

Pure Range-Only Sub-Sea SLAM

P. Newman, J. Leonard
Massachusetts Institute of Technology

Abstract— This paper is about using range-only data to navigate an autonomous underwater vehicle (AUV). We assume the vehicle is equipped with conventional long base line (LBL) transceiver which measures acoustic time of flights (TOFs) between vehicle and small submerged transponders. Using only range data and no prior information other than approximate water column depth, we solve for both transponder location and vehicle trajectory. Results are given using data from a AUV operating in shallow water. A ground truth comparison is made with surveyed transponder locations and trajectory estimates from an on board Doppler/Compass/LBL derived navigation filter.

I. INTRODUCTION AND BACKGROUND

Envision the realistic scenario of an AUV performing an autonomous survey over flat, largely featureless terrain for an extended period of time. A crucial criterion for this mission is reliable navigation with bounded error.

Many AUVs use Doppler Velocity Logs (DVLs) to estimate vehicle position by dead reckoning. The inevitable dilution in accuracy with mission time of this approach can be greatly mitigated (but not overcome) by integrating DVL with an inertial navigation technology [1]. For such systems to be useful great care must be taken in their calibration and alignment. In addition, INS/DVL systems are both expensive and complex to implement – prohibitively so for many platforms.

An alternative is to use acoustic ranging equipment to provide distance measurements to transponders at known locations [2]. Such devices can provide bounded, long-term navigational precision. Conventional approaches rely on GPS-aided calibration of pre-deployed acoustic transponders before the mission begins. Apart from the ensuing fiscal and time expenses, this approach places a hard *a priori* limit

on the operating region of the vehicle.

The goal of the work described in this paper is to relax this constraint and enable the AUV to deduce the transponder locations on the fly, while at the same time using the beacons to navigate. This is the classic Simultaneous Localization and Mapping (SLAM) problem, in which there has been much recent interest and research stemming from the seminal work [3]. Furthermore, we seek a solution without recourse to doppler information even though the approach presented does not in any way preclude its use. The goal is pure, range-only SLAM.

In some ways the problem is made easier by using transponders but in others much harder. Firstly the data-association problem is removed. The correspondences between measurements and features, in this case transponders, are known because each transponder replies to an acoustic interrogation pulse at a unique frequency. However the ranges measured are huge in comparison to the movement of the vehicle between sequential measurements. Hence not only are the transponder locations only partially observable because of the range-only data but also the problem is inherently ill conditioned over small vehicle path lengths.

Regardless of the difficulties, the motivation for this approach is clear. Submerged, on-the-fly calibration of LBL transponders would enable an AUV to lay and extend its own LBL network to fit adaptive mission and navigation criteria. Alternatively, the transponders could simply be dropped from a surface vessel or air-craft with no regard to calibration. With small, cheap acoustic transponders now available, this approach could find application over barren flat terrain and complex reefs alike.

The range-only SLAM problem bears an interesting relationship to the structure-from-motion (SFM) problem in computer vision[4]. In vision, the location of a point in the image plane in effect provides an angle-only measurement to a 3-D scene point. The two problems have a similar mathematical structure, and hence can be considered "duals" of one another.

There are a number of significant differences, however, between the vision and acoustic range-only problems. In vision, the amount of data in a single image is typically very large, with on the order of dozens

P. Newman and J. Leonard are at the Department of Ocean Engineering at MIT email: pnewman/jleonard@mit.edu. Thanks to Sonardyne International for their continued support in sub-sea navigation and Mike Bosse for help before during and after GOATS. This research has been funded in part by NSF Career Award BES-9733040, the MIT Sea Grant College Program under grant NA86RG0074 (project RCM-3), and the Office of Naval Research under grants N00014-97-0202 and N00014-02-C-0210

of matched scene points per image. This contrasts strongly with the paucity of measurements in the sub-sea range-only problem. In addition, measurements are highly noise prone due in large-part to propagation conditions, multi-path, and interference from other acoustic devices operating on or near the vehicle.

McLauchlan has developed the variable state dimension filter (VSDF), an alternate, batch/recursive form of bundle adjustment for structure-from-motion in computer vision [5]. By adjusting the duration of the time-window of measurements, the method can be varied between the two extremes of either a strict recursive solution, using only the most recent measurements, or a full batch algorithm using all the data. Rikoski *et al.* [6], [7] present a related approach, known as delayed decision making (DDM), developed for SLAM using sonar data.

The in-air range-only SLAM problem has been touched on by Kantor [8] using a Kalman filter with linearization of measurement models using some prior of knowledge of transponder locations. The closely related problem of bearing only SLAM with omnidirectional cameras is considered by Deans [9][10]. Similar and inspirational to the approach taken in this paper, this work utilized batch optimization techniques.

Several researchers have considered undersea navigation with respect to a single transponder at an initially unknown location. Vaganay [11] investigated the problem of homing to a single acoustic beacon using range-only measurements. This work uses a two-phase estimation process. During initialization, the vehicle is commanded to execute a circular trajectory, and all the range measurements are used as input into a least-squares optimization to solve for the beacon position and estimated ocean currents in the environment. These results initialize an extended Kalman filter for trajectory estimation for the remainder of the mission. Larsen [12] presents a related concept, called synthetic LBL, in which a single beacon is mapped using range and range-rate measurements.

II. PROBLEM FORMULATION

Unlike many sensors frequently used in SLAM like problems the use of a range-only sensor requires the processing of data from several different vehicle locations 1. The cartesian relationship between two points is clearly unobservable given only the range between them.

In general, a single acoustic time of flight (TOF) between a transceiver at point x_A and an unknown transponder constrains the transponder to lie on a




Fig. 1. A and B represent two poses of an AUV flying at depth D over flat terrain in a water column of depth h with sound velocity c . A single TOF constrains the a transponder to lie on a circle centered on the vehicle. In general two such TOF's constrain the transponder to lie at one of two circle intersections producing a “baseline” ambiguity.

sphere around the current vehicle location. In many cases however we will know the depth of the water column h that the vehicle is operating in and the depth of the vehicle D . Under these mild assumptions the transponder is now constrained to lie in a 2D circle on the sea bed. The geometry of the problem is frequently ‘wide-and-flat’ meaning that the distance R between transponders and vehicle is usually much greater than their vertical separation $h-D$. A consequence of this is that errors in assumed beacon depth h have a small effect on the radius r of the 2D circle as $\frac{dr}{dh} = -h/R \approx 0$; $h \ll R$.

If another TOF is available from a different vehicle location x_B at known displacement from x_A then the transponder must lie at an intersection of the two range circles. The “base-line ambiguity” between the two possible solutions can only be resolved when the vehicle moves to a third location x_C that is not collinear with x_C and x_B .

The problem we wish to solve involves finding not only the location of one or more transponders $x_1, x_2, \dots$ but also the sequence of pose states $x_A, x_B, x_C, \dots$ from which TOF measurements were made. This paper uses a large scale non-linear optimization algorithm to simultaneously solve for transponder locations and vehicle trajectory. We define a state vector $\mathbf{X}$ as the concatenation of unknown vehicle and transponder locations.

$$\mathbf{X} = [x_A^T \ x_B^T \ x_C^T \ \dots \ x_1^T \ x_2^T \ \dots x_i^T]^T$$

The size of this vector is typically large which is the subject of the optimization.

A. Observation Model

A transceiver at point T travelling in a straight line at velocity v transmits an acoustic pulse at time t . The pulse is detected by a stationary transponder at point R at $\frac{|T-R|}{c}$ seconds later, where c is the local in-water speed of sound. The transponder waits for δ_r seconds before transmitting a reply pulse. The transceiver receives the reply pulse at time $t + \Delta_t$ at position $T' = T + v\Delta_t$ where Δ_t is constrained such that:

$$c\Delta_t = |T - R| + |v\Delta_t + T - R| \quad (1)$$

Δ_t is the acoustic TOF. Accurate LBL algorithms use the implicit observation equation 1 to produce high-quality estimates of transceiver position when transponder locations are known. We make the assumption that $|v|$ is small, in which case $\Delta_t = \frac{2}{c} |T - R|$. This assumption is not a prerequisite for what is about to be described and is used here for purposes of simplification. We can write Equation 1 in terms of $\mathbf{X}$ and an observation function h : a TOF observation $z_i(k)$ from pose k to transponder i is a function of vehicle position(s) and transponder locations:

$$z_i(k) = h_{i,k}(\mathbf{X}) \quad (2)$$

The concatenation of all such individual TOF observations is written as Z_t :

$$Z_t = h(\mathbf{X}) \quad (3)$$

The covariance R of Z is a diagonal matrix of uncertainties in each individual observation $R = \text{diag}(\sigma_1^2, \sigma_2^2, \dots)$.

B. Trajectory Model

The vectors $x_A, x_B, x_C, \dots$ are sequential samples of the vehicle trajectory. We expect adjacent poses to be related – for example the position of x_B is likely to be close to x_C . This relationship can be written in terms of $\mathbf{X}$ and a trajectory function f

$$f(\mathbf{X}) = \mathbf{0} \quad (4)$$

In this paper f is used to constrain the vehicle to a path of constant velocity. The acceleration of the vehicle at pose $x_{(\cdot)}$ is modelled as a normally distributed random variable with zero mean and variance Q ie $\ddot{x}_{(\cdot)} \sim N(0, Q)$. Discretizing this constraint leads to the trivial linear relationship between three consecutive poses.

$$x_{(\cdot)-1} - 2x_{(\cdot)} + x_{(\cdot)+1} = 0 \quad (5)$$

Not only does f act as a continuity constraint along the trajectory, it also mitigates lack of observability for each pose. It is unlikely that in a real-world situation each pose will have enough range observations associated with it to allow a solution to be found solely from range-only data.

C. Solving

The observations and various constraint equations are stacked to form a single large observation bundle $Z = h(\mathbf{X})$

$$\begin{bmatrix} Z_t \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} h(\mathbf{X}) \\ f(\mathbf{X}) \end{bmatrix} \quad (6)$$

$$Z = h(\mathbf{X}) \quad (7)$$

We now need to solve equation 6 to find an estimate $\hat{\mathbf{X}}$. A standard non-linear least-squares approach [13] is taken, which repeatedly solves the linearized system of Gauss-Newton equations:

$$\mathbf{H}\delta\mathbf{x} = -\mathbf{J}^T\mathbf{W}\delta\mathbf{z} \quad (8)$$

where $\mathbf{H}$ and $\mathbf{J}$ are the Hessian and Jacobian of h evaluated at the current $\hat{\mathbf{X}}$, W is the observation weight matrix

$$W = \begin{bmatrix} R^{-1} & \mathbf{0} \\ \mathbf{0} & Q^{-1} \end{bmatrix}. \quad (9)$$

The vector $\delta\mathbf{z}$ is the error in predicted observations or error residual $\delta\mathbf{z} = Z - h(\hat{\mathbf{X}})$ and $\delta\mathbf{x}$ is an increment in state space used to calculate a new estimate $\hat{\mathbf{X}}^+ = \hat{\mathbf{X}} + \delta\mathbf{x}$. As each observation depends on only a small subset of $\mathbf{X}$, the Jacobian and Hessian of the problem are extremely sparse. Therefore, finding a solution to Equation 8 is greatly helped by employing sparse techniques.

The acoustic TOF observations from an AUV are noise-prone. Figure 3 shows data from an extreme case in which the combination of a loud acoustic payload, shallow water operation, and noisy actuators caused frequent false detections. The data has an abundance of outliers which cannot reasonably be modelled as a Gaussian process. In particular non-linear least squares techniques are sensitive to outliers – erroneous large components in $\delta\mathbf{z}$ cause large shifts in $\delta\mathbf{x}$. We model the acoustic residuals $\delta\mathbf{z}_t$ as a Cauchy distribution with half maximum width σ – the typical measurement error in ideal conditions.

Accordingly before use in equation 8 each raw acoustic residual ϵ in $\delta\mathbf{z}_t$ is mapped through the Cauchy negative log-likelihood function to produce a



Fig. 2. Caribou the MIT Department of Ocean Engineering’s AUV being launched in the GOATS 2002 Experiment near Italy. The “prong” structure at the front of the vehicle is a Synthetic Aperture Sonar (SAS).

modified residual ϵ' and $\mathbf{J}$ is also modified by an appropriate application of the chain rule.

$$\epsilon' = \text{sgn}(\epsilon) \cdot \ln \left[1 + \frac{\epsilon^2}{\sigma^2} \right] \quad (10)$$

III. RESULTS

This section presents range-only SLAM results using data from the 2002 “GOATS” experiment near the coast of Italy [14]. A network of four small LBL acoustic transponders [15] were deployed in the work area and their surface drop positions surveyed with a high precision differential GPS system. In the shallow operating area the actual seabed location of the transponders is likely to be within a few meters of the drop coordinates. The Odyssey III class vehicle “Caribou” [16] shown in figure 2 was fitted with an acoustic “AVTRAK” transceiver, which transmits to the main vehicle computer the time elapsing between the transmission of an interrogation pulse and detection of a reply pulse from each of the four beacons. The large “prong” structure protruding from the front of the AUV is the receiver array of MIT synthetic aperture sonar (SAS). The SAS is a high power transmit/receive acoustic device and was the primary payload of the GOATS 2002 experiment.

The SAS is an extraordinary payload. It offers remarkable detection capabilities but at the same time substantially changes the vehicle’s hydro-dynamics and controllability. In addition it is a very loud and

rapid (around 3Hz) acoustic source transmitting well into the acoustic ranging frequency band. Previous work using the SAS had scheduled the pinging of both SAS and LBL sensors [14]. This approach although successful needed intricate systems design and integration. The approach taken here is to let the two acoustic devices “run free”. The transmission and reply frequencies of the LBL are known and can be removed from SAS data by appropriate notch filters. The effect of the SAS transmissions on the LBL data is more troublesome. The wide-band chirp of the SAS results in false detections by the “AVTRAK”. The top plot of figure 3 shows the raw TOFs for each of the four transponders. The data must be filtered before insertion into a optimization routine.

A. Data Preparation

We know that the AUV is moving both smoothly and moderately and that the transponders are stationary. It follows that we can immediately cast doubt upon sequences of TOF observations that fluctuate dramatically. Conversely, a sequence of consecutive TOFs that change slowly and have a small sample standard deviation are more likely to correspond to valid measurements. With this in mind, the following filtering strategy is applied to the set of TOFs between the vehicle and each transponder $[t_1, t_2, \dots, t_n]$. A sliding window of width w selects a sequence of measurements $s_i = \{t_{i-w/2}, \dots, t_i, \dots, t_{i+w/2}\}$ around the measurement t_i , which has expected error standard deviation r_i . The median of this set is denoted $\mu_{1/2}^i$ and its sample standard deviation σ^i . A first pass rejects t_i if σ^i is much greater than r_i – for example $> 5r_i$. If however the spread of the sequence is small enough, a scalar η^i is calculated from the median as

$$\eta^i = \frac{1}{r^i} | \mu_{1/2}^i - t^i |$$

which is a robust measure of consensus between t_i and the other members of s_i normalized by the expected observation standard deviation. If η^i is less than a threshold value n (corresponding to a n sigma bound) then t_i is accepted. Applying this scheme to each measurement in turn results in a sparse, conservative collection of accepted measurements grouped in mutually supporting blocks. This collection is referred to as a “data-sketch” and is shown in the middle plot of Figure 3. It is probable that many of the measurements were rejected during the formation of the sketch because adjacent observations offered little or no support. A second pass through the data seeks to recover these lost measurements. The data sketch is filled in




Fig. 3. The effect of filtering and recovering reliable acoustic TOFs. The topmost plot shows the raw acoustic data. The middle plot shows the conservative “data-sketch” which is used to provide support for the recovery of further measurements from the noisy raw data. The lower-most plot shows the result of the process - a set of trusted measurements taken from the raw data which is used in the optimization process.

by linear interpolation resulting in a rational expected data trajectory for each observation time. Then, for each measurement t_i a predicted measurement is calculated using the time-stamp of the measurement and the interpolated data trajectory. The deviation of the real measurement from this expected one is evaluated. If this is within n multiples of the expected observation noise r_i the measurement is reinstated as a valid measurement. The lower plot in Figure 3 shows the resulting cleaned data set after application of this entire process.

Figure 4 shows the result of the non-linear least squares adjustment. During the experiment the AUV navigated using a ten-state Extended Kalman Filter (EKF) with *a priori* knowledge of the transponder locations. The EKF used compass, doppler velocity and LBL data. The complete trajectory of the AUV as estimated by this filter is shown as a solid line in the figure. The last kilometer’s worth of range data was used in the range-only SLAM problem. The AUV’s path was initialized with a 50m long random walk consisting of 150 individual poses. Each transponder was

initialized at the median measured distance from the first vehicle pose at 90-degree intervals. The resulting vehicle poses are drawn as black dots on top of the EKF trajectory in figure 4. The estimated transponder locations can be seen as black dots near the labels T1 to T4. They are in excellent agreement with the DGPS-surveyed locations lying at the center of the square markers. The table in Figure III-A compares the estimated and surveyed baselines. An AUV using this technique would at this point switch navigation to standard LBL navigation using the newly deduced transponder coordinates until a new previously unseen transponder is detected.

The optimization process using only relative measurement between the vehicle and transponders has a gauge deficiency of three – xy-translation and rotation of the whole system. To compare the results to the “ground-truth” trajectory and surveyed transponder locations the solution is translated and aligned such that T1 is coincident with the GPS derived location and the base line between T1 and T2 are co-linear.

The range-only derived trajectory falls short of the




Fig. 4. Result of the range-only optimization. The GPS derived positions of the transponders are shown as squares and the estimated locations as nearby black dots. The estimated vehicle trajectory is plotted on top of that derived from the onboard navigation. The AUV surfaced at the end of the dive and acquired a GPS fix (shown as a star) which is coincident with both the onboard navigation and range only estimated tracks.

Base-Line	Surveyed (m) ± 2	Estimated (m)	Error (m)	Percentage
L[1,2]	278.9	275.2	3.7	1.3 %
L[1,3]	491.2	482.5	8.7	1.8 %
L[1,4]	417.2	412.0	5.3	1.3 %
L[2,3]	438.6	428.6	10.0	2.3 %
L[2,4]	543.9	535.6	8.2	1.5 %
L[3,4]	307.0	302.9	4.1	1.3 %

Fig. 5. A table comparing the surveyed and estimated lengths of the 6 transponder base-lines. The surface drop coordinates of the transponders was recorded with DGPS. The average water column depth was 11m and the transponders were assumed to be lying directly below the drop locations. Note how in general (with the exception of L[2,3]) the longer the base-line the greater the percentage error. This is due to the increased significance of propagation effects on the range measurements over greater distances. The exact cause for larger error in L[2,3] is unclear, however it is the baseline around which fewer measurements were made and the trajectory estimation was poorest.

EKF trajectory at the bottom of the first leg. This can be attributed to two factors. Firstly, the vehicle is operating near to the T2-T3 baseline where T2 and T3, the nearest and most easily detected transponders, provide no cross-baseline information. Secondly the remaining beacons, especially T4, are in the baffles of the vehicle reducing the detection rate of ranges. In this region there are few useful acoustic observations. The interpolating effect of the vehicle trajectory constraint dominates in this region and produces a shortened path filling in between regions possessing denser range measurements.

IV. FUTURE WORK

This work on the sub-sea range-only SLAM problem is ongoing. Several issues need to be addressed before reliable deployment can be achieved. As is so often the case with optimization problems, the choice of the initial conditions is important. Occasionally a combination of randomly guessed transponder and vehicle poses will prevent convergence. A strong baseline hysteresis exists whereby if during the optimization the estimate of a transponder moves erroneously across a base-line between two other transponders the correct solution is rarely obtained. At the time of writing, work is under way to create a more sensible state initialization using a combination of a crude guess of vehicle trajectory using thrust and fin actuation commands and the gradient of range measurement. This process will be greatly helped by allowing the AUV to execute course changes (a 90-degree turn for example) to resolve baseline ambiguities. A series of experiments is planned to test this approach.

A related approach to this problem currently under evaluation is to use odometry (DVL) data and the Delayed Decision Making (DDM) architecture suggested in [6] to initialize transponder locations after only a small change in vehicle position. A possible benefit of this approach is being able to instantiate and remove multiple transponder location hypotheses. By doing so any cross-base-line ambiguity would be naturally resolved as the vehicle changes heading.

REFERENCES

- [1] M. B. Larsen, "High performance doppler-inertial navigation — experimental results," in *IEEE Oceans*, 2000.
- [2] L. L. Whitcomb, D. R. Yoerger, and H. Singh, "Combined Doppler/LBL based navigation of underwater vehicles," in *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, 1999.
- [3] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," in *Autonomous Robot Vehicles*, I. Cox and G. Wilfong, Eds., pp. 167–193. Springer-Verlag, 1990.
- [4] O. Faugeras, *Three-Dimensional Computer Vision: A Geometric Viewpoint*, MIT Press, 1993.
- [5] Philip F. McLauchlan, "A batch/recursive algorithm for 3d scene reconstruction," in *Int. Conf. Computer Vision and Pattern Recognition*, Hilton Head, SC, USA, 2000, vol. 2, pp. 738–743.
- [6] R. Rikoski, J. Leonard, and P. Newman, "Stochastic mapping frameworks," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2002, pp. 426–433.
- [7] J. J. Leonard and R. Rikoski, "Incorporation of delayed decision making into stochastic mapping," in *Experimental Robotics VII*, D. Rus and S. Singh, Eds., Lecture Notes in Control and Information Sciences. Springer-Verlag, 2001.
- [8] George A Kantor and Sanjiv Singh, "Preliminary results in range-only localization and mapping," in *Proceedings of the IEEE Conference on Robotics and Automation (ICRA '02)*, May 2002.
- [9] Matthew Deans and Martial Hebert, "Experimental comparison of techniques for localization and mapping using a bearings only sensor," in *Proc. of the ISER '00 Seventh International Symposium on Experimental Robotics*, December 2000.
- [10] Matthew Deans and Martial Hebert, "Invariant filtering for simultaneous localization and mapping," in *IEEE International Conference on Robotics and Automation*, April 2000, pp. 1042–7.
- [11] J. Vaganay, P. Baccou, and B. Jouvencel, "Homing by acoustic ranging to a single beacon," in *IEEE Oceans*, 2000, pp. 1457–1462.
- [12] M. B. Larsen, "Synthetic long baseline navigation of underwater vehicles," in *IEEE Oceans*, 2000.
- [13] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge University Press, ISBN: 0521623049, 2001.
- [14] J. Edwards, H. Schmidt, and K. LePage, "Bistatic synthetic aperture target detection and imaging with an auv," *IEEE J. Ocean Engineering*, vol. 26, no. 4, pp. 690–699, 2001.
- [15] "http://www.sonardyne.co.uk," .
- [16] "http://www.bluefinrobotics.com," .