

Towards terrain-aided navigation for underwater robotics

STEFAN WILLIAMS*, GAMINI DISSANAYAKE
and HUGH DURRANT-WHYTE

*Australian Centre for Field Robotics, Department of Mechanical and Mechatronics Engineering,
University of Sydney, NSW 2006, Australia*

Received 27 July 2000; accepted 19 November 2000

Abstract—This paper describes an approach to autonomous navigation for an undersea vehicle that uses information from a scanning sonar to generate navigation estimates based on a simultaneous localization and mapping algorithm. Development of low-speed platform models for vehicle control and the theoretical and practical details of mapping and position estimation using sonar are provided. An implementation of these techniques on a small submersible vehicle ‘Oberon’ are presented.

Keywords: Terrain-aided navigation; localization; mapping; uncertainty; autonomous underwater vehicle.

1. INTRODUCTION

Current work on undersea vehicles at the Australian Centre for Field Robotics concentrates on the development of terrain-aided navigation techniques, sensor fusion and vehicle control architectures for real-time platform control. Position and attitude estimation algorithms that use information from scanning sonar to complement a vehicle dynamic model and unobservable environmental disturbances are invaluable in the subsea environment. Key elements of the current research work include the development of sonar feature models, the tracking and use of these models in mapping and position estimation, and the development of low-speed platform models for vehicle control.

While many land-based robots use GPS or maps of the environment to provide accurate position updates for navigation, a robot operating underwater does not typically have access to this type of information. In underwater scientific missions, *a priori* maps are seldom available and other methods for localisation must be considered. Many underwater robotic systems rely on fixed acoustic transponders that are surveyed into the robot's work area [1]. These transponders are then

*To whom correspondence should be addressed. E-mail: stefanw@acfr.usyd.edu.au

interrogated to triangulate the position of the vehicle. The surveying of these transponders can be a costly and time consuming affair — especially at the depths at which these vehicles often operate and their performance can vary with conditions within the water column in which the vehicle is operating.

As an alternative to beacon-based navigation, a vehicle can use its on-board sensors to extract terrain information from the environment in which it is operating. One of the key technologies being developed in the context of this work is an algorithm for Simultaneous Localization and Map Building (SLAM) to estimate the position of an underwater vehicle. SLAM is the process of concurrently building up a feature based map of the environment and using this map to obtain estimates of the location of the vehicle [2–6]. In essence, the vehicle relies on its ability to extract useful navigation information from the data returned by its sensors. The robot typically starts at an unknown location with no *a priori* knowledge of landmark locations. From relative observations of landmarks, it simultaneously computes an estimate of vehicle location and an estimate of landmark locations. While continuing in motion, the robot builds a complete map of landmarks and uses these to provide continuous estimates of the vehicle location. The potential for this type of navigation system for subsea robots is enormous considering the difficulties involved in localization in underwater environments.

This paper presents the results of the application of SLAM to estimate the motion of an underwater vehicle. This work represents the first instance of a deployable underwater implementation of the SLAM algorithm. Section 2 introduces the Oberon submersible vehicle developed at the Centre and briefly describes the sensors and actuators used. Section 3 summarizes the stochastic mapping algorithm used for SLAM, while Section 4 presents the feature extraction and data association techniques used to generate the observations for the SLAM algorithm. In Section 5, a series of trials are described and the results of applying SLAM during field trials in a natural terrain environment along Sydney's coast are presented. Finally, Section 6 concludes the paper by summarizing the results and discussing future research topics as well as on-going work.

2. THE OBERON VEHICLE

The experimental platform used for the work reported in this paper is a mid-size submersible robotic vehicle called Oberon designed and built at the Australian Centre for Field Robotics (see Fig. 1). The vehicle is equipped with two scanning low-frequency terrain-aiding sonars and a color CCD camera, together with bathymetric depth sensors, a fiber optic gyroscope and a magneto-inductive compass with integrated two-axis tilt sensor [7]. This vehicle is intended primarily as a research platform upon which to test novel sensing strategies and control methods. Autonomous navigation using the information provided by the vehicle's on-board sensors represents one of the ultimate goals of the project [8].

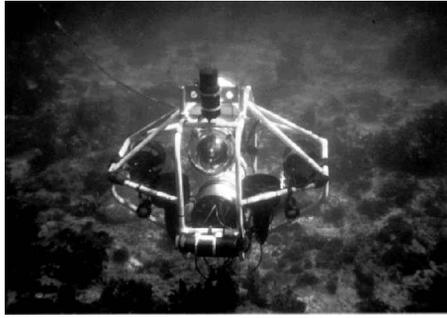


Figure 1. Oberon at sea.

3. FEATURE-BASED POSITION ESTIMATION

This section presents a feature based localisation and mapping technique used for generating vehicle position estimates. By tracking the relative position between the vehicle and identifiable features in the environment, both the position of the vehicle and the position of the features can be estimated simultaneously. The correlation information between the estimates of the vehicle and feature locations is maintained to ensure that consistent estimates of these states are generated.

3.1. The estimation process

The localization and map building process consists of a recursive, three-stage procedure comprising prediction, observation and update steps using an extended Kalman filter (EKF) [3]. The Kalman filter is a recursive, least-squares estimator and produces at time k a minimum mean-squared error estimate $\hat{\mathbf{x}}(k|k)$ of the state $\mathbf{x}(k)$ given a series of observations, $\mathbf{Z}^k = [z(1) \dots z(k)]$:

$$\hat{\mathbf{x}}(k|k) = E[\mathbf{x}|\mathbf{Z}^k]. \quad (1)$$

The filter fuses a predicted state estimate $\hat{\mathbf{x}}(k|k-1)$ with an observation $\mathbf{z}(k)$ of the state $\mathbf{x}(k)$ to produce the updated estimate $\hat{\mathbf{x}}(k|k)$. For the SLAM algorithm the EKF is used to estimate the pose of the vehicle $\mathbf{x}_v(k)$ along with the positions of the N observed features $\mathbf{x}_i(k)$, $i = 1 \dots N$. In the current implementation, the vehicle pose is made up of the two-dimensional position (x_v, y_v) and orientation ψ_v of the vehicle. An estimate of vehicle ground speed, V_v , slip angle, γ_v , and the gyro rate bias, $\dot{\psi}_{\text{bias}}$, is also generated by the algorithm. The ‘slip angle’ γ_v is the angle between the vehicle axis and the direction of the velocity vector. Although the thrusters that drive the vehicle are oriented in the direction of the vehicle axis, the slip angle is often non-zero due to disturbances caused by the deployed tether and currents. A schematic diagram of the vehicle model is shown in Fig. 2.

3.1.1. Prediction. The prediction stage uses a model of the motion of the vehicle to predict the vehicle position, $\hat{\mathbf{x}}_v(k|k-1)$, at instant k given the information

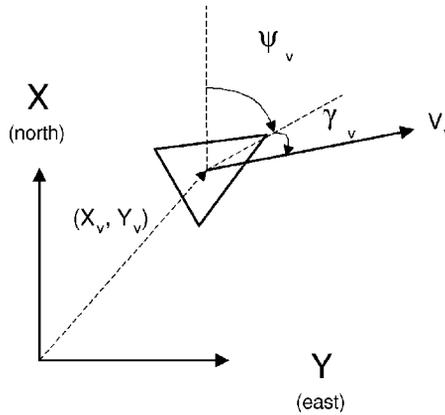


Figure 2. The vehicle model currently employed with the submersible vehicle. The positioning filter estimates the vehicle position, (X_v, Y_v) , orientation, ψ_v , velocity, V_v and slip angle γ_v . The frame of reference used is based on the north-east-down alignment commonly used in aeronautical engineering applications. The x-axis is aligned with the compass generated north reading.

available to instant $k - 1$. A constant acceleration model, shown in (2), is used for this purpose:

$$\begin{aligned}
 \dot{x}_v &= V_v \cos(\psi_v + \gamma_v) + v_x, \\
 \dot{y}_v &= V_v \sin(\psi_v + \gamma_v) + v_y, \\
 \dot{\psi}_v &= \dot{\psi}_{\text{gyro}} - \dot{\psi}_{\text{bias}} + v_\psi, \\
 \dot{V}_v &= v_v, \\
 \dot{\gamma}_v &= v_\gamma, \\
 \dot{\psi}_{\text{bias}} &= v_{\text{bias}},
 \end{aligned} \tag{2}$$

where v_x , v_y , v_ψ , v_v , v_γ and v_{bias} are assumed to be zero-mean, temporally uncorrelated gaussian process noise errors with variance σ_x^2 , σ_y^2 , σ_ψ^2 , σ_v^2 , σ_γ^2 and σ_{bias}^2 respectively. The standard deviations for these noise parameters are shown in Table 1.

The rate of change of vehicle ground speed, \dot{V}_v , and slip angle, $\dot{\gamma}_v$, are assumed to be driven by white noise. The fiber-optic gyroscope measures the vehicle yaw rate and is used as a control input to drive the orientation estimate. Given the small submerged inertia, relatively slow motion and large drag-coefficients induced by the open frame structure of the vehicle and the deployed tether, the model described by (2) is able to capture the motion of the vehicle.

In order to implement the filter, the discrete form of the vehicle model is used to predict the vehicle state $\hat{\mathbf{x}}_v(k|k - 1)$ given the previous estimate $\hat{\mathbf{x}}_v(k - 1|k - 1)$. The discrete, non-linear vehicle prediction equation, \mathbf{F}_v , is shown in (4):

$$\hat{\mathbf{x}}_v(k|k - 1) = \mathbf{F}_v(\hat{\mathbf{x}}_v(k - 1|k - 1), u(k)), \tag{3}$$

Table 1.
SLAM filter parameters

Sampling period	ΔT	0.1 s
Vehicle x process noise SD	σ_x	0.025 m
Vehicle y process noise SD	σ_y	0.025 m
Vehicle heading process noise SD	σ_ψ	0.6°
Vehicle velocity SD	σ_v	0.01 m/s
Vehicle slip angle SD	σ_γ	1.4°
Gyro bias SD	σ_{bias}	$0.3^\circ/\text{s}$
Gyro measurement SD	σ_{gyro}	$0.6^\circ/\text{s}$
Compass SD	σ_{compass}	2.9°
Range measurement SD	σ_R	0.1 m
Bearing measurement SD	σ_B	1.4°
Sonar range		20 m
Sonar resolution		0.1 m

where \mathbf{F}_v is defined by:

$$\begin{aligned}
 \hat{x}_v &= \hat{x}_v + \Delta T \hat{V}_v \cos(\hat{\psi}_v + \hat{\gamma}_v), \\
 \hat{y}_v &= \hat{y}_v + \Delta T \hat{V}_v \sin(\hat{\psi}_v + \hat{\gamma}_v), \\
 \hat{\psi}_v &= \hat{\psi}_v + \Delta T (\dot{\psi}_{\text{gyro}} - \dot{\psi}_{\text{bias}}), \\
 \hat{V}_v &= \hat{V}_v, \\
 \hat{\gamma}_v &= \hat{\gamma}_v, \\
 \hat{\psi}_{\text{bias}} &= \hat{\psi}_{\text{bias}},
 \end{aligned} \tag{4}$$

with the discrete timestamps $(k|k-1)$ and $(k-1|k-1)$ omitted for conciseness.

The features that are tracked in the map are assumed to be stationary over time. This is not a necessary condition given the formulation of the SLAM algorithm using an EKF. However, tracking moving features in the environment is not considered feasible at this time given the available sensors nor is it likely to aid in navigation since accurate models of bodies moving underwater are not likely to be available. This assumption yields a simple feature map prediction equation, \mathbf{F}_m :

$$\hat{\mathbf{x}}_i(k|k-1) = \hat{\mathbf{x}}_i(k-1|k-1). \tag{5}$$

The covariance matrix of the vehicle and feature states, $\hat{\mathbf{P}}(k|k-1)$, is predicted using the non-linear state prediction equation. The predicted covariance is computed using the gradient of the state propagation equation linearized about the current vehicle state estimate, $\nabla \mathbf{F}_v$, and about the control input model, $\nabla \mathbf{F}_u$, the process noise model, \mathbf{Q} , and the control noise model, \mathbf{U} . The filter parameters used in this application are shown in Table 1.

$$\hat{\mathbf{P}}(k|k-1) = \nabla \mathbf{F}_v \hat{\mathbf{P}}(k-1|k-1) \nabla \mathbf{F}_v^T + \nabla \mathbf{F}_u \mathbf{U}(k|k-1) \nabla \mathbf{F}_u^T + \mathbf{Q}(k|k-1), \tag{6}$$

with

$$\mathbf{U}(k|k-1) = \text{diag}[\sigma_{\text{gyro}}^2], \quad (7)$$

and

$$\mathbf{Q}(k|k-1) = \text{diag}[\sigma_x^2 \quad \sigma_y^2 \quad \sigma_\psi^2 \quad \sigma_V^2 \quad \sigma_\gamma^2 \quad \sigma_{\text{bias}}^2]. \quad (8)$$

3.1.2. Observation. There are two types of observations involved in the map building process as implemented on the vehicle. The first is the observation of the orientation from the output of the magneto-inductive compass. The filter generates an estimate of the current yaw of the vehicle by fusing the predicted yaw estimate with the compass output. A shaping state that estimates the yaw rate bias of the gyroscope is also generated. The yaw measurements are incorporated into the SLAM filter using the yaw observation estimate, $\hat{z}_\psi(k|k-1)$, as shown in (9):

$$\hat{z}_\psi(k|k-1) = \hat{x}_\psi(k|k-1). \quad (9)$$

The compass observations are assumed to be corrupted by zero-mean, temporally uncorrelated white noise with variance σ_{compass} .

$$z_{\text{compass}}(k) = \psi(k) + w_{\text{compass}}. \quad (10)$$

There is always a danger that a compass will be affected by ferrous objects in the environment and transient magnetic fields induced by large currents, such as those generated by the vehicle's thrusters. While this may be the case, in practice the compass does not seem to be affected to a large degree by the vehicle's thrusters. In addition, the unit is equipped with a magnetic field strength alarm. When the strength of the magnetic field increases, the alarm is signaled, indicating that the current observation may be in doubt.

Terrain feature observations are made using an imaging sonar that scans the horizontal plane around the vehicle. Point features are extracted from the sonar scans and are matched against existing features in the map. The feature matching algorithm will be described in more detail in Section 4. The observation consists of a relative distance and orientation from the vehicle to the feature. The terrain feature observations are assumed to be corrupted by zero-mean, temporally uncorrelated white noise with variance σ_R and σ_B respectively. The predicted observation, $\hat{z}_i(k|k-1)$, when observing landmark 'i' located at $\hat{\mathbf{x}}_i$ can be computed using the non-linear observation model $\mathbf{H}_i(\hat{\mathbf{x}}_v(k|k-1), \hat{\mathbf{x}}_i(k|k-1))$:

$$\hat{z}_i(k|k-1) = \mathbf{H}_i(\hat{\mathbf{x}}_v(k|k-1), \hat{\mathbf{x}}_i(k|k-1)), \quad (11)$$

where \mathbf{H}_i is defined by:

$$\hat{z}_{iR} = \sqrt{(\hat{x}_v - \hat{x}_i)^2 + (\hat{y}_v - \hat{y}_i)^2},$$

$$\hat{z}_{i\theta} = \arctan\left(\frac{(\hat{y}_v - \hat{y}_i)}{(\hat{x}_v - \hat{x}_i)}\right) - \hat{\psi}_v,$$

with the discrete timestamps $(k|k-1)$ once again omitted for conciseness.

For both types of observation, the difference between the actual observation, $\mathbf{z}(k)$, and the predicted observation, $\hat{\mathbf{z}}(k|k-1)$, is termed the innovation $v(k|k-1)$:

$$v(k|k-1) = \mathbf{z}(k) - \hat{\mathbf{z}}(k|k-1). \quad (12)$$

The innovation covariance, $\mathbf{S}(k|k-1)$, is computed using the current state covariance estimate, $\hat{\mathbf{P}}(k|k-1)$, the gradient of the observation model, $\nabla\mathbf{H}(k|k-1)$, and the covariance of the observation model $\mathbf{R}(k|k-1)$:

$$\mathbf{S}(k|k-1) = \nabla\mathbf{H}(k|k-1)\hat{\mathbf{P}}(k|k-1)\nabla\mathbf{H}(k|k-1)^T + \mathbf{R}(k|k-1), \quad (13)$$

with:

$$\mathbf{R}_\psi(k|k-1) = \text{diag}[\sigma_{\text{compass}}^2], \quad (14)$$

and:

$$\mathbf{R}_i(k|k-1) = \text{diag}[\sigma_{\mathbf{R}}^2 \quad \sigma_{\mathbf{B}}^2]. \quad (15)$$

3.1.3. Update. The state estimate can now be updated using the optimal gain matrix $\mathbf{W}(k)$. This gain matrix provides a weighted sum of the prediction and observation, and is computed using the innovation covariance, $\mathbf{S}(k|k-1)$, and the predicted state covariance, $\hat{\mathbf{P}}(k|k-1)$. This is used to compute the state update $\hat{\mathbf{x}}(k|k)$ as well as the updated state covariance $\hat{\mathbf{P}}(k|k)$:

$$\hat{\mathbf{x}}(k|k) = \hat{\mathbf{x}}(k|k-1) + \mathbf{W}(k|k-1)v(k|k-1), \quad (16)$$

$$\hat{\mathbf{P}}(k|k) = \hat{\mathbf{P}}(k|k-1) - \mathbf{W}(k|k-1)\mathbf{S}(k|k-1)\mathbf{W}(k|k-1)^T, \quad (17)$$

where:

$$\mathbf{W}(k|k-1) = \hat{\mathbf{P}}(k|k-1)\nabla\mathbf{H}(k|k-1)\mathbf{S}^{-1}(k|k-1). \quad (18)$$

4. FEATURE EXTRACTION FOR LOCALIZATION

The development of autonomous map-based navigation relies on the ability of the system to extract appropriate and reliable features with which to build maps. Point features are identified from the sonar scans returned by the imaging sonar and are used to build up a map of the environment.

The extraction of point features from the sonar data is essentially a three-stage process. The range to the principal return must first be identified in individual pings. This represents the range to the object that has produced the return. The principal returns must then be grouped into clusters. Small, distinct clusters can be identified as point features, and the range and bearing to the target estimated. Finally, the range and bearing information must be matched against existing features in the map. This

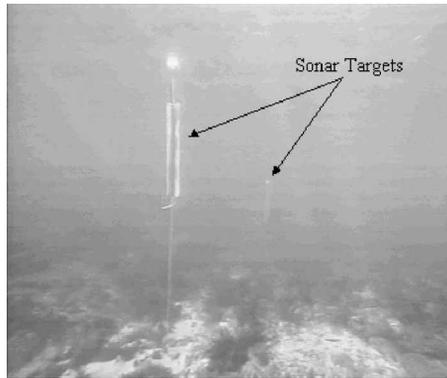


Figure 3. An image captured from the submersible of one of the sonar targets deployed at the field test site.

section provides more details of the feature identification algorithms used to provide observations for the filter.

4.1. Sonar targets

Sonar targets are currently introduced into the environment in which the vehicle will operate (see Fig. 3) in order to obtain identifiable and stable features. A prominent portion of the reef wall or a rocky outcropping might also be classified as a point feature. If the naturally occurring point features are stable they will also be incorporated into the map. Development of techniques to extract terrain aiding information from more complex natural features, such as coral reefs and the natural variations on the sea floor, is an area of active research. The ability to use natural features will allow the submersible to be deployed in a larger range of environments without the need to introduce artificial beacons.

The sonar targets produce strong sonar returns that can be characterized as point features for the purposes of mapping (see Fig. 5a). The lighter sections in the scan indicate stronger intensity returns. In this scan, two sonar targets are clearly visible and can easily be characterized as point features. (The features extracted by the algorithm are shown in Fig. 5b.) More details of the feature extraction algorithms are presented in the following subsections.

4.2. Principal returns

The data returned by the SeaKing imaging sonar consists of the complete time history of each sonar ping in a discrete set of bins scaled over the desired range. The first task in extracting reliable features is to identify the principal return from the ping data. The principal return is considered to be the start of the maximum energy component of the signal above a certain noise threshold. Figure 4 shows a single ping taken from a scan in the field. This return is a reflection from one of the sonar targets and the principal return is clearly visible. The return exhibits very

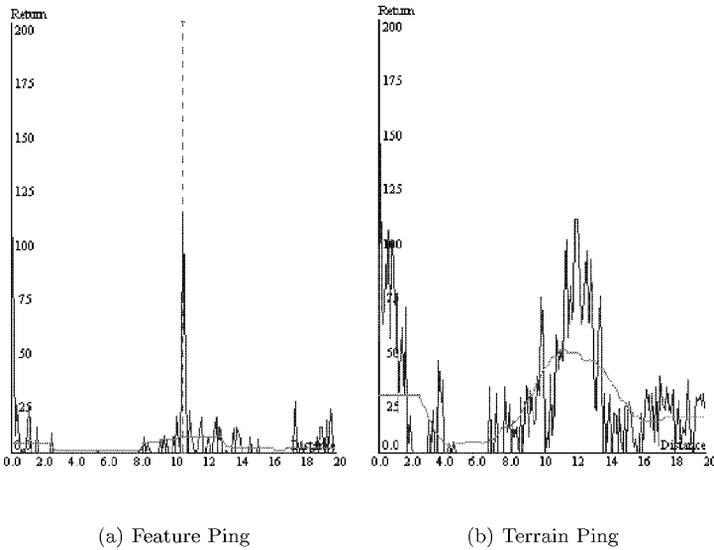


Figure 4. (a) A single SeaKing ping showing the raw ping, the moving average and the computed principal return. This ping is a reflection from one of the sonar targets and shows very good signal to noise ratio. The dashed line marks the principal return. (b) A single SeaKing ping reflected from the reef surrounding the vehicle showing the raw ping and the moving average. The terrain returns are distinguishable from the target returns by the fact that the high energy returns are spread over a much wider section of the ping. The large amplitude return at low range in this ping results from the interface between the oil-filled sonar transducer housing and the surrounding sea water. Large amplitude returns are ignored if they are below 2.0 m from the vehicle.

good signal to noise ratio making the extraction of the principal returns relatively straightforward.

At present the vehicle relies on the sonar targets as its primary source of navigation information. It is therefore paramount for the vehicle to reliably identify returns originating from the sonar targets. Examination of the returns generated by the targets shows that they typically have a large magnitude return concentrated over a very short section of the ping. This differs from returns from other objects in the environment such as rocks and the reef walls that tend to have high energy returns spread over a much wider section of the ping as seen in Fig. 4b.

4.3. Identification of point features

Following the extraction of the principal return from individual pings, these returns are then processed to find regions of constant depth within the scan that can be classified as point features. Sections of the scan are examined to find consecutive pings from which consistent principal return ranges are located. The principal returns are classified as a point feature if the width of the cluster is small enough to be characterised as a point feature and the region is spatially distinct with respect to other returns in the scan [9]. The bearing to the feature is computed using the center

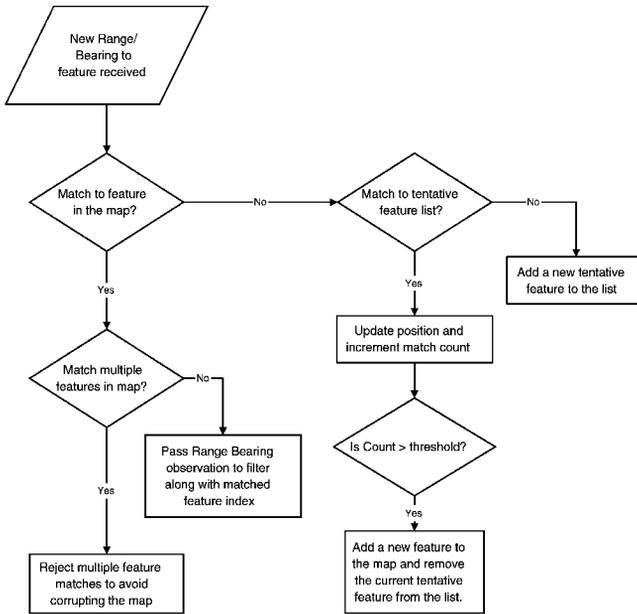


Figure 6. The feature-matching algorithm.

potential feature cannot be associated with any of the tentative features, a new tentative feature is added to the list. Tentative features that are not reobserved are removed from the list after a fixed time interval has elapsed.

5. EXPERIMENTAL RESULTS

The SLAM algorithms have been tested during deployment in a natural environment off the coast of Sydney, Australia. The submersible was deployed in a natural inlet with the sonar targets positioned in a straight line at intervals of 10 m. The vehicle controls were set to maintain a constant heading and altitude during the run. Once the vehicle had reached the end of its tether (approximately 50 m) it was turned around and returned along the line of targets. The slope of the inlet in which the vehicle was deployed meant that the depth of the vehicle varied between approximately 1 and 5 m over the course of the run.

The plot of the final map obtained by the SLAM algorithm (shown in Fig. 7) clearly shows the position of the sonar feature targets along with a number of tentative targets that are still not confirmed as sufficiently reliable. Some of the tentative targets are from the reef wall, while others come from returns off of the tether. These returns are typically not very stable and therefore do not get incorporated into the SLAM map. The absolute location of all the potential point targets identified based on the sonar principal returns are also shown in this map. These locations were computed using the estimated vehicle location at the instant of the corresponding sonar return. The returns seen near the top and bottom of the

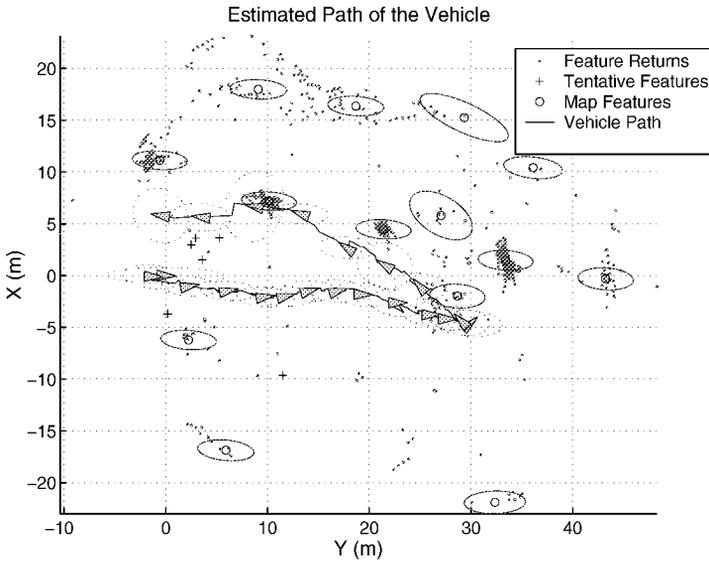


Figure 7. Path of robot shown against final map of the environment. The vehicle position estimates are spaced evenly in time over the run. It is evident that the vehicle speed changes during the run as a function of the tether deployment. The estimated position of the features are shown as circles with the covariance ellipses showing their 95% confidence bounds. Tentative targets that have not yet been added to the map are shown as '+'. The series of tentative targets towards the top of the image occur from the reef wall. These natural point features tend not to be very stable, though, and are thus not incorporated into the map.

map are from the reef walls. As can be seen, large clusters of returns have been successfully identified as targets.

Since there is currently no absolute position sensor on the vehicle, the performance of the positioning filter cannot be measured against ground truth at this time. In previous work, it was shown that the estimator yields consistent results in the controlled environment of the swimming pool at the University of Sydney [10]. To verify the performance of the filter, the innovation sequence can be monitored to check the consistency of the estimates. Figure 8 shows that the innovation sequences are within the covariance bounds computed by the algorithm.

The state estimates can also be monitored to ensure they are yielding sensible estimates. The vehicle is attached to an on-shore command station via a tether. This tether is deployed during the mission and a number of floating buoys keep it from dragging on the ground. The tether catenary creates a force directed back along its length. The effects of this force are evident in the slip angle experienced by the vehicle. When the vehicle executes a large turn, the slip angle tends to change direction. Figure 9 shows the slip angle estimates throughout the run. Shortly after the sharp turn at 360 s, the mean slip angle estimate changes sign — reflecting the fact that the tether has changed its position relative to the vehicle.

In addition to the identified target returns, strong energy returns from the reef walls and the sea floor can also be extracted from the sonar pings. In Fig. 10 these

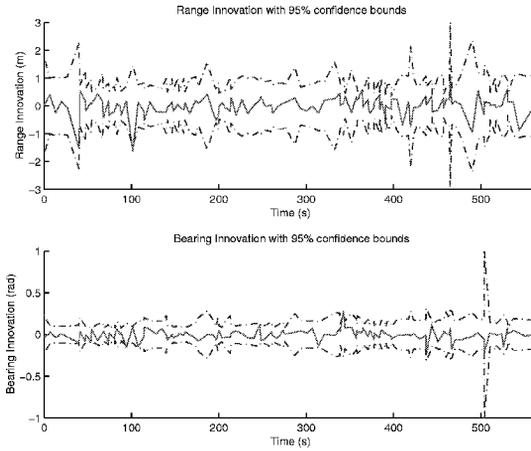
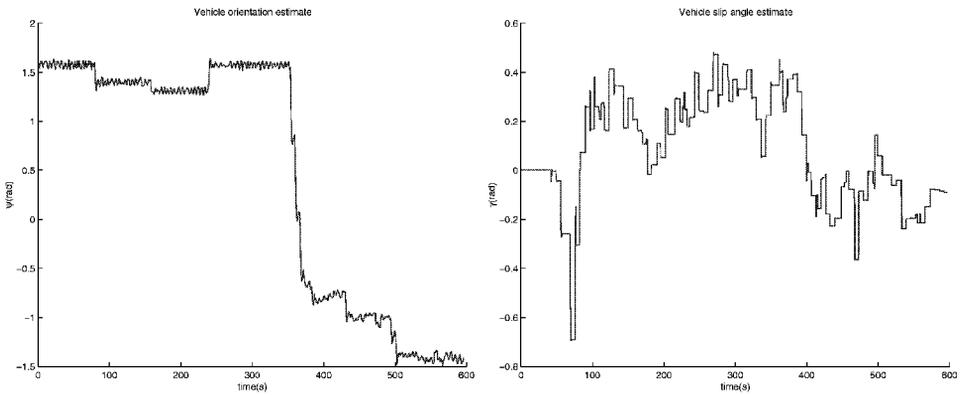


Figure 8. The range and bearing innovation sequences plotted against their 95% confidence bounds. The innovation is plotted as a solid line while the confidence bounds are the dash-dot lines.



(a) Vehicle Orientation ψ_v

(b) Slip Angle γ_v

Figure 9. (a) The vehicle orientation and (b) slip angle computed by the algorithm. The vehicle executes a 180° turn at approximately the 360 s. The mean slip estimate changes from a positive to negative value at this time reflecting the change in the force induced by the tether catenary.

strong returns have been plotted. The return points are color coded to reflect the depth at which the observation was taken. The shape of the inlet can be clearly seen and it is evident that the vehicle is observing the sea floor behind itself as it moves deeper along the inlet. Figure 11 shows a close up view of the same scene showing the vehicle position estimates more clearly.

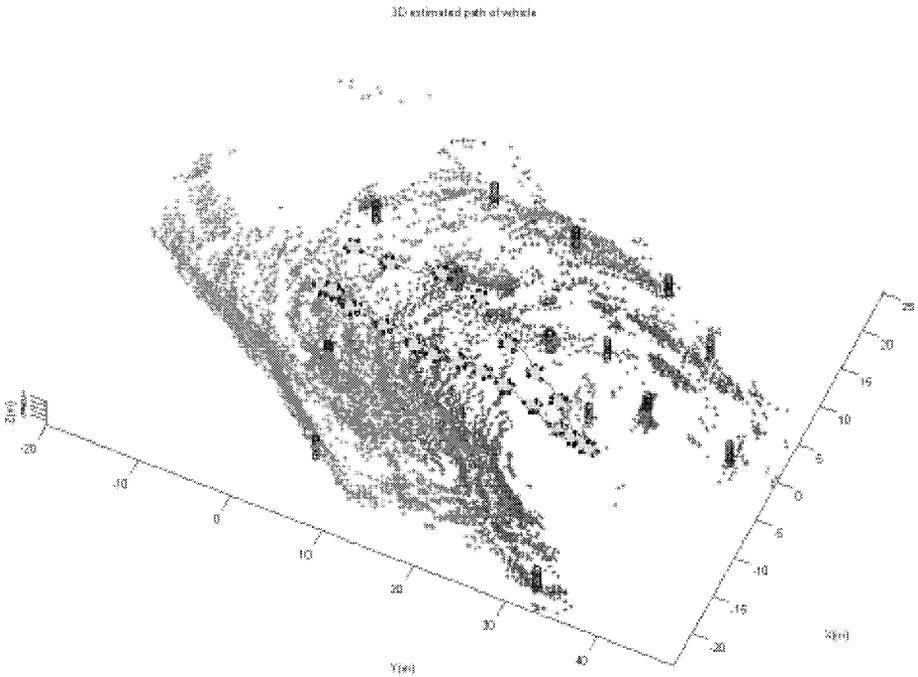


Figure 10. Path of robot shown against final map of the environment. The estimated position of the features are shown as a vertical column of circles. The strong sonar returns are color coded with the depth at which they were observed with a darker point indicating a deeper depth. The shape of the inlet is clearly visible from this plot. The estimated vehicle positions are shown spaced evenly in time.

6. SUMMARY AND CONCLUSIONS

In this paper, it has been shown that SLAM is practically feasible using artificial targets introduced into a natural terrain environment on Sydney's shore-line. By using terrain information as a navigational aid, the vehicle is able to detect unmodeled disturbances in its motion induced by the tether drag and the effect of currents.

The focus of future work is on representing natural terrain in a form suitable for incorporation into SLAM. This will enable the vehicle to be deployed in a broader range of environments without the need to introduce artificial beacons. Another outstanding issue is that of map management. As the number of calculations required to maintain the state covariance estimates increases with the square of the number of beacons in the map, criteria for eliminating features from the map as well as for partitioning the map into submaps becomes important. This is especially true for longer missions in which the number of available landmarks is potentially quite large. Finally, integration of the localisation and map building with mission planning is under consideration. This will allow decisions concerning sensing strategies to be made in light of the desired mission objectives.

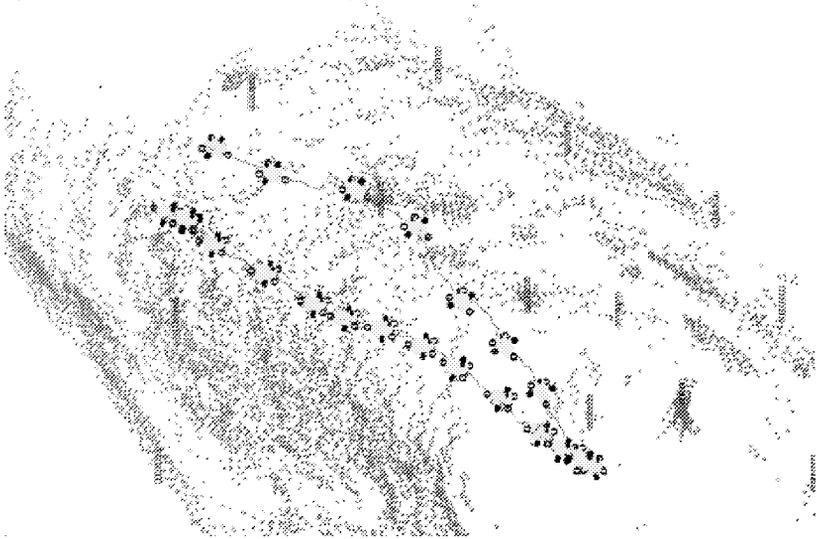


Figure 11. A close-up view of the path of robot shown against final map of the environment. The estimated position of the features are shown as a vertical column of circles. The estimated vehicle positions are shown spaced evenly in time.

REFERENCES

1. L. Whitcomb, D. Yoerger, H. Singh and J. Howland, Advances in underwater robot vehicles for deep ocean exploration: navigation, control and survey operations, in: *Proc. 9th Int. Symp. on Robotics Research*, Snowbird, UT, pp. 346–353 (1999).
2. J. A. Castellanos, J. M. M. Montiel, J. Neira and J. D. Tardos, Sensor influence in the performance of simultaneous mobile robot localization and map building, in: *Proc. 6th Int. Symp. on Experimental Robotics*, Sydney, pp. 203–212 (1999).
3. M. W. M. G. Dissanayake, P. Newman, H. F. Durrant-Whyte, S. Clark and M. Csorba, An experimental and theoretical investigation into simultaneous localisation and map building, in: *Proc. 6th Int. Symp. on Experimental Robotics*, Sydney, pp. 171–180 (1999).
4. H. J. S. Feder, J. J. Leonard and C. M. Smith, Adaptive sensing for terrain aided navigation, in: *Proc. IEEE Oceanic Engineering Society: OCEANS'98*, New York, Vol. 1, pp. 336–341 (1998).
5. J. J. Leonard and H. F. Durrant-Whyte, Simultaneous map building and localisation for an autonomous mobile robot, in: *Proc. IEEE/RSJ Int. Workshop on Intelligent Robots and Systems*, New York, pp. 1442–1447 (1991).
6. W. D. Rencken, Concurrent localisation and map building for mobile robots using ultrasonic sensors, in: *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, New York, Vol. 3, pp. 2192–2197 (1993).
7. S. Williams, P. Newman, S. Majumder, J. Rosenblatt and H. Durrant-Whyte, Autonomous transect surveying of the Great Barrier Reef, in: *Proc. Australian Conf. on Robotics and Automation*, Brisbane, pp. 16–20 (1999).
8. P. Newman and H. F. Durrant-Whyte, Toward terrain-aided navigation of a subsea vehicle, in: *Proc. FSR'97 Int. Conf. on Field and Service Robotics*, Canberra, pp. 244–248 (1997).
9. P. Newman, On the structure and solution of the simultaneous localisation and map building problem, PhD Thesis, University of Sydney (1999).

10. S. Williams, P. Newman, D. Dissanayake and H. Durrant-Whyte, Autonomous underwater simultaneous localisation and mapping, in: *Proc. IEEE Int. Conf. on Robotics and Automation*, San Francisco, CA, Vol. 2, pp. 1793–1798 (2000).
11. R. Bauer and W. D. Rencken, Sonar feature based exploration, in: *Proc. 1995 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Los Alamitos, CA, Vol. 1, pp. 148–153 (1995).
12. R. A. Brooks, A robust, layered control system for a mobile robot, *IEEE J. Robotics Automat.* **2** (1), 14–23 (1986).
13. J. J. Leonard, R. N. Carpenter and H. J. S. Feder, Stochastic mapping using forward look sonar, in: *Proc. FSR '99 Int. Conf. on Field and Service Robotics*, Pittsburgh, PA, pp. 69–74 (1999).
14. P. Maybeck, *Stochastic Models Estimation and Control*, Vol. 1. Academic Press, New York (1982).
15. S. Scheduling, G. Dissanayake, E. Nebot and H. Durrant-Whyte, Slip modelling and aided inertial navigation of an LHD, in: *Proc. IEEE Int. Conf. on Robotics and Automation*, Albuquerque, NM, pp. 1904–1909 (1997).
16. L. Whitcomb, D. Yoerger, H. Singh and Mindell, Towards precision robotic maneuvering, survey, and manipulation in unstructured undersea environments, in: *Robotics Research: Proc. 8th Int. Symp.*, pp. 346–353. Springer-Verlag, London (1998).

ABOUT THE AUTHORS



Stefan B. Williams received the BSc degree (1st class honours) in Systems Design Engineering from the University of Waterloo in Waterloo, ON, Canada in 1997. He is currently completing the requirements for a PhD in Mechatronic Engineering (Field Robotics) at the Australian Centre for Field Robotics in the Department of Mechanical and Mechatronic Engineering at the University of Sydney, Australia. His research work focuses on robotics and artificial intelligence, and concentrates on the development of terrain-aided navigation techniques, sensor fusion and vehicle control architectures for real-time platform control. He has been involved in numerous projects including an autonomous underwater vehicle, an autonomous boat, the design of a modular walking robot and the study of force-feedback haptic systems.



Gamini Dissanayake graduated in Mechanical/Production Engineering from the University of Peradeniya, Sri Lanka. He received his MSc in Machine Tool Technology and PhD in Mechanical Engineering (Robotics) from the University of Birmingham, England in 1981 and 1985, respectively. He was a lecturer at the University of Peradeniya and the National University of Singapore before joining the University of Sydney, Australia, where he is currently a Senior Lecturer attached to the Australian Centre for Field Robotics. His current research interests are in the areas of localization and map building for mobile robots, navigation systems, dynamics and control of mechanical systems, cargo handling, optimization, and path planning.



Hugh Durrant-Whyte received the BSc (Eng) degree (1st class honours) in Mechanical and Nuclear Engineering from the University of London, UK, in 1983, and the MSE and PhD degrees, both in Systems Engineering, from the University of Pennsylvania, USA, in 1985 and 1986, respectively. From 1987 to 1995, he was a Senior Lecturer in Engineering Science, University of Oxford, UK and a Fellow of Oriol College Oxford. Since July 1995 he has been Professor of Mechatronic Engineering at the Department of Mechanical and Mechatronic Engineering, the University of Sydney, Australia, where he leads the Australian Centre for Field

Robotics, a Commonwealth Key Centre of Teaching and Research. His research work focuses on automation in cargo handling, surface and underground mining, defence, unmanned flight vehicles, and autonomous sub-sea vehicles.