

Adaptive Concurrent Mapping and Localization Using Sonar

Hans Jacob S. Feder, John J. Leonard, Chris M. Smith
Department of Ocean Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
{feder,jleonard,cmsmith}@deslab.mit.edu

Abstract

In order to create a truly autonomous vehicle, the task of concurrent mapping and localization (CML) in an a priori unknown environment is an important problem. Traditionally, the task of CML has been separated from the vehicle's motion and sensing strategies. We introduce a method for adaptive concurrent mapping and localization in unknown environments using a scanning sonar sensor. This method maximizes the information gained by the next action of the robot, given the space of available actions. The viability of the approach is shown in simulation and experiments. Results are shown for both adaptive control of motion and adaptive control of motion and sensing. Improved performance is demonstrated in comparison to straight-line motion and random motion.

1 Introduction

The goal of concurrent mapping and localization (CML) is for the mobile robot to build a map of its environment and to use that map to navigate in real time. This problem has been a popular topic in the research community, due to its theoretical challenges and critical importance for many mobile robot applications [11, 14]. CML is an especially important research objective in attempting to provide improved navigation capabilities for autonomous underwater vehicles (AUVs). The objective of this paper is to present a new technique for *adaptively* performing CML, with applicability to AUVs, and to demonstrate these techniques using real sonar data.

A seminal technique for concurrent mapping and localization, called the stochastic map, was published by Smith, Self, and Cheeseman [14]. The stochastic map consists of a single state vector that represents the estimate of the vehicle and feature locations and the error covariance matrix. As the vehicle moves around its en-

vironment taking measurements, the stochastic map is updated using an extended Kalman filter. Moutarlier and Chatila [11], Rencken [12], and Chong and Klee-man [4] have implemented variations of the stochastic map using land robots.

Many important issues remain for future research, such as the data association problem [13], computational complexity issues [7], quantifying the errors introduced by linearization in the extended Kalman filter and handling of dynamic environments [5]. This paper contributes to our knowledge of CML by presenting a technique for adaptively determining the sensing and motion actions of the robot, thereby obtaining more effective CML performance.

The problem of CML becomes particularly challenging when posed for an AUV. Land robots usually operate in passive environments with a given structure, such as office environments. Further, land robots often have relatively accurate localization systems readily available, as well as a variety of sensors. Underwater vehicles cannot use GPS and must operate in unstructured dynamic environments. Sensing options are more limited underwater [15], and positioning systems which use acoustic beacons are expensive and often impractical [18]. Sonar is the principle sensor for AUVs performing CML. As AUVs tend to have more dynamical constraints than land robots, such as loss of controllability at low speeds, the need for adaptively choosing a sensing and motion strategy to obtain the most information about the environment at any time is particularly useful. Such strategies are seen in the way dolphins use sonar to localize and navigate by moving their heads from side to side [1]. Further, due to the low rate of information obtained from a sonar, adaptive strategies are especially beneficial in robotic applications based on sonar sensing.

The structure of this paper is as follows: section 2 outlines the theory of the adaptive concurrent mapping and localization algorithm; section 3 illustrates the approach in simulation for two cases, motion con-

trol with complete scans and motion control with adaptive scanning; section 4 provides experimental results using a single rotating Polaroid sonar mounted on a mobile robot. Finally, section 5 provides concluding remarks and suggestions for future research.

2 Theory

We define the state of the system (robot and map) by the vector $\mathbf{x}[k]$ and the measurements at time step k as the vector $\mathbf{z}[k]$. The measurements are related to the state by a state-to-observation transform $\mathbf{z}[k] = \mathcal{H}(\mathbf{x}[k], \mathbf{d}_z)$, where \mathbf{d}_z is the measurement noise process. The *a posteriori* PDF of $\mathbf{x}[k]$, given a set of measurements $\mathbf{Z}^k \equiv \{\mathbf{z}[k], \dots, \mathbf{z}[1]\}$, can be found from Bayes Theorem as $p(\mathbf{x}[k]|\mathbf{Z}^k) = p(\mathbf{Z}^k|\mathbf{x}[k])p(\mathbf{x}[k])/p(\mathbf{Z}^k)$.

The ‘‘amount’’ of information contained in this distribution can be quantified by the Fisher information, that is [2]:

$$\mathbf{I}_{k|k} \equiv E\{(\nabla_{\mathbf{x}} \ln p(\mathbf{x}|\mathbf{Z}^k))(\nabla_{\mathbf{x}} \ln p(\mathbf{x}|\mathbf{Z}^k))^T\}. \quad (1)$$

In the case that we have an *efficient* estimator for $p(\mathbf{x}[k]|\mathbf{Z}^k)$ the Fisher information is simply given by the inverse of the error covariance of the state $\mathbf{x}[k]$.

In order to perform CML, the state transition (dynamic model) $\mathbf{x}[k+1] = \mathcal{F}(\mathbf{x}[k], \mathbf{u}[k], \mathbf{d}_x)$, in addition to the observation transformation \mathcal{H} , must be known. Here \mathbf{d}_x is the process noise. If the stochastic processes \mathbf{d}_z and \mathbf{d}_x are assumed to be independent, white and Gaussian, and the state transition \mathcal{F} and observation model \mathcal{H} are both linear, the Bayes Least Square (BLS) estimator, $\hat{\mathbf{x}}_{k+1|k} = \mathcal{F}(E\{\mathbf{x}[k]|\mathbf{Z}^k\}, \mathbf{u})$, will be an efficient estimator.

However, in the general problem of CML, neither the dynamic model nor the observation model will be linear. Thus, an efficient estimator cannot be obtained. Further, propagating the Fisher information through nonlinear equations does not guarantee that the statistics will be conserved. Thus, in order to circumvent the problem of transformation of nonlinearities it is assumed that

$$\hat{\mathbf{x}}_{k+1|k} \approx E\{\mathcal{F}(\mathbf{x}[k], \mathbf{u}[k], \mathbf{d}_x)|\mathbf{Z}^k\}. \quad (2)$$

Assuming that this approximation error is small, the system equations are linearized and the Fisher information is transformed through this linearized system, while the state prediction is transformed through the nonlinear system equations. This procedure is that of the extended Kalman filter (EKF) in the information form [3, 9, 17].

2.1 Adaptive Stochastic Map Implementation

In order to illustrate the approach outlined above we have implemented a two-dimensional version of a stochastic map based on the work by Smith *et al.* [14]. In this model, the robot senses features in the environment through range and bearing measurements relative to the robot’s current state (position and orientation). These measurements are used to create a map of the environment, and concurrently localize the robot.

We use $\mathbf{x}[k] = \hat{\mathbf{x}}_{k|k} + \tilde{\mathbf{x}}[k]$ to represent the system state vector $\mathbf{x} = [\mathbf{x}_r^T \ \mathbf{x}_1^T \ \dots \ \mathbf{x}_N^T]^T$, where \mathbf{x}_r and $\mathbf{x}_1 \dots \mathbf{x}_N$ are the robot and feature states, respectively; $\hat{\mathbf{x}}$ is the estimated state vector; $\tilde{\mathbf{x}}$ is the error estimate; k is the time index, where samples are taken every $t = kT$ seconds, for some constant period T . The Fisher information, $\mathbf{I}_{k|k} = E\{\tilde{\mathbf{x}}[k]\tilde{\mathbf{x}}^T[k]\}^{-1}$, of the system takes the form:

$$\mathbf{I}_{k|k}^{-1} = \mathbf{P}_{k|k} = \begin{bmatrix} \mathbf{P}_{rr} & \mathbf{P}_{r1} & \cdots & \mathbf{P}_{rN} \\ \mathbf{P}_{1r} & \mathbf{P}_{11} & \cdots & \mathbf{P}_{1N} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{P}_{Nr} & \mathbf{P}_{N1} & \cdots & \mathbf{P}_{NN} \end{bmatrix}_{k|k} \quad (3)$$

where \mathbf{P} is the inverse of the Fisher information matrix and is identical to the error covariance of the system.

To keep the system model as general as possible, the vehicle dynamics are modeled as a compounding operation [14]. That is, the robot has three degrees of freedom, position (x_r, y_r) and inclination ϕ relative to the internal reference frame. We denote the vehicle’s state by $\mathbf{x}_r = [x_r \ y_r \ \phi]^T$. We define the input to the system $\mathbf{u}[k]$ by a state displacement relative to the current state of the robot. The state transition, \mathcal{F} , is given by:

$$\begin{aligned} \mathbf{x}_r[k+1] &= \mathbf{f}(\mathbf{x}[k], \mathbf{u}[k]) + \mathbf{G}(\mathbf{u}[k])\mathbf{d}_x \\ &= \mathbf{x}_r[k] + \mathbf{T}_\phi[k]\mathbf{u}[k] + \mathbf{G}(\mathbf{u}[k])\mathbf{d}_x, \end{aligned} \quad (4)$$

where $\mathbf{G}(\mathbf{u}[k])$ scales the noise process \mathbf{d}_x as a function of the distance traveled, that is

$$\mathbf{G}(\mathbf{u}[k]) = \begin{bmatrix} \gamma_x \sqrt{u_x^2[k] + u_y^2[k]} & 0 & 0 \\ 0 & \gamma_y \sqrt{u_x^2[k] + u_y^2[k]} & 0 \\ 0 & 0 & \gamma_\phi \end{bmatrix}$$

and

$$\mathbf{T}_\phi[k] = \begin{bmatrix} \cos(\phi[k]) & -\sin(\phi[k]) & 0 \\ \sin(\phi[k]) & \cos(\phi[k]) & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The ϕ in this expression comes from the robot position $\mathbf{x}_r[k]$. The covariance of \mathbf{d}_x is set to the identity matrix without loss of generality.

Equation (4) does not take into account any of the vehicle’s real dynamics. Thus the model is very general. However, if the vehicle’s dynamics were known, they could be used to ensure that the vehicle movements remain realistic.

The observation model \mathcal{H} for the system is given by

$$\mathbf{z}[k] = \mathcal{H}(\mathbf{x}[k], \mathbf{d}_z) = \mathbf{h}(\mathbf{x}[k]) + \mathbf{d}_z, \quad (5)$$

where $\mathbf{z}[k]$ is the observation vector of range and bearing measurements. The observation model, \mathbf{h} , defines the coordinate transformation from state to observation coordinates. The stochastic process \mathbf{d}_z , is assumed to be white, Gaussian, and independent of $\mathbf{x}[0]$ and \mathbf{d}_x , with covariance \mathbf{R} .

The steps of the stochastic map consist of robot displacement, new-feature integration and re-observation of features. Our implementation is very similar to that of Smith *et al.*, and the reader is referred to [14] for a detailed derivation. The difference in our approach is that we include adaption at every time step. This is discussed next.

2.2 The Adaptation Step

At each time step the algorithm seeks to determine the transformation \mathcal{M} from the current information $\mathbf{I}_k|k$ to the resulting information $\mathbf{I}_{k+1}|k+1$ due to an action $\mathbf{u}[k]$ and, from this, infer the optimal action $\mathbf{u}[k]$. Combining the vehicle prediction and EKF update step of the stochastic map, \mathcal{M} is

$$\mathbf{I}_{k+1|k+1} = (\mathbf{F}_x \mathbf{I}_k|k^{-1} \mathbf{F}_x^T + \mathbf{G}(\mathbf{u}[k]) \mathbf{G}(\mathbf{u}[k])^T)^{-1} + \mathbf{H}_x(\hat{\mathbf{x}}_{k+1|k})^T \mathbf{R}^{-1} \mathbf{H}_x(\hat{\mathbf{x}}_{k+1|k}). \quad (6)$$

Where \mathbf{F}_x and \mathbf{H}_x are the Jacobian of \mathbf{f} and \mathbf{h} with respect to \mathbf{x} evaluated at $\hat{\mathbf{x}}_k|k$ and $\hat{\mathbf{x}}_{k+1|k}$ respectively. The first term on the right of Equation (6) represents the previous information of the system, as well as the loss of information that occurs due to the action $\mathbf{u}[k]$. The second term represents the additional information gained by the system due to observations after the action $\mathbf{u}[k]$. As this quantity is a function of $\mathbf{x}[k+1]$, which is unknown, we approximate $\mathbf{x}[k+1]$ by $\hat{\mathbf{x}}_{k+1|k}$ by the assumption of Equation (2). The action that maximizes the information can be expressed as

$$\mathbf{u}[k] = \max_{\mathbf{u}} \arg \mathbf{I}_{k+1|k+1} = \min_{\mathbf{u}} \arg \mathbf{P}_{k+1|k+1}. \quad (7)$$

The information is a matrix and we require a metric to quantify the information. Further, it is desired that this metric have a simple physical interpretation.

We have chosen to define the metric by a cost function $C(\mathbf{P})$, which gives the total area of all the error ellipses, (i.e., highest probability density region) and is thus a measure of our confidence in our map and robot position. That is,

$$C(\mathbf{P}) = \prod_j \sqrt{\lambda_j(\mathbf{P}_{rr})} + \sum_{i=1}^N \prod_j \sqrt{\lambda_j(\mathbf{P}_{ii})}, \quad (8)$$

where $\lambda_j(\cdot)$ is the j -th eigenvalue of its argument. The direction to move can be found by taking the derivative of the cost C with respect to the action \mathbf{u} over the available action space of the robot. This procedure optimizes the information at every time step, and is, therefore, a *local* optimization.

3 Simulation Results

In order to validate the method, the algorithm described above was implemented in simulation. In these simulations, as well as in the experiments, a numerical approximation was performed by evaluating Equation (8) in the action space of the robot, defined by the constraints place on the robot. The robot was constrained to move a distance of 0, 10 or 20 cm at each time step, and could only turn in increments of 22.5° . Further, the vehicle was constrained to not get closer than 40 cm to the features (vertical PVC pipes) as the sonar signal at this range becomes unreliable. In all these simulations, range error was assumed to have a standard deviation of 2 cm while the standard deviation for the bearing was 10° . The standard deviation for the vehicle odometry was set to 5% of the distance traveled. The angle uncertainty was set to 1° . These parameters were chosen as they resemble the situation for the experiments described below. Two different types of simulation are described.

3.1 Adaptive Vehicle Motion

In this simulation, it was assumed that the robot stopped and took a *complete* scan of the environment before continuing. The algorithm chose where to move adaptively. The algorithm takes advantage of the fact that the measurements have different certainty in range and bearing, thus forming an error ellipse. Notice, however, that if the observations of a target have close to equal standard deviation in all directions in space (i.e., a circle), the robot will not move, as the loss of information from odometric error is larger than the information gained by taking a measurement from a different location.

Figure 1 shows the three typical paths of the robot in the presence of two features (8.4 cm radius PVC tubes) as a result of adaption. The robot starts at (0,0) and the paths, solid, dashed, and dashed-dotted lines, occur with approximately equal frequency. The dotted lines around the PVC tubes denote the constraints placed on the robot for how close it can come to the PVC tubes while still obtaining valid sonar returns. (This is a limitation of the standard Polaroid sonar driver circuit). The resulting 95% confidence ellipses are drawn for the middle path, for the estimated position of the center of the features and for the robot's position. The robot's true position is indicated by a '+' while the estimated position is shown by a 'x'. The robot stops moving after about 15 time steps, as more information is lost than gained by moving. Also notice that the robot moves larger distances in the beginning before taking the next scan than at the end of the run; more information is gained by moving in the beginning and taking a measurement from a different location than towards the end of a run.

Figure 2 shows the average over 2000 runs, of the total cost, as defined by Equation (8), of the system under adaption as a percentage of the total cost of the system when moving randomly (solid line), or moving along the negative x-axis (dashed line) under no-adaption. In all the no-adaption runs, the robot move distance was 10 cm. As can be seen, the adaption procedure quickly reaches a high confidence map of the environment, while the random motion slowly starts catching up after the 8th time step. When moving along the negative x-axis, the confidence actually decreases after about the 15th time step, as the robot is so far away from the features that it loses more information by moving than what it gains by sensing at each time step.

3.2 Adaptive Sensing

In using sonar for mapping an environment, one is limited by the relatively slow speed that measurements can be acquired. Thus, we now impose the additional constraint that the sonar can only scan an angle of 15° (i.e., 50% more than the measurement bearing standard deviation) at each time step. The algorithm is required to decide where to direct the attention of the robot. The algorithm thereby adaptively decides where to look as well as where to move. This can be easily implemented in the framework outlined above by adding an additional action u_θ to be controlled and solving Equation (8) given the constraints of the scan angle.

To compare the simulation results to the experi-

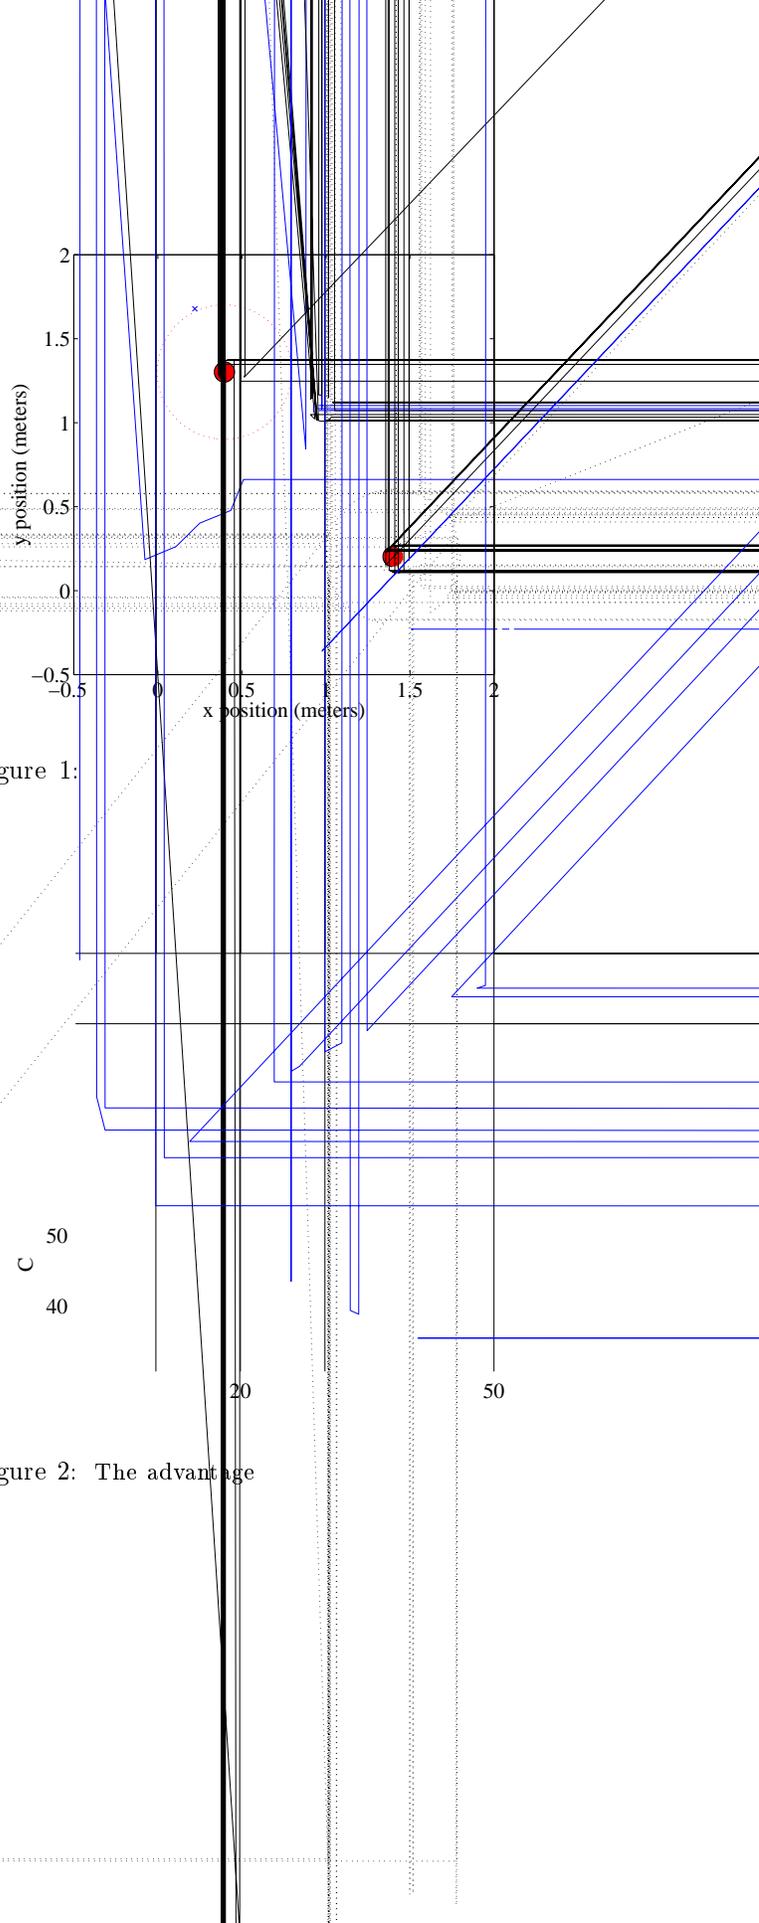


Figure 1:

50
40

20

50

Figure 2: The advantage

randomly under no-adaption, while the dashed line denotes the case when the robot moves along the negative x-axis under no-adaption as a function of sonar returns. The dotted vertical line indicates the point where the adaptive case was terminated as 50 time steps was reached. As can be seen from Figure 3, the adaptive method obtains a map with high confidence after relatively few sonar returns. After a total of 20000 returns, moving randomly and moving along the negative x-axis only achieves 93% and 33% of the confidence level that the adaptive method obtained in 850 sonar returns. Thus, the adaptive strategy requires fewer measurements, and achieves a higher confidence level than any of the no-adaption strategies. The actual vehicle’s motion is very similar to that shown in Figure 1.

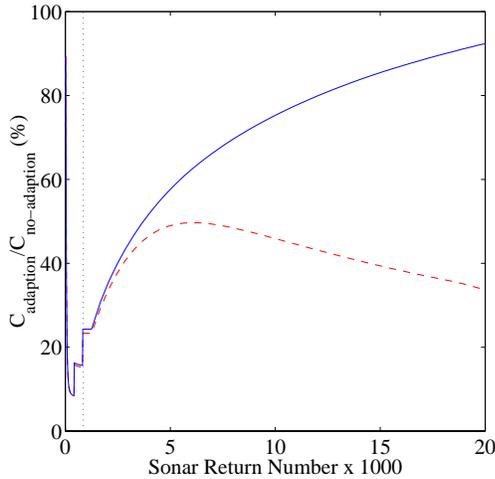


Figure 3: Advantage of adaptive sensing and motion control. The solid line denotes the case when the robot moves randomly under no-adaption, while the dashed line denotes the case when the robot moves along the negative x-axis under no-adaption. The vertical dotted line indicates the point at which the adaptive method had completed its 50 time steps and terminated.

4 Experimental Results

To demonstrate the usefulness of the adaptive method, we implemented the algorithm on a Nomad Scout robot [16]. The robot was equipped with a Polaroid 6500 series ultrasonic sensor operating at 50 kHz mounted on a stepping motor, enabling the sensor to rotate in increments of 0.9 degrees and take measurements, shown in Figure 4. The error in the sensor and the vehicle odometry was assumed to be the same as that used in the simulations. The same constraints

were employed as well. The sonar returns were assumed to be mainly specular, therefore the method of Regions of Constant Depth (RCD) was used to extract features from the scans [8]. The data association problem was solved by a simple nearest neighbor strategy [2]. That is, the closest RCD to an estimated feature was used. This method works well when the features are far apart, as in these experiments. Figure 5 shows the configuration of the robot (triangle) and features (PVC tubes of known radius) in the experimental setup superimposed on a crude map of the room. The dots indicate single sonar returns of the Polaroid sensor when taking a complete scan of the environment. Each complete scan of the environment consisted of 400 sonar returns. In these experiments 15 motion steps were performed in each run.



Figure 4: The Nomad Scout robot with the Polaroid ultrasonic sensor mounted on top.

Figure 6 shows the advantage of adaption over no-adaption for a representative Nomad run, similar to the simulation of Figure 2. As can be seen from this figure, the advantage of performing adaption is clear. Further, the experimental result is well within the one standard deviation bound of the simulated predicted result over 2000 runs.

Figure 7 shows the advantage of performing adaptive motion and sensing over moving in a straight line along the negative x-axis for the Nomad Scout robot (solid line). The simulated result is shown by a dashed line along with the one standard deviation bounds for 2000 simulated runs. As can be seen, the experimental cost ratio is within the one standard deviation of the simulated value. The adaptive strategy produces a high confidence map after relatively few measurements, consistent with the simulations in Figure 3. The non-adaptive method only reaches a confidence level of 50% of the adaptive level, even after taking

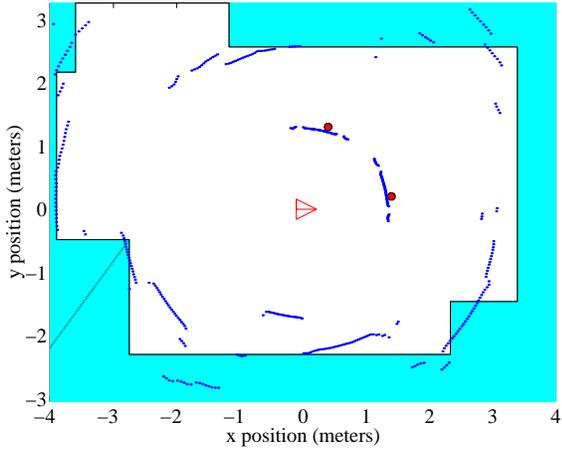


Figure 5: The returns from the Polaroid sensor (dots) with the room model superimposed. The robot is drawn as a triangle, the PVC pipes are indicated by small filled circles. The filled regions represent a crude map of the room.

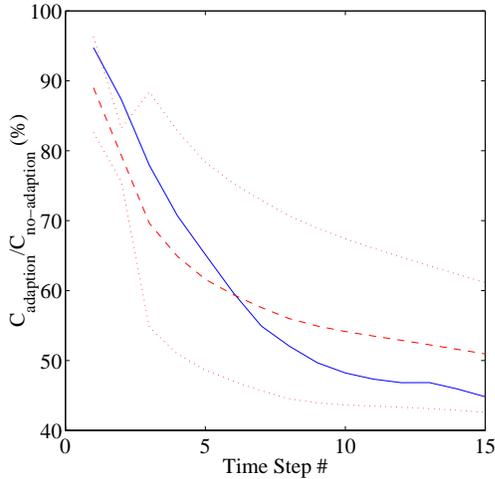


Figure 6: Advantage of adaption over no-adaption in a representative Nomad run. The solid line is the ratio of the cost when performing adaption over the cost of moving in a straight line along the negative x-axis. The dashed line is the average simulation result over 2000 runs, with the dotted line denoting the one standard deviation bounds.

more than 20 times as many measurements.

5 Discussion and Conclusions

We have introduced a method for performing adaptive concurrent mapping and localization (CML) in *a priori* unknown environments for any number of features. The adaptive method was based on choosing

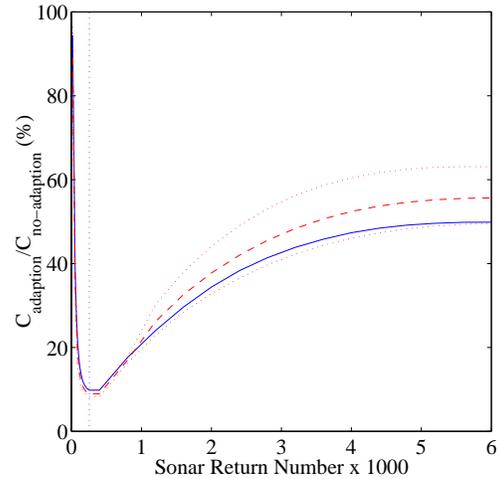


Figure 7: Advantage of adaptive sensing and motion in a representative Nomad experiment. The solid line is the cost of adaptive sensing and motion over the cost of moving in a straight line along the negative x-axis. The dashed line is the average of 2000 simulated runs and the dotted lines show the one standard deviation bound. The dotted vertical line indicates the time where the adaptive case terminated as 15 time steps were completed.

actions that, given the current knowledge, would maximize the information gained in the next measurement. This approach can easily be implemented as an extra step in the stochastic map approach to CML. The validity and usefulness of the approach were verified in both simulation and real world experiments.

Based on our results with real data, we feel confident in the accuracy of the simulation in predicting experimental outcomes. For example, the three typical paths for the robot shown in Figure 1 are representative of both adaptive simulations and real-data experiments, and the plots comparing the advantage of adaptation over no-adaption are similar.

The advantage of performing adaptive CML is notable when only adapting the motion of the vehicle (Figures 2 and 6). However, more substantial gains are obtained when performing adaptive motion *and* sensing. This is apparent from Figures 3 and 7, where the number of sonar returns required to obtain a given confidence level is an order of magnitude fewer than when not performing adaption. This strategy could be especially useful for AUVs, as mechanically scanned sonars are slow, and electronically scanned sonars have a large computational cost under non-adaptive strategies [10].

There are many important and difficult problems not addressed in this paper that must be solved in order to have truly autonomous vehicles, such as, the

data association problem, the validity of linearization in the EKF, feature ontology, and dealing with occlusion and geometric constraints. Future research will concentrate on these issues and on integrating adaptive CML with a dynamic path planner and obstacle avoidance system [6]. Further, integrating and identifying other objects, such as planes, corners and convex objects is being investigated. Work is also in progress to implement this technique with a 675 KHz underwater scanning sonar in a testing tank at MIT, using a large workspace robotic positioning system. This technique provides one starting point for investigating biomimetic echolocation [1] in an underwater robotic system.

Acknowledgment

This research has been funded in part by the Naval Undersea Warfare Center, Newport, RI, U.S.A. H.J.S.F. acknowledges the support of NFR (Norwegian Research Council) through grant 109338/410. J.J.L. acknowledges the support of the Henry L. and Grace Doherty Assistant Professorship in Ocean Utilization. Further, the authors would like to thank Seth Lloyd for fruitful discussions and comments. The authors would also like to thank Jan Meyer, Jong H. Lim, Marlene Cohen, and Tom Consi for their help with the hardware integration.

References

- [1] W. Au. *The Sonar of Dolphins*. New York: Springer-Verlag, 1993.
- [2] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.
- [3] Y. Bar-Shalom and X.-R. Li. *Estimation and Tracking: Principles, Techniques, and Software*. Artech House, 1993.
- [4] K. S. Chong and L. Kleeman. Sonar based map building for a mobile robot. *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 1700–1705, April 1997.
- [5] I. J. Cox and J. J. Leonard. Modeling a dynamic environment using a Bayesian multiple hypothesis approach. *Artificial Intelligence*, 66(2):311–344, April 1994.
- [6] H. J. S. Feder and J.-J. E. Slotine. Real-time path planning using harmonic potentials in dynamic environments. *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 874–881, April 1997.
- [7] J. J. Leonard and H. F. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, pages 1442–1447, Osaka, Japan, 1991.
- [8] J. J. Leonard and H. F. Durrant-Whyte. *Directed Sonar Sensing for Mobile Robot Navigation*. Boston: Kluwer Academic Publishers, 1992.
- [9] J. S. Manyika and H. F. Durrant-Whyte. *Data Fusion and Sensor Management: A decentralized information-theoretic approach*. New York: Ellis Horwood, 1994.
- [10] M. Medeiros and R. Carpenter. High resolution array signal processing for AUVs. In *AUV 6*, pages 10–15, 1996.
- [11] P. Moutarlier and R. Chatila. Stochastic multi-sensory data fusion for mobile robot location and environment modeling. *5th Int. Symposium on Robotics Research*, pages 207–216, 1989.
- [12] W. D. Rencken. Concurrent localisation and map building for mobile robots using ultrasonic sensors. In *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, pages 2192–2197, Yokohama, Japan, 1993.
- [13] C. M. Smith and J. J. Leonard. A multiple hypothesis approach to concurrent mapping and localization for autonomous underwater vehicles. In *International Conference on Field and Service Robotics*, Sydney, Australia, 1997.
- [14] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. Cox and G. Wilfong, editors, *Autonomous Robot Vehicles*. Springer-Verlag, 1990.
- [15] W. K. Stewart. Remote sensing issues for intelligent underwater systems. In *International Conference on Computer Vision*, pages 230–235, 1991.
- [16] N. Technologies. www.robots.com, 1998.
- [17] J. Uhlmann. *Dynamic Map Building and Localization: New Theoretical Foundations*. PhD thesis, University of Oxford, 1995.
- [18] J. Vaganay, J. G. Bellingham, and J. J. Leonard. Outlier rejection for autonomous acoustic navigation. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 2174–2181, April 1996.