

Scan Matching for Flooded Subterranean Voids

David Silver, David Bradley and Scott Thayer
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA

Abstract— This paper presents a new technique for scan matching with sparse, noisy sensors such as sonar range finders. This method allows an underwater vehicle to estimate its position within a 2-D plane. When combined with other sensors typically used on such vehicles, full 6 degree of freedom pose estimation is possible in flooded subterranean environments. Our approach uses a particle filter and an approximation of the likelihood of sensor readings, based on nearest neighbor distances, to approximate the probability distribution over possible poses. We present simulated results based on field collected data, demonstrating our approaches advantages over traditional scan matching techniques.

1. INTRODUCTION

The navigation and mapping of subterranean voids presents both an opportunity and a significant challenge for the field of robotics. One of the most common and important instances of such voids are abandoned mines. In the U.S. alone, tens of thousands of abandoned mines threaten nearby surface structures, ecosystems, and ongoing mining operations [1]. Proactive measures can be taken to prevent these threats from manifesting into disastrous situations; however, prevention relies upon knowledge of mine layout and condition. Mine maps are the primary source of such information. In the case of abandoned mines, these maps, when they exist, are often inaccurate, if legible.

Robots are sensible tools for mapping mines and similar voids since they can potentially navigate areas too small or too dangerous for people to enter. In the case where the abandoned mine is flooded, access to the mine is often only possible via boreholes. Current plans call for building robots that can deploy via boreholes as narrow as 8 inches in diameter. Constructing a robot that can be deployed through such a small borehole, and still perform useful mapping functions, is a significant challenge. The borehole size significantly constrains, amongst other things, the number and quality of sensors such a robot can use. Since the entire purpose of deploying the robot is to obtain an accurate map, new techniques are needed for accurate positioning and mapping in the presence of such limited sensing.

Previous mine mapping robots [2] used laser range finders both for pose estimation and mapping. These sensors provided highly accurate and dense measurements. The analogous sensor for underwater environments is the rotary sonar, which involves a single sonar transducer spinning around a single axis. Current technological limitations further constrain such a robot to use a small imaging rotary sonar for horizontal scanning. Such a device has approximately one fifth the azimuth resolution of a laser scanner, and produces a full scan once every few seconds,




Fig. 1. An initial design for the Minefish, a borehole deployable AUV

as opposed to dozens of times per second. Also, the range data it produces contains much more noise.

This paper describes a new technique for scan matching that is suitable to the limited sensing capabilities present in flooded subterranean environments. This new method uses a particle filter to estimate the robots pose within a 2 dimensional plane. When this estimate is combined with other sensors available to underwater vehicles, the robot's 3-D position and orientation can be estimated.

In the next section, related work in underwater pose estimation is described. Section 3 describes related work in scan matching. Section 4 describes our new approach to scan matching. Experimental results are presented in Section 5, and we conclude with a discussion and directions for future work in Section 6.

2. UNDERWATER POSE ESTIMATION

Several approaches have been developed for performing pose estimation with an Autonomous Underwater Vehicle (AUV). While underwater environments provide many challenges for positioning, they also provide some advantages. The most important of these is the ability to measure depth directly using a pressure sensor. Such depth sensors, when combined with a standard two axis tilt sensor, provide direct measurement of the roll, pitch, and Z coordinate of an AUV's pose. The remaining challenge is determining the yaw, X, and Y values of the AUV's pose. Traditional inertial navigation systems, combined with thruster modeling, provide short term estimates of these values. The main focus of underwater pose estimation

is therefore often to provide long term stability of these values. For an overview of AUV navigation techniques, see [3].

One of the most common techniques in use with commercial AUVs is Long Base Line (LBL) positioning, which uses acoustic time of flight to beacons at known locations to accurately determine AUV position. LBL can also be used when the initial beacon locations are unknown [4]. However, for LBL to be applicable, the AUV must have acoustic line of site to the beacons currently in use. While this is usually the case for an open water AUV, it is quite the opposite in an enclosed flooded environment. Also, deployment of such beacons to subterranean voids requires borehole access at every beacon location.

Another technology in use for AUV positioning is the Doppler Velocity Log (DVL). A DVL measures the Doppler shift in multiple acoustic signals to determine an AUV's velocity in the XY plane. When combined with other traditional sensors, a DVL can be used to compute a dead-reckoning position with comparatively little drift over long periods of time [5]. Unfortunately, at this time there are no commercially available DVLs that meet our size constraints. Until they can be further miniaturized, DVLs are not a viable option. An alternative method for direct velocity estimation is by tracking visual features, using a system such as the one described in [6]. Again, environmental factors specific to subterranean voids add difficulty. The corridors of such voids are often lacking in the distinct visual features necessary for robust tracking. Also, such voids are often zero light environments, requiring a drain on an AUV's limited power supplies to illuminate an area.

Traditional landmark based Simultaneous Localization and Mapping (SLAM) [7] techniques are also applicable to AUVs. In particular, using an extended Kalman filter with landmarks [8] extracted from sonar range finders is a common approach. The extracted landmarks may be artificial [9] or natural [10][11]. As with artificial beacons, artificial landmarks are not practical in this context due to the difficulty of deployment. Natural landmarks also present problems. Again, one issue is with line of sight. There may not always be distinctive natural landmarks visible due to the limited line of sight. Also, even when natural landmarks are visible, robust detection with a noisy range sensor is a challenge

3. SCAN MATCHING

A common method of pose estimation for indoor mobile robots is scan matching. By solving for the rigid transformation between consecutive scans from a range sensor the robot's motion in the time period between the scans can be inferred. The sensor utilized is most often a scanning laser range finder.

One of the most popular approaches for scan matching is the Iterative Closest Point (ICP) [12] algorithm. In ICP, the transformation between scans is found iteratively by assuming that every point in the first scan corresponds to its closest point in the second scan, and solving closed form for the transformation using these correspondences. The algorithm is seeded by an initial odometry estimate, and iterated until convergence. This process can be viewed as an optimization




Fig. 2. ICP error is computed by summing the squared distance between each point in the red set and its closest point in the blue set

problem. In this context, the error between scans is the sum of squared distance between pairs of closest points, hereby referred to as the ICP error (Figure 2). At each step, the transformation that minimizes this error is computed, and the correspondences recalculated until a local minimum of error is reached. Variations on this approach based on different methods for computing point correspondences have been proposed [13].

The ICP algorithm will always converge to at least a local minimum. Given enough accurate data points and a sufficient seed transformation, this local minimum can be arbitrarily close to the actual robot motion. However, sparse and noisy data, such as that from a imaging sonar, can cause ICP to fail. A single noisy reading can significantly affect the computed transformation, causing the estimated robot pose to drift over time.

Other scan matching approaches, such as [14], are based on extracting feature points from each scan, and then matching features between scans. A mapping between features in different scans then leads directly to the rigid transformation of the robot between scans. However such feature extraction techniques are not always robust enough to be used with higher noise sensors such as sonar.

4. PARTICLE SCAN MATCHING

One of the issues with ICP scan matching using sonar sensors is the presence of outlier readings. Since ICP is minimizing the sum of squared distance, the noise is also squared, meaning a few outliers can easily corrupt the estimated pose. Subsequent pose estimates will then suffer because of this preexisting error. This problem can be alleviated to some degree by matching new scans to a history of previous scans, rather than just the last scan. However, this is not always enough to overcome the noise inherent in sonar sensors.

Maintaining a record of the uncertainty inherent in the pose estimation process is essential for minimizing the effects of noise. A common approach to approximating such uncertainty estimates is to maintain a set of samples from the probability

distribution being measured. Given enough samples, any distribution can be accurately approximated. This approach is often called particle filtering. Previous AUV localization systems have also made use of particle filters. In particular, [15] and [16] use a particle filter to estimate an AUV's position in the XY plane based on its depth, and a previously existing seabed elevation map. The approach used is based on the Monte Carlo Localization (MCL) algorithm [17] [18]. In this approach, a Bayesian filter is implemented by approximating each of the necessary distributions with samples drawn from them. The distributions in question are the pose prior:

$$P(X_{t-1}|X_{t-2}, U_{t-1}, \dots, U_0, Y_{t-1}, \dots, Y_0) \quad (4.1)$$

the motion model

$$P(X_t|X_{t-1}, U_t) \quad (4.2)$$

and the likelihood of the sensor readings

$$P(Y_t|X_t) \quad (4.3)$$

where X_t, U_t, Y_t are the pose, control inputs, and sensor readings at time t . The likelihood of the sensor readings is used to approximate the probability of a given pose, since from Bayes rule

$$P(X_t|Y_t) \propto P(Y_t|X_t) \quad (4.4)$$

The filter is implemented by drawing a sample from the motion model for each sample in the prior set of particles. The particles are then weighted by the likelihood of the sensor measurements. This likelihood is computed based on a preexisting map of the environment. Finally, the particles are resampled with replacement. The probability of selecting a particle is proportional to its weight.

Previous approaches that combined particle filters and scan matching, such as [19], have used scan matching as a replacement for odometry in the prediction phase of a particle filter, depending on landmark based techniques for the update phases. However, landmarks are hard to detect reliably with a noisy range sensor. Also, such an approach uses a prior fixed model of the uncertainty introduced in scan matching, as opposed to observing it explicitly at each cycle.

Our approach to combining particle filtering and scan matching considers ICP scan matching from a probabilistic viewpoint. Given two sensor scans of a static environment taken at different times, any difference in the scans can be attributed to either robot motion or sensor noise. Transforming both scans into the same coordinate frame, based on their respective pose estimates, makes any remaining differences between the scans attributable to either sensor noise or to positioning error. Therefore, the smaller these differences, the more likely both the pose estimate and the sensor reading (eq. 4.4). Since ICP error is a quantitative measurement of the differences between two scans, it can be used as an approximation of the likelihood of the sensor readings involved. By minimizing ICP error, ICP scan matching finds the pose that maximizes the likelihood of the sensor reading in question.

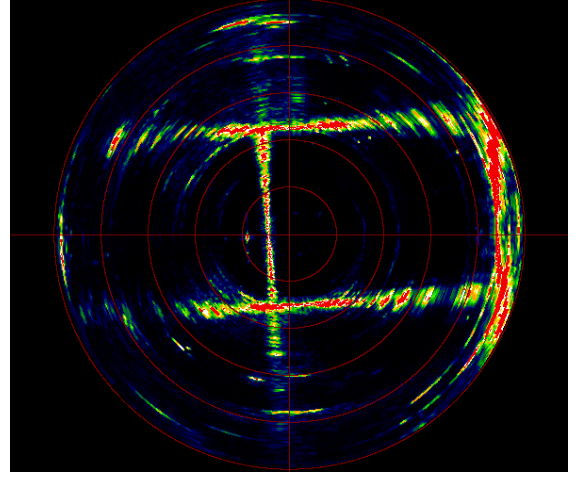


Fig. 3. Output from an Imagenex 881A rotary sonar, tested in the Carnegie Mellon swimming pool

Armed with this method for approximating likelihood, we use the following approach to pose estimation: A set of particles S_0 that represent the initial belief of the AUV pose at time 0, $P(X_0)$, are initialized. The particles are then updated by control inputs U_t and sensor measurements Y_t as follows:

- 1) For each particle i in S_{t-1} , sample from the AUV's motion model $P(X_t|S_{t-1}^i, U_t)$ to generate a new sample for \hat{S}_t , a set of temporary particles.
- 2) Weight each particle \hat{S}_t^i by the ICP error of the most recent sensor measurements Y_t , transformed into world coordinates by \hat{S}_t^i . The ICP error is calculated with respect to previous sensor measurements, transformed by their respective pose estimates.
- 3) Sample with replacement from \hat{S}_t , with probability *inversely* proportional to the weights, to form S_t .

This approach is quite similar to MCL. The difference is the lack of a preexisting map. Instead, the map is built online by transforming each scan with the latest pose estimate. In this sense our approach is a combination of traditional scan matching and MCL.

One issue with 2-D scan matching in a 3-D environment is the incorrect assumption that the AUV is confined to a 2-D plane. While other sensors may measure changes in roll, pitch, and depth, accounting for the effect of these changes on sensor measurements that are believed to lie within the XY plane is more difficult. In this aspect, imaging sonars offer an advantage. While such sonars have an azimuth resolution of a few degrees, the resolution in declination of such systems is usually on the order of 20 or 30 degrees. The result of this lowered resolution is to make the effect of small changes in roll, pitch, and depth appear just as a small increase in noise, which particle scan matching can account for. Larger changes in these values over longer periods of time are accommodated by using an expiring history of sensor readings.




Fig. 4. 4 consecutive laser scans showing a cave-in in the Mathies mine. The sonar scans simulated from these laser scans are shown in green.

5. EXPERIMENTAL RESULTS

A. Simulated Data

The current target vehicle for particle scan matching, the Minefish, is still under development. Consequently, experiments were run in simulation. The data used for the simulation was collected by Groundhog [2], a robot for mapping dry mines, in an autonomous run in the abandoned Mathies coal mine near Pittsburgh, PA.

Laser scans logged at approximately 1 Hz were used as the basis for simulated sonar data. Data collected from an Imagenex 881A rotary sonar (Figure 3) was used as a reference for determining what behavior the simulated sonar should exhibit. Each laser scan was downsampled, to provide an azimuth resolution of 2.5 degrees from the initial 0.5 degrees. To simulate the slower scan rate of an imaging sonar, only a 45 degree section of each 180 degree scan was used. Consecutive slices were used from consecutive scans. Finally, 25 cm² variance Gaussian noise was added to simulate the lesser accuracy of sonar compared to laser, and multipath noise was added to approximately one tenth of all readings. Figure 4 shows an example of this process.

Proprioceptive pose estimates, corresponding to the output from an IMU filtered with control inputs, were simulated from groundtruth poses. Groundtruth was used to generate relative changes in position and orientation between simulated sonar scans. Gaussian noise was added to these deltas to simulate noisy output. The noise added to position changes had a variance of 0.25 cm², and 0.0025 radians² to changes in orientation. These numbers result in much greater drift than is typical in such a system, and were chosen to test the limits of particle and ICP scan matching.

B. Simulated Results

Both ICP and particle scan matching were run on the the combination of simulated initial pose estimates and sensor

readings. Particle scan matching was run with between 10 and 5000 particles at a time. A history of the last 100 scans was used for matching against the current scan. Such a long history is computationally feasible due to the small size of each scan (Figure 5). The final pose estimates from both approaches were used in combination with the simulated sonar readings to produce occupancy grids[20]. These maps are shown in Figure 7, along with maps built using the groundtruth and simulated IMU poses. The groundtruth map was created using the original laser data rather than the sonar data simulated from it.

Figure 7 demonstrates the difficulty that ICP encounters when given sparse and noisy sensor readings. ICP corrects for some of the noise in the simulated IMU, but remains too dependent on the initial pose that it is given. The result is that the poses produced by ICP scan matching still appear similar to a random walk.

Figure 7 also demonstrates the ability of particle scan matching to cope with this same high level of noise. The pose estimates generated using 100 pose particles are a close approximation of ground truth, exhibiting marginal drift over the approximately 175 meters of traverse. In simulation, particle scan matching with 100 particles ran at approximately 3 Hz. Given an expected decrease in the computational burden in moving from simulation to an optimized implementation, this is more than fast enough to allow particle scan matching to run in real time without interfering with the other computation necessary on an AUV.

Increasing the number of pose particles does result in slight improvements. The pose estimates generated using 1000 pose particles result in a map that exhibits even less drift from groundtruth. The relationship between accuracy and the number of particles is visible in Figure 6, which plots the experimentally recorded variance from the true orientation deltas against the number of particles used. This variance exhibits rapid decay, indicating that the pose uncertainty is sufficiently approximated with a small number of particles.




Fig. 5. Particle scan matching in action. The blue points represent sonar readings transformed into world coordinates using previous pose estimates. The green points represent the current sonar readings, and the current set of pose particles is shown in red.

This is possible with so few particles because the change in pose between cycles is known to be small. In contrast, particle based systems for solving global localization or kidnapped robot style problems require far more particles.

6. CONCLUSION

We have presented a new approach to scan matching with sparse, noisy range sensors. Our approach combines concepts from Monte Carlo Localization and ICP scan matching to produce a 2-D pose estimate. This 2-D estimate can be combined with other sensors for full 3-D AUV pose estimation. Simulations based on field collected data demonstrate the advantages of particle scan matching in the presence of high levels of noise.

Future work will focus on two aspects of particle scan matching. The first will be to apply particle scan matching to the Minefish prototype upon its completion to demonstrate full 3-D pose estimation. There is also the possibility of additional improvement to the algorithm itself. In particular, alternative error functions based on nearest neighbor distances will be explored. One possibility is to use a distance metric based on a model of sensor noise. Another is to use an M-estimator [21] or other robust statistic, in order to lessen the negative effects of outlier readings. Techniques for adapting the number of particles to reduce computational overhead, such as KLD-sampling [22], will also be explored.

Another possible improvement is related to the map used online for approximating the likelihood of sensor readings. Currently, this map is built by transforming each sensor scan into world coordinates using the pose estimate produced by processing that scan. This approach does not allow the map to reflect the known pose uncertainty at the time that portion of the map was built. Also, once a sensor reading is added to the map, it remains unchanged until it expires. Future research will focus on ways to address both of these issues.




Fig. 6. The variance of the error in orientation changes, recorded for various numbers of pose particles

ACKNOWLEDGMENTS

Chuck Whittaker, Zachary Omohundro, Mike Sokolsky, and Chris Atwood, for their efforts in developing the Minefish.

George Kantor and David Wettergreen, for continued feedback and commentary.

John Hatch, for use of the Mathies mine.

The Mine Safety and Health Administration and Pennsylvania Department of Environmental Protection, for their support and expertise.

Senator Arlen Specter, for his continued support of this work.

REFERENCES

- [1] J. Belwood and R. Waugh, "Bats and mines: Abandoned does not always mean empty," *Bats*, vol. 9, no. 3, 1991.
- [2] C. Baker, A. Morris, D. Ferguson, S. Thayer, C. Whittaker, Z. Omohundro, C. Reverte, W. Whittaker, D. Hähnel, and S. Thrun, "A Campaign in Autonomous Mine Mapping," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, New Orleans, LA, 2004.
- [3] L. Whitcomb, "Underwater robotics: Out of the research laboratory and into the field," in *IEEE International Conference on Robotics and Automation*, 2000.
- [4] P. Newman and J. Leonard, "Pure range-only sub-sea slam," in *IEEE International Conference on Robotics and Automation*, September 2003.
- [5] L. Whitcomb, D. Yoerger, and H. Singh, "Advances in doppler-based navigation of underwater robotic vehicles," 1999.
- [6] M. Dunbabin, P. Corke, and G. Buskey, "Low-cost vision-based auv guidance system for reef navigation," in *Proceedings of the IEEE International Conference on Robotics and Automation*, New Orleans, LA, 2004.
- [7] J. Leonard and H. Durrant-Whyte, "Simultaneous Map Building and Localization for an Autonomous Mobile Robot," in *IEEE/RSJ International Workshop on Intelligent Robots and Systems*, May 1991, pp. 1442–1447.
- [8] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," *Autonomous Robot Vehicles*, pp. 167–193, 1990.
- [9] S. B. Williams, P. Newman, M. W. M. G. Dissanayake, and H. Durrant-Whyte, "Autonomous underwater simultaneous localisation and map building," in *Proceedings of IEEE International Conference on Robotics and Automation*, 2000.




Fig. 7. Occupancy grid maps given different pose estimates: Groundtruth(left), Simulated IMU(top), ICP Scan Matching(bottom) and Particle Scan Matching(right). The pose estimates used for creating each map are shown in blue. The groundtruth poses are superimposed in red over the particle based estimates, for comparison.

- [10] J. Leonard, R. Carpenter, and H. Feder, "Stochastic mapping using forward look sonar," in *International Conference on Field and Service Robotics*, 1999, pp. 69–74.
- [11] S. Williams and I. Mahon, "Simultaneous localisation and mapping on the great barrier reef," in *Proceedings of the IEEE International Conference on Robotics and Automation*, New Orleans, LA, 2004.
- [12] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 2, pp. 239–256, 1992.
- [13] F. Lu and E. Milios, "Robot pose estimation in unknown environments by matching 2d range scans," *Journal of Intelligent and Robotics Systems*, vol. 18, pp. 249–275, 1997.
- [14] M. Tomono, "A scan matching method using euclidean invariant signature for global localization and map building," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2004.
- [15] R. Karlsson and F. Gustafsson, "Particle filter for underwater terrain navigation," in *IEEE Workshop on Statistical Signal Processing*, September 2003, pp. 526–529.
- [16] A. Bachmann and S. Williams, "Terrain aided underwater navigation - a deeper insight into generic monte carlo localization," in *Proceedings of the Australasian Conference on Robotics and Automation*, 2003.
- [17] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte carlo localization for mobile robots," in *IEEE International Conference on Robotics and Automation*, May 1999.
- [18] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte carlo localization: Efficient position estimation for mobile robots," in *Proc. of the National Conference on Artificial Intelligence*, 1999.
- [19] D. Hahnel, 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 *Proceedings of IROS2003*, 2003, pp. 206–211.
- [20] H. Moravec and A. Elfes, "High resolution maps from wide angle sonar," in *IEEE Int. Conf. on Robotics and Automation*, 1985.
- [21] P. Huber, *Robust Statistics*. New York: John Wiley and Sons, 1981.
- [22] D. Fox, "Adapting the sample size in particle filters through kld-sampling," 2003.