Large-scale path planning for Underwater Gliders in ocean currents

Dushyant Rao & Stefan B. Williams

School of Aerospace, Mechanical and Mechatronic Engineering,
University of Sydney, Australia
Email: drao7945@uni.sydney.edu.au, stefanw@acfr.usyd.edu.au

Abstract

Underwater gliders are a class of AUVs designed for high endurance over long distances, but their reduced velocity makes them more susceptible to ocean currents during deployment. Thus, feasible paths need to be generated through the ocean current field. This paper proposes a method for determining energy-optimal paths that account for the influence of ocean currents. The proposed technique is based on Rapidly-Exploring Random Trees (RRTs). Using real ocean current and bathymetry data, results produce comparable paths to grid based methods, and offer an improvement in terms of avoiding high-energy shallow regions. Future work will focus on heuristically biasing the RRT growth to further improve the generated paths, and implementation of the algorithms on a glider platform.

1 Introduction

Autonomous Underwater Vehicles (AUVs) have been used extensively in oceanography of late, to model ocean currents and tides, predict weather patterns and perform other marine research tasks. They provide a means by which to continuously measure ocean parameters such as temperature and salinity, as they can spend long periods of time underwater, and at significantly lower operating costs than manned research vessels [Blidberg, 2001].

Underwater gliders are a class of AUVs that are buoyancy-driven, reducing and expanding their volume to dive and climb through the ocean [Graver, 2005]. A pair of wings mounted on the body then converts this vertical motion into horizontal propulsion. As a product of this propulsion mechanism, their velocities are typically lower than for motor-driven vehicles (herein referred to as AUVs), reaching a maximum of around 0.5 m/s. Since drag is proportional to the square of velocity, gliders experience quadratically less drag, and can offer high endurance over long ranges. They can achieve distances on the order of thousands of kilometres per deployment, and have considerably longer deployment periods than other AUVs, lasting weeks or months depending on the platform design.

However, this velocity, while improving endurance and range, means that they can often be adversely affected by ocean currents. In a large area off the East coast of Australia used for this study, ocean currents average around 0.3 m/s and reach a maximum of 1.7 m/s [BlueLink, 2009]. Accordingly, glider trajectories need to account for these current magnitudes to ensure that the mission goals are realisable.

Path planning for AUVs is an area that has been extensively explored [Carroll *et al.*, 1992, Xu *et al.*, 2007, Zhang *et al.*, 2008]. However, gliders can be far more affected by ocean currents than an AUV. While the AUV can often overcome these currents if necessary, the glider may drift significantly from its planned trajectory. Accordingly, path planning for a glider is more critical than that for a regular AUV. While path planning for AUVs is typically aimed at minimising energy consumption, glider path planning is necessary just to ensure that the destination is accessible. Nevertheless, previously implemented AUV planning algorithms may be adaptable for a glider application.

One past approach involves partitioning the ocean field into a non-uniform square grid and implementing the A* graph search algorithm to generate a path through the field [Carroll et al., 1992]. However, ocean currents are not directly considered in the path cost and are merely used as guidelines in the planning process. This grid A* application has since been extended to incorporate ocean currents, finding the minimum energy cost path for an AUV through the field [Garau et al., 2005]. While useful for an AUV, this energy optimisation does not necessarily extend well to a glider application, and a planning technique tailored to a glider would be ideal. Constraining the vehicle to move in an 8-connected grid also means that optimality is compromised in the graph discretisation alone, and a continuous technique is desirable.

Continuous planning techniques have also been explored for both AUVs and gliders alike. Past work investigates optimal path planning in 3D for an underwater glider [Mahmoudian *et al.*, 2007], but ocean currents are not considered, collapsing the optimal solution into a kinodynamic technique using Dubins curve segments. More recent work models ocean currents as 3D B-spline functions and produces near-optimal paths in terms of time and energy efficiency [Zhang *et al.*, 2008].

Another study which considers ocean currents in path planning applies a parallel swarm parameterisation

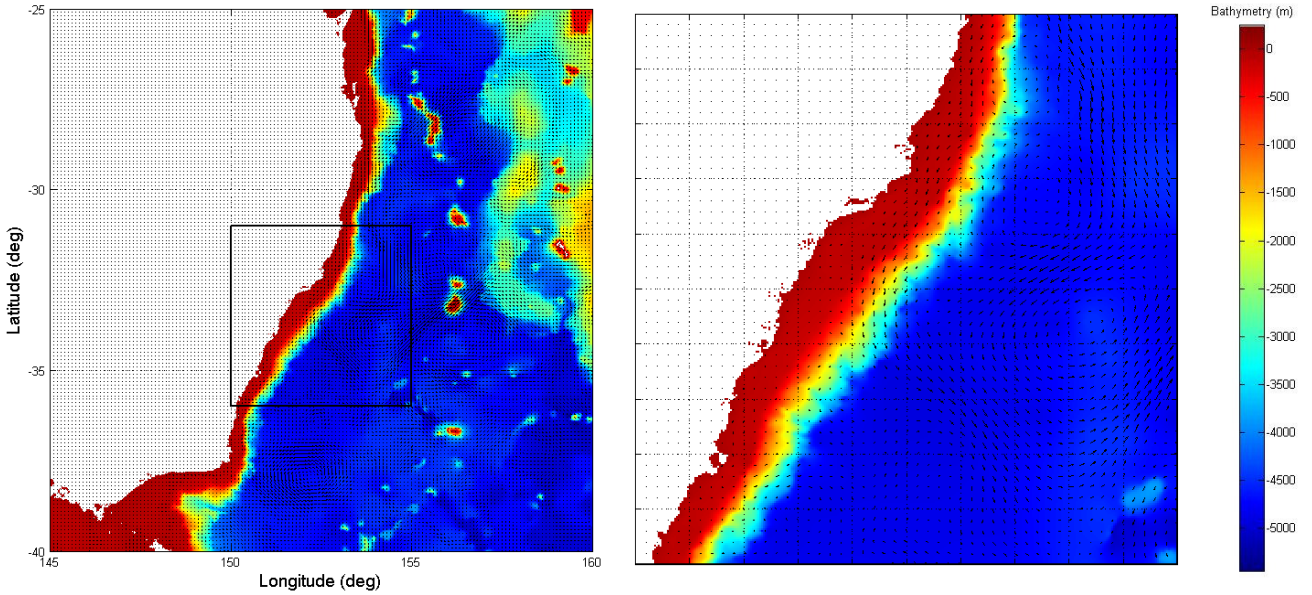

Figure 1: BlueLink Ocean current data overlaid on GEBCO bathymetry data encoded by colour, with red and blue representing shallowest and deepest regions respectively. Left: currents in Pacific Ocean off the coast of NSW, Right: inset (zoomed)

to optimise a non-linear cost surface [Witt & Dunbabin, 2008]. The cost considers energy expenditure, obstacles, shallow water and time, and their planner is robust in handling variable vehicle thrust and time-variant currents.

For this application however, some of these considerations can be neglected. The current fields change very slowly with time, and gliders vary their velocity less than AUVs – they are less equipped to use this as a degree of freedom [Graver, 2005]. Thus, it is a useful exercise to explore other options and develop a simpler solution fitted to the requirements of a glider.

Rapidly-Exploring Random Trees, or RRTs, [LaValle, 1998] have been used to plan paths for non-holonomic ground vehicles [LaValle & Kuffner, 2001], but this has yet to be applied to planning in ocean currents, or any such vector cost field. This randomised approach allows a non-uniform graph to be generated by means of the RRT, and heuristic biases can be applied to direct the RRT growth [Simmons & Urmson, 2003]. Interestingly, RRTs have been applied to AUV path planning [Tan *et al.*, 2004] but in this work, ocean currents were not considered. Instead, the main advantage of RRTs was their ability to quickly explore high-dimensional search spaces. However, this feature is of less value for an underwater glider – since longitudinal motion is largely devoted to executing the dive cycles required for propulsion, the main scope for path planning is in the 2D lateral plane.

Nonetheless, RRTs may be used for path planning in the lateral plane. To the knowledge of the authors, RRTs have yet to be applied to planning in ocean currents.

This paper looks to use RRTs to develop a path planning solution tailored to an underwater glider in the presence of ocean currents. These currents are often greater in magnitude than the glider velocity, and so the focus is on feasible paths, while striving for energy optimality.

For the path planning implementation in this paper, a number of assumptions will be made:

- Ocean current and bathymetry data is assumed to be constant over a glider mission. Since mission durations are far greater than computation time, paths can easily be regenerating during a mission to account for the dynamic ocean current field.
- Distances between degrees of latitude and longitude are assumed to be fixed, such that the ocean surface is approximated by a 2D plane. This approach is sufficient for the purposes of gauging the effectiveness of different path planning techniques.
- Dive depth is assumed to be half of the local ocean depth. This may not accurately reflect dive characteristics, but serves to demonstrate the effectiveness of the algorithm in avoiding shallow regions.

The structure of the remainder of this paper is as follows. Section 2 outlines the ocean model and ocean current data. Section 3 outlines the A* search algorithm as applied to a uniform grid and the development of time and energy cost functions. Section 4 describes the use of Rapidly-Exploring Random Trees (RRTs) to generate non-uniform heuristically biased networks to replace the uniform grid. Section 5 discusses simulation results for the system and finally, Section 6 draws some preliminary conclusions and provides a detailed list of future work.

2 Ocean model

The ocean current data used in this investigation is obtained from BlueLink Ocean Forecasting Australia

[BlueLink, 2009], a joint effort between the Bureau of Meteorology, the Royal Australian Navy and CSIRO. Due to the relatively slow temporal variation in the field, forecasting data is provided on a daily basis.

The ocean current vectors $\overrightarrow{v_c}$ are provided as North / East / Down vector components (u, v, w) at each (x, y, z) position in the ocean. Vertical ocean current components are generally more difficult to measure so datasets are more sparse and prone to error. However, these vertical components are generally negligible in comparison to lateral currents, so only the 2D lateral ocean surface current data is used here, with Northing and Easting components (u, v) at (x, y).

Global ocean data for a particular day is stored in a NetCDF data format, and is downloaded to MATLAB via a java frontend. The current simulation implementation utilises a historic dataset from December 2008, but the system can easily be extended to periodically update the ocean data available. To minimise transfer time, only the required region of the data is imported.

Gridded bathymetry data is obtained from the British Oceanographic Data Centre [BODC, 2009]. Ocean depths are given in metres below sea level z for each (x, y), and are provided in the same NetCDF format.

2.1 Ocean environment

The area of interest is the region of the Pacific Ocean off the coast of New South Wales. The currents in this region form the Southern end of the East Australian Current and the separation zone leads to a number of eddies off the coast near Sydney [Bluelink, 2009]. The regional ocean currents in the dataset used reach a maximum of around 1.68 m/s and have a mean of 0.32 m/s. With a typical glider speed of around 0.25 m/s, short of the 0.5 m/s maximum, even the average ocean currents are stronger than the glider velocity. This planning scenario is quite demanding, and the need to consider ocean currents in path planning is apparent.

3 A* Graph Search optimisation

This section examines grid-based planning with an A* graph search optimisation and their applicability to planning for ocean currents.

3.1 Uniform grid formulation

The original A* path planning implementation [Hart *et al.*, 1968] is generalised to a series of nodes $\{n_i\}$ connected by line segments $\{e_{ij}\}$, where $\{e_{ij}\}$ connects nodes $\{n_i\}$ and $\{n_j\}$. In the case of a uniform square grid of size $p \times q$, the nodes are represented as grid cells (i, j) for $0 \leq i < p$ and $0 \leq j < q$.

A path from a start cell s to a destination cell d is defined by a series of connected grid cells / nodes:

$$\Gamma = \{s, x_1, \ldots, x_k, x_{k+1}, \ldots, d\} \quad (1)$$

The nodes are assumed to be 8-connected, such that each grid cell is connected by an edge to each of its 8 immediate grid neighbours – vertical, horizontal and diagonal.

The next stage is to determine the cost and heuristic used to achieve optimal planning. The next section looks at a path cost and heuristic to determine time-optimal paths, and the following section extends this to achieve energy optimality.

3.2 Time-based cost and heuristic

If we represent our 2D underwater setting by this uniform discretisation, with grid spacing Δx, we have an ocean current vector $\overrightarrow{v_c}$ and a possible glider velocity vector $\overrightarrow{v_v}$ at each grid node. $\overrightarrow{v_c}$ is fixed based on the downloaded ocean current data, and we have $v_v = 0.25\ m/s$ but can choose the direction of the glider velocity, θ_v. The net velocity is therefore given by

$$\overrightarrow{v_{tot}} = \overrightarrow{v_c} + \overrightarrow{v_v} \quad (2)$$

where the net velocity has magnitude v_{tot} and heading θ_{tot}.

From each cell, the heading θ_{cn} to each neighbouring cell occurs in increments of 45°. In order to move towards a neighbour cell then, the net velocity must be in the cell's direction, giving $\theta_{tot} = \theta_{cn}$. Separating the total velocity into x and y components:

$$v_{tot} \cos\theta_{tot} - v_c \cos\theta_c = v_v \cos\theta_v \quad (3)$$
$$v_{tot} \sin\theta_{tot} - v_c \sin\theta_c = v_v \sin\theta_v \quad (4)$$

Summing the squares of (3) and (4) and simplifying, we have a quadratic expression for the unknown v_{tot}:

$$v_{tot}{}^2 - 2v_c \cos(\theta_c - \theta_{tot})\ v_{tot} + (v_c{}^2 - v_v{}^2) = 0 \quad (5)$$

The solutions to this equation represent the vector configurations that yield the appropriate heading vector. Any negative or complex solutions are impossible and can be discarded. If two positive real solutions remain, we can simply choose the larger v_{tot}.

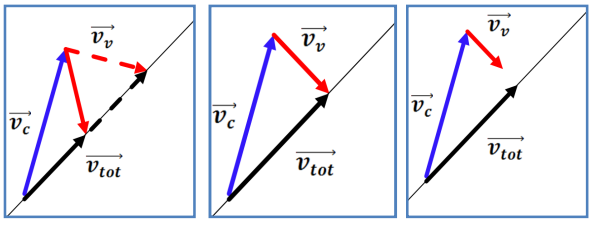

Figure 2: Possible vector configurations: 2 positive real solutions (left), 1 solution (centre), no solutions (right).

If the traversal is possible, the time cost of that line segment can be determined using the edge distance:

$$t_{cell} = \frac{\Delta x\sqrt{m^2+n^2}}{v_{tot}} \tag{6}$$

The original A* framework indicates that the node cost $f(x)$ at x_k is a sum of an accumulated path cost $g(x)$ from s to x_k and a heuristic cost estimate $h(x)$ from x_k to d. For any node x_k, the path cost can be projected from its parent node x_{k-1} as $g(x_k) = g(x_{k-1}) + t_{cell}$. The heuristic $h(x)$ needs to be *admissible* for A* to find optimal paths, and for the original implementation, the Euclidean distance between the node and goal node was used. In this implementation, the time-based heuristic utilises the minimum distance between $x_k = (i_k, j_k)$ and $d = (i_d, j_d)$ along an 8-connected grid, which can be expressed as:

$$\begin{aligned} width &= min(|i_d - i_k|, |j_d - j_k|) \\ length &= max(|i_d - i_k|, |j_d - j_k|) \\ d_{heur} &= \Delta x\big[\,(\sqrt{2}-1)width + length\big] \end{aligned} \tag{7}$$

The velocity is estimated using the global ocean current maximum, $v_{heur} = v_{c_{max}} + v_v$, and the time heuristic found by $h(x) = \frac{d_{heur}}{v_{heur}}$. Since d_{heur} and v_{heur} are both optimistic, the time heuristic is admissible. The 8-connected distance is slightly less optimistic than the Euclidean distance, but is still admissible, meaning that the algorithm expands fewer nodes but still returns optimal results.

3.3 Energy-based path cost and heuristic

Realistically, energy consumption is a more critical factor than time optimality. If we consider energy optimality instead of time optimality, the longevity of glider deployments may be considerably improved. Fortunately, low energy paths are typically close to time optimal paths, but account for regions that are demanding on energy consumption, such as coastal and reef areas.

The primary source of energy consumption (60-80%) on a glider mission is the actuation required to execute the longitudinal dive cycles through the ocean [Friedman, 2008]. Consequently, the energy consumption per dive cycle can be modelled as the sum of three major components:

$$W_{yo} = W_{b_{yo}} + W_{m_{yo}} + W_{\epsilon_{yo}} \tag{8}$$

Here, $W_{b_{yo}}$ represents the ballast piston responsible for changing the vehicle's buoyancy, $W_{m_{yo}}$ denotes the moving mass actuation required to achieve pitch control, and $W_{\epsilon_{yo}}$ is the 'hotel load' which is a sum of energy consumption from other sources.

For a fixed glide path angle ξ and dive depth z_{yo}, the number of dive cycles in a fixed time t can be expressed as [Rao, 2009]:

$$n_{yo} = \frac{t}{t_{yo}} = \frac{v_v \tan\xi}{2z_{yo}}\, t \tag{9}$$

The dive depth, as mentioned earlier, is taken to be half of the local ocean depth.

Thus, for a travel time between cells of t_{cell}, the corresponding energy cost can be expressed as:

$$W_{cell} = \left(W_{b_{yo}} + W_{m_{yo}} + W_{\epsilon_{yo}}\right)\frac{v_v \tan\xi}{2z_{yo}}\, t_{cell} \tag{10}$$

As before, the path cost for any node x_k is then projected from its parent node x_{k-1} as $g(x_k) = g(x_{k-1}) + W_{cell}$.

When converting to an energy metric, we need to ensure that the heuristic $h(x)$ is still admissible. If the heuristic time cost is denoted $h_t(x)$, the heuristic energy cost can be found using the global maximum depth to provide an optimistic estimate:

$$h_W(x) = \left(W_{b_{yo}} + W_{m_{yo}} + W_{\epsilon_{yo}}\right)\frac{v_v \tan\xi}{2z_{max}}\, h_t(x) \tag{11}$$

3.4 Limitations

As with previous grid-based implementations however, the resulting paths are only optimal over the grid discretisation, and limiting the vehicle motion to 8 headings immediately reduces the solution space. This is even more severe in the face of ocean currents - if $\overrightarrow{v_c}$ is noticeably larger than $\overrightarrow{v_v}$, a situation may arise where none of the 8 directions are achievable. This erroneously suggests that the vehicle is trapped in the cell, whereas a path can always be found outside one of the eight directions by following $\overrightarrow{v_c}$.

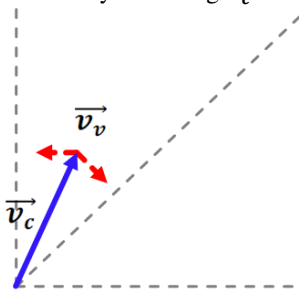

Figure 3: Erroneous 'trapped' scenario, where no vehicle heading can propel the vehicle in one of the 8 directions.

To eliminate this problem, a smoother non-uniform graph can be generated in place of the square grid, as outlined in the following section.

4 RRT Graph generation

This section details the construction of Rapidly-exploring Random Trees (RRTs) to construct the graph for search with A*. RRTs were first introduced as a randomised approach to kinodynamic planning which, unlike traditional kinodynamic techniques, also achieved

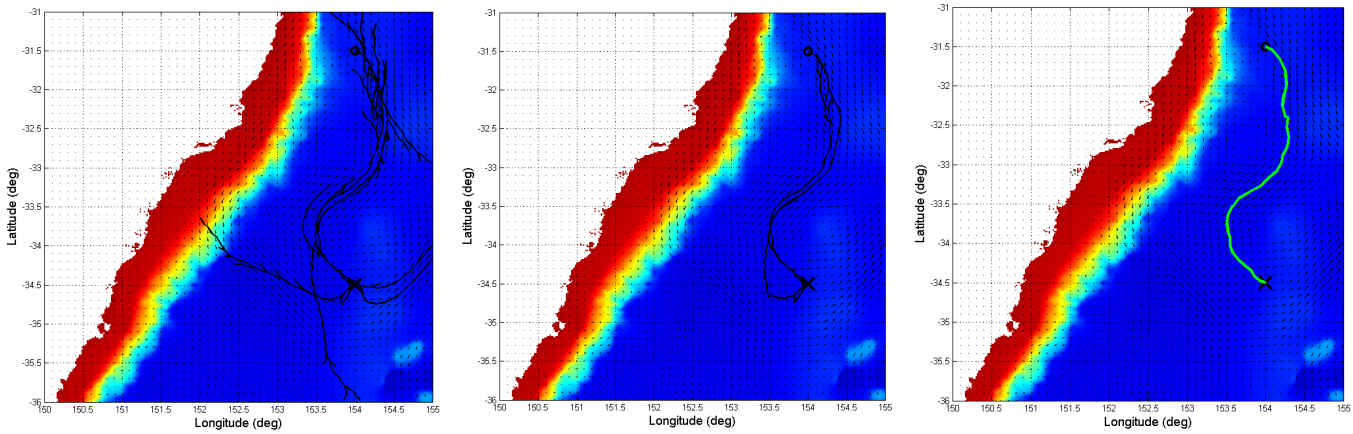

Figure 4: Bidirectional RRT generation (left); pruned RRT network (centre), optimal path from A* graph search (right). The start and destination are denoted by 'o' and 'x' respectively.

obstacle avoidance [LaValle and Kuffner, 2001]. Consider a vehicle with state space control law $\boldsymbol{x_{k+1}} = f(\boldsymbol{x_k}, \boldsymbol{u})$, starting at point (0,0) in the 2D plane. We begin by entering this first state $\boldsymbol{x_1}$ as a vertex in the tree. The tree is then expanded as follows. A random state $\boldsymbol{x_{rand}}$ is then selected based on a uniform distribution within the vehicle's configuration space. The nearest neighbour to $\boldsymbol{x_{rand}}$ within the tree is found to be $\boldsymbol{x_k}$, and we apply a control $\boldsymbol{u}$ which moves the state $\boldsymbol{x_k}$ towards $\boldsymbol{x_{rand}}$ over a fixed timestep Δt. This process is then repeated to grow the tree.

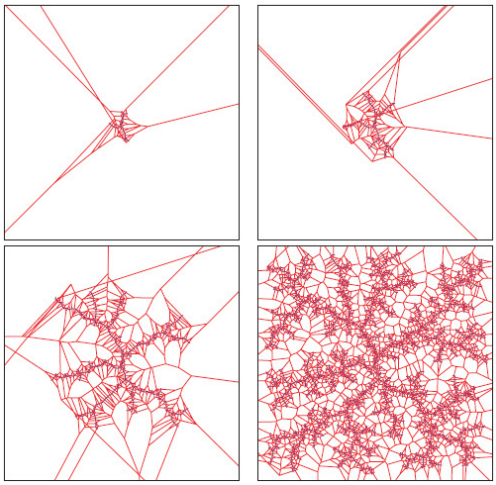

Figure 5: Voronoi diagram for RRT nodes during growth [LaValle and Kuffner, 2001]

The benefit of this method is that it is innately biased towards exploration, and looks to fill the configuration space. Observing the Voronoi diagram for the tree nodes (Figure 5), the Voronoi regions corresponding to outermost states are larger, so a random state chosen from a uniform distribution is more likely to expand an outer state. With this inherent Voronoi bias, the tree rapidly explores the workspace.

For this application, we are working with large distance scales on the order of tens and hundreds of kilometres, over which the non-holonomic constraints of the vehicle can safely be ignored. Given that planning is performed on the 2D plane, the vehicle can then be modelled as a point (x, y) that can easily move in any direction. The control applied to move state $\boldsymbol{x_k} = (x_k, y_k)$ towards $\boldsymbol{x_{rand}} = (x_{rand}, y_{rand})$ is then the fixed glider velocity $v_v = 0.25\ m/s$ over a fixed timestep Δt in the direction $\theta_v = \tan^{-1}\left(\frac{y_{rand}-y_k}{x_{rand}-x_k}\right)$ pointing towards $\boldsymbol{x_{rand}}$. The next state can then be found by summing the glider and ocean current velocities, as

$$\boldsymbol{x_{k+1}} = \boldsymbol{x_k} + (\overrightarrow{\boldsymbol{v_v}} + \overrightarrow{\boldsymbol{v_c}})\Delta t \qquad (12)$$

4.1 Path planning with RRTs

Having specified the nature of RRT growth, it is possible to create a network of possible solutions by generating RRTs bidirectionally, one from the start and one from the destination node. The latter is generated in reverse to simulate the actual glider motion, with $\boldsymbol{x_{k+1}} = \boldsymbol{x_k} - (\overrightarrow{\boldsymbol{v_v}} + \overrightarrow{\boldsymbol{v_c}})\Delta t$. In other words, the forward tree simulates the glider moving from $\boldsymbol{x_k}$ to $\boldsymbol{x_{k+1}}$, while the backward tree simulates the glider moving from $\boldsymbol{x_{k+1}}$ to $\boldsymbol{x_k}$. This also simplifies the process of integrating the two trees, as outlined in the following paragraphs.

A noticeable characteristic of the growth of the RRT is the lack of spatial uniformity. In the original case the tree explores the search space evenly, but in this application, the growth is heavily biased towards ocean currents. This is a product of the Voronoi bias – branches in the direction of the currents are propagated further, they are therefore the outermost points, and have larger Voronoi regions than branches travelling against the current.

Following growth of both trees, states common to both are connected. States $\boldsymbol{x_k}$ and $\boldsymbol{x_k'}$ in the forward and backward trees are considered common if their Euclidean distance is within some threshold ϵ:

$$\sqrt{(x_k - x_k')^2 + (y_k - y_k')^2} < \epsilon \qquad (13)$$

These are then considered as connection points between the two trees.

To simplify the graph, branches of nodes that do not lead to a connection point are 'pruned'. This is achieved by searching for each node which has no adjoining 'daughter nodes' ('branch end'). The algorithm then removes nodes from the tree in reverse until it reaches a node with multiple connections ('split branch'). By repeating this process for each 'branch end', the unconnected nodes are removed, and we are left with a network of nodes, each with a position and pointers to its connected nodes.

This non-uniform network is then searched for an optimal path using the same A* approach as before. Since the trees are generated with propagation over fixed timestep Δt, the time cost between successive nodes is constant at $t_{cell} = \Delta t$. However, the time heuristic has to be adjusted slightly because the previous heuristic is only admissible over an 8-connected grid. Instead, we now use the Euclidean distance given in Equation 12.

$$d_{heur} = \sqrt{(x_d - x_k)^2 + (y_d - y_k)^2} \quad (14)$$

The velocity estimate is the same as that for the grid case, and the time heuristic is again given by $h_t(x) = \frac{d_{heur}}{v_{heur}}$. The energy heuristic can then be found from Equation 11.

4.2 Biasing RRT growth

Improved results can be obtained by heuristically biasing the growth of the RRTs.

The probability distribution for any RRT implementation is a critical aspect [Simmons & Urmson, 2003], even for the uniform distribution of the original case. The growth of the tree needs to be extensive enough to ensure optimal solutions are not overlooked, but needs to be small enough to avoid expanding into unnecessary areas. This is largely satisfied by applying the adjusted Iterative k-Nearest RRT (IkRRT) algorithm [Simmons & Urmson, 2003] which weights the Voronoi regions of each node by some heuristic cost, and adjusts the RRT implementation to consider the k-nearest neighbours rather than a single neighbour.

The cost of each node in the tree is denoted as the energy cost heuristic, the sum of Equations 10 and 11. When selecting an RRT node to expand, the algorithm finds the k-nearest neighbours to $\boldsymbol{x_{rand}}$, sorts them by node cost, and iterates through each node starting with the lowest cost node. A node is selected probabilistically, such that low cost nodes are favoured over high cost nodes [Simmons & Urmson, 2003]. The comparative results with and without energy-based cost biasing are shown in Figure 6.

As mentioned, the path is naturally biased towards following ocean currents. To ensure the path does attempt to reach the destination, a 20% destination bias is applied. In other words, every fifth $\boldsymbol{x_{rand}}$ on average is the destination node (for the forward tree) and the start node (for the backward tree).

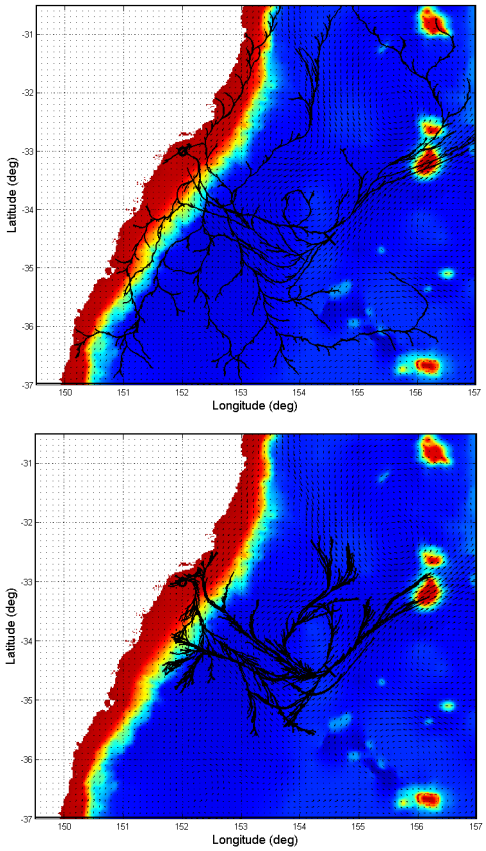

Figure 6: RRT growth without biasing (top) and IkRRT heuristic cost biasing (bottom)

Rapid changes in path heading are undesirable as they create a fluctuating path which sacrifices time and energy efficiency. Over smaller distance scales, the vehicle model rectifies this, only allowing smooth turns to change heading. However, over long ranges, other methods need to be implemented to produce sufficiently smooth paths. In the current implementation, this is achieved by limiting the change in θ_v between successive nodes to $\pm 20°$. This value was selected because it sufficiently removed rapid fluctuations but did not excessively constrain the growth of the path.

5 Results

Results have been obtained with a fixed RRT size of 10000 nodes over both trees. However, since RRTs are probabilistic, they are susceptible to variation over multiple runs. Preliminary results suggest a maximum variation of approximately 9% in energy cost over 5 consecutive runs (Figure 8). This inconsistency will need to be considered in future work, but can be somewhat alleviated by finding the best RRT path over multiple runs of the algorithm. Given the long timescale of glider missions, computation time is not a driving factor, and the algorithm can easily be run multiple times to yield improved results. The results shown here are the best of five consecutive runs.

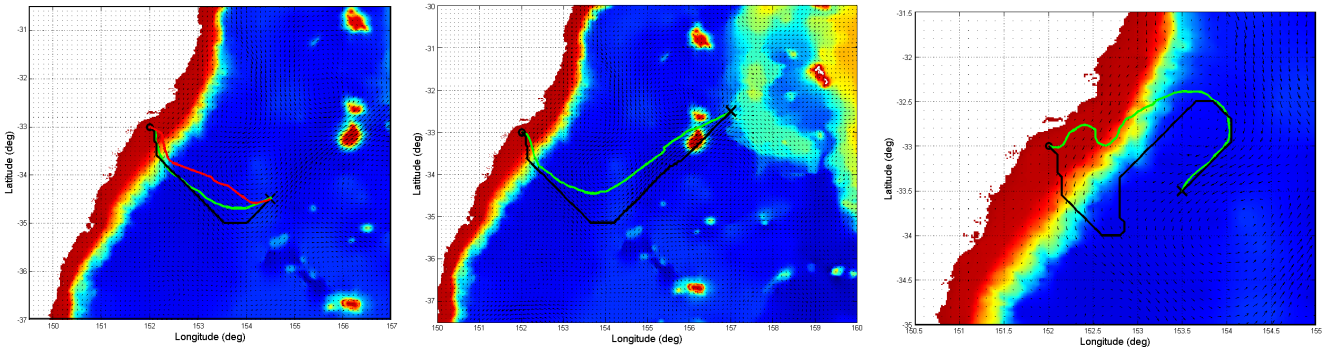

Figure 7: Path results for mission scenarios A, B and C (left to right)

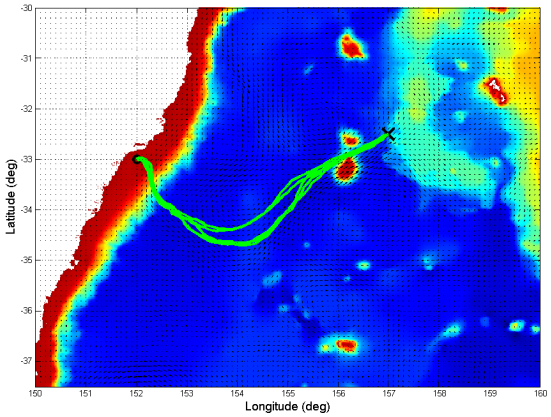

Figure 8: Path variation over multiple runs

To evaluate the performance of the current algorithms, results are compared with those for a basic heading controller, where the vehicle heading is set to the destination at every point on the path. Results in Figure 7 show that both the grid and RRT based methods are successful in producing feasible paths through the field. Over the three missions A, B and C, the energy consumption has been computed as the sum of Equation 10 over all path nodes. The results are shown in Table 1.

Mission	**Energy Cost (kJ)**		
	Heading controller	**RRT**	**Grid**
A	289.1	263.7	267.8
B	N/A	417.5	432.9
C	N/A	475.5	431.7

Table 1: Energy cost results

Mission A represents a simple planning scenario for a short distance mission. In this case the destination is downstream from the start, and the direct method yields satisfactory results. However, since it does not consider path costs, its energy expenditure is greater than both RRT and grid paths. The RRT path yields a marginal 2% improvement over the grid-based path.

Mission B is a more difficult mission as it spans over a larger range and has shallow regions close to the destination point. For this less simplistic planning scenario, a heading controller is no longer sufficient to reach the destination. Instead, the vehicle reaches a steady state position where its propulsion cannot overcome ocean currents. The drawbacks of the grid discretisation are also evident, as the grid path cannot avoid the adjacent shallow region due to lack of discretisation freedom. On the other hand, the RRT path travels around the shallow areas and offers a 6% energy improvement over the grid path.

However, for a highly difficult mission against strong ocean currents (Mission C), the grid path consumes 9% less energy than the RRT path. This scenario reveals some potential downsides of the RRT approach. With RRTs, paths that are against the Voronoi bias (against ocean currents) are less likely to be found, and the final path can fluctuate unnecessarily due to its probabilistic nature.

6 Conclusions and Future Work

This paper discusses the application of Rapidly-Exploring Random Trees to underwater glider path planning in an ocean current field. Biasing the growth of these trees allows us to generate a network which provides smoother paths than typical grid partitioned methods.

Early results suggest that RRTs have strong potential in underwater planning applications because of the Voronoi bias, which innately drives the trees in the direction of ocean currents. This can then be exploited to generate paths that make full use of the ocean current velocity advantage. A comparison between RRT and grid techniques indicates that both methods may be suitable for glider path planning in different mission scenarios. RRT methods can overcome the drawbacks in the grid discretisation and thereby avoid high energy shallow regions, while the grid based planner can handle highly difficult missions against strong ocean currents, where RRTs may struggle to find path solutions against the Voronoi bias.

However, one point to note is that the grid paths are optimal, whereas the RRT paths could potentially be improved if the RRT growth is improved accordingly. Thus, future work will examine additional biasing methods by which improved RRT paths can be generated.

Additional future work will look at the following:

- Real-time update of ocean current and bathymetry data as they become available.
- Modelling ocean surface as a sphere rather than a 2D plane, to more accurately represent distances between latitude / longitude co-ordinates.
- Extension of the algorithm to produce strong results in a non-deterministic environment.
- Employing the algorithms on a glider deployment, in order to obtain empirical data.

Nevertheless, this paper has successfully applied Rapidly-Exploring Random Trees to an underwater application and demonstrated numerically the potential of these techniques, providing a foundation for future research in the area.

Acknowledgments

This work is supported by the ARC Centre of Excellence programme, funded in part by the Australian Research Council (ARC) and the New South Wales State Government. The authors would like to acknowledge Surya Singh and Sisir Babu Karumanchi for assistance with RRTs, as well as Asher Bender, Daniel Steinberg and Ariell Friedman for development of the glider model.

References

[Blidberg, 2001] D. R. Blidberg. The Development of Autonomous Underwater Vehicles (AUV): A brief summary. *IEEE International Conference on Robotics and Automation*, 2001.

[BlueLink, 2009] Bureau of Meteorology, Royal Australian Navy, CSIRO. BlueLink Ocean Forecasting Australia. *http://www.marine.csiro.au/bluelink*, 2009

[BODC, 2009] British Oceanographic Data Centre. Gridded Bathymetric Datasets. https://www.bodc.ac.uk/data/online_delivery/gebco/select, 2009

[Carroll *et al.*, 1992] K.P. Carroll, S.R. McClaran, E.L. Nelson, D.M. Barnett, D.K. Friesen, and G.N. Williams. AUV path planning: an A* approach to path planning with consideration of variable vehicle speeds and multiple, overlapping, time-dependent exclusion zones. In *Proceedings of the 1992 Symposium on Autonomous Underwater Vehicle Technology*, pages 79–84, 1992

[Friedman, 2008] A. L. Friedman. Design and Development of the Buoyancy Engine and Pitch Control System for an Autonomous Underwater Glider. Undergraduate Thesis, School of Aeronautical, Mechanical and Mechatronic Engineering, 2008.

[Garau *et* al., 2005] B. Garau, A. Alvarez, G. Oliver. Path Planning of Autonomous Underwater Vehicles in Current Fields with Complex Spatial Variability: an A* Approach. *IEEE International Conference on Robotics and Automation*, 2005.

[Graver, 2005], J. G. Graver. Underwater Gliders: Dynamics, Control and Design. PhD Thesis, Mechanical and Aerospace Engineering, Princeton University, 2005.

[Hart *et al.*, 1968] P. Hart, N. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. IEEE Transactions on Systems Science and Cybernetics, 4(2), pages 100–107, 1968.

[Inanc *et al.*, 2005] T. Inanc, S. C. Shadden, J. E., Marsden. Optimal Trajectory Generation in Ocean Flows. *American Control Conference*, 2005.

[LaValle, 1998] S. M. LaValle. Rapidly-exploring random trees: A new tool for path planning. *Iowa State University*, 1998.

[LaValle & Kuffner, 2001] S. M. LaValle, J. J. Kuffner. Randomized Kinodynamic Planning. *The International Journal of Robotics Research*, 20, pages 378-385, 2001.

[Mahmoudian *et al.*, 2007] N. Mahmoudian, C. A. Woolsey, J. Geisbert. Steady Turns and Optimal Paths for Underwater Gliders. *AIAA Guidance, Navigation and Control Conference and Exhibit*, 2007.

[Rao, 2009] D. Rao. Path Planning for an Underwater Glider in Ocean Currents. Undergraduate Thesis, School of Aeronautical, Mechanical and Mechatronic Engineering, 2009.

[Simmons & Urmson, 2003] R. Simmons, C. Urmson. Approaches for heuristically biasing RRT growth. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1178 – 1183, 2003.

[Tan *et al*., 2004] C. S. Tan, R. Sutton, J, Chudley. An incremental stochastic motion planning technique for autonomous underwater vehicles. In *Proceedings of IFAC Control Applications in Marine Systems Conference*, pages 483–488, 2004.

[Witt & Dunbabin, 2008] J. Witt, M. Dunbabin. Go With the Flow: Optimal AUV Path Planning in Coastal Environments. *Australian Conference on Robotics and Automation*, 2008.

[Xu *et al.*, 2007] Z. D. Xu, R. Zhang, C. X. Ziying. Research on Global Path Planning in the Marine Environment for AUV. *International Conference on Intelligent Systems and Knowledge Engineering (ISKE)*, 2007.

[Zhang *et* al., 2008] W. Zhang, T. Inanc, S. Ober-Blöbaum, J. Marsden. Optimal Trajectory Generation for a Glider in Time-Varying 2D Ocean Flows B-spline Model. *IEEE International Conference on Robotics and Automation*, 2008.