

AN EXPERIMENTAL INVESTIGATION OF COOPERATIVE SLAM

Matthew Walter, John Leonard

*Massachusetts Institute of Technology
Cambridge, MA, USA
{mwalter, jleonard}@mit.edu*

Abstract: This paper considers the problem of cooperative navigation and mapping by a heterogeneous team of multiple autonomous underwater vehicles (AUVs). We address a form of cooperative Simultaneous Localization and Mapping (SLAM) in which only one vehicle is responsible for maintaining estimates of the map and poses for each robot. By combining inter-vehicle measurements with observations of the environment made by each vehicle, the result is a better knowledge of the poses of each robot in the group. We present results of an experiment in which three land robots perform cooperative SLAM in a large, ambiguous environment. Comparisons with ground truth are provided.

Keywords: Navigation, Mapping, Multi-Robot systems

1. INTRODUCTION

Due in large part to the spatial benefits, a great deal of attention of late in the robotics community has been focused on developing networks of coordinated autonomous vehicles. There are a wealth of applications for air, surface, and sub-sea systems (Fiorelli *et al.*, 2003) which would benefit from such an approach. Among the components necessary for achieving autonomy is the ability to navigate reliably and efficiently in environments for which there is little or no *a priori* information.

Motivated by this requirement, much effort has been spent on addressing the problem of single vehicle localization by posing it in a probabilistic state estimation framework. A number of techniques have been presented for both local navigation as well as the so-called global “kidnapped robot” problem of localization in a completely unknown environment, notably the work detailed in (Borenstein *et al.*, 1996; Konlige, 1999). Localization for the case of a team of robots, though, is a relatively new field. For vehicles operating in an *a priori* known environment, in (Fox *et al.*, 2000) the authors use Monte Carlo particle filters to represent the probability distributions for each of

the vehicles in a group. Inter-vehicle observations are incorporated to improve state estimation, though the corresponding vehicle correlation is ignored, resulting in beliefs which may suffer from overconfidence. Alternatively, in (Roumeliotis and Bekey, 2002) the authors present a method which relies upon a single Kalman Filter distributed among the vehicles in the group. Upon the observation of one robot by another, it is shown that by exchanging only individual estimates of pose and covariance, inter-vehicle correlation can still be maintained.

1.1 AUV Navigation

Due to the unique characteristics of the ocean environment, many of the technologies used for land and air-based vehicles (*e.g.* GPS and laser range finders) are not available options underwater. Traditionally, position is estimated from acoustic time-of-flight measurements, requiring the deployment and calibration of a network of transponders. Long baseline (LBL) (Milne, 1983) functions at distances of up to 10 km from the beacon array though the period between measurements can be large (10 seconds) and the accuracy

is generally limited to the order of 1 meter. It is possible to achieve millimeter accuracy using higher frequency variations on LBL but for distances no greater than 100 meters (Whitcomb *et al.*, 1998).

The pose estimate can be improved through the incorporation of dead reckoning sensors including gyros, Inertial Navigation Systems (INS), and Doppler Velocity Logs (DVL) (Larsen, 2000). The corresponding measurements of heading, angular rates, and vehicle velocity provide a dead reckoned pose estimate (Whitcomb *et al.*, 1998) which is prone to error drift. Research has recently been presented to reduce the localization error using vision-based mosaicing (Fleischer, 2000; Eustice *et al.*, 2004). While such implementations offer a promising alternative to the need to deploy beacon arrays, they do not address the case of cooperative localization.

Our approach to moving baseline navigation utilizes a heterogeneous team of vehicles to achieve improved localization accuracy without requiring an LBL net or that each vehicle be equipped with a plethora of sensors. Through the use of a logical extension to a delayed state Kalman Filter (Leonard *et al.*, 2002), we consider a state estimation scheme based upon inter-vehicle measurements and communication. With one or more vehicles capable of performing Simultaneous Localization and Mapping (SLAM), the corresponding navigation precision can then be exploited to obtain more accurate estimates of pose for each vehicle in the group.

By formulating moving baseline navigation in a Kalman Filter framework, our work is closely aligned with that of (Roumeliotis and Bekey, 2002). The approaches differ in that we directly incorporate SLAM capabilities which utilize inter-vehicle measurements as well as observations of the environment made by each robot in the group. Notably, our work is unique in a number of areas including the use of

- multiple vehicle delayed state kalman filter for cooperative SLAM
- perceptual grouping and feature tracking based upon measurements made concurrently by multiple robots
- joint compatibility branch and bound for multiple robot data association

In the remainder of the paper, we present the general structure of the filter as well as the methods we have chosen for data association and the incorporation of measurement data for state estimation. Results are then presented for a land-based experimental validation of our algorithm for a team of three vehicles operating in an indoor environment.

2. PROBLEM FORMULATION

Consider a heterogeneous group comprised of M vehicles, categorized as either *slaves* or *masters*. The so-called *slave* robots are those which are equipped

with a limited set of sensors both for dead reckoning as well as observing the environment. Additionally, their dynamics may be poorly understood resulting in models which are relatively inaccurate. As a result of these deficiencies, *slave* vehicles are thus prone to error when navigating individually in large environments.

The *master* vehicles, on the other hand, are better suited for navigation, having acceptable dead reckoning capabilities as well as an accurate sensor package for environmental observations. Additionally, they may be able to perform SLAM or implement one of the aforementioned vision-based navigation algorithms. Critical to our moving baseline algorithm, we also require that *master* vehicles have a means of making inter-vehicle measurements which, in the case of AUV's, may be provided by the ability to acoustically interrogate the other members of the group. It is also assumed that there exists a communications channel over which sensor data is shared.

2.1 Simultaneous Localization and Mapping

We specifically address the case in which one of the *master* vehicles benefits from improved navigational accuracy as a result of the ability to perform SLAM. There are several approaches for implementing SLAM (Montemerlo *et al.*, 2002; Thrun, 1998) and in this paper we consider the feature-based formulation which has previously been implemented for underwater vehicles (Williams *et al.*, 2000; Newman *et al.*, 2003). The addition of inter-vehicle measurements together with shared observations of the environment made by the other robots in the group provides a context for performing cooperative SLAM. Figure 1 provides a simple sketch of a team of three robots in which one of the *master* vehicles is able to initialize a point feature using observations made by the other vehicles in conjunction with inter-vehicle measurements.

3. MOVING BASELINE NAVIGATION

Consider a group of $M = 3$ vehicles, one *slave* and two *masters*. Let the state, \mathbf{x} , include a pose (position and orientation) history for each of the vehicles together with the mapped features. Represented by a Gaussian distribution, the joint pose and map posterior is then described by the mean, $\hat{\mathbf{x}}$, and covariance, \mathbf{P}

$$\hat{\mathbf{x}}[k] = \begin{bmatrix} \hat{\mathbf{x}}_{v_1}^{(n)}[k] \\ \hat{\mathbf{x}}_{v_2}^{(n)}[k] \\ \hat{\mathbf{x}}_{v_3}^{(n)}[k] \\ \hat{\mathbf{x}}_{f_1}[k] \\ \vdots \\ \hat{\mathbf{x}}_{f_i}[k] \end{bmatrix} \quad (1a)$$

$$\mathbf{P}[k|k] = \begin{bmatrix} \mathbf{P}_{vv}^{(n)}[k|k] & \mathbf{P}_{vf}^{(n)}[k|k] \\ \mathbf{P}_{fv}^{(n)}[k|k] & \mathbf{P}_{ff}[k|k] \end{bmatrix} \quad (1b)$$

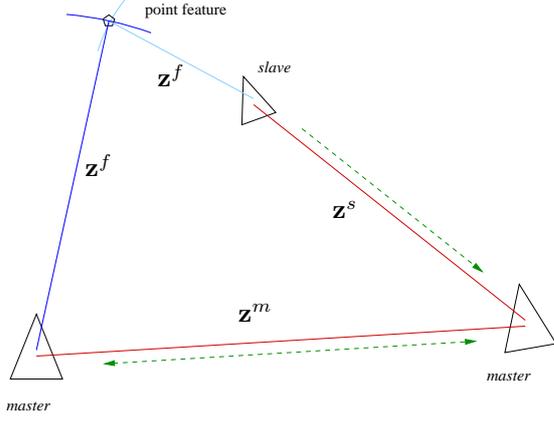


Fig. 1. Moving baseline representation of a group of two *master* vehicles and one *slave*. The *master* vehicle at the right is able to map the point feature using the individual observations made by the other members of the group together with the inter-vehicle measurements.

where, $\hat{\mathbf{x}}_{v_i}^{(n)}[k]$ represents the current and past n poses of the vehicle's trajectory, estimated based upon all information available up to time k .

$$\hat{\mathbf{x}}_{v_i}^{(n)}[k] = \begin{bmatrix} \hat{\mathbf{x}}_{v_i}^{k-n}[k] \\ \hat{\mathbf{x}}_{v_i}^{k-(n-1)}[k] \\ \vdots \\ \hat{\mathbf{x}}_{v_i}^k[k] \end{bmatrix} \quad (2)$$

The sub-blocks of the covariance matrix correspond to the smoothed vehicle and feature uncertainty and cross-correlation. More specifically, as shown in Equation 3, $\mathbf{P}_{vf}^{(n)}[k|k]$ includes the correlation between features and the pose history for each vehicle in the group. Similarly, $\mathbf{P}_{vv}^{(n)}[k|k]$ incorporates the covariance estimate for each vehicle as well as the cross-correlation for each state in the trajectory history.

$$\mathbf{P}_{vf}^{(n)}[k|k] = \begin{bmatrix} \mathbf{P}_{v_1^{k-n} f_1}[k|k] & \dots & \mathbf{P}_{v_1^{k-n} f_l}[k|k] \\ \vdots & \ddots & \vdots \\ \mathbf{P}_{v_2^k f_1}[k|k] & \dots & \mathbf{P}_{v_2^k f_l}[k|k] \\ \mathbf{P}_{v_2^{k-n} f_1}[k|k] & \dots & \mathbf{P}_{v_2^{k-n} f_l}[k|k] \\ \vdots & \ddots & \vdots \\ \mathbf{P}_{v_3^k f_1}[k|k] & \dots & \mathbf{P}_{v_3^k f_l}[k|k] \\ \mathbf{P}_{v_3^{k-n} f_1}[k|k] & \dots & \mathbf{P}_{v_3^{k-n} f_l}[k|k] \\ \vdots & \ddots & \vdots \\ \mathbf{P}_{v_3^k f_1}[k|k] & \dots & \mathbf{P}_{v_3^k f_l}[k|k] \end{bmatrix} \quad (3)$$

3.1 Kalman Filter

The standard extended Kalman filter (Gelb, 1982) time prediction and measurement update steps are used to manage the joint prior over pose history and features. The *master* vehicle performing SLAM propagates the state estimate for each robot in time via

a nonlinear discrete time model with $\boldsymbol{\mu}_{v_i}[k+1]$ the communicated inputs. Errors in the dynamics are represented by white Gaussian noise, $\mathbf{w}_{v_i} \sim N(\mathbf{0}, \mathbf{Q}_{v_i})$.

$$\mathbf{x}_{v_i}^{k+1}[k+1] = \mathbf{f}(\mathbf{x}_{v_i}^k[k], k) + \mathbf{g}(\boldsymbol{\mu}_{v_i}[k+1], \mathbf{w}_{v_i}) \quad (4)$$

The mean and covariance are then propagated with the standard extended Kalman filter time propagation step. The appropriate Jacobians are used as presented in (Smith *et al.*, 1990) to manage the correlation between both pose history and features.

$$\hat{\mathbf{x}}_{v_i}^{k+1} = \mathbf{f}(\hat{\mathbf{x}}_{v_i}^k[k], k) + \mathbf{g}(\boldsymbol{\mu}_{v_i}[k+1], \mathbf{0}) \quad (5a)$$

$$\mathbf{P}_{(vv)_i}^{k+1} = \mathbf{F}_x(\hat{\mathbf{x}}) \mathbf{P}_{(vv)_i}^k \mathbf{F}_x^T(\hat{\mathbf{x}}) + \mathbf{G}_w \mathbf{Q}_{v_i} \mathbf{G}_w^T \quad (5b)$$

Both inter-vehicle and feature observations made either by the SLAM robot or another member of the group are used to update the state estimate. Vehicle and feature measurements, represented as \mathbf{z}^v and \mathbf{z}^f , respectively, are assumed to be nonlinear in the state and corrupted by zero mean white Gaussian noise, $\mathbf{v}_{v_i}^{f,v} \sim N(\mathbf{0}, \mathbf{R}_{v_i}^{f,v})$, with the uncertainty varying with the vehicle.

$$\mathbf{z}^v = \mathbf{h}^v(\mathbf{x}_{v_i}, \mathbf{x}_{v_j}) + \mathbf{v}_{v_i}^v \quad (6a)$$

$$\mathbf{z}^f = \mathbf{h}^f(\mathbf{x}_{v_i}, \mathbf{x}_f) + \mathbf{v}_{v_i}^f \quad (6b)$$

The mean and covariance are then updated in the standard manner

$$\mathbf{S} = \mathbf{H}_x \mathbf{P}_{vf}^{t_n}[k+1|k] \mathbf{H}_x^T + \mathbf{R}_{v_i}^{f,v}$$

$$\mathbf{K} = \mathbf{P}_{vf}^{t_n}[k+1|k] \mathbf{H}_x^T \mathbf{S}^{-1}$$

$$\hat{\mathbf{x}}[k+1] = \hat{\mathbf{x}}[k+1|k] + \mathbf{K} (\mathbf{z}^{f,v} - \hat{\mathbf{z}}^{f,v}) \quad (7a)$$

$$\mathbf{P}_{vf}^{t_n}[k+1|k+1] = \mathbf{P}_{vf}^{t_n}[k+1|k] - \mathbf{K} \mathbf{S} \mathbf{K}^T \quad (7b)$$

3.2 Feature Initialization

With an estimate of recent pose history for each member of the group, the SLAM robot builds the map of the environment using observations made by each vehicle. Measurements from multiple vantage points are particularly useful in the case where the sensors provide limited observability as is the case with range-only sonar used commonly for underwater vehicles.

Consider that a new feature is to be added to the map using measurements made by two vehicles at different points in time. Using the corresponding vehicle pose estimates, the SLAM robot initializes the feature based upon a nonlinear spatial relationship such as those presented in (Leonard *et al.*, 2002) for the case of range-only measurements.

$$\hat{\mathbf{x}}_{f_{i+1}}[k] = \mathbf{g}(\hat{\mathbf{x}}_{v_i}[k_i], \hat{\mathbf{x}}_{v_j}[k_j], \mathbf{z}_i, \mathbf{z}_j)$$

The mean state vector and covariance matrix are amended, maintaining the cross-correlation between feature state and vehicle poses.

$$\hat{\mathbf{x}}[k] \rightarrow \left[\hat{\mathbf{x}}^T[k] \quad \hat{\mathbf{x}}_{f_{i+1}}^T \right]^T \quad (8a)$$

$$\mathbf{P}[k|k] \rightarrow \begin{bmatrix} \mathbf{P}[k|k] & \mathbf{B}^T \\ \mathbf{B} & \mathbf{A} \end{bmatrix} \quad (8b)$$

$$\mathbf{B} = \mathbf{G}_x \mathbf{P}[k|k]$$

$$\mathbf{A} = \mathbf{G}_x \mathbf{P}[k|k] \mathbf{G}_x^T + \mathbf{G}_z \mathbf{R}_{i,j} \mathbf{G}_z^T$$

where $\mathbf{R}_{i,j}$ is a block-diagonal matrix representing the measurement noise strength.

3.3 Data Association

In order to perform the filter update step with either inter-vehicle or feature measurements, a means of performing data association is necessary. Posing the problem as a search in correspondence space, the SLAM robot must match a set of m measurements, $\mathcal{Z}_m = \{\mathbf{z}_1^f(v_i), \mathbf{z}_2^v(v_j), \dots, \mathbf{z}_m^f(v_k)\}$, possibly made by different vehicles, with the vehicles and known features while identifying spurious measurements for perceptual grouping. The search is then for the hypothesis

$$\mathcal{H} = \{\alpha_1, \alpha_2, \dots, \alpha_m\}$$

which correctly defines the correspondence

$$\{(\mathbf{z}_1^f(v_i), f_{\alpha_1}), (\mathbf{z}_2^v(v_j), v_{\alpha_2}), \dots, (\mathbf{z}_m^f(v_k), f_{\alpha_m})\}$$

While there are a number of algorithms for estimating this correspondence, we use the joint compatibility branch and bound (JCBB) test (Niera and Tardós, 2001). In considering the correlation between each of the measurement pairings, JCBB is well suited for moving baseline navigation, as the compatibility of observations made by more than one robot must be taken into account. At the cost of added computational complexity, JCBB is more robust to spurious pairings, particularly in the presences of increased pose uncertainty.

Joint compatibility generates a hypothesis incrementally, measuring the feasibility of each candidate pairing with an *implicit measurement function*

$$\xi_{i\alpha_i}(\mathbf{x}, \mathbf{h}_i^{f,v}) = 0$$

relating the state and measurement. The corresponding innovation is then given by $\nu_{i\alpha_i} = \xi_{i\alpha_i}(\hat{\mathbf{x}}, \mathbf{z}_i^{f,v})$. The compatibility of the candidate pairing, α_i , with the current hypothesis, $\mathcal{H} = \{\alpha_1, \alpha_2, \dots, \alpha_{i-1}\}$, is ensured using a vector-valued representation of the implicit function and innovation

$$\xi_{\mathcal{H}_i}(\mathbf{x}, \mathcal{Z}_i) = \begin{bmatrix} \xi_{\mathcal{H}_{i-1}}(\mathbf{x}, \mathcal{Z}_{i-1}) \\ \xi_{i\alpha_i}(\mathbf{x}, \mathbf{z}_i) \end{bmatrix} \quad (9a)$$

$$\nu_{\mathcal{H}_i} = \begin{bmatrix} \nu_{\mathcal{H}_{i-1}} \\ \nu_{i\alpha_i} \end{bmatrix} = \begin{bmatrix} \xi_{\mathcal{H}_{i-1}}(\hat{\mathbf{x}}, \mathcal{Z}_{i-1}) \\ \xi_{i\alpha_i}(\hat{\mathbf{x}}, \mathbf{z}_i) \end{bmatrix} \quad (9b)$$

The Mahalanobis metric provides a measure of the accuracy of the current hypothesis and is required to fall below a Chi-squared threshold.

$$D_{\mathcal{H}_i}^2 = \nu_{\mathcal{H}_i}^T \mathbf{C}_{\mathcal{H}_i}^{-1} \nu_{\mathcal{H}_i} < \chi_{d,\alpha}^2 \quad (10)$$

The innovation covariance for the joint hypothesis, $\mathbf{C}_{\mathcal{H}_i}$, serves as a weighting factor and can be generated incrementally with the hypothesis.

$$\mathbf{C}_{\mathcal{H}_i} = \mathbf{H}_{\mathcal{H}_i} \mathbf{P}[k|k] \mathbf{H}_{\mathcal{H}_i}^T + \mathbf{G}_{\mathcal{H}_i} \mathbf{R} \mathbf{G}_{\mathcal{H}_i}^T \quad (11)$$

$$\mathbf{H}_{\mathcal{H}_i} = \begin{bmatrix} \mathbf{H}_{\mathcal{H}_{i-1}} \\ \mathbf{H}_{i\alpha_i} \end{bmatrix} \quad \mathbf{G}_{\mathcal{H}_i} = \begin{bmatrix} \mathbf{G}_{\mathcal{H}_{i-1}} \\ \mathbf{G}_{i\alpha_i} \end{bmatrix}$$

In this manner, JCBB performs a depth-first search for the correct hypothesis, only incorporating candidate pairings which satisfy (10) and are thus compatible with the current hypothesis.

4. EXPERIMENTAL RESULTS

In order to evaluate the performance of moving baseline navigation, experiments were performed using a team of three land-based vehicles. Two B21 mobile robots served as the *masters* while a single ATRV-Jr functioned as the *slave*, both of which are manufactured by iRobot. Each vehicle was equipped with wheel encoders for measuring odometry and a SICK laser scanner for observing the environment. The *slave's* scanner had a maximum range of 8 meters while the range for the two *masters* was 25 meters. An additional laser scanner with a maximum range of 8 meters was mounted on each B21 for the purpose of making inter-vehicle measurements, as shown in Figure 2. Pole markers placed on each vehicle served both as an identification aid as well as a means of inferring the relative pose between robots. The environment consisted of 64 hurdles laid out systematically on four adjacent indoor tennis courts, providing a convenient method for measuring the feature ground truth. The vehicles were driven manually and the data was post-processed.



Fig. 2. Photograph of the experimental setup.

In line with the discussion, one *master* was designated as the SLAM robot, maintaining estimates of the pose history for each robot over a window of $n = 30$ time steps. The vehicle dynamics (4) were represented using simple kinematic models driven by the measurements from the wheel encoders. The second *master* was prone to a significant degree of heading drift which was not included in the model. Inter-vehicle

measurements were manually limited to a frequency of roughly 1Hz. Additionally, while observations of the *slave* provided a measure of relative vehicle pose, most observations of one *master* by the other yielded relative position but not orientation.

The map generated by the moving baseline navigation algorithm is shown in Figure 3 together with the ground truth hurdle locations. The estimated trajectories for each vehicle are superimposed. The apparent discontinuity in the trajectory of the second *master* corresponds to a period during which the vehicle was out of observation range. Tracking was then performed solely using dead reckoning which, as a result of the error in the heading dynamics, caused the state to diverge. The SLAM robot was able to relocate the vehicle by means of hypothesis tracking using a combination of inter-vehicle observations.

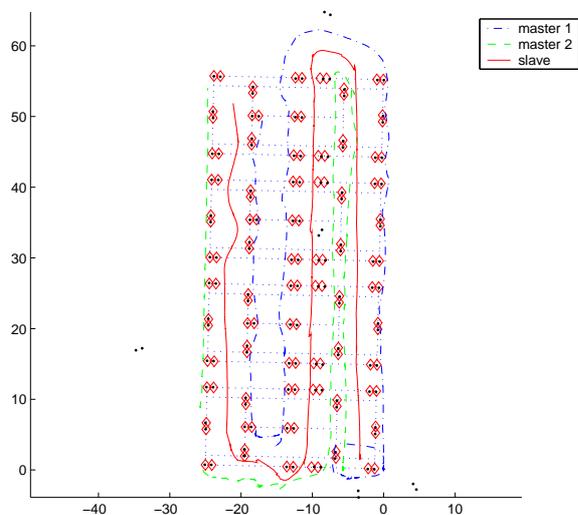


Fig. 3. Overlay of map and ground truth with diamonds designating hand-measured feature location and dots the estimated pose. The estimated trajectories for the three vehicles are included.

To better understand the benefits of moving baseline navigation, the data for the second *master* as well as the *slave* were each processed separately using a single-vehicle SLAM algorithm. Due in part to the aforementioned heading model error, the filter for the second *master* failed shortly into the run. The *slave* vehicle, on the other hand, was able to perform SLAM for roughly two-thirds of the experiment before diverging. As indicated in Figure 4, moving baseline navigation provides a pose estimate for the *slave* vehicle which is significantly more accurate than that which would be attained had the vehicle operated on its own.

Animations of the results of the experiment are available at www.mit.edu/~mwalter/Research/

5. CONCLUSION

We have presented a framework for performing navigation and mapping with a heterogeneous team of

vehicles using a logical extension to the delayed state Kalman filter. By requiring that only a few vehicles in the group be well suited for navigation, it is shown that a combination of inter-vehicle measurements together with data sharing allow each vehicle in the group to navigate with improved accuracy. We have specifically considered the case in which one of the vehicles is able to perform SLAM and have discussed the incorporation of multiple vantage point observability offered by the team of robots in both building a map of the environment as well as tracking each of the vehicles. Results of a land-based experiment involving three vehicles operating in a large, ambiguous environment were presented. A comparison with the results from one of the robots performing SLAM independently reveals the improvement in localization accuracy offered by moving baseline navigation.

Future work will address the further limitations imposed by the underwater environment, particularly with regard to communications bandwidth and the limited observability offered by range-only measurements. We are currently in the process of testing moving baseline navigation using an AUV in conjunction with two surface craft.

ACKNOWLEDGEMENTS

This research has been funded by the Office of Naval Research under grants N00014-02-C-0210, N00014-03-1-0879, and N00014-97-0202. M. Walter gratefully acknowledges the support of a National Defense Science and Engineering Graduate Fellowship.

The authors would like to thank Michael Bosse, Brad Hasegawa, Sung Joon Kim, Paul Newman, Edwin Olson, and Andrew Patrikalakis for their help in collecting the data processed in this paper. Thanks also to Jerome Vaganay for many helpful discussions related to the cooperative navigation of multiple AUVs.

REFERENCES

- Borenstein, J., B. Everett and L. Feng (1996). *Navigating Mobile Robots: Systems and Techniques*. A.K. Peters, Ltd., Wellesley, MA.
- Eustice, R., O. Pizarro and H. Singh (2004). Visually augmented navigation in an unstructured environment using a delayed state history. In: *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*.
- Fiorelli, E., P. Bhatta, N.E. Leonard and I. Shulman (2003). Adaptive sampling using feedback control of an autonomous underwater glider fleet. In: *Proceedings of the 13th International Symposium on Unmanned Untethered Submersible Technology*. Durham, NH.
- Fleischer, S. (2000). Bounded-Error Vision-Based Navigation of Autonomous Underwater Vehicles. PhD thesis. Stanford University.

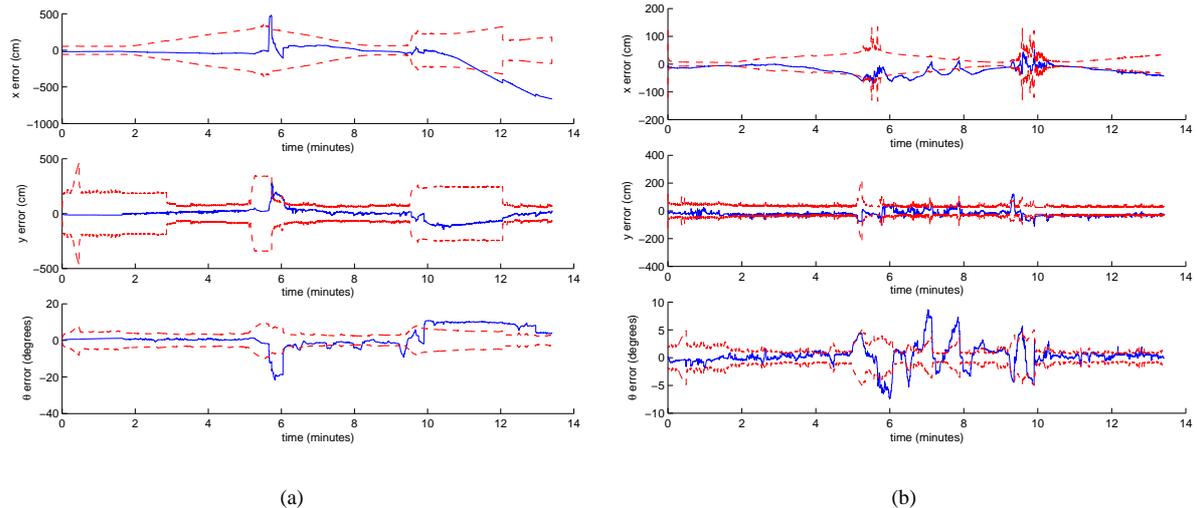


Fig. 4. Pose error together with the 3σ uncertainty bounds for the *slave* vehicle corresponding to (a) the vehicle performing SLAM independently and (b) the results of moving baseline navigation.

Fox, D., W. Burgard, H. Kruppa and S. Thrun (2000). A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots*.

Gelb, A., Ed. (1982). *Applied Optimal Estimation*. MIT Press. Cambridge, MA.

Konlige, K. (1999). Markov localization using correlation. In: *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Larsen, M. B. (2000). High performance doppler inertial navigation — experimental results. In: *IEEE Oceans*.

Leonard, J., R. Rikoski, P. Newman and M. Bosse (2002). Mapping partially observable features from multiple uncertain vantage points. *International Journal of Robotics Research*. To Appear.

Milne, P.H. (1983). *Underwater Acoustic Positioning Systems*. Spon Ltd.. New York.

Montemerlo, M., S. Thrun, D. Koller and B. Wegbreit (2002). FastSLAM: A factored solution to the simultaneous localization and mapping problem. In: *Proceedings of the AAAI National Conference on Artificial Intelligence*. AAAI. Edmonton, Canada.

Newman, P.M., J.J. Leonard and R. Rikoski (2003). Towards constant-time SLAM on an autonomous underwater vehicle using synthetic aperture sonar. In: *International Symposium on Robotics Research*. Sienna.

Niera, J. and J.D. Tardós (2001). Data association in stochastic mapping using the joint compatibility test. *IEEE Transactions on Robotics and Automation* **17**(6), 890–897.

Roumeliotis, S.I. and G.A. Bekey (2002). Distributed multi-robot localization. *IEEE Transactions on Robotics and Automation* **18**(5), 781–795.

Smith, R., M. Self and P. Cheeseman (1990). Estimating uncertain spatial relationships in robotics. In: *Autonomous Robot Vehicles* (I. Cox and G. Wilfong, Eds.). pp. 167–193. Springer Verlag.

Thrun, S. (1998). Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence* **99**(1), 21–71.

Whitcomb, L.L., D.R. Yoerger, H. Singh and D.A. Mindell (1998). Towards precision robotic maneuvering, survey, and manipulation in unstructured undersea environments. In: *Robotics Research – Eighth International Symposium* (Y. Shirai and S. Hirose, Eds.). Springer-Verlag. London. pp. 45–54.

Williams, S.B., P. Newman, M.W.M.G. Dissanayake and H.F. Durrant-Whyte (2000). Autonomous underwater simultaneous localisation and map building. In: *Proceedings of the IEEE Conference on Robotics and Automation*. San Francisco CA, USA. pp. 22–28.