

An Efficient FastSLAM Algorithm for Generating Maps of Large-Scale Cyclic Environments from Raw Laser Range Measurements

Dirk Hähnel¹ and Wolfram Burgard¹ and Dieter Fox² and Sebastian Thrun³

¹University of Freiburg, Department of Computer Science, Freiburg, Germany

²University of Washington, Computer Science and Engineering, Seattle, WA, USA

³Stanford University, Computer Science Department, Stanford, CA, USA

Abstract

The ability to learn a consistent model of its environment is a prerequisite for autonomous mobile robots. A particularly challenging problem in acquiring environment maps is that of closing loops; loops in the environment create challenging data association problems [9]. This paper presents a novel algorithm that combines Rao-Blackwellized particle filtering and scan matching. In our approach scan matching is used for minimizing odometric errors during mapping. A probabilistic model of the residual errors of scan matching process is then used for the resampling steps. This way the number of samples required is seriously reduced. Simultaneously we reduce the particle depletion problem that typically prevents the robot from closing large loops. We present extensive experiments that illustrate the superior performance of our approach compared to previous approaches.

I. Introduction

Learning maps with mobile robots is one of the fundamental problems in mobile robotics. In the literature, the mobile robot mapping problem is often referred to as the *simultaneous localization and mapping problem (SLAM)* [5], [6], [9], [12], [13], [19]. This is because mapping includes both, estimating the position of the robot relative to the map and generating a map using the sensory input and the estimates about the robot's pose.

One of the hardest problems in robotic mapping is that of loop closure. As a robot traverses a large cycle in the environment, it faces the hard data association of correctly connecting to its own map under large position errors. This problem has long been acknowledged for its hardness, and a number of approaches have addressed it [4], [9], [20]. Recently, Murphy and colleagues have presented Rao-Blackwellized particle filters [8], [17] as an effective way of representing alternative hypotheses on robot paths and associated maps. Montemerlo *et al.* [14], [15] extended this idea to efficient landmark-based SLAM using Gaussian representations to and were the first to successfully realize it on real robots.

In this paper we present a highly efficient approach to simultaneous localization and mapping with laser scans.

As previously proposed by Murphy [17], our approach applies a Rao-Blackwellized particle filter to estimate a posterior of the path of the robot, in which each particle has associated to it an entire map. This differs from work in [20], where only a single map is retained. To scale to large-scale environments, we transform sequences of laser range-scans into odometry measurements using range-scan registration techniques [10]. This way our system can deal with significantly larger environments than Murphy's approach [17], since the scan matching yields odometry estimates that are an order of magnitude more accurate than the raw wheel encoder data. Simultaneously, the transformation of sequences of scans into odometry measurements reduces the well-known particle deprivation problem [21], since the number of resampling operations is significantly reduced. By using a learned model of the residual errors of the range registration our approach can correctly integrate the corrected odometry into the particle filtering process. As a result, we obtain a drastic reduction in the number of particles needed to build large-scale maps, or, put differently, an improved ability to map large environments. This is demonstrated in our experimental results section, in which we compare our approach to previous techniques.

This paper is organized as follows. In the following section, we will discuss techniques for incremental probabilistic mapping and localization. In Section III, we describe our approach to integrate scan matching with Rao-Blackwellized particle filters to achieve a robust approach for simultaneous mapping and localization. Section IV presents several experiments illustrating that our approach can successfully learn accurate maps with range scanners in large-scale environments. Additionally, we present experiments illustrating that our technique outperforms existing approaches.

II. Incremental Probabilistic Mapping and Localization

In probabilistic terms the goal of map learning is to find the map and the robot positions which yield the best interpretation of the data d_t gathered by the robot [19]. Here the data $d_t = \{u_{0:t-1}, z_{1:t}\}$ consists of a stream

of odometry measurements $u_{0:t-1}$ and perceptions of the environment $z_{1:t}$. The mapping problem can be phrased as recursive Bayesian filtering for estimating the robot positions along with a map of the environment:

$$p(x_{1:t}, m \mid z_{1:t}, u_{0:t-1}) = \alpha \cdot p(z_t \mid x_t, m) \cdot \int p(x_t \mid x_{t-1}, u_{t-1}) p(x_{1:t-1}, m \mid z_{1:t-1}, u_{0:t-2}) dx_{1:t-1} \quad (1)$$

In probabilistic mapping and localization it is typically assumed that the odometry measurements are governed by a so-called probabilistic motion model $p(x_t \mid x_{t-1}, u_{t-1})$ which specifies the likelihood that the robot is at x_t given that it previously was at x_{t-1} and the motion u_{t-1} was measured. On the other hand, the observations follow the so-called observation model $p(z \mid x)$, which defines for every possible location x in the environment the likelihood of the observation z .

Unfortunately, estimating the full posterior in Equation 1 is not tractable in general. One popular approach is to restrict observations to landmark detections, and represents robot positions by Gaussians [6]. In this context, the Bayes filter can be approximated efficiently by an EKF for which the state consists of the robot positions along with positions of the landmarks. Other researchers attempted to overcome the restrictions to landmark observations by using laser range-finders and incremental scan matching [2], [18], [22]. The general idea of these approaches can be summarized as follows (see also [19]). At any point $t-1$ in time, the robot is given an estimate of its pose \hat{x}_{t-1} and a map $\hat{m}(\hat{x}_{1:t-1}, z_{1:t-1})$. After the robot moved further on and after taking a new measurement z_t , the robot determines the most likely new pose \hat{x}_t such that

$$\hat{x}_t = \underset{x_t}{\operatorname{argmax}} \{ p(z_t \mid x_t, \hat{m}(\hat{x}_{1:t-1}, z_{1:t-1})) \cdot p(x_t \mid u_{t-1}, \hat{x}_{t-1}) \}. \quad (2)$$

It does this by trading off the consistency of the measurement with the map (first term on the right-hand side in (2)) and the consistency of the new pose with the control action and the previous pose (second term on the right-hand side in (2)). The map is then extended by the new measurement z_t , using the pose \hat{x}_t as the pose at which this measurement was taken.

Whereas this approach has the advantage that it yields accurate results and can be implemented efficiently if the registration is performed with respect to a global map or with respect to only a fixed number of scans. Its major disadvantage lies in the greedy maximization step. When the robot has to close larger loops, this approach suffers from registration errors during loop closures and therefore tends to fail in large environments. To overcome this problem, extensions of this approach have been developed which maintain a posterior about the position of the vehicle [9], [19]. The key idea of these techniques is to delay the maximization until the robot detects that a loop has been closed. This is usually done by identifying that the robot enters an already known area from an




Fig. 1. Graphical model of concurrent mapping and localization as filtering process.

unknown area and simultaneously observes a high likelihood of its observations for potential positions that are under consideration. If the registration of the vehicle in its map can be done with high likelihood, the previous poses are corrected backwards in time according to the pose correction that is necessary to properly close the loop. The position posterior is then replaced by a Dirac distribution which has its mode at the most likely position when the robot closes the loop. Furthermore, subsequent backwards corrections are stopped when the robot reaches this node. Whereas this approach can reliably close even large loops, it has the disadvantage that it under-estimates the uncertainty in the robot pose when closing loops.

More recently, Murphy and colleagues [17], [8] have presented Rao-Blackwellized particle filters as an efficient way to represent the full posterior of the robot pose. Figure 1 depicts a graphical model of Rao-Blackwellized simultaneous mapping and localization. The key idea of this approach is to solve the recursive Bayes filter update by the following equation:

$$p(x_{1:t}, m \mid z_{1:t}, u_{0:t-1}) = p(m \mid x_{1:t}, z_{1:t}, u_{0:t-1}) p(x_{1:t} \mid z_{1:t}, u_{0:t-1}) \quad (3)$$

Here, a particle filter is used to represent robot trajectories $x_{1:t}$ and a different map is conditioned on each sample of the particle filter. The importance weights of the samples are computed according to the likelihoods of the observations in the maximum likelihood map constructed using exactly the positions this particular particle has taken. The key advantage of this approach is that the samples approximate at every point in time the full posterior over robot poses and maps. The first successful realization on real robots of an extended version of this technique has been presented by Montemerlo *et al.* [15].

However, particle filters are known to be subject to major approximation errors. One of these errors is known as the particle depletion problem [21]. This problem can lead to a divergence of the filter and can result in the lack of particles in the vicinity of the correct state. In the SLAM context this can prevent the robot from closing a given loop. There are two parameters that have a major influence on the approximation error. First, the number of particles needs to be high enough to represent the posterior. However, too many samples can prevent




Fig. 2. Parameters of the probabilistic motion model.

the filtering process from being fast enough for online-processing. Furthermore, the number of resampling steps needs to be limited in order to avoid that the samples converge too quickly to the maximum likelihood state which is undesirable especially in ambiguous situations. On the other side, too few resampling steps could result in a divergence of the filter since many samples are wasted on unlikely states and the uncertainty typically introduced by robot motions would exceed the certainty gained by incorporating observations of the environment.

In the following section we will describe our solution to this problem. This approach transforms sequences of laser measurements into odometry measurements using a scan matching procedure and utilizes the remaining laser scans for map estimation.

III. Combining Laser-based FastSLAM with Scan Matching

Our approach to reduce the problems described above is to use a scan matching routine to correct the odometry and to use this corrected path information as input for the sampling step in the Rao-Blackwellized particle filter.

The 2d scan matching we apply, which is described in detail in [10], aligns a scan relative to the previous scans by computing an occupancy grid map [16] from the previous measurements. To avoid the time consuming ray-tracing operation required to compute the likelihood of a measurement $p(z | x)$ we apply an approximation which considers only the endpoint of a beam [11], [19]. This way, $p(z | x)$ can be computed efficiently using fast look-up operations. To also be able to incorporate maximum range measurements, our system assumes that the cell 20cm in front of that in which the maximum range measurement ends must be unoccupied.

A key question when combining a scan matching routine with a probabilistic technique is how to estimate the uncertainty of the scan matching process so as to correctly incorporate the corresponding uncertainty during the prediction step of the sampling procedure. In our current system we use a parametric model of the odometry error and learn the parameters of this model using data acquired during experiments. To learn the parameters of the model used for the experiments described here we performed an experiment in which we generated a statistics of alignment




Fig. 3. Sample densities obtained with the models for the raw odometry (left image) and for scan matching (right image) for ten incremental movements of a real robot.

errors after convergence of scan matching. Using a data set recorded in the Intel Research Lab Seattle, we applied our system equipped with a manually designed motion model. We then took the resulting map (see Figure 9) as ground truth and compared the raw odometry and the results of the scan matching with the positions corrected by the routine. The error model we use has three parameters as it assumes that in every single movement there are three errors involved (see also Figure 2). First, whenever the robot starts to move, it makes a small rotational error $\alpha' - \alpha$. Second, the robot introduces a certain error $d' - d$ to the distance between the final location and the starting position. Finally, the true final orientation differs by a certain amount from the measured orientation which is expressed by a non-zero difference $\beta - \beta'$. The means and the variances of the relative errors in these three parameters were learned by comparing the approximated displacement after convergence of the scan matching routine with the (estimated) ground truth information. Alternative models of odometry errors and corresponding techniques for parameter estimation have been proposed by Borenstein and Feng [3], Doh *et al.* [7], as well as Bengtsson and Baerveldt [1].

Figure 3 plots the resulting sample densities obtained when relying on pure odometry (left image) and the densities obtain with the error model for the scan matching process (right image). As the figure shows, the samples are much more focused if the scan matching routine is used. This leads to the desired effect that the variance of the posterior is reduced, that fewer samples can be used, and that larger loops can be closed.

A graphical model of our approach to integrate results from the scan matching process into the Rao-Blackwellized sampling routine is depicted in Figure 4 (c.f. 1). The key idea is to compute every k steps a new odometry measurement u'_j out of the $k - 1$ previous observations z and the k most recent odometry readings. The k -th observation is then used to compute the weights of the samples in the particle filter. Note that this clear separation between laser scans used for odometry and laser scans used for map estimation ensures that all information is used only once.

One important aspect when using Rao-Blackwellized particle filters for mapping is the efficient update of the maps of the individual particles. Montemerlo *et al.* [15]




Fig. 4. Graphical model of the integration of scan matching and probabilistic mapping.

proposed a tree-structure to efficiently update the map. In the system described here, we only use a limited number of scans gathered by the robot to update the map of a particle. The scans chosen need to intersect with the area visible according to the pose of the corresponding particle. This way, the update of the map of every particle can be achieved in constant time. Although this is an approximation only, we never found any evidence that the quality of the resulting maps was decreased significantly.

IV. Experimental Results

The approach described above has been implemented and tested using different robotic platforms and in different environments as well as in extensive simulation runs. In all experiments, we found out that the system can operate online and can also robustly close large and nested loops.

A. Mapping Large-Scale Environments with Multiple Cycles

The first experiment was carried out using a Pioneer 2 robot equipped with a SICK LMS laser range-finder in the Intel Research Lab, Seattle, WA. The size of this environment is 28m \times 28m. The robot traveled 491m with an average speed of 0.19m/s. Figure 5 shows the map generated based on the raw odometry data provided by the robot. As can be seen from the figure, the robot suffers from serious errors in odometry so that the resulting map is useless without any correction. Figure 6 (left) shows the map created with our scan matching technique. Although local structures of the map appear to be very accurate, the map is globally inconsistent. For example many structures like walls, doors etc. can be found twice and with a small offset between them. Finally, the right image of Figure 6 shows the resulting map obtained with our system. Although the sharpness of this map is not as high as that of the map created only with scan matching, they are globally consistent. The map was created in real-time, i.e. the computation time needed to process the data did not exceeded the time to record them. We used 100 samples, a number we found to yield satisfactory results on all data sets. Figure 9 shows a map created using




Fig. 5. Mapping of the Intel Research Lab with the raw odometry data.




Fig. 6. Map of the Intel Research Lab after scan matching (left) and obtained in real-time with 100 samples (right).

500 particles. Whereas this map is more accurate and has a similar crispness as the scan matching map, the time to compute this map was several hours. Figure 7 visualizes the trajectories of all samples shortly before and after closing the major loop in this data set. As the left image illustrates, the robot is quite uncertain about its position relative to the starting position upon its return. However, after a few resampling steps the uncertainty has been reduced drastically (right image).

A second example map obtained with our approach is depicted in Figure 8. The map shows the fourth floor of the 50m \times 12m large Sieg Hall of the University of Washington. As can be seen from the figure, the robot went several times around the circle and still successfully learned a consistent map. This map was generated in real-time using 100 samples. The grid resolution was 10cm.




Fig. 7. Trajectories of all 100 samples shortly before (left) and after (right) closing the loop.




Fig. 8. Map of the Sieg Hall at the University of Washington created in real-time.

B. Comparison to Previous Online Techniques

The second experiment is designed to show the advantage of our integrated technique over previous approaches that represent a posterior over poses in a single map only [19], [9]. For this experiment we used a data set generated for the Wean Hall of the Carnegie Mellon University using our B21r simulator. The size of this environment is $32\text{m} \times 10\text{m}$. In the simulation the robot moved 251m with an average speed of 0.78m/s. To obtain realistic data, we added a serious amount of noise to the ground truth data provided by the simulation system. The resulting input trajectory is depicted in Figure 10 (left). Please note, that pure scan matching again failed to correctly close the loop using this data set. We implemented the particle filter strategy presented by Thrun *et al.* [20], [19]. In our system we achieve this by using only that sample with the highest likelihood during the resampling process whenever the robot closes a loop. After this, we continue with the normal resampling procedure described above. The middle image in Figure 10 shows the result of this procedure. The point in time when the system discovered that it closed the loop and the resulting inconsistencies are also labeled. The inconsistencies are a consequence of the fact that only the particle with the highest importance factor survives at the time the robot closes the loop. Since this particle does not always correspond to the correct position of the robot (as it is the case in this example) and since the motion model cannot compensate for this error, the resulting map contains errors. In contrast to that, our approach, that integrates scan matching with a Rao-Blackwellized particle filter, yields a consistent map of the environment (see right image of Figure 10) since it provides accurate predictions and simultaneously maintains the robot pose uncertainty in the posterior.

Finally, we analyzed whether the standard Rao-Blackwellized particle filter without odometry correction by scan matching provides the same performance as our approach. For this purpose we run the standard procedure using the input data for the Intel Research Lab. We used shorter resampling steps (three times more often than in the other runs to avoid a fast divergence) and 200 samples which was the maximum number of samples that allowed updates in real-time on our 1.8GHz Pentium IV PC. Since the standard procedure was not able to learn a consistent map, we repeated this experiment with increasing numbers of samples. It turned out that under 1000 samples, which was the maximum number our PC equipped with 768MB of main memory could handle, we




Fig. 9. Map of the Intel Research Lab obtained offline with 500 samples.



Fig. 11. Map created with the standard Rao-Blackwellized particle filtering technique in real-time using 100 samples and based on the raw odometry data.

could not observe a case in which the standard algorithm converged. An example map typically obtained using the standard algorithm is depicted in Figure 11. The fact that our algorithm reliably converges with 100 samples indicates that the integration of the scan matching routine yields an enormous improvement.

V. Conclusions

In this paper we presented a highly efficient algorithm for simultaneous mapping and localization using laser scans that combines a scan matching procedure with Rao-Blackwellized particle filtering. The scan matching routine is used to transform sequences of laser measurements into odometry measurements. The corrected odometry and the remaining laser scans are then used for map estimation in the particle filter. The lower variance in the corrected odometry reduces the number of necessary resampling steps and this way decreases the particle depletion problem. In practical experiments we demonstrated that our approach allows to learn maps of large-scale environments in real-time with as few as 100 samples. Simultaneously, it outperforms previous approaches with




Fig. 10. (left) Map obtained from the raw input data in the simulation experiment. In the middle, the resulting map is shown if the mapping process is continued with the maximum likelihood sample after closing the loop at the marked place. The inconsistencies in the right part of the map show that the loop is not correctly closed. The map on the right was built using our approach.

respect to robustness and efficiency.

Acknowledgments

This work has partly been supported by the EC under contract number IST-2000-29456 and by the German Science Foundation (DFG) under contract number SFB/TR8-03. It has also been sponsored by DARPA's MARS, CoABS, MICA, and SDR Programme (contract numbers N66001-01-C-6018, NBCH1020014, F30602-98-2-0137, F30602-01-C-0219, and NBCHC020073) and by the NSF under grant numbers IIS-9876136, IIS-9877033, and IIS-0093406.

VI. REFERENCES

- [1] O. Bengtsson and A. Baerfeldt. Localization in changing environment - estimation of a covariance matrix for the idc algorithm. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1931–1937, 2001.
- [2] P. Besl and N. McKay. A method for registration of 3d shapes. *Trans. Patt. Anal. Mach. Intell.* 14(2), pages 239–256, 1992.
- [3] J. Borenstein and Feng. L. Measurement and correction of systematic odometry errors in mobile robots. *IEEE Journal of Robotics and Automation*, 12(6):869–880, 1996.
- [4] M. Bosse, P. Newman, M. Soika, W. Feiten, J. Leonard, and S. Teller. An atlas framework for scalable mapping. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2003.
- [5] J.A. Castellanos, J.M.M. Montiel, J. Neira, and J.D. Tardós. The SPMAP: A probabilistic framework for simultaneous localization and map building. *IEEE Transactions on Robotics and Automation*, 15(5):948–953, 1999.
- [6] G. Dissanayake, H. Durrant-Whyte, and T. Bailey. A computationally efficient solution to the simultaneous localization and map building (SLAM) problem. In *ICRA'2000 Workshop on Mobile Robot Navigation and Mapping*, 2000.
- [7] N. Doh, H. Choset, and W. K. Chung. Accurate relative localization using odometry. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2003.
- [8] A. Doucet, J.F.G. de Freitas, K. Murphy, and S. Russell. Rao-Blackwellised particle filtering for dynamic Bayesian networks. In *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2000.
- [9] J.-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proc. of the IEEE Int. Symp. on Computational Intelligence in Robotics and Automation (CIRA)*, 1999.
- [10] D. Hähnel, D. Schulz, and W. Burgard. Mapping with mobile robots in populated environments. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.
- [11] K. Konolige. Markov localization using correlation. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1999.
- [12] J.J. Leonard and H.J.S. Feder. A computationally efficient method for large-scale concurrent mapping and localization. In *Proc. of the Ninth Int. Symp. on Robotics Research (ISRR)*, 1999.
- [13] F. Lu and E. Miliotis. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
- [14] M. Montemerlo and S. Thrun. Simultaneous localization and mapping with unknown data association using FastSLAM. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2003.
- [15] M. Montemerlo, S. Thun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to simultaneous mapping and localization. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 2002.
- [16] H.P. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, pages 61–74, Summer 1988.
- [17] K. Murphy. Bayesian map learning in dynamic environments. In *Neural Info. Proc. Systems (NIPS)*, 1999.
- [18] T. Röfer. Using histogram correlation to create consistent laser scan maps. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.
- [19] S. Thrun. A probabilistic online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 20(5):335–363, 2001.
- [20] S. Thrun, W. Burgard, and D. Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2000.
- [21] R. van der Merwe, A. Doucet, N. de Freitas, and E. Wan. The unscented particle filter. Technical Report CUED/F-INFENG/TR 380, Cambridge University, Department of Engineering, 2000.
- [22] G. Weiß, C. Wetzler, and E. von Puttkamer. Keeping track of position and orientation of moving indoor systems by correlation of range-finder scans. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 595–601, 1994.