

Recovering Particle Diversity in a Rao-Blackwellized Particle Filter for SLAM After Actively Closing Loops

Cyrril Stachniss[†]

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

Giorgio Grisetti^{†‡}

[‡] *Dipartimento Informatica e Sistemistica
Università “La Sapienza”
I-00198 Rome, Italy*

Wolfram Burgard[†]

Abstract—Acquiring models of the environment belongs to the fundamental tasks of mobile robots. Approaches addressing the problem of simultaneous localization and mapping (SLAM) typically process the perceived sensor data and do not influence the motion of the mobile robot. In this paper, we present an approach to actively closing loops during exploration. It applies a Rao-Blackwellized particle filter to maintain multiple hypotheses about potential trajectories of the robot and corresponding maps. To prevent the particle filter from becoming overly confident, we present a technique to recover the particle diversity after successfully closing a loop. This way the particle depletion problem is avoided. The combination of our approach with the active loop closing strategy allows to deal with multiple nested loops. Experimental results presented in this paper illustrate the advantage of our method over previous approaches to mapping with Rao-Blackwellized particle filters.

I. INTRODUCTION

Simultaneous localization and mapping belongs to one of the fundamental problems in mobile robotics. Robots that are able to concurrently acquire a model of their environment and to localize themselves relatively to this model are regarded as fulfilling a major precondition of truly autonomous mobile vehicles. Recently, Rao-Blackwellized particle filters (RBPF) have been introduced as an effective means for solving the SLAM problem with occupancy grid maps [4, 17]. The key idea of this technique is to use a particle filter in which each particle carries its own map. The individual maps are computed based on the trajectory of the corresponding particle.

Whenever robots build maps of unknown terrain autonomously, the question arises where to move next to acquire useful sensor data. As we demonstrated in a previous work [19], the quality of a map constructed by a mobile robot depends on its trajectory during data acquisition. This is typically due to the fact that the vehicle needs to re-localize itself during exploration to build an accurate environmental model. A good pose estimation is necessary to make the correct data association, i.e., to determine if the current measurement fits into the map built so far. Avoiding repeated visits of the same place reduces the probability of making correct data associations and therefore increases the risk that the filter does not converge to the correct solution. Therefore, exploration strategies, which always guide the robot to unknown areas,

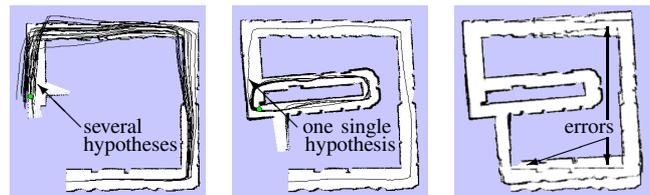


Fig. 1. This figure illustrates that a loss of particle diversity introduced by repeated loop closing can lead to a wrong solution in the context of mapping with a Rao-Blackwellized particle filter.

perform typically worse compared to a system considering place re-visiting actions.

However, especially if a Rao-Blackwellized particle filter is used to solve the SLAM problem, re-visiting known areas can be problematic. Consider, for example, a robot that moves from place A to place B and then repeatedly observes B . While it is mapping B it does not get any further information about A . Since each particle represents a whole trajectory of the robot, also hypotheses representing ambiguities about A will vanish while reducing potential uncertainties about B . The loss of uncertainty about A can later result in filter divergence if the robot returns to A . In practice, such situations occur when the robot has to map environments with nested loops. If the vehicle repeatedly traverses an inner loop, it can lose particles that are necessary to correctly close an outer loop (see Figure 1, for example). This phenomenon is known as the particle depletion problem [21].

In this paper, we consider the problem of actively closing loops during exploration. In our recent work [19], the robot is able to detect loops and to actively follow the previously traversed path to reduce its pose uncertainty. Whereas this approach has been demonstrated to be very effective and to yield accurate maps, its major drawback lies in the heuristic stopping criterion that aborts the loop closing process. Due to the risk of particle depletion, the robot should spend only a limited amount of time in an inner loop and has to leave an inner loop whenever the uncertainty drops below a certain threshold computed from the uncertainty when it entered the loop [19].

The contribution of this paper is a new technique to recover the uncertainty of the Rao-Blackwellized particle filter in

the context of nested loops. Our approach determines an approximation of the posterior represented by the particles at the entry of a loop and propagates its uncertainty through the loop. This way, hypotheses needed to close an outer loop are maintained. The major advantage of this approach is that the robot can, in principle, stay arbitrary long in an inner loop without losing information necessary to close outer loops.

This paper is organized as follows. After the discussion of related work, we explain the idea of Rao-Blackwellized mapping in Section III. Section IV presents our exploration technique with active loop closing. Section V then describes how to recover the diversity of a particle filter when the robot leaves a loop. Finally, Section VI contains experimental results carried out on real robots as well as in simulation.

II. RELATED WORK

In the context of exploration, many techniques presented so far focus on generating motion commands that minimize the time needed to cover the whole terrain [2, 11, 22]. Most of these approaches, however, assume that an accurate position estimation of the robot is given during exploration. In contrast to this, Ko *et al.* [10] presented a multi-robot exploration system based on a mapping technique similar to that of Gutmann and Konolige [7] which is able to deal with unknown relative start locations of the robots. In contrast to our approach, their robots consider rendezvous instead of place re-visiting actions to re-localize.

In the area of SLAM, the vast majority of papers has focused on the aspect of state estimation as well as belief representation and update [3, 6, 4, 5, 7, 8, 16, 17, 20]. These techniques, however, are passive and only process incoming sensor data without explicitly generating control commands. Recently, some techniques have been proposed which actively control the robot during SLAM. For example, Makarenko *et al.* [15] introduced an utility function which trades-off the cost of exploring new terrain with the utility of selected positions with respect to a potential reduction of the pose uncertainty. Sim *et al.* [18] presented an approach in which the robot follows a parametric curve to explore the environment. Both techniques integrate the pose uncertainty into the decision process of where to move next. However, they rely on the fact that the environment contains landmarks that can be uniquely determined during mapping. In contrast to this, our approach makes no assumptions about distinguishable landmarks and uses raw laser range scans to compute accurate grid maps.

Only very few works address the problem of revoking a previously made decision in the SLAM context. Hähnel *et al.* [9] maintain a data association tree in which each branch represents a sequence of associations. Whenever a branch becomes more likely than the current best one their approach switches to the alternative data association sequence. Their work can be regarded orthogonal to our technique for recovering the uncertainty of a particle filter. In fact, both approaches can be combined.

Our approach presented here extends our previous work [19] and presents a way to recover particle diversity when apply-

ing a Rao-Blackwellized particle filter to solve the SLAM problem. Our algorithm enhances the ability to correctly close loops, especially, in the context of nested loops.

III. RAO-BLACKWELLIZED MAPPING

According to Murphy [17], the key idea of solving the SLAM problem with a Rao-Blackwellized particle filter (RBPF) is to estimate a posterior $p(x_{1:t} | z_{1:t}, u_{0:t-1})$ about potential trajectories $x_{1:t}$ of the robot given its observations $z_{1:t}$ and its odometry measurements $u_{0:t-1}$. This distribution is then used to compute a posterior over maps and trajectories:

$$p(x_{1:t}, m | z_{1:t}, u_{0:t-1}) = p(m | x_{1:t}, z_{1:t})p(x_{1:t} | z_{1:t}, u_{0:t-1}) \quad (1)$$

Eq. (1) can be solved efficiently, since the quantity $p(m | x_{1:t}, z_{1:t})$ can be computed analytically once $x_{1:t}$ and $z_{1:t}$ are known. To estimate $p(x_{1:t} | z_{1:t}, u_{0:t-1})$ over the potential trajectories, Rao-Blackwellized mapping uses a particle filter in which an individual map is associated to each sample. Each of those maps are constructed given the observations $z_{1:t}$ and the trajectory $x_{1:t}$ represented by the corresponding particle. During resampling, the weight ω of each particle is proportional to the likelihood $p(z_t | m, x_t)$ of the most recent observation z_t given the map m associated to this particle and its pose x_t . Throughout this paper we use a highly optimized variant of the original algorithm for mapping with RBPFs. An efficient implementation is necessary due to the online requirement needed for autonomous exploration.

IV. EXPLORATION WITH ACTIVE LOOP-CLOSING

The goal of an exploration task is to minimize the uncertainty of the robot about the world. The uncertainty of a posterior can be determined by the entropy \mathcal{H} . In the following we derive how to compute the entropy of a Rao-Blackwellized particle filter for mapping. To improve the readability, we use d to refer to all sensor and odometry measurements and x to represent the whole trajectory of the vehicle. For the entropy holds

$$\mathcal{H}(p(m, x | d)) = \mathcal{H}(p(m | x, d)p(x | d)) \quad (2)$$

$$= - \int_x \int_m p(m | x, d)p(x | d) \cdot \log(p(m | x, d)p(x | d)) dm dx \quad (3)$$

$$= - \int_x \int_m p(m | x, d) dm \cdot p(x | d) \log(p(x | d)) dx - \int_x \int_m p(m | x, d)p(x | d) \log(p(m | x, d)) dm dx \quad (4)$$

$$= - \int_x p(x | d) \log(p(x | d)) dx - \int_x p(x | d) \int_m p(m | x, d) \log(p(m | x, d)) dm dx \quad (5)$$

$$= \mathcal{H}(p(x | d)) + \int_x p(x | d) \mathcal{H}(p(m | x, d)) dx. \quad (6)$$

Eq. (5) is obtained from Eq. (4) since $\int_m p(m | x, d) dm = 1$. This derivation shows that in the context of Rao-Blackwellized

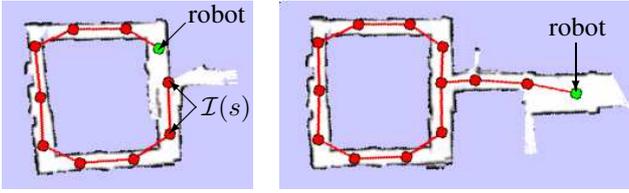


Fig. 2. The red/gray circles and lines in these two image represent the nodes and edges of $\mathcal{G}^{[s]}$. In the left image, $\mathcal{I}(s)$ contained two nodes. After closing the loop the robot continues to acquire new terrain (right image).

particle filters the entropy can be divided into two components. Whereas the first term represents the entropy of the posterior about the trajectory of the robot, the second term corresponds to the uncertainty about the map weighted by the likelihood of the corresponding trajectory. Thus, to minimize the robot's overall uncertainty, one needs to reduce the map uncertainty of the individual particles as well as the pose uncertainty.

Throughout this paper we mainly focus on the first component of Eq. (6), although a reduction of the uncertainty in the posterior about the trajectory of the robot typically assigns small weights to inconsistent maps, which leads to a reduction of the second term too. To minimize the uncertainty about the trajectory of the robot our approach identifies opportunities to actively close loops during terrain acquisition. This approach actively re-enters known areas of the map and follows a previously traversed path. As a result, unlikely path hypotheses are eliminated from the particle set or get a low importance weight so that the overall uncertainty is reduced.

To determine whether there exists a possibility to close a loop we consider two different environmental representations. Each particle s maintains an occupancy grid map $m^{[s]}$ and a topological map $\mathcal{G}^{[s]}$ during the exploration task. The vertices in $\mathcal{G}^{[s]}$ represent positions visited by the robot and the trajectory of particle s corresponds to the edges in $\mathcal{G}^{[s]}$. New nodes are created and added to the graph structure whenever the robot moved for a certain distance or it cannot observe any previously created node from its current location.

There exist techniques to combine topological maps with other kind of spacial representations. This is typically done to handle large-scale maps or to simplify the loop-closing problem [1, 12, 13]. Those approaches can attach detailed local maps to the nodes of the topological map. Building up such a hierarchy is not intended by our work, since we only use the topological map to detect loops in the environment.

Figure 2 shows such a graph for one particular particle during different phases of an exploration task. In both images the topological map $\mathcal{G}^{[s]}$ is depicted on top of metric map $m^{[s]}$. To motivate the idea of our active loop-closing algorithm, consider the left image of Figure 2. Here the robot is almost closing a loop. This can be detected by the fact that the length of the shortest path between the current pose of the robot and previously visited locations in the topological map $\mathcal{G}^{[s]}$ is large, whereas it is small in the grid map $m^{[s]}$. Thus, to determine whether or not a loop can be closed we compute for each sample s the set $\mathcal{I}(s)$. This set of positions of interest contains all nodes that are close to the current pose $x_t^{[s]}$ of

particle s based on the grid map $m^{[s]}$ but are far away given the topological map $\mathcal{G}^{[s]}$

$$\mathcal{I}(s) = \{x_{t'}^{[s]} \in \text{nodes}(\mathcal{G}^{[s]}) \mid \text{dist}_{m^{[s]}}(x_{t'}^{[s]}, x_t^{[s]}) < c_1 \wedge \text{dist}_{\mathcal{G}^{[s]}}(x_{t'}^{[s]}, x_t^{[s]}) > c_2\}. \quad (7)$$

Here $\text{dist}_m(x_1, x_2)$ is the length of the shortest path from x_1 to x_2 given the grid map and $\text{dist}_{\mathcal{G}}(x_1, x_2)$ the shortest path based on the topological map. The terms c_1 and c_2 are constants that must satisfy the constraint $c_1 < c_2$.

If $\mathcal{I}(s) \neq \emptyset$, there exist so-called shortcuts from $x_t^{[s]}$ to the positions in $\mathcal{I}(s)$. These shortcuts represent edges that would generate a cycle in the graph structure of $\mathcal{G}^{[s]}$ and therefore constitute an opportunity to close a loop (compare Figure 2). To determine the most likely movement that guides the robot around the loop, one in principle has to integrate over all particles and consider all potential outcomes of that particular action. Since this would be too time consuming for online-processing we consider only the particle s^* with the highest accumulated logarithmic importance weight

$$s^* = \underset{s}{\operatorname{argmax}} \sum_{t'=1}^t \log \omega_{t'}^{[s]} \quad (8)$$

Here $\omega_t^{[s]}$ is the weight of sample s at time step t . If $\mathcal{I}(s^*) \neq \emptyset$, we choose the node x_{t_e} from $\mathcal{I}(s^*)$ which is closest to $x_t^{[s^*]}$:

$$x_{t_e} = \underset{x \in \mathcal{I}(s^*)}{\operatorname{argmin}} \text{dist}_{m^{[s^*]}}(x_t^{[s^*]}, x) \quad (9)$$

In the sequel, x_{t_e} is denoted as the *entry point* at which the robot has the possibility to close a loop. The term t_e corresponds to the last time the robot was at the node x_{t_e} .

Once the robot detects a loop, it approaches x_{t_e} and then follows the path taken after previously arriving at x_{t_e} . During this process the uncertainty in the pose of the vehicle typically decreases, because the robot is able to localize itself in the map built so far and unlikely particles vanish.

Furthermore, we have to define a criterion for deciding when the robot actually has to stop following a loop. As illustrated in the introduction (see Figure 1), this criterion can have a major influence on the quality of the resulting map. Therefore, one has to find an adequate criterion which allows the filter to eliminate unlikely hypotheses but at the same time makes the convergence to a wrong solution unlikely. The termination criterion of the loop-closing process is the main weakness of our previous approach [19] in which the robot has to leave an inner loop whenever the uncertainty drops below a certain threshold computed from the uncertainty when it entered the loop. In the following we will introduce an alternative criterion which takes into account how well the current sensor information is suited to identify unlikely hypotheses within the particle set.

To estimate how well the current set of N particle represents the true posterior Liu [14] introduced the effective number of particles N_{eff} (also called effective sample size):

$$N_{eff} = \frac{1}{\sum_{s=1}^N (\omega^{[s]})^2} \quad (10)$$

The idea behind this measure is to determine the variance in the importance weights of the particles. Liu uses N_{eff} to resample in an intelligent way but it is also very useful in the context of active loop-closing. We monitor the change of N_{eff} over time, which allows to analyze how the new acquired information affects the filter. If N_{eff} stays constant the new information does not help to identify unlikely hypotheses represented by the individual particles. In that case, the variance in the importance weights of the particles does not change over time. If, in contrast, the value of N_{eff} decreases over time, the new information can be used to identify that some particles are less likely than others. This is exactly the criterion we need to decide whether or not the loop-closing should be aborted. As long as new information helps to identify unlikely particles we follow the loop. As soon as the observations do not provide any new knowledge about the environment we continue to explore new terrain. In the experimental section we will show how N_{eff} behaves when closing loops.

As long as no loop is detected, we use a frontier-based exploration strategy [2, 22] to choose target points for the robot. In our current system we determine frontiers based on the map of the most likely particle s^* . According to Yamauchi [22], a frontier is any known cell that is an immediate neighbor of an unknown, unexplored cell.

V. RECOVERING PARTICLE DIVERSITY AFTER LOOP CLOSURE

In addition to the active loop-closing technique described above, we need a way to recover hypotheses vanished from a particle filter during the repeated traversal of an inner loop. Even if the stopping criterion based on N_{eff} makes particle depletion unlikely, the vanishing of important hypotheses and the resulting problem of filter divergence remains. Note that the risk of particle depletion increases with the size of the environment. Also, the smaller the number of particles, the higher is that risk.

Consider, as an example, a robot that has accurately mapped an inner loop. In such a case the particle filter will have converged to a very peaked distribution and typically only one hypothesis present at the entry point will have survived. Thus, it is not guaranteed that this hypothesis is the one which perfectly closes the outer loop. In principle, a robot therefore has to maintain a sufficient variety of particles allowing it to perform the next loop closure. Since the robot does not know in advance how many loops it will find in the environment this problem cannot be solved in general with a finite number of particles only.

If one knew the starting point of such an inner loop in advance, one solution would be to suspend the particle filter and to start for each particle an individual RBPF initialized with the current state of that particle. After the convergence of all filters one can then attach their solutions to the corresponding particles in the suspended filter. Apart from the fact that a loop cannot be recognized in advance this approach is not feasible for online tasks like exploration since the amount

of computational resources needed grows exponentially in the number of loops.

The recovering technique described in this section is an approximation of this approach. The key idea is to simulate this process as soon as the robot detects a loop. Given the current set of particles, the robot computes the posterior at the entry point of the loop given the particles in its current belief. In this approximative particle set the states and weights are computed according to

$$\tilde{x}_{t_e}^{[s]} = \text{pred}_{t_e}(x_t^{[s]}) \quad (11)$$

$$\tilde{\omega}_{t_e}^{[s]} = \omega_t^{[s]}. \quad (12)$$

Here $\text{pred}_{t_e}(x_t^{[s]})$ is the state of the ancestor of $x_t^{[s]}$ at time t_e . Whenever the robot stops the loop closing behavior it uses this posterior to propagate the variety of the particles through the loop. In probabilistic terms this corresponds to rewriting the term $p(x_{1:t} | z_{1:t}, u_{0:t-1})$ in Eq. (1) in the following way:

$$p(x_{1:t} | z_{1:t}, u_{0:t-1}) = p(x_{t_e+1:t} | x_{1:t_e}, z_{t_e+1:t}, u_{t_e:t-1})p(x_{1:t_e} | z_{1:t_e}, u_{0:t_e-1}) \quad (13)$$

In our current implementation this posterior is approximated by importance sampling from $p(x_{t_e+1:t} | x_{1:t_e}, z_{t_e+1:t}, u_{t_e:t-1})$. The trajectory drawn from this posterior is attached to each particle in $p(x_{1:t_e} | z_{1:t_e}, u_{0:t_e-1})$. This process propagates the different hypotheses from the entry point into the current belief before leaving the loop. If the robot then has to close a second loop it has a higher chance to find hypotheses closing this loop accurately.

Note that in general a mapping system has to maintain a stack of saved states especially in environments with several nested loops. Due to the fact that we control the robot actively and never start a second loop-closing process before completing the current one, we only have to maintain a single saved state at each point in time.

As we demonstrate in the experiments this technique is a powerful tool to recover vanished hypotheses without restarting the mapping algorithm from scratch. It only needs to attach a local trajectory to each particle which can be done within a few seconds (on a 2.8GHz Pentium IV).

VI. EXPERIMENTS

Our approach has been implemented and evaluated in real world and in simulation. The experiments described here are designed to illustrate the benefit of our active loop closing technique with the ability to recover the diversity of the particles after loop closing. We also demonstrate how N_{eff} evolves during exploration and why it is an useful criterion to stop the loop-closing process. Furthermore we discuss the advantages of our approach compared to our previous algorithm described in [19].

A. Recovering the Particle Diversity

This experiment is designed to show the effect of our technique to recover the particle variety when the robot leaves a loop. The environment used to carry out this experiment is



Fig. 3. This figure shows the same experiment as depicted in Figure 1, but using our recovering technique. In the left image the robots saves the set of approximated particles at time step t_e and later on recovers the vanished hypotheses (middle image). This allows the robot to correctly close the outer loop (right image).

depicted in the right image of Figure 3. The robot started in the outer loop, entered the inner loop, and stayed there for a long period of time. As shown in Figure 1, without our recovering technique the filter can converge to a wrong solution. The reason is that at the time when the robot leaves the loop only one hypothesis of the original particle set at the entry point has survived. Accordingly, the robot lacks an appropriate particle to properly close the outer loop. Using our algorithm, however, the robot can recover the hypotheses at the entry point and can propagate them through the loop (see left and middle image of Figure 3). The most likely map of the posterior after closing the outer loop is shown in the right image.

To provide a more quantitative analysis we mapped the environment 30 times without the capability of restoring the filter and 30 times with this option. The standard technique was able to build a correct map in only 40% of all runs. In all other cases the algorithm did not produce an accurate map. In contrast to this, our algorithm yielded a success rate of 93%. We repeated this experiment in different environments and got similar results. Figure 4 shows two (partial) maps of the Killian Court at the MIT. The left map has been built without the recovering technique using 40 particles and shows inconsistencies due to vanished hypotheses. The right map has been constructed using our recovering technique in which the correct hypothesis has been restored. The average success rate of our approach was 55% whereas the standard approach found the correct data association in only 5% of all 40 runs.

This shows that our recovering technique is a powerful extension to autonomous exploration with mapping systems based on RBPFs especially in the context of (multiple) nested loops. Note that in general the success rate of the standard approach increases with number of particles used. Since each particle carries its own map, it is of utmost importance to keep this value as small as possible. Therefore our approach also can be regarded as a contribution to limit the number of particles during Rao-Blackwellized mapping.

Additionally, we analyzed in our experiments the approximation error obtained by retrospectively recovering the particles at the entry point of a loop. Using this system we observed that in our experiments typically around 75% of the particles in the filter at time step t_e had a successor in the current set and were therefore saved. In principle, this value must drop for loops of increasing length. However, in our experiments we found similar values in several different

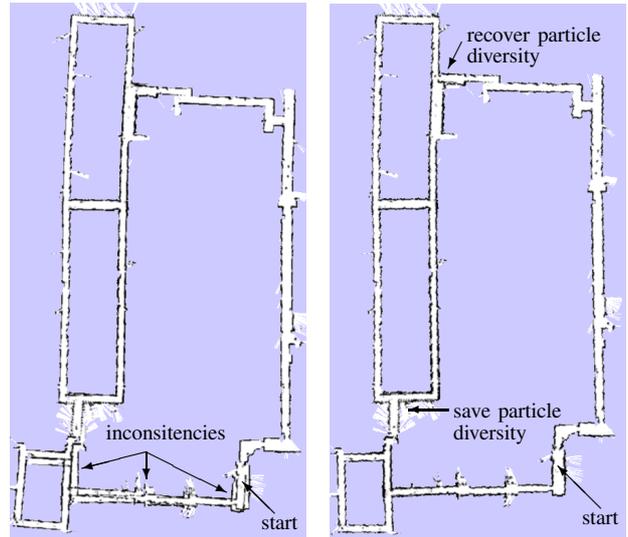


Fig. 4. This figure shows two maps of the Killian Court at the MIT. The size of the environment is $150\text{m} \times 80\text{m}$. The left map was constructed with the standard RBPF approach. If, in contrast, the robot is able to recover hypotheses the map becomes more accurate (right image).

environments. The Kullback-Leibler distance (KL-distance) between the recovered particle set and the true one at time step t_e was between 1.0 and 1.5 compared to a value around 13 in the situation in which only a single hypothesis survived. This again illustrates the advantage of our new technique over the previous approach [19]. Using the technique described here the robot can move arbitrarily long through a nested loop without losing the capability to close an outer loop.

B. Using N_{eff} to Stop the Loop-Closing Procedure

In this experiment we analyze the constraint that terminates the active loop-closing behavior. The process is stopped whenever N_{eff} stays constant for a certain period of time.

One typical evolution of N_{eff} is depicted in the left image of Figure 5. The robot started at position A and in the first part of the experiment explored unknown terrain (between the positions A and B). As can be seen, N_{eff} decreases over time. After the loop has been closed correctly and unlikely hypotheses had partly been removed by the resampling action (position B) the robot processed the inner loop and N_{eff} stayed more or less constant. This indicates that acquiring further data in this area has only a very small effect on the relative likelihood of the particles and the system could not determine which hypotheses represented unlikely configurations. In such a situation, it therefore makes more sense to focus on new terrain acquisition and to not continue the loop-closing process.

If the robot takes into account the evolution of N_{eff} to stop the loop-closing procedure, it moves through this loop only as long as it can acquire useful data to identify unlikely particles. It collects enough data to make the correct data association. Unless the number of particles is not too small or the size of the loop is not too big, the filter will typically not converge to a wrong hypothesis. One example for an exploration task done with 30 particles in an environment with several loops is depicted in Figure 6. The loop closing procedure was executed

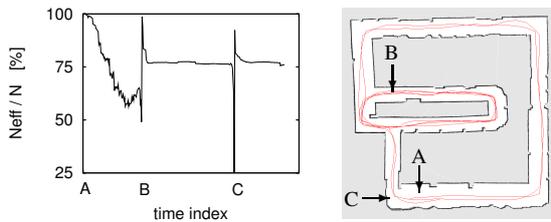


Fig. 5. The graph plots the evolution of the N_{eff} function over time during an experiment in the environment shown in the right image. The robot started at position A. The position B corresponds to the closure of the inner loop, and C corresponds to closure of the outer loop.

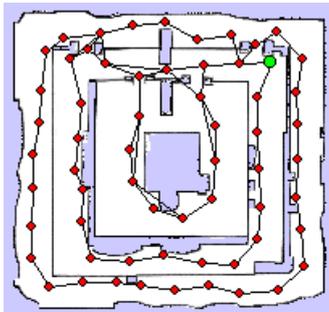


Fig. 6. The map acquired by autonomous exploration in an environment with 3 loops and a size of $40\text{m} \times 40\text{m}$. The most likely grid map is depicted in the background whereas the topological map is shown in the foreground.

four times during this exploration task. As can be seen, all loops have been closed accurately.

Furthermore, we analyzed the length of the trajectory traveled by the robot. Due to the active loop-closing our technique generates longer trajectories compared to a purely frontier-based exploration strategy. We performed several experiments in different environments and measured the average overhead. Typically the overhead was less than 10%, but it obviously depends on number of loops in the environment and their length.

VII. CONCLUSION

In this paper we presented a novel approach to maintaining particle diversity during actively closing loops in a mapping system based on a Rao-Blackwellized particle filter. When closing a loop our approach determines an approximation of the particle set at the time the robot entered the loop. It uses this posterior to propagate the particle diversity through the loop. Compared to previous approaches this allows the robot to traverse a nested loop for an arbitrary amount of time without depleting important particles.

The approach has been implemented and tested on real robot data as well as in simulation. As experimental results demonstrate, we obtain a robust exploration algorithm that produces more accurate maps compared to standard combinations of SLAM algorithms with exploration techniques, especially, in the context of nested loops.

ACKNOWLEDGMENT

This work has partly been supported by the German Science Foundation (DFG) under contract number SFB/TR-8, by the

Marie Curie program under contract number HPMT-CT-2001-00251, and by the EC under contract number FP6-004250-CoSy. Furthermore, we would like to thank John Leonard and Mike Bosse for providing us the Killian Court dataset.

REFERENCES

- [1] M. Bosse, P.M. Newman, J.J. Leonard, and S. Teller. An atlas framework for scalable mapping. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2003.
- [2] W. Burgard, M. Moors, D. Fox, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2000.
- [3] 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.
- [4] A. Doucet, J.F.G. de Freitas, K. Murphy, and S. Russel. Rao-blackwellized particle 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 simultaneous localization and mapping without predetermined landmarks. In *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2003.
- [6] U. Frese, P. Larsson, and T. Duckett. A multilevel relaxation algorithm for simultaneous localisation and mapping. *IEEE Transactions on Robotics*. to appear.
- [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] J. Ko, B. Stewart, D. Fox, K. Konolige, and B. Limketkai. A practical, decision-theoretic approach to multi-robot mapping and exploration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2003.
- [11] S. Koenig and C. Tovey. Improved analysis of greedy mapping. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2003.
- [12] B. Kuipers and Y.T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and Autonomous Systems*, 1981.
- [13] B. Lisen, D. Silver D. Morales, G. Kantor, I.M. Rekleitis, and H. Choset. Hierarchical simultaneous localization and mapping. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2003.
- [14] J.S. Liu. Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statist. Comput.*, 6:113–119, 1996.
- [15] A.A. Makarenko, S.B. Williams, F. Bourgoult, and F. Durrant-Whyte. An experiment in integrated exploration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.
- [16] 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.
- [17] K. Murphy. Bayesian map learning in dynamic environments. In *Neural Info. Proc. Systems (NIPS)*, 1999.
- [18] R. Sim, G. Dudek, and N. Roy. Online control policy optimization for minimizing map uncertainty during exploration. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2004.
- [19] C. Stachniss, D. Hähnel, and W. Burgard. Exploration with active loop-closing for FastSLAM. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2004.
- [20] S. Thrun. An online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 2001.
- [21] 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.
- [22] B. Yamauchi. Frontier-based exploration using multiple robots. In *Proc. of the Second Int. Conf. on Autonomous Agents*, 1998.