucet [3] the optimal choice of the proposal distribution with respect to the variance of the particle weights and under the Markov assumption is

$$p(x_t \mid m_{t-1}^{(i)}, x_{t-1}^{(i)}, z_t, u_t) = \frac{p(z_t \mid m_{t-1}^{(i)}, x_t) p(x_t \mid x_{t-1}^{(i)}, u_t)}{\int p(z_t \mid m_{t-1}^{(i)}, x') p(x' \mid x_{t-1}^{(i)}, u_t) dx'}. \quad (3)$$

In several particle filter applications [1, 14] the odometry motion model $p(x_t \mid x_{t-1}, u_t)$ has been chosen as the proposal distribution. When modeling a mobile robot equipped with a laser range finder, this choice can be suboptimal, since the accuracy of the laser range finder leads to extremely peaked likelihood functions. In such a case, the likelihood function $p(z_t \mid m_{t-1}^{(i)}, x_t)$ dominates the product $p(z_t \mid m_{t-1}^{(i)}, x_t) p(x_t \mid x_{t-1}^{(i)}, u_t)$ within the meaningful area of this distribution (see Fig. 1). In our current system we therefore approximate $p(x_t \mid x_{t-1}^{(i)}, u_t)$ by a constant k within the interval $L^{(i)}$ given by

$$L^{(i)} = \{x \mid p(z_t \mid m_{t-1}^{(i)}, x) > \epsilon\}. \quad (4)$$

Under this approximation Eq. (3) turns into

$$p(x_t \mid m_{t-1}^{(i)}, x_{t-1}^{(i)}, z_t, u_t) \simeq \frac{p(z_t \mid m_{t-1}^{(i)}, x_t)}{\int_{x' \in L^{(i)}} p(z_t \mid m_{t-1}^{(i)}, x') dx'} \quad (5)$$

We furthermore locally approximate the distribution around the maximum of the likelihood function by a Gaussian:

$$p(x_t \mid m_{t-1}^{(i)}, x_{t-1}^{(i)}, z_t, u_t) \simeq \mathcal{N}(\mu_t^{(i)}, \Sigma_t^{(i)}) \quad (6)$$

With this approximation, we obtain a closed form which is suitable for sampling. For each particle i the parameters $\mu_t^{(i)}$ and $\Sigma_t^{(i)}$ can be determined by evaluating the likelihood function for a set of points $\{x_j\}$ sampled around the corresponding local maximum found by the scan-matching process

$$\mu_t^{(i)} = \frac{1}{\eta} \sum_{j=1}^{K} x_j p(z_t \mid m_{t-1}^{(i)}, x_j) \quad (7)$$

$$\Sigma_t^{(i)} = \frac{1}{\eta} \sum_{j=1}^{K} p(z_t \mid m_{t-1}^{(i)}, x_j)(x_j - \mu_t^{(i)})(x_j - \mu_t^{(i)})^T. \quad (8)$$

Here $\eta = \sum_{j=1}^{K} p(z_t \mid m_{t-1}^{(i)}, x_j)$ is a normalizer. Observe that the computation of $\mu_t^{(i)}$ and $\Sigma_t^{(i)}$ as well as the scan-matching process are carried out for each particle. The $\{x_j\}$ are chosen to cover an area dependent on the last odometry reading uncertainty $x_j \in \{x_t \mid p(x_t \mid x_{t-1}, u_t) > \chi\}$, and with a density depending on the grid map resolution. In our current system we apply a scan-matching routine similar to that of Hähnel *et al.* [10].

Furthermore, we have to specify how the importance weights are computed under this proposal distribution. We can approximate the i-th particle importance weight $w^{(i)}$ by:

$$\begin{aligned} w_t^{(i)} &= w_{t-1}^{(i)} p(z_t \mid m_{t-1}^{(i)}, x_{t-1}^{(i)}, u_t) \\ &= w_{t-1}^{(i)} \int p(z_t \mid m_{t-1}^{(i)}, x') p(x' \mid x_{t-1}^{(i)}, u_t) dx' \\ &\simeq w_{t-1}^{(i)} k \int_{x' \in L^{(i)}} p(z_t \mid m_{t-1}^{(i)}, x') dx' \\ &\simeq w_{t-1}^{(i)} k \sum_{j=1}^{K} p(z_t \mid m_{t-1}^{(i)}, x_j) \\ &= w_{t-1}^{(i)} k \eta \end{aligned} \quad (9)$$

The proposal distribution based on our likelihood function enables a robot equipped with a laser range finder to draw samples in an very accurate way (see Fig. 2). For example, in

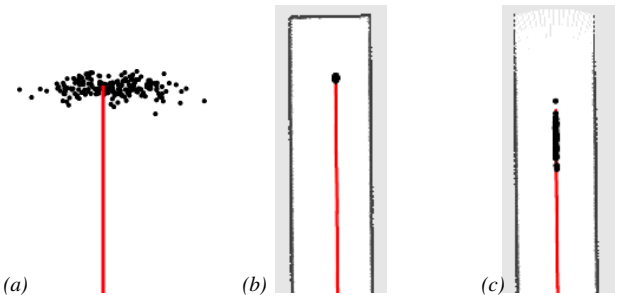

Fig. 2. Proposal distributions typically observed during mapping. In a featureless open space the proposal distribution is the raw odometry motion model *(a)*. In a dead end corridor the particles the uncertainty is small in all of the directions *(b)*. In an open corridor the particles distributes along the corridor *(c)*.

Fig. 3. **Intel Lab,** The robot starts in the upper part of the circular corridor, and runs several times around the loop, before entering the rooms. On the left, map generated with 15 particles. On the right a particular showing the accuracy of the map in the loop closure point, the resolution is 1 cm.

a corridor the samples are typically spread along its direction. The resulting densities have a much lower uncertainty than in the situation in which the odometry motion model is utilized. Accordingly our technique allows to greatly decrease the number of particles. The approach of Hähnel *et al.* [8], which uses a fixed proposal distribution obtained from the residual error of the scan-matching process, in principle needs a conservative proposal distribution to cover the worst case and therefore typically needs more particles than our algorithm.

B. Selective Resampling

A further aspect that has a major influence on the performance of a particle filter is the resampling step. During resampling, the particles with a low importance weight $w^{(i)}$ are typically replaced by samples with a high weight. On the one hand, resampling is necessary since only a finite number of particles are used. On the other hand, the resampling step can delete good samples from the sample set, causing particle depletion [19]. Accordingly, it is important to find a criterion when to perform a resampling step. Liu [11] introduced the so-called effective number of particles N_{eff} to estimate how well the current particle set represents the true posterior. This quantity is computed as

$$N_{\mathit{eff}} = \frac{1}{\sum_{i=1}^{N} \left(w^{(i)}\right)^2}. \tag{10}$$

The intuition behind N_{eff} is as follows. If the samples were drawn from the true posterior, the importance weights of the samples would be equal to each other. The worse the approximation the higher the variance of the importance weights. Since N_{eff} can be regarded as a measure of the dispersion of the importance weights, it is a useful measure to evaluate how well the particle set approximates the true posterior. Our approach follows the one proposed by Doucet [3] to determine whether or not a resampling should be carried out. We resample each time N_{eff} drops below a given threshold of $N/2$ where N is the number of particles. In our experiments we found that this approach drastically reduces the risk of replacing good particles, because the number of resampling operations is reduced and resampling operations are only performed when needed.

V. Experiments

The approach described above has been implemented and tested on real robot data. An on line implementation worked on a Pioneer2 AT, equipped with a Sick PLS range finder. The experiments carried out in a variety of environments have shown the effectiveness of our approach in indoor and outdoor environments. The quality of the maps was extremely good, allowing in some cases to generate map of 1cm of resolution, without observing considerable inconsistencies. We never required more than 100 particles, even for environments with a size of 250 $\times$ 250 meters. Our algorithm can be executed in real time with 50 particles on a Pentium IV-2.8GHz, and the robot traveling at 0.5 m/s. In the next section we discuss the behavior of the filter in three real world datasets. More maps generated with our approach are available on the Web.[1] Furthermore, we give a quantitative analysis of the performance of the presented approach.

A. Mapping Results

The data sets discussed here have been recorded at the Intel Research Lab in Seattle, at the Killian Court at MIT, and on the campus of the University of Freiburg. The maps of these environments are depicted in Figures 3, 4, and 5.

a) Intel Lab: The size of the Intel Lab is 28 $\times$ 28 meters. The data set has been recorded with a Pioneer II robot equipped with a SICK sensor. To process this data set our algorithm needed only 15 particles. As can be seen in Fig. 3, the quality of the final map is so high that the map can be magnified up to 1 cm of resolution without an observable significant error.

b) Freiburg Campus: The second data set has been recorded outdoors at the Freiburg Campus. Our approach needed only 30 particles to produce a good quality map such as the one shown in Fig. 4. Note that this environment partly violates the approximation that the environment is planar. Additionally, there were objects like grass and bushes as well

[1] http://www.informatik.uni-freiburg.de/∼stachnis/research/rbpfmapper/

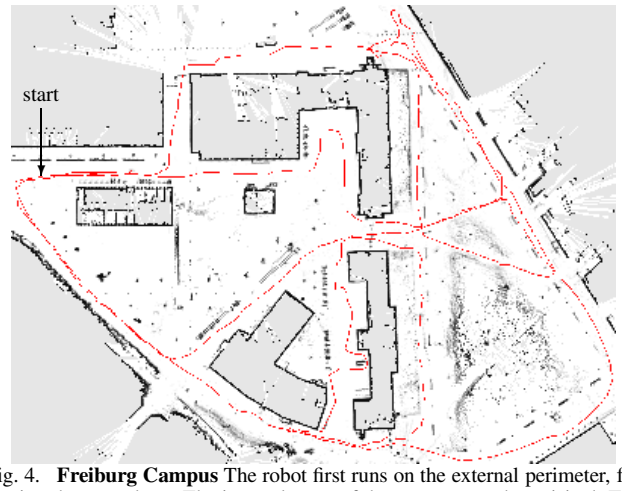

Fig. 4. **Freiburg Campus** The robot first runs on the external perimeter, for closing the outer loop. The internal parts of the campus are then visited. The area is 250 × 250 m, and the overall trajectory is 1.7 km. Map generated with 30 particles.

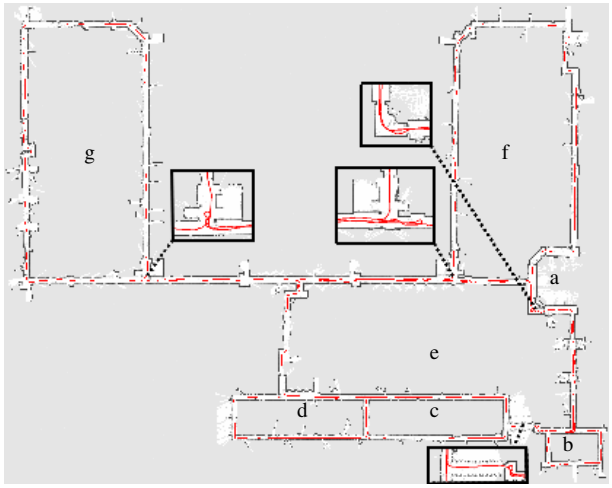

Fig. 5. **MIT Killian Court** The robot starts from the point labeled with a, then goes for the first loop b. It then moves to c, d, and comes back to a and b, prior to visit f and g. The environment is of 250 × 215 m and the robot traveled for 1.9 km. This map has been generated with 80 particles. The rectangles show magnifications of several parts of the map.

as moving objects like cars and people. Despite the resulting spurious measurements our algorithm was able to generate an accurate map.

c) MIT Killian Court: The third experiment was performed with a dataset acquired at the MIT Killian court (see Fig. 5). This dataset is extremely challenging for a Rao-Blackwellized mapper since there are several nested loops, which can cause the algorithm to fail due to particle depletion. In this data set our selective resampling procedure turned out to be extremely useful since it maintains hypotheses about potential trajectories. A consistent, topologically correct map can be generated with 60 particles. However, the resulting maps sometimes show artificial double walls. By employing 80 particles it is possible to achieve good quality maps.

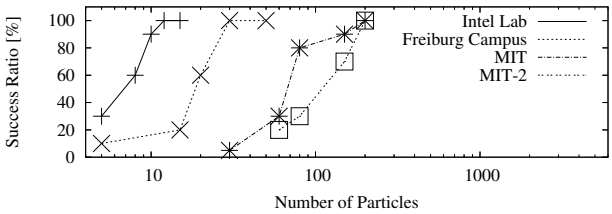

Fig. 6. Success rate of 10 runs of the algorithm in different environments. For the experiment MIT-2 we disabled the adaptive resampling.

B. Quantitative Results

In order to measure the improvement in terms of a number of particles, we have compared the performances of our informed proposal distribution and the uninformed proposal used in [8]. The following table summarizes the number of particles needed by a RBPF for providing a topologically correct map in at least 60% of all applications of our algorithm.

Proposal Distribution	**Intel**	**MIT**	**Freiburg**
our approach	8	60	20
approach of [8]	40	400	400

It turns out that in all of the cases the number of particles required by our approach was approximately one order of magnitude less than the ones required by the other. Moreover, the resulting maps are better due to our improved sampling process that takes the last reading into account.

Figure 6 summarizes results about the success ratio of our algorithm in the environments considered here. The plots show the percentage of correctly generated maps depending on the number of particles used. The quality of the maps has been evaluated by visual inspection. As a measure of success we used the topological correctness.

C. Effects of Improved Proposals and Adaptive Resampling

The increased performance of our approach is due to the interplay of two factors, namely the improved proposal distribution, which allows to generate samples with an high likelihood, and the adaptive resampling controlled by monitoring N_{eff}. For proposals that do not consider the whole input history it has been proved that N_{eff} can only decrease (stochastically) over time [3]. Only after a resampling operation N_{eff} recovers its maximum value. It is important to notice that the behavior of N_{eff} strictly depends on the proposal: the worse the proposal, the faster N_{eff} drops.

In our experiments we found that evolution N_{eff} under our proposal distribution shows three different behaviors depending on the information obtained from the robot's sensors. Whenever the robot moves through unknown terrain, N_{eff} typically drops slowly. This is because the proposal distribution becomes less peaked and the likelihood of observations often differ slightly. The second behavior can be observed when the robot moves through a known area. In this case, due to the improved proposal distribution, each particle keeps localized within its own map and the weights are more

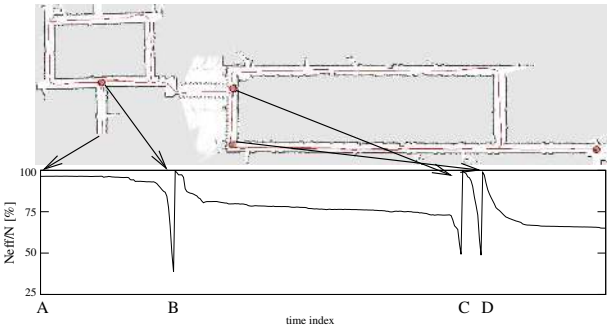

Fig. 7. The graph plots the evolution of the N_{eff} function over time during an experiment in the environment shown in the right image. The time index labeled with B corresponds to the closure of the small loop, and C, D corresponds to the big loop closure.

or less equal, depending on the amount of inconsistencies detected among the acquired scan and maps. This results in a constant behavior of N_{eff}. Finally, when closing a loop, some particles are correctly aligned with their map and some other are not. The correct particles have a high weight, while the wrong ones have a low weight. Thus the variance of the importance weights increases and N_{eff} substantially drops.

Accordingly, our threshold criterion on N_{eff} typically forces a resampling action when the robot is closing a loop. In all other cases the resampling is avoided, this way keeping the necessary variety in the particle set. As a result, the particle depletion problem is seriously reduced. To analyze this we performed an experiment in which we compared the success rate of our algorithm to that of a particle filter which resamples at every step. As Figure 6 illustrates our approach more often converged to the correct solution (MIT curve) for the MIT data set compared to the particle filter with the same number of particles and a fixed resampling strategy (MIT-2 curve).

VI. Conclusions

In this paper we presented an improved approach to learning grid maps with Rao-Blackwellized particle filters. Based on the likelihood model of a scan-matching process for last laser range scan our approach computes a highly accurate proposal distribution. This allows to seriously reduce the number of required samples. Additionally we apply a selective resampling strategy based on the number of effective particles. This approach reduces the number of resampling steps in the particle filter and thus limits the particle depletion problem.

The approach has been implemented and evaluated on data acquired with different mobile robots equipped with laser range scanners. Tests performed with our algorithm in different large-scale environments have demonstrated its robustness and the ability of generating high quality maps. In these experiments the number of particles needed by our approach often was by one order of magnitude smaller compared to previous approaches.

Acknowledgment

This work has partly been supported by the Marie Curie program under contract number HPMT-CT-2001-00251 and by the German Science Foundation (DFG) under contract number SFB/TR-8 (A3), and by the EC under contract number FP6-004250-CoSy. The authors also acknowledge Mike Bosse and John Leonard for providing the dataset of the MIT Killian Court.

References

[1] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte carlo localization for mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1998.

[2] 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.

[3] A Doucet. On sequential simulation-based methods for bayesian filtering. Technical report, Signal Processing Group, Departement of Engeneering, University of Cambridge, 1998.

[4] A. Doucet, J.F.G. de Freitas, K. Murphy, and S. Russel. Rao-blackwellized partcile filtering for dynamic bayesian networks. In *Proc. of the Conf. on Uncertainty in Artificial Intelligence (UAI)*, 2000.

[5] A. Eliazar and R. Parr. DP-SLAM: Fast, robust simultainous localization and mapping without predetermined landmarks. In *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2003.

[6] U. Frese and G. Hirzinger. Simultaneous localization and mapping - a discussion. In *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2001.

[7] J.-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proc. of the International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 2000.

[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 Int. Conf. on Intelligent Robots and Systems (IROS)*, 2003.

[9] D. Hähnel, W. Burgard, B. Wegbreit, and S. Thrun. Towards lazy data association in slam. In *Proc. of the International Symposium of Robotics Research (ISRR)*, 2003.

[10] D. Hähnel, D. Schulz, and W. Burgard. Map building with mobile robots in populated environments. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.

[11] J.S. Liu. Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statist. Comput.*, 6:113–119, 1996.

[12] M. Montemerlo, S. Thrun D. Koller, and B. Wegbreit. FastSLAM 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges. In *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2003.

[13] M. Montemerlo and S. Thrun. Simultaneous localization and mapping with unknown data association using FastSLAM. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2003.

[14] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to simultaneous localization and mapping. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 2002.

[15] H.P. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, pages 61–74, Summer 1988.

[16] K. Murphy. Bayesian map learning in dynamic environments. In *Neural Info. Proc. Systems (NIPS)*, 1999.

[17] M.K. Pitt and N. Shephard. Filtering via simulation: auxilary particle filters. Technical report, Department of Mathematics, Imperial College, London, 1997.

[18] S. Thrun. An online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 2001.

[19] R. van der Merwe, N. de Freitas, A. Doucet, and E. Wan. The unscented particle filter. Technical Report CUED/F-INFENG/TR380, Cambridge University Engineering Department, August 2000.