

Cooperative AUV Navigation using a Single Surface Craft

Maurice F. Fallon, Georgios Papadopoulos and John J. Leonard

Abstract Maintaining accurate localization of an autonomous underwater vehicle (AUV) is difficult because electronic signals such as GPS are highly attenuated by water making established land-based localization systems, such as GPS, useless underwater. Instead we propose an alternative approach which integrates position information of other vehicles to reduce the error and uncertainty of the on-board position estimates of the AUV. This approach uses the WHOI Acoustic Modem to exchange vehicle localization estimates — albeit at low transmission rates — while simultaneously estimating inter-vehicle range. The performance capabilities of the system were tested using Oceanserver’s Iver2 and the MIT Scout kayaks.

1 Introduction

Localization or navigation of vehicles using only onboard local sensors, such as a Doppler Velocity Logger (DVL) or Inertial Measurement Unit (IMU), are certain to experience accumulated positioning error. One can, of course, utilize more precise sensors to reduce the rate of accumulated error — DVL units with error accumulation rates as low as 0.2% are commercially available. However this approach may not be satisfactory due to practical, power or financial limitations.

Regardless of the platform used, the accumulation of error and uncertainty is simply slowed, rather than bounded. The result of this is that an AUV surveying the ocean floor or a land robot building a street map must be halted on occasion so as to reset the position uncertainty — either by surfacing for a GPS fix or by repositioning at a known location. This procedure wastes both energy and time, requires a human interface and may be unacceptable in many operating environments.

The standard approach for bounding error underwater is Long Baseline (LBL). Two or more beacons are deployed at known locations — either as buoys on the

Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139.
mfallon, gpapado, jleonard@mit.edu

water surface or moored on the seabed. The AUV transmits an acoustic query to the beacons which reply in a manner which allows the AUV to estimate the beacon/AUV range and to then improve its own position estimate. Recent improvements to this system have removed the need for round-trip timing (Synced LBL) and also allowed for estimation of both range and angle using an array of receiving sensors (USBL).

While these technologies are now all commercially available, the mobility of the AUV is restricted as typical coverage is limited to an area within a few kilometers of the beacon. To relax this restriction an alternative approach considers a system in which a surface vehicle (with access to GPS) or a submerged vehicle (with accurate dead reckoning instrumentation) communicates with a fleet of much less accurately localized vehicles so as to improve the positioning of the latter. One example of this approach is the Moving Long Base Line (MLBL) navigation proposed by Vaganay *et al.* [9], in which typically two surface vehicles serve as mobile beacons for one or more AUVs. Other related recent research has been performed by Bahr *et al.* [1], Eustice *et al.* [4] and Maczka *et al.* [5]. It should also be recognized that multi-AUV navigation falls within the wider problem of multi-robot cooperative localization, see [6] for a more general introduction to the field.

In this paper, we describe experiments that extend the MLBL approach to situations in which a single surface vehicle is used to estimate the position of a submerged AUV using range-only measurements. In Section 2 the basic framework of this technique is discussed. Our algorithm is outlined in Section 3 followed by a number of modifications which improve performance. Section 4 presents the results of a combination of simulation and realistic experiments to illustrate the concept. Finally conclusions drawn from the experiments and the directions of future work are presented in Section 5.

2 Cooperative Localization Under Water

This paper retains the framework for underwater localization previously introduced in [1] and also used in [4]. We shall assume there to be one surface vehicle providing the submerged fleet of vehicles with position information while perhaps operating as a communications moderator — in the dual role of a Communications and Navigation Aid (CNA). Each of the autonomous underwater vehicles maintains a dead reckoning filter, drawing upon measurements of velocity, heading and depth. Finally, communication through the water channel is possible using the WHOI Acoustic Modem — at transmission rates of the order of 32 bytes per 10 seconds — in a process which also yields a time-of-flight measurement which can be used to estimate the inter-vehicle range.

There are a number of methods which could be used to integrate the received position information. Our earlier work, [1], proposed an algorithm which utilized the on-board dead reckoning estimate of the AUV and a pair of CNA range estimates to form a complete estimate of the AUV state vector.

The seabed, the water surface and deep sea thermoclines within the water body have the ability to cause significant multi-path signal interference and the receipt of a substantial amount of infeasible outlier measurements. A typical dataset was illustrated for a regular Long Baseline systems in [8]. For these reasons it would be reasonable to assume that the received measurement set obtained from the WHOI modem would contain substantial multi-modality, thereby motivating this approach.

However the advanced processing within the WHOI modem decoder has the ability to suppress the bulk of these effects, such that the received measurements decoded by the modem contain only a moderate amount of noise. For this reason the proposed approach instead uses an implementation of the Extended Kalman Filter.

A particle filtering approach [3] could also have been considered as this would have more accurately incorporated the non-linearity of the correction step, however because we will maintain full control of the CNA's motion this issue can broadly be avoided.

Previous *proof-of-concept* experiments illustrated that the range variance is broadly independent of range itself, however detailed examination of this was not carried out [2]. The modem transducer was then directly clamped to the underside of the kayak. Our more recent experiments have instead hung the transducer 2-3 meters below the kayak hull. We expect less noise interference from the kayak motor and less reflections from the water surface in this configuration.

Figure 1 illustrates WHOI modem range data plotted versus GPS-derived 'ground truth', as measured in the Charles River adjacent to MIT recently. Because the ground truth distance between the two vehicles was determined using imprecise GPS measurements, it is difficult to precisely estimate the distribution of the range measurements. Other issues, such as the position of the GPS sensor relative to the modem on the kayak must also be recognized. In the absence of precise ground truth, we estimate the range variance to be between 4–8m.

3 Single Surface Craft Cooperative Navigation

The configuration we will consider in this work will be of a single CNA supporting N underwater vehicles¹. Each AUV will maintain an estimate of its own position and uncertainty. This estimate will be propagated using the usual Kalman prediction step so as to integrate heading, forward and starboard velocity measurements.

As mentioned above, this estimate will be corrected using range and position information relative to an CNA using the WHOI acoustic modem. At present the 32 byte packet transmitted from the CNA shall contain latitude, longitude, depth and heading as well as a UNIX time-stamp. Transmission of a packet consists of two stages: first a *mini packet* is transmitted to initiate the communication sequence. The inter-vehicle range can be estimated using this mini packet. Following this, the information packet is transmitted in a process which lasts approximately 5-6 sec-

¹ Subsequent research will aim to relax the necessity of a dedicated surface vehicle

onds. In all, it is prudent to reserve 10 seconds per transmission. Similarly the AUV will transmit a message containing its own position estimate as well the associated covariance matrix which can be used to help the CNA plan its own supporting motion — also requiring 10 seconds per transmission.

It is envisaged that the MLBL will be integrated within a multi-AUV setup in which use of the communication channel is shared between many communicating processes. As a result the transmission rate of a position/range pair is likely to be substantially below one measurement per 10 seconds. Furthermore only a portion of transmitted messages will actually be received. For these reasons it is prudent to optimize the location from which the ASC transmits so as to maximize the benefit achieved from the correction step. Although a basic zig-zag motion plan was adopted in this work, future work will consider more elaborate motion planning for the CNA.




Fig. 1 Analysis of range estimates derived from the WHOI Modem. Upper Left: Comparison of modem range estimate (red dots) and range derived from GPS ‘ground truth’ (blue crosses) for each *fully successful* 10 second transmission period. Lower Left: Illustration of the frequency of successful transmissions. Category 0 represents an entirely failed transmission; Category 1: successful range transmission; Category 2: successful range and packet transmission. Category 2 corresponds to the modem ranges in upper left plot. Right: Histogram of range error (using estimated range versus GPS ‘ground truth’ range), also illustrated is a normal distribution fitted to the data (red, $\bar{r} = 0.66\text{m}$, $\sigma_r = 7.5\text{m}$) and the normal distribution used in the experiments in Section 4 with (cyan, $\bar{r} = 0\text{m}$, $\sigma_r = 5\text{m}$). This range data corresponds to Experiment 1.

3.1 Utilizing Partial Messages

As illustrated in Figure 1, a significant proportion of the (range) mini packets are received without the information packet — meaning that the usual correction step cannot be made²

By linearly predicting the CNA position using previous position estimates, an estimate of the CNA at this time can be formed. This estimate can then be used with the previously orphaned range measurement to allow another correction step to occur. While post processing of the data from the experiments presented in Section 4 in this manner reduced the average error by approximately one meter, in future we propose to introduce redundancy into the transmitted messages so as to avoid this scenario. See Section 5 for more discussion.

3.2 Online Compass Bias Correction

A Bayesian filter - such as a Kalman filter or particle filter - assumes that measurements are formed using unbiased estimators. Heading is however particularly difficult measurement to estimate properly. Compass accuracy can be effected by the characteristics of the local region, the magnetism of the vehicle itself and magnetic declination. It is particularly severe for imprecise sensors used aboard the CNA platform. As a result, the compass used in the experiments presented in Section 4 is a dominant source of navigation error. Typically compass bias is corrected using a calibration process which can be both complex and time consuming. In this scenario, the EKF corrections garnered using the CNA range and position can be used to estimate the compass bias and to remove its effect.

Between successive corrections of the EKF, the filter will be predicted according to the dynamical model. The frequency of the prediction step will be much higher than the correction step. The distance between the posterior estimate of a correction step at time k_1 and the predicted position at time k_2 is the estimated relative distance




Fig. 2 Compass Bias Correction Example: MLBL position estimate (blue) is corrected towards the ground truth (red) in a consistent direction. The angular correction of the 4 correction steps, $\theta_{1:4}$, can be used to form an estimate of the bias angle, which is then removed. Note that multiple iterations of the prediction step take place between each correction step.

² For a typical mission in the open ocean inter-vehicle ranges of the order of 1km are expected, making this an even more significant issue.

traveled in that time

$$\Delta \mathbf{x}_{k_2|k_1}^- = \bar{\mathbf{x}}_{k_2} - \mathbf{x}_{k_1}. \quad (1)$$

where $\mathbf{x}_{k_1} = [x_{k_1}, y_{k_1}]$ represents the state vector at time $k-1$. The CNA position and range measurement are then integrated to correct the posterior position estimate

$$\Delta \mathbf{x}_{k_2|k_1} = \mathbf{x}_{k_2} - \mathbf{x}_{k_1} \quad (2)$$

If the sensors contributing to the measurement, $\mathbf{z}_{k_2}$, are unbiased the expected value of the update will be zero. However if there exists a compass bias, the EKF will act to correct the filter in the direction opposite to the bias

$$\theta_{k_2|k_1} = \arccos \left(\frac{\Delta \mathbf{x}_{k_2|k_1}^- \cdot \Delta \mathbf{x}_{k_2|k_1}}{|\Delta \mathbf{x}_{k_2|k_1}^-| |\Delta \mathbf{x}_{k_2|k_1}|} \right) \quad (3)$$

Figure 2 illustrates the issue for a sequence of MLBL corrections for a biased compass. It can be seen that the angle of the correction is consistently in the hypothesized bias direction. However as the CNA consistently maneuvers relative to the AUV, a closed form expression for the bias angle cannot be formed.

Instead we will propose to successively estimate the bias until this effect is removed. Consider the net angular correction set of N successive corrections, $(\theta_{k-N+1}, \dots, \theta_k)$. We assume that the median of this set, given by $\tilde{\theta}_k$, will be in the direction of, but less than, the bias angle, i.e. $[0 \leq \tilde{\theta}_k < \theta_{\text{bias}}]$.

This value is assumed to be an initial estimate of the bias and used to correct the heading estimate subsequently. After the next N corrections, any remaining bias is again estimated and added to the running bias estimate. Eventually the bias will be assumed to be known and can be removed.

4 Experiments

A number of experiments were carried out in the Charles River, adjacent to MIT, to demonstrate the concept of Moving Long Baseline using the Surface Crafts for Oceanographic and Undersea Testing (SCOUT) kayaks designed in MIT and the low-cost Iver2 from Oceanserver (see Figure 4). Each of the kayaks was equipped with a WHOI modem, a compass and a GPS sensor while the Iver's basic sensor suite consisted of only a compass and a WHOI modem. The Iver2's only velocity estimate was a **constant** value of 1.028 m/s (2 knots) specified by the mission plan.

Each vehicle's onboard computer ran an implementation of the MOOS software platform [7]. Maintaining an accurately synchronized clock is essential for the estimation of inter vehicle ranges; to do so the Iver2 utilized a precisely synchronized timing board developed by Eustice et al. [4] while the SCOUT kayaks used the Plus-Per-Second (PPS) contained within its received GPS data messages.

Experiment 1: A single SCOUT kayak designated as the 'AUV' completed a survey-type mission while another kayak maintained a zig-zag pattern behind the

‘AUV’ — taking on the CNA role. The onboard GPS sensor was used to determine the ground truth position as well as to simulate forward and starboard velocities. Measurements drawn from the CNA transmissions were used by the ‘AUV’ to reduce its uncertainty. The designated ‘AUV’ carried out 1.5 circuits of a rectangle, covering approximately 1800 metres in total over a period of 37 minutes.

Note the large increase in the error of the position measurement between 22–26 minutes. This was caused by a combination of poor CNA position estimation (caused by visibility of just 4 GPS satellites) and the CNA moving close, yet parallel, to the AUV. It is envisaged that this could have been avoided with the use of a more accurate GPS unit or by forbidding the CNA from taking such a trajectory.

The following are a number of metrics for this test: mean error 12.5m, mean ‘AUV’ velocity 0.82m/s, mean CNA velocity 1.08m/s. There were 205 transmissions of which 130 were fully successful, 63 resulted in a failed packet transmission but a successful range estimate while 12 resulted in complete transmission failure. The algorithm can be seen to bound the error of the position estimate to approximately 20m.

Experiment 2: In a second fully realistic experiment, the Iver2 carried out a pre-defined ‘lawnmower’ pattern running at a depth of 2.4m while again the SCOUT kayak supported by transmitting its GPS position to the AUV via the WHOI modem. In addition the Iver2 transmitted its own MLBL position estimate, which was received by the CNA and used to plan locations from which to transmit.

Figure 6(b) illustrates the path taken by the vehicles. The test lasted 28 minutes and in total the Iver2 travelled 2 km. The AUV surfaced twice as a safety precaution. After 9 minutes the AUV first surfaced and received a GPS fix at (-201.6 -242.0) as shown as a red cross, at that time the front seat filter estimated a position of (-258.7, -276.5) while the MLBL filter estimate (-208.9, -238.1) giving an error of 66.7m and 8.3m error respectively. When the Iver surfaced for the second time (after 19 minutes), the corresponding errors were 53.7m and 14.1m. As well as estimating the AUV position with error, both of these MLBL filter estimates were within a 95% confidence interval upon surface. Note that after each surface the AUV transited




Fig. 3 Error (left) and 95% confidence (right) for the MLBL algorithm (blue) and the dead reckoning alone (green) for Experiment 1 where we have defined %90 confidence in terms of the largest eigenvalue of the covariance matrix.

from the GPS location back to its planned location on the mission path before diving and continuing the mission.

It should be mentioned that between 4–8 and 12–18 minutes no packets were successfully received by the AUV and as a result no MLBL corrections were possible (See Figure 5). This can be attributed to a number of factors

- The CNA was positioned behind the AUV and as a result churned water from the AUV propeller is likely to have reduced communication capabilities.
- With each failed transmission the AUV/CNA range grew until about 225m which is considered long for this experimental river environment ³.
- The presence of a tourist cruise ship nearby.

In future tests, precautions will be taken to avoid these issues.




Fig. 4 Vehicles Used: OceanServer Iver2 (left) and the MIT Scout kayak (right)




Fig. 5 Results for Experiment 2. Left: Modem range estimates with successful packet transmission (red dots) and modem range estimates but failed packet transmission (black crosses). Right: 95% confidence for the MLBL algorithm (blue) and the dead reckoning along (green). Note the two long portions of the run in which ranges were determined but no packet was successfully transmitted and the resultant growth in position uncertainty.

³ Note that the maximum range of the WHOI modem in the open ocean is estimated to be of the order of 4–5 times greater than the river environment.


5 Future Work and Conclusions

The concept of a single surface vehicle supporting the localization of an AUV has been outlined. Full experimental results with a single CNA supporting an Iver2 were presented. The resultant position estimate was shown to be substantially more accurate than the vehicle's own onboard navigation filter. Future work will focus on extending this framework for testing with three Iver2 vehicles and eventually towards the scenario in which a set of heterogeneous vehicles are continuously submerged with only a single vehicle occasionally surfacing to access the GPS.


Secondly, the performance of the algorithm is directly determined by the quality and frequency of received measurements. We will consider the optimization of the transmitted messages (and the re-transmission of failed data) so as to reduce the proportion of useless or partial messages received by the AUV. In this work the path taken by the CNA was an arbitrary zig-zag behind the AUV. Motion planning of the CNA's path — so as to transmit messages from the most advantageous location — will also be carried out in future.

References

1. Alex Bahr and John J. Leonard. Cooperative localization for autonomous underwater vehicles. In *Proceedings of the 10th International Symposium on Experimental Robotics (ISER)*, Rio de Janeiro, Brasil, July 2006.
2. Joseph Curcio, John J. Leonard, Jerome Vaganay, Andrew Patrikalakis, Alexander Bahr, David Battle, Henrik Schmidt, and Matthew Grund. Experiments in moving baseline navigation using autonomous surface craft. In *Proceedings of MTS/IEEE Oceans*, volume 1, pages 730–735, September 2005.
3. Arnaud Doucet, Nando de Freitas, and Neil Gordon, editors. *Sequential Monte Carlo methods in practice*. Springer-Verlag, 2000.
4. Ryan M. Eustice, Louis L. Whitcomb, Hanumant Singh, and Matthew Grund. Experimental results in synchronous-clock one-way-travel-time acoustic navigation for autonomous underwater vehicles. In *IEEE International Conference on Robotics and Automation (ICRA)*, Rome, Italy, April 2007.
5. Darren K. Maczka, Aditya S. Gadre, and Daniel J. Stilwell. Implementation of a cooperative navigation algorithm on a platoon of autonomous underwater vehicles. *Oceans 2007*, pages 1–6, 2007.
6. Anastasios I. Mourikis and Stergios I. Roumeliotis. Performance analysis of multirobot cooperative localization. *IEEE Transactions on Robotics*, 22(4):666–681, 2006.
7. Paul M. Newman. MOOS - a mission oriented operating suite. Technical Report Tech. Rep. OE2003-07, 2003.
8. Edwin Olson, John J. Leonard, and Seth Teller. Robust range-only beacon localization. *IEEE Journal of Oceanic Engineering*, 31(4):949–958, October 2006.
9. Jerome Vaganay, John J. Leonard, Joseph A. Curcio, and J. Scott Willcox. Experimental validation of the moving long base line navigation concept. pages 59–65, June 2004.



(a)



(b)

Fig. 6 Paths taken by the AUV and CNA during Experiment 1 (upper) and 2 (lower), see Section 4 for more details. CNA measurements were transmitted from the black dots. Note that the final 500m of Experiment 1 has been omitted as it overlaps with what is shown.