Robust weighted scan matching with quadtrees

Visser, A.; Slamet, B.A.; Pfingsthorn, M.

Publication date
2009

Document Version
Final published version

Published in
Proceedings of the Fourth International Workshop on Synthetic Simulation and Robotics to Mitigate Earthquake Disaster (SRMED 2009), 6th July 2009, Graz, Austria

Robust Weighted Scan Matching with Quadtrees

Arnoud Visser[1], Bayu A. Slamet[1], and Max Pfingsthorn[2]

[1] Universiteit van Amsterdam, Science Park 107, Amsterdam, The Netherlands
[2] Jacobs University Bremen, Campus Ring 1, Bremen, Germany

Abstract. This paper presents the improvement of the robustness and accuracy of the weighted scan matching algorithm matching against the union of earlier acquired scans. The approach allows to reduce the *correspondence error*, which is explicitly modeled in the weighted scan matching algorithm, by providing a more complete and denser frame of reference to match new scans. By making use of the efficient quadtree data structure, earlier acquired scans can be stored with millimeter accuracy for environments with dimensions larger than 100x100 meter. This can be realized with the preservation of real-time performance. In our experiments we illustrate the significant gains in robustness and accuracy that can be the result with this approach.

1 Introduction

Robot motion estimation is a central topic in probabilistic robotics research. Reliable robot motion estimation is essential prerequisite to be able to build maps of unknown environment, such as encountered in search and rescue missions. The general idea is that 'raw' motion estimates provided by odometry sensors or inertial navigation systems (INS) accumulate significant errors and therefore preclude an accurate spatial interpretation of sensor information, e.g. for the purpose of mapping, planning or navigation. By establishing correspondences between current and past observations these motion estimates can be corrected. As laser range scanners usually provide rather accurate measurements of a robot's physical surroundings they are a key enabler for more accurate displacement estimation based on such a correspondence analysis.

Algorithms that estimate displacement by comparing the current laser scan with one or more previously acquired laser scans are typically referred to as *scan matchers*. While scan matching algorithms usually indeed yield more accurate motion estimates, they do not resolve the fact that all robot observations remain correlated through the robot's trajectory. This implies that every estimate builds on the previous one. Hence, the estimates may still diverge incrementally from the robot's true motions in the long run.

This fundamental issue lies at the core of *simultaneous localization and mapping* (SLAM) research. As such, many advanced techniques that aim to detect and correct error accumulation have been put forward by SLAM researchers.

Particle filters [1], explicit loop closure strategies [2], Kalman filters [3] and information filters [4], elaborate exploration strategies that explicitly take the localization and mapping accuracy into account [5] and highly advanced data structures [6] to facilitate these algorithms are just some examples, see also [7] for a survey on several of these techniques and [8] for a more in-depth discussion.

Although these SLAM techniques have proved to be very effective in achieving their objective, they usually do a global minimalization of the localization error for all registered viewpoints. In this paper a pure incremental approach is described, with only a local minimalization of the localization error of a new viewpoint. We aim to boost the robustness of the scan matcher itself, so that there are significantly more accurate position estimates to start with. In a later stage a global minimalization of the localization error can be applied with SLAM techniques, for instance as a post-processing step.

Scan matching is a research field with a long history [9] and many of the techniques applied in this article are applied before. Kd-trees, quadtrees and octrees are well known techniques to index scan points [10]. Matching against local submaps is a known to improve the robustness of the match. Matching against all past scans is less common, typically the past scans are accumulated in evidence grids [11]. Weighted Scan Matching is a technique which by its advanced way to characterize uncertainty is able to achieve high accuracy [12]. This article is innovative in the combination of those known techniques, which further improves the robustness of the Weighted Scan Matching algorithm.

The way to characterize uncertainty proposed by Pfister [13] is based on an analysis of the physical phenomena that challenge scan matchers. The analysis reveals three sources for uncertainty. The measurement noise and sensor bias are the well known sources of uncertainty. In [14] it is demonstrated that the accuracy of the scan matching improves considerably when the remaining *correspondence error* is explicitly taken into account as third source of uncertainty. Our approach focuses on further reducing the correspondence error by providing the algorithm more points by involving all relevant scans observed in the past in the matching process. The reduced correspondence error also allows tightening the thresholds of the algorithm, which gives an additional, although slight, increase of the accuracy. Yet, the major improvement is not the accuracy, but the increased robustness in situation where not many scan points are available. The additional points for the algorithm are provided by a neighborhood search on all previous scans. Quadtrees [15] are used to make this feasible in real time.

2 Related Work

The idea of matching against accumulated scans (instead of single scans) was already introduced during the development of the original Iterative Closest Point algorithm (e.g. [9]). In the Computer Vision community many variants of the basic ICP concept have been proposed, but not all variants are directly applicable to the problem of a laser scanner on a mobile robot. In the mobile robot case, there are typically many viewpoints with a few measurement points, while in the case of range images (generated with structured light) there are many measurement points for a few viewpoints.

Although commercial laser scanners as the SICK LMS 200 can give accurate distance measurements over a range of 80 meters, the angular resolution is limited, resulting in sparse sampling of surfaces further away. To get a detailed estimate the profile of the real surface at such a large distance, measurements of several scans have to be combined, and non-sampled parts have to be interpolated [16].

In [14] we used the *local submaps* introduced in [6] that also aim at providing a richer frame of reference for matching new scans. In its original form the reference scan (scan A in Fig. 1) was extended by *all* measurements from a *few* scans in the neighborhood. The extent of the local submap is limited to a few scans because the computational effort of the scan matching algorithm grows quadratically with the number of points considered. The technique presented in this paper could be seen as an efficient variant of the local submap that considers as matching candidates only a *few* measurements in the neighborhood derived from a data structure that stores *all* previous scans.

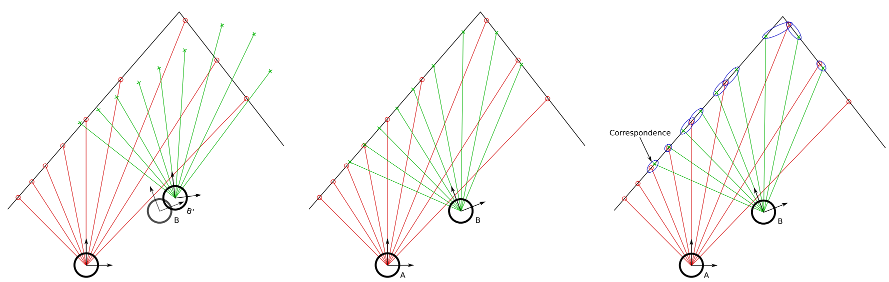

(a) Odometry Estimate (b) Scan Matcher Estimate (c) Correspondence Error

Fig. 1. Two scans A (red) and B (green). a) The second scan B' at an (erroneously) estimated location, B shows the true location of the second scan. b) The scans describe the same geometry after matching. c) The residual correspondence error after matching.

In [17] the authors present a technique for specific use with really sparse sensor readings that is based on *multiscans*. A single multiscan then accumulates sparse sensor readings until the integrated scan is dense enough. This scan is then matched against features in the environment and a particle filter is used to maintain multiple hypotheses. In our approach the quadtree also integrates multiple scans into a denser frame of reference, but the matching process is point based (not dependent on the presence of particular features in the environment). Also, we present results that did not (yet) take advantage of particle filtering.

Quadtrees seem a natural extension to occupancy grid maps as they efficiently index the occupancy information and thereby enable their fast retrieval. Notice that with a quadtree one is not forced to set a lower limit on the dimension of the cell (which is typically set to 15 centimeters for occupancy grid maps). This allows to store efficiently measurements with millimeter precision. However, only few robotics applications (e.g. [18]) benefit from quadtrees in two dimensions. The performance boost possible with scan matching is already recognized for higher dimensions. Octrees and advanced kd-tree search algorithms are for instance used in [19] for 3D point clouds. Octrees enable [20] to maintain a particle filter with 5000 particles in real-time for an underwater vehicle.

3 Weighted Scan Matching using Quadtrees

As a robot explores a planar environment there are several aspects that affect the range sensing process and thereby complicate the scan matching [13]. Matching two scans A and B is equivalent to estimating the relative translation τ and rotation ρ between the poses θ_A and θ_B where the scans were recorded. This can be done by projecting all individial scan beams p_B on the local coordinate frame of pose θ_A, but more sophisticated algorithms maintain subsets $c_A \subset p_A$ and $c_B \subset p_B$ of points that are deemed correlated. Matching the correlated subsets c_A and c_B is finding the estimate displacement $\Delta\hat{\theta}_{AB} = (\hat{\tau}, \hat{\rho})$ that minimizes the error ϵ defined as:

$$\epsilon = c_A - \hat{R} \cdot c_B - \hat{T} \tag{1}$$

where $\hat{R}$ is the rotation matrix of rotation $\hat{\rho}$ and $\hat{T}$ the displacement vector of translation $\hat{\tau}$. The error ϵ can be decomposed in several components. The *bias error* refers to the measurement bias that may be inherent to the used scanning device and the *measurement error* is the uncertainty attributed to the noise involved in the measurement process. For every point correspondence pair, the measurements c_A and c_B in Eq. (1) can be decomposed into the terms:

$$c_A^i = r_A^i + b_A^i + \delta c_A^i \tag{2}$$

where r_A^i is the 'true' measurement and b_A^i and δc_A^i respectively denote the bias error and noise error. Note that in this paper the bias error will further be ignored as for many datasets the bias of the used sensor is unknown, and as demonstrated in [13] it has minor impact when compared to the other sources for uncertainty.

The *correspondence error* is the consequence from the fact that a range sensor measures distances at discrete angles, which means that the samples taken from two different viewpoints will not coincide. The difference in sample distance on the surface will be small at close distance, but is more prominent for surfaces at larger distances. Even when both the measurement error δc_A^i and bias error b_A^i are ignored, and the distance to the real surface r_A^i is acquired, this measurement can still be halfway two other scan points r_B^j and r_B^{j+1}, making an maximum error of $||r_B^i - r_B^{j+1}||/2$, which is a function of the distance and incident angle towards that surface. In the weighted scan matching algorithm this quantity is actively estimated. This phenomenon is illustrated in Fig. 1(c) where the ellipses indicate the associated correspondence error.

Ignoring the bias error b_A^i then gives the following error ξ^i per pair of corresponding points:

$$\xi^i = \underbrace{\left(r_A^i - R \cdot r_B^i - \hat{\tau}\right)}_{Correspondence\ Error} + \underbrace{\left(\delta c_A^i - R \cdot \delta c_B^i\right)}_{Measurement\ Error} \tag{3}$$

The error ξ^i per pair can be used in Eq. (1) to find the displacement that minimizes the overall error function.

The sub-problem of the measurement error was handled in [13] by modeling the uncertainty of every single scan beam with a Gaussian distribution. Also the modeling of the correspondence error was addressed in [13], but this error was estimated for matching two scans from two poses A and B. Here we demonstrate that the correspondence error can be reduced by matching the current scan A against the integration of nearby measurements from many previous scans $B_{1...N}$. The correspondence error is not longer a function of the distance, but of the density of nearby measurements at that point.

The central idea is to use a quadtree [15] to keep an index of all relevant scans observed in the past. Assume that a particular scan s_t acquired at time t was found as being relevant. Then all global projections of the points $p \in s_t$ will be inserted into the quadtree Q that is maintained. The relevance of a scan is determined by simple thresholds on the estimated displacement returned by the scan matcher. Given the displacement estimate $\Delta\hat{\theta}$ for a new scan then the scan is considered relevant if either the translation threshold τ_{max} or the rotation threshold ρ_{max} is exceeded. Hence, these thresholds are no probabilistic measures and only serve as a simple constraint to reduce redundancy in the quadtree.

Key to our approach is that the weighted scan matcher matches the new scan s_t against the quadtree, which means that s_t is matched simultaneously against *all past relevant scans*. This is feasible in real-time due to the excellent performance of quadtrees on nearest neighbor finding. The benefit of matching a scan simultaneously against multiple previous scans is illustrated in Fig. 2 where a robot zigzags through a corridor. In the example the robot is configured with a laser range scanner that has a field of view of 180 degrees, which is indicated using different shades of gray. Assume that the quadtree stores the scans acquired at time steps t_1 and t_2, then two alternatives for t_3 illustrate the difference. The robot could proceed with traditional incremental scan matching and only match the new scan against the one from t_2, or it can benefit from the quadtree and match simultaneously against *both* previous scans. The latter has two immediate benefits: the quadtree provides a *denser* and *more complete* frame of reference for matching the new scan. The density gain involves parts of the wall marked in blue that are covered by both scans of t_1 and t_2. The gain in completeness involves the part of the wall marked in red that is not covered by the scan of t_2 but *is* covered by the one of t_1.

When considering the density and completeness gains illustrated in Fig. 2 one may indeed expect the correspondence error of Eq. (3) to be reduced. This is demonstrated in Fig. 4 of Sec. 4.

4 Experiments and Results

In our experiments we will use two implementations of the Weighted Scan Matcher (WSM) [13]. The point-correlation procedure of the original implementation was replaced with a nearest neighbor search algorithm of a quadtree. We did not make any additional modifications to the internal workings of the scan matchers, so we refer the interested reader to prior research [14, 21] and the original papers for further details.

The experiments will investigate the improvements that can be gained from using quadtrees for weighted scan matching. The visualizations were created with the standard occupancy rendering techniques from [14]. All presented results are strictly based on scan matching, the SLAM algorithm was purely incremental. As

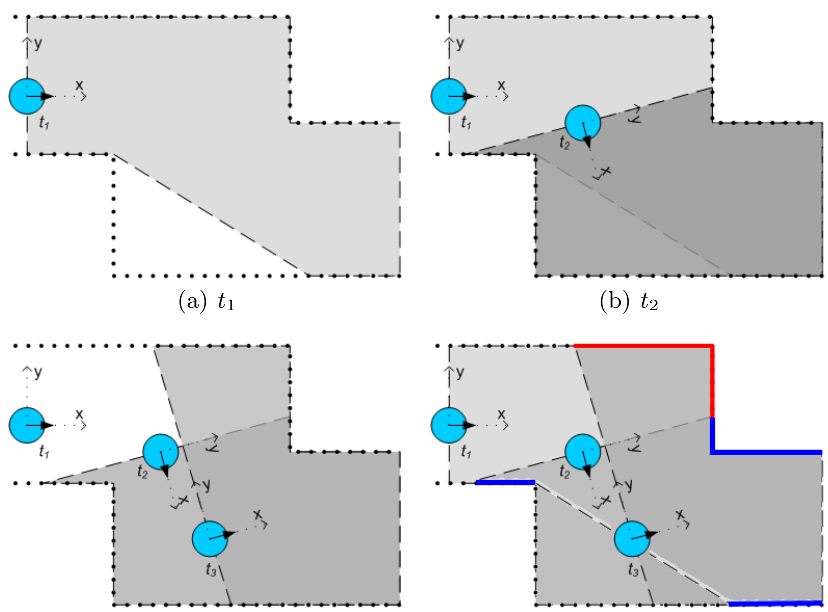

(a) t_1 (b) t_2

(c) t_3 incremental (d) t_3 using quadtree

Fig. 2. The benefit of using a quadtree that indexes all scans observed in the past. After the scans acquired at time-steps t_1 and t_2 the quadtree has a denser and more complete reference frame to offer for matching the scan of t_3. The walls shaded with blue were covered by both scans of t_1 and t_2 and illustrate the density gain. The part of the wall shaded with red marks reference data not present in the scan of t_2 but which *is* covered by the quadtree that also holds the scan of t_1.

long as the scan matching routine reported good correspondence, the result was used for localization. When the correspondence dropped below a threshold, the map was extended. When the correspondence was bad, the measurements from the dead reckoning were used. The initially found locations are not corrected afterwards in a global optimalization routine.

In the remainder of this section we will refer to the original implementation of WSM as the *incremental* version and we use Q-WSM to refer to the version that benefits from quadtrees.

4.1 Performance Benchmarks

For these experiments we used the high-fidelity simulator [22] that is also used during the Virtual Robots competition that is part of the RoboCup Rescue World Championships [3].

Figure 3 shows the results acquired on a dataset that was recorded in 'The Park'[4]. The challenge posed by this dataset is that often the range scanner only observes one or two trees and that the robot moves on slippery grass. The dataset covers an area of approximately 80 by 40 meters where the robot started in the bottom-right corner and traverses the park in clockwise direction. The robot's path is shaded with gray for clarity and should describe a single closed loop from tip to tail.

The incremental version of WSM accumulates a significant error, leaving a gap of several meters. The Q-WSM version closes the loop implicitly. The error of the incremental version can mainly be found in the upper right corner, where a sharp point turn is made[5] at a location without clear features. For the majority of the scans the results of both versions are equally good, the additional robustness of the Q-WSM version is mainly visible at the difficult situations.

For a more detailed comparison of both algorithms, the average correlation distances that remain after matching are plotted in Fig. 4. Q-WSM finds matches with less residual correspondence error almost throughout the dataset. Over the whole dataset the average correlation distance reduces from 10.20 mm to 5.62 mm.

To see if the reduced correlation distance also improves the robustness of the scan matcher the uncertainty measures returned by the algorithms were analyzed. This uncertainty measure is the full 3-by-3 covariance matrix of the Gaussian distribution over the displacement estimate and does not lend itself well for charting in its original form. Therefore we took the on-diagonal elements that describe the independent uncertainty in x and y direction and combined these into a Euclidean distance measure. The trace of the covariance matrix acquired with Q-WSM is in most cases (94%) smaller than the trace of incremental version. The average uncertainty of the Q-WSM algorithm reduced to 36% of the average uncertainty of the incremental version.

[3] USARSim source code and manual are available online at http://www.sourceforge.net/projects/usarsim

[4] Outdoor area of the DM compWorldDay1 map

[5] Data collected by novice driver according to [23]

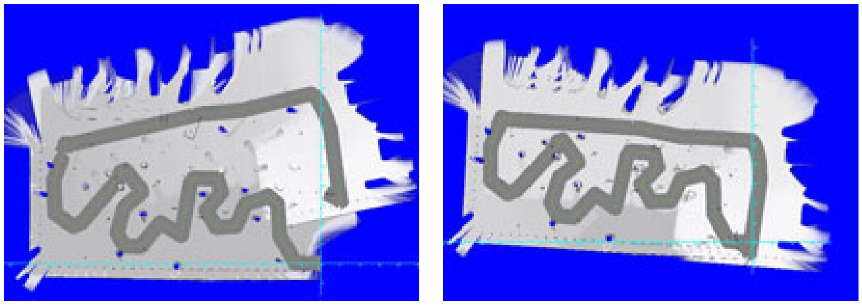

(a) Incremental WSM (b) Q-WSM

Fig. 3. Comparison of the two versions of WSM for a drive through a park, with poor odometry and sparse range scans.

The used laser range scanner was configured at 1 Hz with a resolution of 181 beams over a field of view of 180 degrees. The time necessary to match a scan ranged from 40 to 411 milliseconds for Q-WSM. The average time was 128 milliseconds. These results were acquired on a three-year old notebook (1.9GHz CPU and 1GB RAM) and clearly illustrate that the quadtree approach does not preclude real-time performance for this dataset.

4.2 Comparison with SLAM approaches

To validate the results acquired in the simulator on data that suffers from real-world odometric errors and sensor noise we use some dataset published on Radish [6] from [24]. The results in Fig. 5 serve to illustrate the general applicability of the presented approach. For this purpose we employed Q-WSM on several publicly available robotics datasets, some of which are also often referred in SLAM research. During the workshop also the comparison to other datasets can be made, such as the CMU Newell Simon Hall, Intel Lab Oregon or AP Hill dataset.

The 'Edmonton Convention Center' concerns quite a large area with lots of open spaces and 'Stanford Gates' is a huge dataset where the robot traverses several hundreds of meters. For comparison, the maps generated with CARMEN[7] are given. The "MIT CSAIL building" was used as example by [23] to illustrate how difficult SLAM could be. The map was generated by an Extended Kalman Filter, including two manual corrections. Especially note the amount of detail that is preserved by the Q-WSM algorithm for these datasets, like table legs and other small obstacles. Those small obstacles are not visible on the original maps, due to the bigger gridsize. Still, the claim is not made that SLAM is outperformed. For instance in the 'Stanford Gates' map it is clear that loop closing

[6] Available on http://radish.sourceforge.net
[7] http://carmen.sourceforge.net/

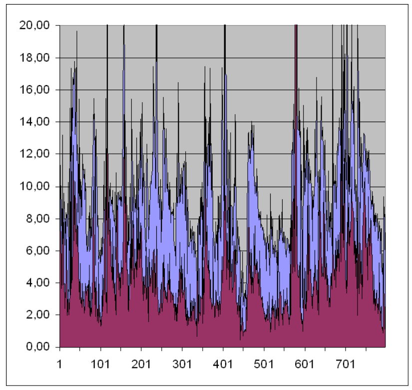

Fig. 4. Detailed comparison of the performance of the quadtree-based and incremental WSM algorithm for a drive through a park. The figure gives the residual correlation distance, Q-WSM in front of incremental WSM. The x-axis represents the time line defined by the 800 scans that comprised this dataset and the y-axis plots for each scan the average correlation distance (mm) that remains after applying the estimated displacement.

could help to get rid of the offset of 20 centimeters visible in the corridor at the right. The comparison with these SLAM approaches is made, to demonstrate how far one can come without SLAM and how easy it would be for a SLAM algorithm to make the final corrections with this initial match. Last but not least, this map is generated on the fly and directly available to an operator of for instance a rescue mission. It is clear that the maps are of enough quality to be used for navigation.

5 Discussion

It is clear that while driving around the number of points stored in the quadtree grows. The depth of the tree is limited by setting a minimal cell size, but this minimal cell size is rather small (1 millimeter), well below the accuracy of the laser scanner. When the number of points in the quadtree grows, searching for nearest neighbors stays efficient, a well known characteristic of quadtrees. Yet, another well known characteristic of quadtrees is that the construction time increases when the tree grows in size. Fortunately searching is performed more often than registration of scan points on the map. No significant performance drop was experienced on the datasets of Radish, with map sizes of several hundred meters.

6 Conclusion

This paper presents an improvement of the weighted scan matcher which targets the correspondence error by providing a denser and more complete frame of

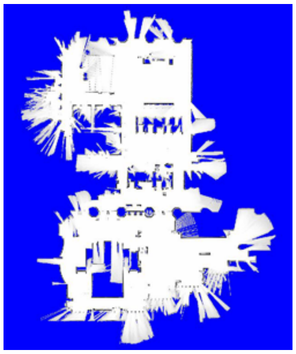

(a) Edmonton Convention Center - courtesy from Roy

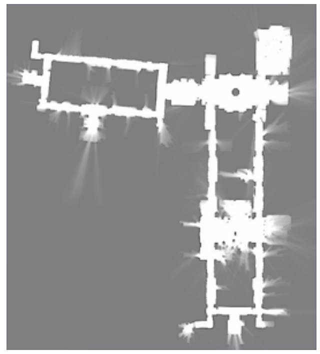

(b) Stanford Gates - courtesy from Gerkey

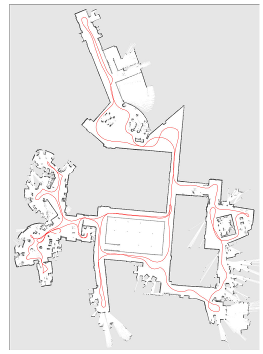

(c) MIT CSAIL building - courtesy from Roy [23]

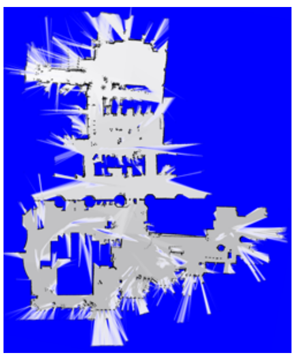

(d) Edmonton Convention Center - Q-WSM

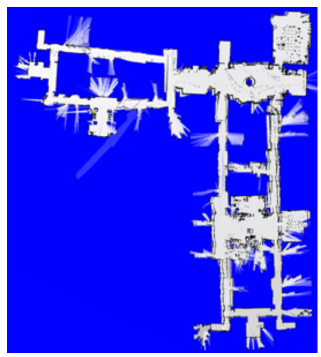

(e) Stanford Gates - Q-WSM

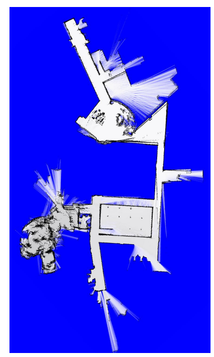

(f) MIT CSAIL building - Q-WSM

Fig. 5. Three real measured datasets from Radish. Here we present results that are acquired just by running the scan matcher Q-WSM as an iterative process, without any global optimization at the end of the process.

reference for matching new scans. This is realized by storing all relevant scans observed in the past in a quadtree, which enables the weighted scan matcher to preserve real-time performance.

In our experiments we showed that our approach can reduce the residual correspondence error and lead to increased robustness and accuracy. In several tests on real-robot data the applicability of our approach was illustrated. This paper demonstrates a different approach to robotic mapping under real-time constraints.

Acknowledgements

The research reported here is performed in the context of the Interactive Collaborative Information Systems (ICIS) project, supported by the Dutch Ministry of Economic

Affairs, grant nr: BSIK03024. The authors thank Olaf Booij and Zoran Zivkovic for their useful comments.

References

1. Montemerlo, M., Thrun, S., Koller, D., Wegbreit, B.: Fastslam: A factored solution to the simultaneous localization and mapping problem. In: Proceedings of the AAAI National Conference on Artificial Intelligence, Edmonton, Canada, AAAI (2002)
2. Olson, E., Leonard, J., Teller, S.: Fast Iterative Optimization of Pose Graphs with Poor Initial Estimates. In: Proceedings of the IEEE International Conference on Robotics and Automation. (2006) 2262–2269
3. Smith, R.C., Cheeseman, P.: On the representation and estimation of spatial uncertainty. International Journal of Robotics Research **5** (1986) 56–68
4. Thrun, S., Liu, Y., Koller, D., Ng, A., Ghahramani, Z., Durrant-Whyte, H.: Simultaneous localization and mapping with sparse extended information filters. International Journal of Robotics Research **23(7-8)** (2004)
5. Stachniss, C., Hähnel, D., Burgard, W., Grisetti, G.: On actively closing loops in grid-based fastslam. Advanced Robotics **19** (2005) 1059–1080
6. Howard, A., Sukhatme, G.S., Matarić, M.J.: Multi-robot mapping using manifold representations. Proceedings of the IEEE - Special Issue on Multi-robot Systems (2006)
7. Thrun, S.: Robotic Mapping: A Survey. In Lakemeyer, G., Nebel, B., eds.: Exploring Artificial Intelligence in the New Millenium, Morgan Kaufmann (2002)
8. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics (Intelligent Robotics and Autonomous Agents). The MIT Press (2005)
9. Chen, Y., Medioni, G.: Object modeling by registration of multiple range images. In: Proceedings of IEEE International Conference on Robotics and Automation (ICRA). Volume 3. (1991) 2724–2729
10. Nüchter, A., Surmann, H., Hertzberg, J.: Automatic model refinement for 3d reconstruction with mobile robots. In: Proceedings of the 4th IEEE International Conference on Recent Advances in 3D Digital Imaging and Modeling (3DIM '03), IEEE Computer Society Press (2003) 394 – 401
11. Bosse, M., Zlot, R.: Map Matching and Data Association for Large-Scale Two-dimensional Laser Scan-based SLAM. The International Journal of Robotics Research **27** (2008) 667–691
12. Censi, A.: Scan matching in a probabilistic framework. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Orlando, Florida (2006) 2291–2296
13. Pfister, S.T., Kriechbaum, K.L., Roumeliotis, S.I., Burdick, J.W.: Weighted range sensor matching algorithms for mobile robot displacement estimation. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA). Volume 2. (2002) 1667– 1674
14. Pfingsthorn, M., Slamet, B.A., Visser, A.: A Scalable Hybrid Multi-Robot SLAM Method for Highly Detailed Maps. In: RoboCup 2007: Robot Soccer World Cup XI. Volume 5001 of Lecture Notes in Artificial Intelligence., Berlin Heidelberg, Springer-Verlag (2008) 457–464
15. Samet, H.: The quadtree and related hierarchical data structures. ACM Comput. Surv. **16** (1984) 187–260

16. Rutishauser, M., Stricker, M., Trobina, M.: Merging range images of arbitrarily shaped objects. In: Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR). (21-23 Jun 1994) 573–580
17. Beevers, K.R., Huang, W.H.: SLAM with sparse sensing. In: Proceedings of IEEE International Conference on Robotics and Automation (ICRA). (2006) 2285–2290
18. Yahja, A., Singh, S., Stentz, A.: An efficient on-line path planner for outdoor mobile robots operating in vast environments. Robotics and Autonomous Systems **33** (2000) 129–143
19. Nüchter, A., Lingemann, K., Hertzberg, J.: 6D SLAM with Cached kd-tree Search. In: Robot Navigation. Number 06421 in Dagstuhl Seminar Proceedings, Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany (2007)
20. Fairfield, N., Kantor, G., Wettergreen, D.: Real-time slam with octree evidence grids for exploration in underwater tunnels. Journal of Field Robotics **24** (2007) 3–21
21. Slamet, B.A., Pfingsthorn, M.: ManifoldSLAM: a Multi-Agent Simultaneous Localization and Mapping System for the RoboCup Rescue Virtual Robots Competition. Master's thesis, Universiteit van Amsterdam (2006)
22. Balakirsky, S., Scrapper, C., Carpin, S., Lewis, M.: Usarsim: providing a framework for multi-robot performance evaluation. In: Proceedings of PerMIS 2006. (2006)
23. Kollar, T., Roy, N.: Trajectory Optimization using Reinforcement Learning for Map Exploration. The International Journal of Robotics Research **27** (2008) 175–196
24. Hähnel, D., Fox, D., Burgard, W., Thrun, S.: A highly efficient fastslam algorithm for generating cyclic maps of large-scale environments from raw laser range measurements. In: Proceedings of the Conference on Intelligent Robots and Systems (IROS). (2003)