Terrain-Aided Navigation for Long-Range AUVs in Dynamic Under-Mapped Environments

Georgios Salavasidis
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
geosal@noc.ac.uk

Andrea Munafò
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
andmun@noc.ac.uk

Davide Fenucci
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
davfen@noc.ac.uk

Catherine A. Harris
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
cathh@noc.ac.uk

Thomas Prampart
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
thoram@noc.ac.uk

Robert Templeton
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
robert.templeton@noc.ac.uk

Micheal Smart
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
micsmart@noc.ac.uk

Daniel T. Roper
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
danrope@noc.ac.uk

Miles Pebody
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
miles.pebody@noc.ac.uk

Povl E. Abrahamsen
Polar Oceans
British Antarctic Survey
Cambridge, CB3 0ET, UK
epab@bas.ac.uk

Stephen D. McPhail
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
sdm@noc.ac.uk

Eric Rogers
Electronics and Computer Science
University of Southampton
Southampton, SO17 1BJ, UK
etar@ecs.soton.ac.uk

Alexander B. Phillips
Marine Autonomous and Robotic Systems
National Oceanography Centre
Southampton, SO14 3ZH, UK
abp@noc.ac.uk

Abstract

Deploying long-range Autonomous Underwater Vehicles (AUVs) mid-water column in the deep ocean is one of the most challenging applications for these submersibles. Without external support and speed over the ground measurements, Dead-Reckoning (DR) navigation inevitably experiences an error proportional to the mission range and the speed of the water currents. In response to this problem, a computationally feasible and low-power Terrain-Aided Navigation (TAN) system is developed. A Rao-Blackwellized Particle Filter (RBPF) robust to estimation divergence is designed to estimate the vehicle's position and the speed of water currents. To evaluate performance, field data from multi-day AUV deployments in the Southern Ocean are used. These form a unique test case for assessing the TAN performance under extremely challenging conditions. Despite the use of a small number of low-power sensors and a Doppler Velocity Log (DVL) to enable TAN, the algorithm limits the localisation error to within a few hundreds of metres, as opposed to a DR error of 40 km, given a 50 m resolution bathymetric map. To evaluate further the effectiveness of the system under a varying map quality, grids of 100 m, 200 m and 400 m resolution are generated by sub-sampling the original 50 m resolution map. Despite the high complexity of the navigation problem, the filter exhibits robust and relatively accurate behaviour. Given the current aim of the oceanographic community to develop maps of similar resolution, the results of this work suggest that TAN can enable AUV operations of the order of months using global bathymetric models.

Abbreviations

ALR1500/6000	Autosub Long Range 1500/6000	LL	Latitude and Longitude
AUV	Autonomous Underwater Vehicle	MEMS	Micro-Electro-Mechanical System
BR	Bottom-Relative	MMSE	Minimum Mean Square Error
CTD	Conductivity Temperature Depth	MXX	Mission XX
DEM	Digital Elevation Model	NED	North-East-Down
DT	Distance Travelled	PF	Particle Filter
DVL	Doppler Velocity Log	RBPF	Rao-Blackwellized Particle Filter
DynOPO	Dynamics of the Orkney Passage Outflow	RMSE	Root Mean Square Error
EKF	Extended Kalman Filter	RRS	Royal Research Ship
GPS	Global Positioning System	SBL	Short Baseline
IA	Interval Analysis	SIS	Sequential Importance Sampling
IMU	Inertial Measurement Unit	SISR	SIS with Resampling
INS	Inertial Navigation System	TAN	Terrain-Aided Navigation
JCR	James Clark Ross	USBL	Ultra Short Baseline
KF	Kalman Filter	WR	Water-Relative
LBL	Long Baseline		

1 Introduction

Conventional techniques to sample seawater properties require access to a research ship, for which the operational window can be constrained by factors such as weather condition, limited ship manoeuvrability and inability to access certain areas (e.g. regions covered by a thick layer of sea ice), in addition to the prohibitively high cost of scientific ship time. However, recent advances in maritime robotics suggest that safe and cost-effective techniques for acquiring subsea measurements in large areas will soon be available. An example platform that currently enables sampling of large water masses is the underwater glider (Bachmayer et al., 2004). However, the lack of precise control and navigation, the saw-tooth gliding behaviour and the restricted payload mean that gliders are not suitable for all applications. Examples include science-driven missions where submersibles are required to operate for a prolonged period of time at a relatively fixed depth,

following the surface ice or seabed at a near-constant range, observing slow-moving targets mid-water column (e.g. tracking zooplankton and/or descending particle aggregates), autonomously detecting and tracking hot spots of biological activity for days to months as they move in an unpredictable and complicated trajectory, etc.

To address these needs, several new platforms have recently been developed opening up a diverse range of new applications in the open ocean (Jakuba et al., 2018; Yoerger et al., 2018; Hobson et al., 2018). Similarly, in response to ever-increasing demand for conducting research activities at even larger temporal and spatial scales, the National Oceanography Centre (NOC) in the UK has developed long-range propeller-driven Autonomous Underwater Vehicles (AUVs), such as the Autosub Long Range 6000 (ALR6000; Furlong et al. 2012) and the Autosub Long Range 1500 (ALR1500; Roper et al. 2017). These platforms combine the extreme range of subsea gliders with the increased manoeuvrability, position control and payload capacity of large flight-style AUVs. By using high energy density lithium primary cells and minimising the energy consumption within all sub-systems, ALR vehicles can achieve a multi-month endurance (Salavasidis, 2019; Salavasidis et al., 2020). This capability coupled with the ample flooded space available and the increasing availability of low-power sensors, developed primarily for subsea gliders (Hayes et al., 2007), have provided such platforms with a variety of payload options and use cases. For example, the extreme range of these vehicles can facilitate the attempt of previously impossible tasks, such as a continuous mid-water column biogeochemical survey from Svalbard (Norway) to Point Barrow (Alaska, USA) under the Arctic sea ice (Salavasidis, 2019) or a long-term monitoring of the mesopelagic zone (200-1000 m below the surface), which hosts a diverse biological community of unique organisms and includes sharp variations in seawater properties of high scientific interest (e.g. temperature, salinity, etc.).

Despite the considerable effort directed to underwater navigation, a self-contained low-power solution is not yet available, precluding multi-month AUV missions in Global Positioning System (GPS) denied environments (e.g. under the sea ice). The problem becomes further complicated when deploying AUVs mid-water column and navigating relative to the surrounding water, particularly in areas with high water flow intensity. AUV navigation in the absence of estimates of speed over the ground is currently the weakest (Medagoda et al., 2016), making mid-water navigation the next frontier for maritime robotics.

Survey-class AUVs typically rely on simple Dead-Reckoning (DR) techniques coupled with Doppler effect-based velocity sensors (Salavasidis et al., 2019), such as a Doppler Velocity Log (DVL). Although this approach provides acceptable accuracy in short oceanographic missions, the error in estimating the vehicle's position is cumulative, as each estimate is relative to the previous one, hence DR experiences unbounded error growth. The rate of the error increase depends on the quality of the employed motion sensors and the physical environment where the AUV operates. The typical error rate for survey-class AUVs operating in a straight line is of the order of 0.1-2% of the Distance Travelled (DT), given the velocity over the ground (Paull et al., 2014). Otherwise, when only water-relative velocity is available, the error is mainly proportional to the speed of the water currents. Also, the error rate can be further reduced to less than 0.05% of DT by using a high-grade Inertial Navigation System (INS) coupled with a DVL using a continuous bottom-tracking regime. However, such a system can be costly, bulky and power-demanding, making it incompatible with the energy constraints of long-range AUVs.

When operating in open ocean, a widely used solution to the navigation problem is regular surfacing to obtain GPS position fixes. However, surfacing results in a trade-off between navigational accuracy and the effective use of the vehicle, which is therefore an important consideration for deep-water missions. Moreover, surface access for certain applications can be limited or even impossible, e.g. beneath the Arctic ice cap. Alternatively, AUVs can rely on external acoustic-based support systems to limit the navigation error (Paull et al., 2014). Three schemes are the most frequently used (Vickery, 1998): a) Long Baseline (LBL), b) Short Baseline (SBL), and c) Ultra Short Baseline (USBL). In the case of the LBL, the vehicle can self-localise with respect to a set of moored beacons or, in case of SBL and USBL, a support-ship or an unmanned surface vehicle can be used (Mandt et al., 2001; Phillips et al., 2018).

Acoustic-based navigation is a well-established technique, but it has the disadvantage of constraining the

AUV's operational window, as the vehicle is required to remain within range of the external support system (posing a major limitation for long-range AUVs). Furthermore, acoustic-based systems can be expensive and require a considerable time and effort to be calibrated, deployed, surveyed and recovered (resulting in further cost increases and delays until the mission can begin). If surface support is provided, then the operational cost of a dedicated ship can be impractical for some applications. Assistance from the surface can also be restricted due to, for example, a thick layer of sea ice.

To overcome these limitations and enable long-endurance underwater missions, long-range AUVs require a self-contained navigation solution that uses only on-board sensors and does not rely on external support or periodic surfacing. Navigation techniques that exploit the geophysical properties of Earth appear to offer a viable solution (Teixeira, 2007). This type of navigation is fully implementable on-board an AUV, thus particularly appealing for long-range vehicles.

Terrain-Aided Navigation (TAN) is a form of geophysical-based localisation where observable terrain features are matched to a pre-existing bathymetric map to infer information about the state of the vehicle (Melo and Matos, 2017; Salavasidis et al., 2019), with the first applications to the underwater domain and AUVs dating back to the 1990s (Bergem, 1993; Di Massa, 1997; Di Massa and Stewart, 1997). To perceive water depth and enable TAN, an AUV needs to be equipped with one or more sonar sensors. By combining range measurements with the operating depth of an AUV, computed by converting the hydrostatic pressure into depth, bathymetric observations can be formed. It has been shown, for example in Meduna (2011), that the performance of TAN depends on the accuracy of the employed sensors (motion and bathymetric), the number of bathymetric observations available at a time (i.e. whether the vehicle is equipped with multi- or single-beam sonars) and also the resolution of the employed bathymetric map (including the terrain variability itself). A comprehensive review of underwater TAN and relevant background literature, including sensor technologies, can be found in Melo and Matos (2017).

Despite the demonstrated effectiveness of TAN in aerial applications, e.g. as a navigation method for cruise missiles (Siouris, 2004), the application of these techniques in the unstructured underwater environment is relatively immature with outstanding challenges to be addressed (Melo and Matos, 2017). Implementations are extremely limited in number and the majority of them rely on platforms equipped with power hungry sensors (e.g. a high-grade INS and a 127-beam echo-sounder; Anonsen and Hallingstad 2006) and require bathymetric maps with resolutions varying from sub-metres to up to only a few metres. However, large scale high-resolution terrain maps are rarely available, particularly in remote deep oceans, and the reliance on power-hungry sensors severely limits the AUV endurance. Therefore, real-world applications are mostly short missions, ranging from hundreds of metres to few kilometres in shallow areas for which detailed bathymetric surveys have been undertaken.

To reduce the on-board power consumption and enable longer AUV missions, recent studies have focused on developing TAN techniques for platforms with limited sensor suites (Salavasidis, 2019; Meduna, 2011; Meduna et al., 2008). AUV operations near the ocean bottom require no additional ranging sensors as the downward-facing DVL, which is typically included in the standard sensor suite for AUV, can provide range-to-bottom measurements (Salavasidis et al., 2019; Donovan, 2011; Meduna et al., 2010). However, challenging conditions can occur when AUVs are required to operate far outside of the DVL's bottom-tracking range while, for example, collecting biogeochemical and environmental DNA samples mid-water column in deep oceans. No suitable (in size and energy consumption) off-the-shelf echo-sounders are commercially available with the ability to ping the seabed from thousands of metres away. For the ALR vehicles, the NOC has developed a bespoke 4000 m range 12 kHz single-beam echo-sounder to enable TAN during mid-water operations in deep oceans. This sensor complies with the low-power budget of the ALR vehicles by functioning at a low-repetition rate (of the order of minutes) while being optimised to have a low quiescent power demand (< 0.5 W) and having a minimum possible volume to easily fit into the ALR vehicles.

Furthermore, despite the promising TAN results achieved using low-power sensor suites, the impact of using low-resolution bathymetric maps has, as yet, seen relatively little investigation. Instead, there are surprisingly very few studies that specify the resolution of the utilised bathymetric map (as also stated in Meduna 2011;

Meduna et al. 2009) and examine its contribution to the overall navigation performance.

Even though an inverse relationship between map coarseness and TAN accuracy can be expected, the rate of degradation and its dependencies is currently unclear. Most of the reported field experiments use terrain maps with grid cell size significantly smaller than 50 m, orders of resolution that are rarely available in remote deep oceans. For example, the currently available digital bathymetric model of the world's oceans (GEBCO_2014; Weatherall et al. 2015) is gridded at a regular interval of 30 arc-seconds (approximately 920 m resolution at the equator) and the majority of the grid cell values are subject to interpolation because of a lack of actual bathymetric data. TAN techniques using such coarse bathymetric maps have only been tested in computer simulations. In particular, preparatory studies for crossings of the Arctic Ocean using ALR1500 have investigated the potential of using TAN given an approximately 500 m resolution bathymetric map (Salavasidis et al., 2016, 2018).

In response to the shortcomings just discussed, this work develops and evaluates a TAN algorithm that relies on basic low-power motion sensors, a downward-facing DVL treated as a four-beam echo-sounder and low-resolution terrain maps. To make the navigation filter computationally feasible in real-time for vehicles using processing boards with limited computational resources, the filter estimates the 2-D vehicle's position and the 2-D speed of the local water currents (to withstand the drift caused by the water flow while operating out of the DVL's bottom-tracking range). The state estimation is performed by the application of the Rao-Blackwellized Particle Filter (RBPF; Schon et al. 2005).

The performance of the developed navigation system is evaluated using field data collected during three multi-day deployments of ALR6000 in the Southern Ocean. These experiments comprise a unique test case given their range of up to $\approx$ 200 km and depths over 3000 m while traversing areas with low topographic information. To imitate operations in the mid-water ocean, the available bottom-tracking velocity is replaced by the water-relative velocity derived from the rotation speed of the vehicle's propeller, thus introducing a drift caused by water currents (which can be in excess of 0.6 m/s). To further investigate whether the navigation filter can cope with coarse bathymetric maps, typically available in remote deep oceans, the quality of the original 50 m resolution map of the area is degraded via a sub-sampling process and three new maps of 100 m, 200 m and 400 m resolution are generated. To deal with these challenging conditions, the navigation filter is augmented with techniques to escape from local minima caused by terrain similarities.

Aiming to develop robust TAN techniques that will enable multi-month AUV deployments at any possible mission depth, this work extends that in Salavasidis et al. (2019) by navigating AUVs in areas with strong water currents and using low-resolution maps. The remainder of the paper is structured as follows: Section 2 introduces the problem of navigating AUVs mid-water column in deep oceans using field data collected during past deployments in the Southern Ocean. Section 3 presents the main challenges for enabling long-range TAN. Section 4 formulates TAN as a recursive Bayesian estimation problem and described the RBPF algorithm, which is specifically designed for long-range AUVs. Section 5 specifies parameters required to enable the RBPF algorithm and constructs the experimental environment where the proposed navigation system is evaluated. Section 6 evaluates the *a-posteriori* performance of the TAN algorithm, while Section 7 concludes the paper, summarising the most important outcomes of this work and providing directions for future research.

2 Deployments of ALR6000 in the Southern Ocean

Science expeditions as part of the Dynamics of the Orkney Passage Outflow (DynOPO) research programme (Naveira Garabato and Forryan, 2017; British Antarctic Survey, 2017) were undertaken in April 2017 at the Orkney Passage (see Figure 1). This is a submarine valley of the Southern Ocean, north-east of the Antarctic Peninsula, that connects the Weddell Sea and the Scotia Sea with water depths in excess of 5300 m, steep terrain morphology and the presence of strong time-varying water currents.

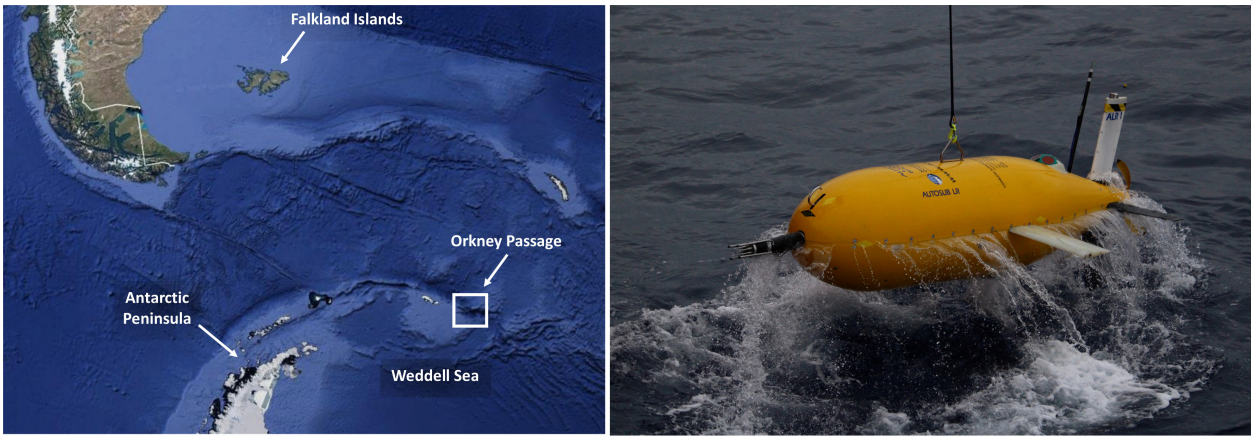

Figure 1: Left: ALR6000 deployments at the Orkney Passage, a narrow gap in the ocean-floor ridge that extends north-east of the Antarctic Peninsula (Google Earth, 2017). Right: ALR6000 recovery from a deep multi-day deployment in the Orkney Passage.

The main objective of the expedition was to map the flow of the deep water masses and investigate the causes of the warming of the deep ocean. Hence, ALR6000 performed three deep long-range missions (hereafter referred to as Missions 41, 42 and 44) and measured seawater properties (e.g. temperature, salinity, etc.) together with the speed of the water currents and the water turbulence close to the ocean bottom. The information obtained using ALR6000 has provided scientists with an invaluable set of data for understanding better the deep water circulation and its role in regulating the Earth's climate system (Naveira Garabato et al., 2019). For the deployment and recovery of the vehicle, the UK's ice-strengthened Royal Research Ship (RRS), James Clark Ross (JCR), was used.

2.1 Dead-Reckoning Navigation Using Water-Relative Velocity

To minimise the energy consumption and enable long-range missions, ALR6000 is typically equipped with a small number of low-power navigation sensors (Salavasidis, 2019; Salavasidis et al., 2020). These include a GPS module to obtain position fixes when at the surface; a 6-axis magnetic compass/attitude module to determine the vehicle's direction; a Conductivity, Temperature and Depth (CTD) probe to calculate the vehicle's operating depth; and two DVLs, one facing upwards and the other downwards. The downward-facing DVL is used to estimate the vehicle velocity (bottom-relative or water-relative) and has a maximum reliable range of approximately 180 m for bottom-lock, whereas the upward-facing DVL is primarily used to avoid collision with ice while operating under the polar sea ice. To level the velocity vector in a local horizontal reference frame, the downward-facing DVL has been integrated with tilt sensors (Micro-Electro-Mechanical System (MEMS) accelerometers).

To meet the science requirement of data-collection near the ocean bottom, ALR6000 utilised an altitude following control scheme with 90 m altitude demand and constant 0.7 m/s forward speed over the ground. In terms of navigation, the vehicle self-localised using GPS position fixes when on the surface and dead-reckoning when submerged (Salavasidis et al., 2019). The vehicle's path was pre-specified using waypoints along bathymetric contour lines (for the vehicle's safety). The track following algorithm given in McPhail and Pebody (1998) was used to guide the vehicle among these waypoints. The maximum depth achieved was approximately 3700 m while operating for nearly 200 km, a range achieved in approximately three days. The availability of the on-site research ship enabled the position of the vehicle to be intermittently recorded for mission safety and post-processing purposes using the ship's USBL head and an acoustic transponder available on the ALR6000. These USBL measurements are hereafter treated as the vehicle's baseline position.

As reported in Salavasidis et al. (2019), the ALR6000's real-time DR navigation was sufficiently accurate to

complete the demanded tasks during the DynOPO missions. Despite the use of basic low-power motion sensors, high navigation performance was achieved mainly due to the successful use of online sensor calibration and error compensation techniques while operating within the DVL's bottom-tracking range. The real-time DR navigation accumulated a Root Mean Square Error (RMSE) of 1.2 km at the surfacing location after a mission of approximately 195 km, which corresponds to an error growth of less than 0.65% of DT. Moreover, the largest proportion of the error appeared to occur when the vehicle operated under the influence of strong water currents (i.e. while descending to a depth greater than 1100 m and ascending from a depth over 3500 m), whereas the error increase during the bottom-tracking regime was relatively low (Salavasidis et al., 2019).

Many of the current oceanographic science missions focus on biogeochemical processes taking place mid-water column in remote deep oceans. To imitate such deployments in realistic conditions and demonstrate the extent of the navigation challenge whilst operating in areas with strong water currents, the available Bottom-Relative (BR) velocity from the three DynOPO missions is replaced by Water-Relative (WR) velocity estimates. One option is to use the WR velocity provided by the downward-facing DVL. However, previous field operations have shown that DVLs are likely to give poor quality WR velocity due to a sporadic lack of particulate matter to provide back-scatter in clear deep oceanic water. An alternative option is to measure and convert the propeller's rotation speed into forward speed using the following first order mapping function:

$$v_k^f = 0.00742 \cdot Pr_k \cdot \cos(\theta_k), \tag{1}$$

where θ is the pitch displacement of the vehicle; Pr denotes the propeller's Revolutions Per Minute (RPM); and v_k^f corresponds to the horizontally levelled forward speed. The coefficient used in Equation (1) is calculated by calibrating the mapping function against high-quality speed estimates over the ground. Moreover, the starboard speed is set to zero due to the lack of actuation in this direction.

Figure 2 shows the WR velocity computed using Equation (1) in comparison to the water and bottom-relative velocities obtained from the downward-facing DVL for a random time-interval of Mission 44 (M44) and a depth greater than 2500 m. The DVL's water-relative velocity appears to be noisy[1], whereas the velocity computed using the propeller RPM appears to be reliable in the steady-state phase of the vehicle (little or no acceleration/deceleration). By including the BR velocity in Figure 2, information regarding the speed of the local water currents can be inferred. Note also the drop-outs (zero forward speed) that the downward-facing DVL experiences for various reasons (e.g. steep terrain, clear oceanic water, etc.).

Figure 3 gives the estimated path of the vehicle during the three DynOPO missions as calculated in real-time by the on-board DR navigation (Salavasidis et al., 2019) and in post-processing using WR speed computed via Equation (1). For notational convenience, a superscript asterisk will be used to indicate that DR navigation uses the propeller-derived WR velocity (DR*), as opposed to DR (without asterisk) navigation when using BR velocity. Moreover, the ship-based USBL measurements are included in Figure 3 as localisation references. It is assumed that each mission starts when GPS position fixes become unavailable, whereas the end of each mission corresponds to the surfacing location. Note also that ALR6000 failed to complete M42 because a backup system (independent of the mission management computer) activated the primary abort mechanism. Therefore, the end of M42 is assumed to be the location where the vehicle aborted the mission. In addition, Table 1 lists navigation details for each mission and quantifies the navigation error with respect to the last recorded USBL measurement.

Despite the real-time navigation achieving a sufficient level of accuracy, DR* experiences substantial positioning drift due to the strong water currents in the area. In particular, the localisation error is approximately 28.4% of DT for M42 and 21.4% of DT for M44, which corresponds, respectively, to an RMSE of 33 km and 41 km when compared to the last recorded USBL measurement. These results clearly reveal the need for an effective self-contained navigation solution enabling long-range missions mid-water column in the deep oceans.

[1]Note that the noise in the WR velocity provided by the DVL can be reduced to a certain extent by combining velocities derived from several DVL depth cells (Medagoda et al., 2016).

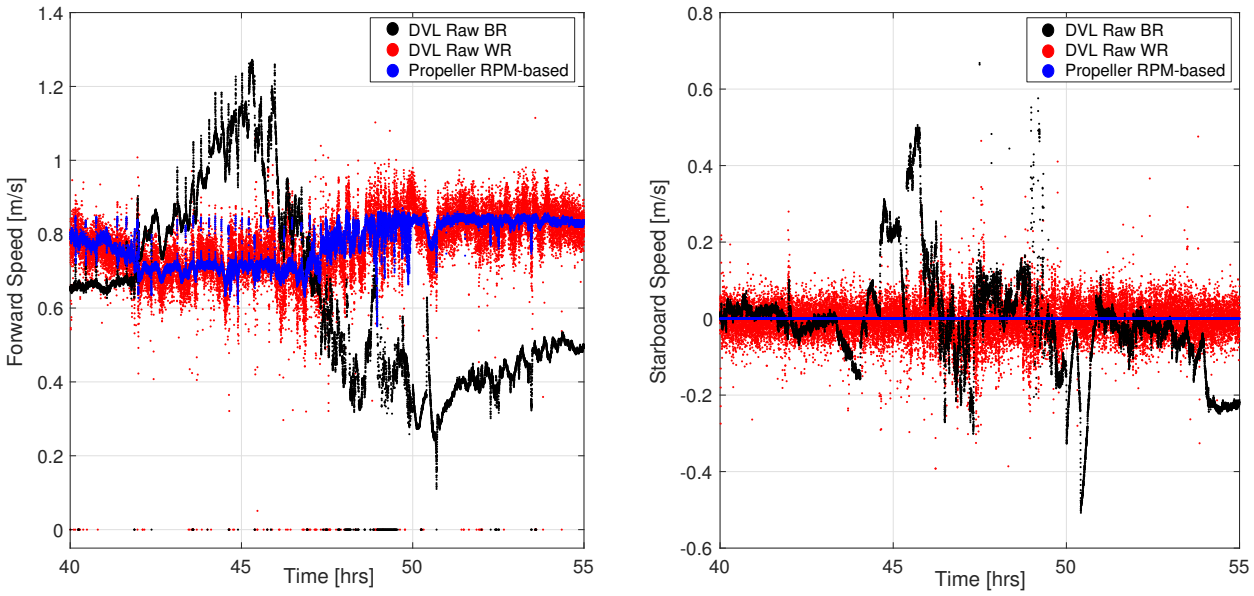

Figure 2: Comparison between raw speed estimates provided by the DVL (BR and WR speed) and the forward speed computed using the propeller RPM from M44. For the DVL's water-relative speed, the third depth cell is visualised (bin length 8 m).

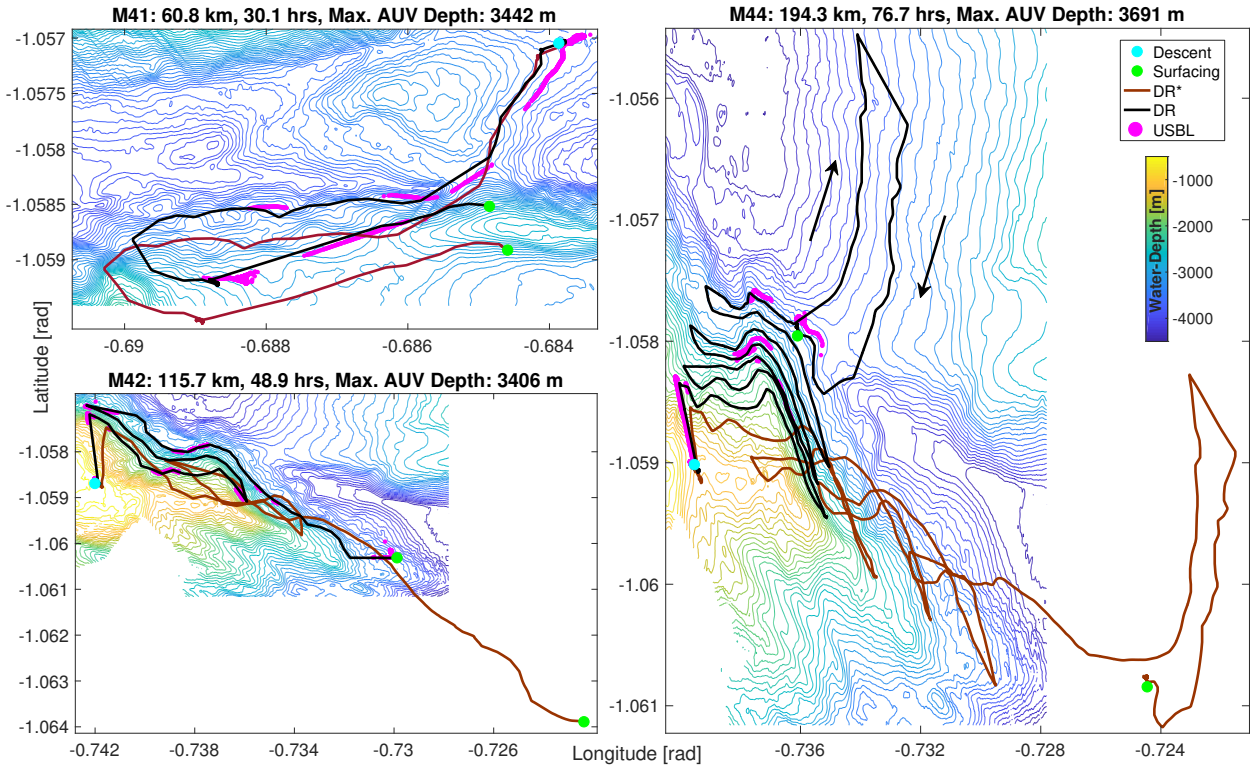

Figure 3: ALR6000 position estimates via DR navigation. The black dashed lines correspond to the real-time navigation (using BR velocity when available) whilst the red dashed lines are computed using speed estimates obtained from the propeller RPM. USBL measurements are shown in magenta as localisation references.

Table 1: Details of the three DynOPO missions

Parameters	M41	M42	M44	Units
Total duration	30.1	48.9	76.7	hrs
Total Distance Travelled (DT)	60.8	115.7	194.3	km
Maximum operating depth	3442	3406	3691	m
Descent water depth	3483	478	1136	m
Ascent water depth	2920	-	3505	m
Number of USBL measurements	4248	2529	4077	-
DT until the last USBL	51.6	115.7	191.7	km
DR error (w.r.t. the last USBL)	1.1	0.3	1.2	km
DR error (w.r.t. DT)	2.1	0.27	0.6	%
DR* error (w.r.t. the last USBL)	2.9	33	41	km
DR* error (w.r.t. DT)	5.7	28.4	21.4	%

w.r.t.: with respect to.

3 Challenges to Enable Long-Range Terrain-Aided Navigation

TAN is a self-contained localisation technique with demonstrated effectiveness in bounding the localisation error during short AUV missions in areas with known bathymetry (Melo and Matos, 2017). However, the application of these techniques during multi-month long missions can be challenging because of: a) the lack of accurate large-scale bathymetric maps; and b) the demand for energy conservation and hence the use of a limited sensor suite and a low-power processing board with low computational power.

3.1 Bathymetric Maps

A bathymetric map is a terrain grid model—typically a uniform one—that represents the true seabed to some varying degree of accuracy, with the deviation of each grid cell from the reality constituting an error. The size of the grid cells is commonly referred to as resolution of the grid model, with a higher size corresponding to a lower resolution. The accuracy of a grid model has been shown to have a positive correlation with the model's resolution (Li, 1992), therefore high-resolution models can represent the potentially complex underwater terrain with higher fidelity.

In shallow waters, a small scale high-resolution terrain map can be generated by performing extensive surveys using precise mapping systems installed on accurately georeferenced and stabilised platforms. Conversely, the construction of reliable large scale terrain maps in remote deep oceans can be costly and challenging (if not impossible). This is because these areas may be inaccessible (e.g. due to a thick sea ice), or the surface-made maps can be subject to substantial errors due to the large footprint of the ship-mounted sonars (Jensen et al., 2011). For example, currently available bathymetric maps of deep areas are gridded to hundreds of metres (per side) because of the sparse and often unreliable depth data (Mayer et al., 2018). Such terrain models may also contain large areas lacking actual depth measurements, which are typically filled via interpolation. As a result, TAN algorithms must be sufficiently robust when used for long-range remote applications to deal with areas where only sparse[2] bathymetric maps are available.

Although large scale maps are currently of low-quality, within the *GEBCO Seabed 2030* project (Mayer et al., 2018) the oceanographic community aims to develop a new refined Digital Elevation Model (DEM) of the world's oceans. This work will combine old bathymetric data of acceptable quality with new surveys scheduled for the future to generate new DEMs of resolutions varying with depth by 2030. In particular, 72.6% of oceans are expected to be mapped with 400 m cell sizes, 11% with 200 m cell size and 13.7%

[2]In the content of this work, the sparsity/coarseness of a bathymetric map refers to both low-resolution and a lack of actual depth data over large areas (where the depth values of the grid cells are computed using interpolation techniques).

with 100 m cell size. The remaining 2.7% corresponds to very deep waters (5750-11000 m) where such high-resolution would be impractical and hence will be mapped with 800 m cell size.

3.1.1 Bathymetric Error Sources

Sources of error (and thus of uncertainty) in bathymetric maps are typically grouped into three wide categories (Desmet, 1997; Hare et al., 2011): *Type 1* - depth measurement errors (e.g. uncompensated biases and positional inaccuracies) and temporal changes (ageing of sounding data); *Type 2* - errors introduced when processing available depth measurements with gridding algorithms to create grid models (e.g. numerical errors, errors introduced by averaging/smoothing and interpolating); and *Type 3* - errors due to the terrain variability within each grid cell and between the available depth measurements. Therefore, the uncertainty associated with each grid cell depends on that in the depth measurements and the gridding techniques utilised, together with the topographic variability. High-resolution DEMs are primarily affected by *Type 1* errors, whereas low-resolution DEMs containing large regions where depth values are subject to interpolation can be massively affected by errors of *Type 2* and *Type 3*.

3.1.2 Bathymetric Uncertainty Modeling

A common and widely used approach to quantifying the vertical bathymetric uncertainty is to use a single global (i.e. space-invariant) measure of deviation, such as the RMSE or standard deviation statistics (Wechsler, 2007). Then, assuming error normality and identical distribution for each grid cell, the map uncertainty is often represented by a zero-mean Gaussian distribution with standard deviation σ_g, denoted by $\mathcal{N}(0, \sigma_g^2)$. This simplified approach can be sufficient in shallow waters and areas with low topographic variability, where high-quality depth measurements can be obtained and the gridding algorithm does not introduce large errors. However, when these conditions do not hold, this approach can become crude or even invalid for capturing the overall error effect (see Monckton 1994). Quantifying the uncertainty in coarse bathymetric maps in areas with a strong presence of errors of *Type 2* and *Type 3* is a very challenging task (Wechsler, 2007). The reason is that coarse maps are built based on an extremely small number of depth measurements and uncertainty models usually require information that is often not available (e.g. metadata associated to existing depth samples). Lacking definitive information regarding the spatial structure of the map error, the RMSE can be the only statistic available for use.

To expand the above error model to deep water areas, the vertical uncertainty model can be augmented by associating an additional zero-mean Gaussian noise to each depth value of the grid. The intensity of this noise model is calculated as a function of the depth at each location of interest $\boldsymbol{x}$, which can mathematically be expressed as:

$$\sigma_h(\boldsymbol{x}) = 0.5\sqrt{1 + \Big(0.023 \cdot h(\boldsymbol{x})\Big)^2}, \tag{2}$$

where $h(\boldsymbol{x})$ corresponds to the depth value (expressed in metres) at a location of interest $\boldsymbol{x}$ (typically expressed by latitude and longitude) and $\sigma_h(\boldsymbol{x})$ is the computed vertical standard deviation in metres. This last equation assumes that the error in the depth values is the combination of uncorrelated errors in travel time (one metre) and sound speed (expressed as 2.3% of the depth value) of the sonar measurements used to create the bathymetric map.

Equation (2) was derived from the 95% confidence level of the maximum acceptable vertical uncertainty in sounding data for deep bathymetric measurements, according to the 5th edition of S-44 hydrographic survey standards and *Order 2* values (International Hydrographic Organization, 2008). Despite the fact that the formula was defined as quality control for bathymetric measurements and is not meant to quantify the uncertainty in bathymetric maps, the functional form of the error is useful for modeling uncertainties in sonar measurements and DEMs (Marks and Smith, 2008; Claus and Bachmayer, 2015). Therefore, it can be used to quantify the uncertainty due to the errors of *Type 1*.

3.2 Low-Power Sensor Suite

In addition to the lack of reliable large-scale terrain maps, the requirement for minimum energy consumption limits long-range AUVs to low-power motion sensors and sonars with low perceptual ability. This, in turn, complicates further the estimation process since it has been shown that the use of a high-grade INS in conjunction with an information-rich sonar (e.g. a multi-beam echo-sounder) would enable fast detection of unique terrain features that would allow the navigation filter to reliably converge to the true solution (Nygren and Jansson, 2004). To further save energy, long-range AUVs typically use a low-power processing board with limited computational resources (Phillips et al., 2012, 2017). This imposes additional constraints regarding the on-board computational power available for the navigation system.

4 Terrain-Aided Navigation for Long-Range AUVs

The use of a small number of low-power motion sensors unavoidably causes a rapid localisation error growth (e.g. see Section 2). Despite the challenges described in the previous sections, this work develops a TAN algorithm and investigates whether the localisation error can be bounded whilst using low-power sensors and coarse bathymetric maps during long duration AUV missions.

The AUV navigation task is formulated as a non-linear Bayesian estimation problem. This is then addressed by employing a Particle Filter (PF) and, in particular, the Sequential Importance Sampling with Resampling (SISR) algorithm (Salavasidis et al., 2019; Arulampalam et al., 2002), which is a recursive sample-based estimator with the ability to represent an arbitrary state posterior probability density function. However, while such an estimator can effectively deal with non-linearities, it may require significant processing power, which is a limiting factor for the applicability of these techniques on AUVs with low processing power. The latter is especially true for high-dimensional estimation problems, as the PF would require a larger number of samples to ensure that the state space is sufficiently covered.

To develop a navigation system with low computational requirements, the TAN algorithm concentrates on modeling and estimating a minimum number of parameters. Even though an AUV operates in the 3-D world, where three parameters are required to fully determine the vehicle's position, two parameters are essentially unknown and not directly observable. These unknowns constitute the easting and northing coordinates of some horizontal datum, or the Latitude and Longitude (LL) angles of the Earth's spherical coordinate system. The third dimension (operating depth) can be computed by measuring the hydrostatic pressure, using the CTD probe which contains a pressure sensor, and converting it into depth by applying the UNESCO pressure-to-depth formula (Fofonoff and Millard, 1983). Furthermore, given the slow dynamics of the water currents, a common assumption made in the oceanographic community (Visbeck, 2002), estimating the speed of the water currents becomes possible. Therefore, the final navigation filter is considered to be four dimensional (4-D), with the state vector consisting of the 2-D position of the vehicle and the 2-D speed of the water currents near the vehicle (i.e. velocity biases).

To reduce further the computational complexity of the 4-D particle filter, the RBPF can be used for state estimation (Schon et al., 2005; Teixeira, 2007). Based on the Rao-Blackwell theorem and Rao-Blackwellization techniques (Nordlund, 2002), the performance of the SISR algorithm can be enhanced (Doucet et al., 2000) whilst reducing the computational demand. This can be achieved by decomposing/marginalising the state vector into linear and non-linear sub-states. The linear part of the state is then estimated optimally using a Kalman Filter (KF) whilst the non-linear parameters are estimated using a PF. As the dimension of the PF decreases, fewer particles are required, resulting in a more computationally efficient filtering algorithm.

4.1 Process Model

The vehicle's motion can be modeled as a discrete Markov process that relies on 2-D kinematics. This propagates prior position estimates in time by integrating the vehicle's velocity expressed in the north and east directions, $\boldsymbol{u}_k = [u_k^N, u_k^E]^T$, which is obtained by rotating the horizontally-levelled body frame velocity using a 2-D rotation matrix (parametrized by yaw). Considering the spherical approximation of the Earth, the vehicle's position is expressed in LL. To include the water flow in the estimation process (assuming that only WR velocity is available), the state vector can be augmented with a 2-D velocity vector representing the water currents in the north and east directions. Given the slow dynamics (Visbeck, 2002), water currents can be modeled as random walk with a low-intensity Gaussian step. The process model is thus influenced by the vehicle's WR velocity and the speed of the local water currents, as the following discrete-time model shows:

$$\begin{gathered}\left[\begin{array}{c}\boldsymbol{x}_{k+1}\\ \hdashline \boldsymbol{c}_{k+1}\end{array}\right] = \left[\begin{array}{c:c}I_{2\times2} & F(\boldsymbol{x}_k)\\ \hdashline 0_{2\times2} & I_{2\times2}\end{array}\right]\left[\begin{array}{c}\boldsymbol{x}_k\\ \hdashline \boldsymbol{c}_k\end{array}\right] + \left[\begin{array}{c:c}F(\boldsymbol{x}_k) & 0_{2\times2}\\ \hdashline 0_{2\times2} & 0_{2\times2}\end{array}\right]\left[\begin{array}{c}\boldsymbol{u}_k\\ \hdashline 0_{2\times1}\end{array}\right] + \left[\begin{array}{c}\boldsymbol{\eta}_k^{\boldsymbol{x}}\\ \hdashline \boldsymbol{\eta}_k^{\boldsymbol{c}}\end{array}\right]\\ \text{with: } F(\boldsymbol{x}_k) = \Delta t \cdot G(\boldsymbol{x}_k) = \Delta t \cdot \left[\begin{array}{cc} R_e^{-1} & 0\\ 0 & \left(R_e \cdot \cos(x_k^{Lat})\right)^{-1}\end{array}\right]\end{gathered} \tag{3}$$

where $\boldsymbol{x}_k = [x_k^{Lat}, x_k^{Lon}]^T$ corresponds to the Markovian state sub-vector representing the vehicle LL position at time k (expressed in radians); $\boldsymbol{c}_k = [c_k^N, c_k^E]^T$ refers to the state sub-vector containing the speed of the local water currents in the north and east directions (expressed in m/s); Δt is the position propagation period, during which the vehicle's velocity and the local water currents are assumed to be constant; $F(\boldsymbol{x}_k)$ is a matrix used to convert the north-east displacement, caused by the vehicle's motion and the water currents during Δt, to the respective displacement in LL; R_e is the radius of Earth in metres; and the vectors $\boldsymbol{\eta}_k^{\boldsymbol{x}} = [\eta_k^{Lat}, \eta_k^{Lon}]^T$ and $\boldsymbol{\eta}_k^{\boldsymbol{c}} = [\eta_k^N, \eta_k^E]^T$ compose the process noise. $\boldsymbol{\eta}_k^{\boldsymbol{x}}$ represents the noise in position, representing sensor inaccuracies and unmodeled dynamics. $\boldsymbol{\eta}_k^{\boldsymbol{x}}$ is expressed in radians and follows normal distribution $\mathcal{N}\big(0, G(\boldsymbol{x}_k)Q_k^{\boldsymbol{x}}G(\boldsymbol{x}_k)^T\big)$. As the position noise intensity $Q_k^{\boldsymbol{x}}$ is expressed in square metres, a conversion to square radians is performed by $G(\boldsymbol{x}_k)$. Similarly, $\boldsymbol{\eta}_k^{\boldsymbol{c}}$ is the noise accounting for temporal and spatial variations of the speed of the water currents and follows $\mathcal{N}(0, Q_k^{\boldsymbol{c}})$. The intensity of the noise for the water currents is typically low and expressed in square m/s. Note that to obtain $F(\boldsymbol{x}_k)$, $G(\boldsymbol{x}_k)$ is scaled by Δt. This is required to convert velocities into the associated displacement in radians for a given Δt. For notational simplicity, the matrix parametrization by latitude is hereafter omitted $\big(F(\boldsymbol{x}_k) = F_k$, $G(\boldsymbol{x}_k) = G_k\big)$.

Several studies have also included errors in the vertical direction in the estimation process (e.g. see Anonsen and Hallingstad (2006)). This allows the navigation filter to compensate for biases in depth and range sensors, in addition to errors in estimating the tidal variation. However, to maintain a relatively low-dimension state vector, these errors can be assumed to be negligible when compared to the errors introduced by coarse DEMs. For example, the typical tidal range in the open ocean is up to one metre (Stammer et al., 2014).

4.2 Observation Model

Given the current energy budget of ALR6000, long-range TAN can be enabled only via the use of low-power low-informative sonars (Salavasidis, 2019). Such a sensor is the NOC's single-beam echo-sounder (Salavasidis et al., 2020) or the downward-facing DVL already available on the vehicle. Depending on the employed sonar, an on-board measurement can be either a scalar value r_k, when a single range-to-bottom measurement is available at a time, or a vector containing up to N_r ranges, $\boldsymbol{r}_k = [r_{k,1}...r_{k,N_r}]^T$, e.g. obtained by the $\{N_r = 4\}$-beam DVL. To compensate for the vehicle's attitude and the sonar's beam orientation, each range-to-bottom measurement $r_{k,i}$ can be projected into the 3-D space using the following expression:

$$\boldsymbol{\chi}_{r_{k,i}}^{NED} = \left[\begin{array}{c}\chi^N\\ \chi^E\\ \chi^D\end{array}\right]_{r_{k,i}} = \mathfrak{R}(\psi_k, \phi_k, \theta_k) \cdot \hat{\boldsymbol{q}}_{r_i} \cdot r_{k,i}, \tag{4}$$

where: $\mathfrak{R}(\psi_k, \phi_k, \theta_k)$ is the 3-D rotation matrix parametrized by the vehicle attitude at time k; $\hat{\boldsymbol{q}}_{r_i}$ is the direction unit vector of the i-beam; and $\boldsymbol{\chi}_{r_{k,i}}^{NED}$ is the 3-D location where the i-beam hits the seabed with respect to a sonar-attached North-East-Down (NED) reference frame. As $r_{k,i}$ is expressed in metres, $\boldsymbol{\chi}_{r_{k,i}}^{NED}$ will be also in metres.

To compute the absolute location where the i-beam hits the terrain, the north-east component of the projected range, $\begin{bmatrix} \chi_{r_{k,i}}^{N} \\ \chi_{r_{k,i}}^{E} \end{bmatrix}$, needs to be added to the vehicle's position $\boldsymbol{x}_k$. Note that as $\boldsymbol{x}_k$ is expressed in LL, the north-east offset needs to be first converted into the corresponding offset in LL:

$$\mathbf{x}_{k,i} = \begin{bmatrix} \mathbf{x}_{k,i}^{Lat} \\ \mathbf{x}_{k,i}^{Lon} \end{bmatrix} = \begin{bmatrix} x_k^{Lat} \\ x_k^{Lon} \end{bmatrix} + G_k \begin{bmatrix} \chi_{r_{k,i}}^{N} \\ \chi_{r_{k,i}}^{E} \end{bmatrix}, \tag{5}$$

where $\mathbf{x}_{k,i}$ is the LL location where the i-th beam hits the terrain.

The water depth at $\mathbf{x}_{k,i}$ can be calculated by adding the AUV's operating depth to the vertical component $\chi_{r_{k_i}}^{D}$ of the projected range vector of the i-beam:

$$z_{k,i} = \chi_{r_{k,i}}^{D} + d_k + \omega_{k,i}, \tag{6}$$

where: $z_{k,i}$ is the on-board i-th bathymetric observation expressed in metres; d_k is the operating depth expressed in metres; and $\omega_{k,i}$ is the observation noise assumed to be described by $\mathcal{N}(0, \sigma_{\omega_{k,i}}^2)$. The additive noise aims to represent errors in both sonar measurements and the derived operating depth.

Bathymetric observations, $\boldsymbol{z}_k = [z_{k,1}, ..., z_{k,N_r}]^T$, can be modeled using a function that relates $\boldsymbol{z}_k$ to the water depth at the location where the sonar beams hit the seabed:

$$z_{k,i} = \bar{h}(\mathbf{x}_{k,i}) + \omega_{k,i}, \tag{7}$$

where $\bar{h}(\mathbf{x}_{k,i})$ is a function representing the true water depth at $\mathbf{x}_{k,i}$. However, the true terrain $\bar{h}(\cdot)$ is unknown and only a terrain model can be available for use:

$$h(\mathbf{x}) = \bar{h}(\mathbf{x}) + m(\mathbf{x}), \tag{8}$$

where $m(\mathbf{x})$ is a spatially varying noise, assumed to follow $\mathcal{N}\big(0, \sigma_m^2(\mathbf{x})\big)$, representing errors in the vertical direction of the employed grid map.

To calculate the water depth in any inter-grid location of the bathymetric model, a bi-linear interpolation method is used. Even though higher order methods can also be used, such as bi-cubic or spline interpolation, they are likely to give a small performance gain at the expense of a large computational effort (Rees, 2000). Therefore, most TAN implementations have adopted relatively simple interpolation schemes, avoiding also overfitting and overestimating the terrain variability (Meduna, 2011).

Taking into account the statistical independence between the errors relating to the map, sonar and depth, an observation function can be formed as the following discrete-time model:

$$z_{k,i} = h(\mathbf{x}_{k,i}) + \zeta_{k,i}, \tag{9}$$

where $h(\cdot)$ is a non-linear function that relates the water depth value of the bathymetric map at $\mathbf{x}_{k,i}$ to the on-board observation $z_{k,i}$. The additive noise $\zeta_{k,i}$ represents the overall impact of the errors in the bathymetric map, together with those in the sensors used to obtain $z_{k,i}$. Therefore, $\zeta_{k,i}$ typically has a spatially variant noise intensity, depending on the quality of the employed DEM and the accuracy of bathymetric observations. Assuming the noise to be represented by $\mathcal{N}(0, \sigma_{z_{k,i}}^2)$ and taking into account the noise independence between bathymetric observations and the DEM, $\sigma_{z_{k,i}}$ is computed as:

$$\sigma_{z_{k,i}} = \sqrt{\sigma_{\omega_{k,i}}^2 + \sigma_m^2(\mathbf{x}_{k,i})}. \tag{10}$$

Further details concerning Equation (10) are given in Section 5.2, where the uncertainty in both bathymetric observation and map are quantified. Note that while more detailed error models can also be considered, unmodeled sources of error (e.g. errors in the vehicle's attitude, sensor misalignments, etc.) are expected to have a negligible effect when compared to the errors in the bathymetric map and thus can be ignored.

Given the present form of the observation function and assuming statistical independence between each bathymetric observation, the measurement likelihood function takes the following form:

$$p(\boldsymbol{z}_k|\boldsymbol{x}_k) = \alpha \cdot exp\left(-\frac{1}{2}\sum_{i=1}^{N_r} \frac{1}{\sigma_{z_{k,i}}}\Big(z_{k,i} - h(\mathbf{x}_{k,i})\Big)^2 \right), \tag{11}$$

where α is a normalisation constant.

4.3 Rao-Blackwellized Particle Filter (RBPF)

The vehicle's position, $\boldsymbol{x}_k$, is used as an input to the observation function $h(\cdot)$ in Equation (9). Given that the terrain below the vehicle can vary in a highly non-linear way (thus making $h(\cdot)$ a non-linear function), no unique mapping always exists between bathymetric observations and the vehicle's position. Hence, due to the non-linearity introduced by $h(\cdot)$, a non-linear estimator is required to approximate $\boldsymbol{x}_k$. In contrast, the speed of the water currents $\boldsymbol{c}_k$ is modeled as a linear function with additive Gaussian noise. Using Bayes' theorem, water currents can therefore be marginalised out from the state vector and can be optimally estimated using a KF. More precisely, the posterior density $p(\boldsymbol{x}_k, \boldsymbol{c}_k|\boldsymbol{z}_k)$ can be decomposed into $p(\boldsymbol{c}_k|\boldsymbol{x}_k, \boldsymbol{z}_k) \times p(\boldsymbol{x}_k|\boldsymbol{z}_k)$, where the linear and the non-linear sub-states are estimated using, respectively, a KF and a PF.

4.3.1 Sequential Importance Sampling Resampling

Suppose that N particles are used and $\boldsymbol{x}_k^i$ corresponds to the i-th particle with associated weight w_k^i at time k. The SISR particle filter is a recursive estimation algorithm that consists of two phases: prediction phase and update phase. The implementation of these two phases is given below.

SISR Prediction Phase: Particles are propagated forward in time using the decomposed state vector derived from Equation (3):

$$\boldsymbol{x}_{k+1}^i = \boldsymbol{x}_k^i + F_k^i \boldsymbol{u}_k + F_k^i \boldsymbol{c}_k^i + \boldsymbol{\eta}_k^{i,\boldsymbol{x}}, \tag{12}$$

where F_k^i is the conversion matrix, defined in Equation (3), parametrized by the latitude of the i-th particle. Note that the state transition model for the i-th particle uses as input the speed of the water currents $\boldsymbol{c}_k^i$, as estimated by the i-th KF. To compactly represent the noise in position, the uncertainty caused by the water currents is incorporated in $\boldsymbol{\eta}_k^{i,\boldsymbol{x}}$ as:

$$\boldsymbol{\eta}_k^{i,\boldsymbol{x}} \sim \mathcal{N}\big(0, F_k^i P_{k|k-1}^{\boldsymbol{c}} (F_k^i)^T + G_k^i Q_k^{\boldsymbol{x}} (G_k^i)^T\big), \tag{13}$$

where $P_{k|k-1}^{\boldsymbol{c}}$ is the water current error covariance at time k as predicted by the i-th KF and G_k^i is a matrix parametrized by the latitude of the i-th particle, see Equation (3).

SISR Update Phase: Once a bathymetric observation becomes available, particles are weighted using the likelihood function $p(\boldsymbol{z}_k|\boldsymbol{x}_k^i)$, derived from Equation (9), which provides a statistical measure of conformity (or matching) between bathymetric observations and the *a priori* terrain map. The SISR algorithm weights each particle recursively as:

$$w_k^i = w_{k-1}^i \cdot p(\boldsymbol{z}_k|\boldsymbol{x}_k^i). \tag{14}$$

where w_k^i is the weight of the i-th particle at time k. This expression is the direct result of using the prior density $p(\boldsymbol{x}_k|\boldsymbol{x}_{k-1})$ as the importance density in the SISR algorithm (Arulampalam et al., 2002).

Resampling: To mitigate the degeneracy effect (Carpenter et al., 1999), where after a few iterations only a very small number of particles contribute to the final state estimation, a resampling step has been introduced (Doucet et al., 2000). The concept behind resampling is to replace particles that have a small weight with others that have larger weight. However, frequently performing resampling can lead to the so-called particle collapse[3], where particles are concentrated in a very small region (Arulampalam et al., 2002). Therefore, to determine whether to trigger resampling, the effective sample size has been introduced as a measure of degeneracy (Nordlund, 2002):

$$N_{eff} = \frac{1}{\sum_{i=1}^{N} (w_k^i)^2} \tag{15}$$

where w_k^i is the normalised weight of the i-th particle. According to Nordlund (2002), an indication for severe degeneracy is when the effective sample size decreases below $N_{th} = \frac{2N}{3}$. In this work, the Systematic Resampling (SR) is utilised, as it is computationally more efficient than many other approaches (Li et al., 2015).

4.3.2 Kalman Filter

Each particle maintains an estimate of the speed of the water currents local to the vehicle by employing a KF. The KF model is given by:

$$\begin{aligned} \boldsymbol{c}_{k+1} &= \boldsymbol{c}_k + \boldsymbol{\eta}_k^{\boldsymbol{c}} \\ \boldsymbol{z}_k^c &= F_k \boldsymbol{c}_k + F_k \boldsymbol{u}_k + \boldsymbol{\eta}_k^{\boldsymbol{x}} \end{aligned}$$

where the process model is derived from Equation (3) and $\boldsymbol{z}_k^{\boldsymbol{c}}$ is the KF observation model corresponding to the position displacement during Δt, expressed in LL. According to Schon et al. (2005), the predicted $\boldsymbol{x}_{k+1}$ carries information about the speed of the water currents, therefore it is possible to interpret the residual $\boldsymbol{x}_{k+1} - \boldsymbol{x}_k$ as a KF observation ($\boldsymbol{z}_k^c$ is a fictitious measurement).

Kalman Filter Prediction The Kalman filter is a recursive Bayesian estimation technique that alternates between prediction and update phases. By modeling the speed of the water currents as a Gaussian random walk, the prediction for the i-th KF is given by:

$$\begin{aligned} \boldsymbol{c}_{k+1}^i &= \hat{\boldsymbol{c}}_{k+1|k}^i \\ &= \hat{\boldsymbol{c}}_{k|k}^i, \end{aligned} \tag{16}$$

and the predicted error covariance by:

$$P_{k+1|k}^{i,\boldsymbol{c}} = P_{k|k}^{i,\boldsymbol{c}} + Q_k^{\boldsymbol{c}}. \tag{17}$$

Kalman Filter Update The assumption that each particle represents the vehicle's true position provides the filter with the capability to estimate water currents using a KF. The employed i-th KF uses as observation $\boldsymbol{z}_k^i$ the LL distance travelled of the i-th particle during the Δt time-interval. This observation is then used to update the predicted speed of the water currents as (Teixeira, 2007; Schon et al., 2005):

$$\begin{aligned} \hat{\boldsymbol{c}}_{k|k}^i &= \hat{\boldsymbol{c}}_{k|k-1}^i + K_k^i \boldsymbol{y}_k^i \\ P_{k|k}^{i,\boldsymbol{c}} &= (I - K_k^i F_k^i) P_{k|k-1}^{i,\boldsymbol{c}} \end{aligned} \tag{18}$$

where $\boldsymbol{y}_k^i$ is the Kalman innovation and K_k^i is the Kalman gain that are computed, respectively, by Equations (19) and (20):

$$\boldsymbol{y}_k^i = \boldsymbol{z}_k^{i,\boldsymbol{c}} - (F_k^i \hat{\boldsymbol{c}}_{k|k-1}^i + F_k^i \boldsymbol{u}_k), \tag{19}$$

[3]Note that the particle collapse can be addressed by applying additional noise (so-called jitter) immediately after executing resampling, as reported in Salavasidis et al. (2016); Claus and Bachmayer (2015).

$$K_k^i = P_{k|k-1}^{i,\boldsymbol{c}}(F_k^i)^T \left[F_k^i P_{k|k-1}^{\boldsymbol{c}}(F_k^i)^T + G_k^i Q_k^{\boldsymbol{x}}(G_k^i)^T\right]^{-1}. \tag{20}$$

The innovation vector contains the difference between the actual/measured i-th particle displacement, $\boldsymbol{z}_k^{i,\boldsymbol{c}}$, and the predicted displacement, $\hat{\boldsymbol{z}}_k^{i,\boldsymbol{c}}$, using $\hat{\boldsymbol{c}}_{k|k-1}^i$.

4.4 RBPF Point Estimates

While the RBPF algorithm computes the state posterior probability distribution, an AUV typically needs an estimate of the vehicle's position that, for example, can be used by the control unit. Therefore, the navigation system is required to calculate a single vector, known as a point estimate, which serves as a 'best' estimate of the vehicle's state. Position and water current estimates can be obtained by taking the Minimum Mean Square Error (MMSE) estimate:

$$\begin{bmatrix} \hat{\boldsymbol{x}}_k \\ \hat{\boldsymbol{c}}_k \end{bmatrix} = \sum_{i=1}^{N} w_k^i \begin{bmatrix} \boldsymbol{x}_k^i \\ \boldsymbol{c}_k^i \end{bmatrix}, \tag{21}$$

where w_k^i is the normalised particle weight. Covariances associated with the estimates of both the vehicle's position and the speed of the water currents can be computed, respectively, by Equations (22) and (23):

$$P_k^{\hat{\boldsymbol{x}}} = \sum_{i=1}^{N} w_k^i (\boldsymbol{x}_k^i - \hat{\boldsymbol{x}}_k)(\boldsymbol{x}_k^i - \hat{\boldsymbol{x}}_k)^T, \tag{22}$$

$$P_k^{\hat{\boldsymbol{c}}} = \sum_{i=1}^{N} w_k^i \big(P_{k|k}^{i,c} + (\boldsymbol{c}_k^i - \hat{\boldsymbol{c}}_k)(\boldsymbol{c}_k^i - \hat{\boldsymbol{c}}_k)^T\big). \tag{23}$$

4.5 RBPF Initialisation

The recursive nature of the RBPF algorithm requires the filters to be initialised. This includes the initialisation of both the PF and the KF. In deep-water missions where a long descent is expected, the filter initialisation can be performed on the arrival of the first bathymetric observation to minimise the use of the on-board computer (Salavasidis et al., 2019). However, to properly initialise the filter an estimate of the vehicle's position with the associated uncertainty ($\{\hat{\boldsymbol{x}}_k, P_k^{\hat{\boldsymbol{x}}}\}$) is required. Therefore, a surface-acquired GPS position with the associated uncertainty is propagated in time (until initialisation is triggered) using a process equivalent to performing the KF prediction phase:

$$\begin{aligned} \hat{\boldsymbol{x}}_{k+1} &= \hat{\boldsymbol{x}}_k + F_k(\hat{\boldsymbol{x}}_k)\boldsymbol{u}_k, \\ P_{k+1}^{\hat{\boldsymbol{x}}} &= S_k P_k^{\hat{\boldsymbol{x}}} S_k^T + G_k Q_k^{\boldsymbol{x}} G_k^T, \end{aligned} \tag{24}$$

where $P_{k+1}^{\hat{\boldsymbol{x}}}$ is the time-updated position error covariance and S_k is the Jacobian matrix of the position propagation model with respect to the state vector $\hat{\boldsymbol{x}}$ evaluated at $\hat{\boldsymbol{x}}_k$. Note that the position propagation model above is a simplified version of Equation (3), as terms relating to water currents are eliminated. Once a bathymetric observation becomes available, the particle weight is initialised to $1/N$ and particles are spread using a normal distribution centred at $\hat{\boldsymbol{x}}_k$ with covariance $P_k^{\hat{\boldsymbol{x}}}$.

In terms of the water currents, KFs can be initialised using a zero-mean normal distribution with either an arbitrary (but realistic) noise intensity or intensity calculated *a priori* based on spatio-temporal data of the speed of the water currents (e.g. using ocean models or historical data).

4.6 RBPF Simplification

The number of KFs required to implement the RBPF in the current formulation is equal to the number of particles, N, used in the SISR algorithm. This dependence arises from the requirement to track the

vehicle's position considering the spherical shape of the Earth during long-range missions, instead of using a local static NED reference frame as a navigation origin. The latter approach ignores the Earth's curvature and hence introduces a navigation error that is a function of the deviation in latitude (the range between the vehicle's position and the origin of the navigation frame), the length of the circle of latitude (thus the latitude itself) and the vehicle's velocity (which for AUVs travelling slowly can be ignored). For short-range applications, this error is negligible provided the vehicle remains within a short distance from the navigation origin and avoids extreme latitudes. However, the error becomes substantial during long-range missions and thus the navigation system needs to take into account the Earth's approximately spherical shape.

One simple method to address this problem is by moving the origin of the NED navigation frame along the vehicle and keeping track the frame's origin in LL, which would coincide with the vehicle's position. This concept is exploited in Equation (12), which propagates each particle using the computed displacement in a local moving NED reference frame attached to its prior position. This propagation is then converted through F_k into the corresponding displacement in LL. Considering a minimum error in latitude (as the origin of the NED frame coincides with prior particle position) and the insignificant displacement during Δt, compared to the radius of Earth, the error introduced by locally flattening the spherical shape of Earth can be considered negligible even in high latitudes.

One disadvantage of dealing directly with the approximately spherical shape of Earth is the non-linearity introduced via the matrix F_k in the process model, see Equation (3). This non-linearity leads the RBPF to employ as many KFs as the number of particles the SISR algorithm uses. However, a simplification can still be made. Instead of placing the origin of each particle's reference frame at the prior position of the particle, the origin is now attached to the prior point estimate $\hat{\boldsymbol{x}}_k$. Then, all NED reference frames are collapsed into a single frame. This has an effect on the computational complexity of the RBPF as it allows the RBPF to use a single KF for all particles (since F_k will be identical for all particles).

The technique described above, which avoids the use of multiple KFs, inevitably re-generates errors related to the use of a NED navigation frame, although the frame moves along with the particles. This error depends on the spatial distribution of the particles (particles' deviation from the origin of the navigation frame) and the particle position in latitude. However, considering that F_k varies slowly with time, the frame-related error will remain insignificant even at high latitude zones and thus can easily be captured by increasing the process noise intensity. Simple calculations show that for the extreme scenario of operating at 85° latitude, a particle 50 km from the origin of the navigation frame will introduce a propagation error of less than 4.5 cm in the east direction ($\Delta t = 1$ sec).

4.7 RBPF Algorithm

The required steps to perform an iteration of the RBPF are shown in the pseudocode given in Algorithm 1. At each iteration, the algorithm receives as input: the particles (with the corresponding weights) and the speed of the water currents from the preceding iteration; the AUV north-east speed $\boldsymbol{u}_k$; and the observation z_k (if available). The number of particles N and the resampling threshold N_{th} are constants and hence not treated as input arguments. Outputs of the algorithm are the posterior particles with the associated weights, the KF estimates and the filter point estimates. All steps listed in the algorithm are as described in this section.

4.8 RBPF Reset Capability

A major disadvantage of the current form of the RBPF algorithm is its inability to recover from false convergence. The filter can diverge for various reasons, where the three most common are: a) the use of an unreliable bathymetric map; b) fusing low-accuracy bathymetric observations; and c) traversing areas with low topographic information. If position estimates suddenly deviate from the vehicle's true position (e.g. the filter becomes stuck in a local minimum), there is no functionality that would enable re-convergence. The

Algorithm 1 Rao-Blackwellised Particle Filter

Inputs: $\{\boldsymbol{x}_{k-1}^i, \boldsymbol{c}_{k-1}^i, w_{k-1}^i\}_{i=1}^N, P_{k-1|k-2}^{\boldsymbol{c}}, \boldsymbol{u}_k, \boldsymbol{z}_k$
Outputs: $\{\boldsymbol{x}_k^i, \boldsymbol{c}_k^i, w_k^i\}_{i=1}^N, P_{k|k-1}^{\boldsymbol{c}}, \hat{\boldsymbol{x}}_k, \hat{\boldsymbol{c}}_k, P_k^{\hat{\boldsymbol{x}}}, P_k^{\hat{\boldsymbol{c}}}$

function RBPF
  if Initialisation Needed **then**
    $[\{\boldsymbol{x}_{k-1}^i, \boldsymbol{c}_{k-1}^i, w_{k-1}^i\}_{i=1}^N\}] \leftarrow$ **Initialise RBPF** ▷ Section 4.5
  end if
  $[\{\boldsymbol{x}_k^i\}_{i=1}^N] \leftarrow$ **SISR Prediction Phase** ▷ Equation (12)
  $[\{\hat{\boldsymbol{c}}_{k-1|k-1}^i\}_{i=1}^N, P_{k-1|k-1}^{\boldsymbol{c}}] \leftarrow$ **KF Update Phase** ▷ Equation (18)
  if $\boldsymbol{z_k}$ available **then**
    $[\{w_k^i\}_{i=1}^N] \leftarrow$ **SISR Update Phase** ▷ Equation (14)
    $w_k^i \leftarrow \frac{w_k^i}{\sum_{l=1}^N w_k^l}$, for $i = 1 \rightarrow N$ ▷ Weight Normalization
  end if
  $[\hat{\boldsymbol{x}}_k, \hat{\boldsymbol{c}}_k, P_k^{\hat{\boldsymbol{x}}}, P_k^{\hat{\boldsymbol{c}}}] \leftarrow$ **Point Estimates** ▷ Section (4.4)
  $N_{eff} \leftarrow$ **Effective Sample Size** ▷ Equation (15)
  if $N_{eff} < N_{th}$ **then**
    $[\{\boldsymbol{x}_k^i, \boldsymbol{c}_k^i\}_{i=1}^N\}] \leftarrow$ **Systematic Resampling** ▷ See Li et al. (2015)
    $w_k^i \leftarrow 1/N$, for $i = 1 \rightarrow N$ ▷ Particle Weight Reset
  end if
  $[\boldsymbol{c}_k^i, P_{k|k-1}^{\boldsymbol{c}}] \leftarrow$ **KF Prediction Phase** ▷ Equations (16) & (17)
end function

detection of the filter's divergence and enabling re-convergence are two challenging and application-dependent problems, resulting in the unavailability of a general solution.

A common approach to partly address PF failures is the injection of extra particles (Thrun et al., 2005), where new particles are generated when a performance measure indicates suspicious filter behaviour. The RBPF performance can be monitored to some extent via the likelihood function $p(\boldsymbol{z}_k|\boldsymbol{x}_k)$ or, alternatively, through the summation of the non-normalised particle weights, $W_k = \sum_1^N w_k^i$ (Teixeira et al., 2017). If this measure drops below a certain threshold, or violates certain conditions, new particles are generated either using a uniform distribution over the state space or drawing particles at locations dictated by the likelihood function.

This work uses concepts similar to those described above for monitoring the behaviour of the filter but the full RBPF reset is considered (rather than injecting new particles). To form the condition that triggers a reset, the summation of the non-normalised particle weights (W_k) is low-pass filtered (smoothed) in time using two distinct smoothing factors (α). This forms the fast-varying term (W_k^f) and the slow-varying term (W_k^s), where W_k is smoothed using a high smoothing factor (α_f) and a low smoothing factor (α_s), respectively. This smoothing process can be expressed mathematically as:

$$W_k^{\{s,f\}} = W_{k-1}^{\{s,f\}} + \alpha_{\{s,f\}}\Big(W_k - W_{k-1}^{\{s,f\}}\Big), \tag{25}$$

where $\{s, f\}$, as a superscript for W or a subscript for α, represents the fact that the equation holds for both s and f.

Provided that the smoothing factors are selected such that $0 < \alpha_s \ll \alpha_f < 1$, the filter reset is triggered when the following condition becomes active (Thrun et al., 2005):

$$\beta - \frac{W_k^f}{W_k^s} > 0 \tag{26}$$

where β is a pre-defined constant ($\beta \leqslant 1$). By reducing β, more conservative reset schemes are considered and hence the RBPF becomes more tolerant to bathymetric mismatches. In reality, the selection of β must take into account the likelihood function. If σ_z becomes very high, both W_k terms become smooth and insensitive, which will result in occasional resets. The opposite can happen when the likelihood function is very narrow/sharp and W_k^f becomes very reactive causing resets to be more likely.

For the RBPF reset itself, water currents (Kalman filter) are initialised by a zero-mean normal distribution with intensity as discussed in Section 4.5. In terms of the SISR initialisation, particles are initialised taking into consideration the existing particle spatial distribution, which is approximated by a normal distribution with first and second moments defined by the filter's point estimates, see Equations (21) and (22). Moreover, the estimated error covariance is stretched by a constant factor, which allows the re-initialised particles to cover a larger area, increasing the probability of enclosing the true position of the vehicle.

In general, it may be preferable to perform regular filter resets when operating in uncertain environments (e.g. low-quality bathymetric map and/or unknown environmental conditions). This helps the filter to avoid becoming stuck in a local minimum, for which a late reset may be insufficient for recovery (considering the particular reset mechanism described above). Conversely, missions well-suited for TAN can benefit from using a lower β that will reduce the chance of triggering a reset, after which the filter may require some time to re-converge to a solution. Moreover, frequent filter resets can also be risky as false convergence may occur. In turn, this may trigger repeatedly further resets, causing large discontinuities in estimating the vehicle's position. Therefore, precautionary safety measures need to be taken to avoid multiple consecutive resets, such as setting a minimum elapsed time between filter re-initialisations.

5 Water-Relative Navigation: Experimental Environment & RBPF Configuration Setup

To challenge the developed TAN algorithm and explore its capabilities, this section forms a test environment where the algorithm's *a-posteriori* performance can be quantified in realistic and demanding conditions. For this purpose, it is believed that the field data obtained from the DynOPO expedition comprise a unique test case given the mission range and the associated challenges (e.g. the motion along bathymetric contour lines which results at a low-informative along-track terrain morphology).

To further increase the complexity, it is assumed that the ALR6000 navigates relative to the water flow, which is achieved by replacing the BR velocity with speed estimates derived from the propeller RPM. Section 2 has already demonstrated that the ALR6000's on-board navigation system can experience errors exceeding 40 km in only a three-day duration mission in the Orkney Passage, where water currents can be of the same order of magnitude as the ALR6000's forward speed.

Motivated by the current focus of the oceanographic community to develop bathymetric maps of world's oceans primarily of 400 m resolution (Mayer et al., 2018), this section also assesses the TAN performance relative to the map resolution. Therefore, three maps of 100 m, 200 m and 400 m resolution have been generated by removing bathymetric data and degrading the accuracy of the original 50 m resolution map.

5.1 Generating Coarse Bathymetric Maps of the Orkney Passage

The original bathymetric maps of the Orkney Passage were generated using data collected on 12 scientific cruises on RRS JCR, RV Polarstern, and BIO Hesperides between 1989 and 2017. The data were collected during bathymetry surveys using Kongsberg EM120 (Hesperides and JCR up to 2011), EM122 (JCR after 2011), and Atlas Hydrosweep (Polarstern) multibeam sonars. The grids and links to the source data are available through the UK Polar Data Centre (Abrahamsen, 2020). Due to the steep-terrain morphology and depths sometimes exceeding 5000 m, a surface-made high-resolution map is not possible (Jensen et al.,

2011; Mayer et al., 2018). The generated map first has been gridded into squares of approximately 50 m per side ($\approx 0.052°/0.025°$ in longitude and latitude, respectively) with the centre of each grid cell to take the mean value of all multi-beam readings that fell within that particular grid cell. In areas lacking depth measurements, gaps of up to ten grid cells wide have been filled using spline interpolation. Note that the resolution of the constructed map can be sometimes higher than the resolution that the ship's multi-beam sonar can actually provide. For example, for a beam width $1° \times 1°$ (transmit $\times$ receive) and depths over 5000 m, the radius of the (approximately circular) footprint is over 40 m.

Even though the resolution of the original map is significantly lower compared to the resolution of the typical maps used for TAN in the literature (Melo and Matos, 2017), it is of better quality when compared to the currently available bathymetric models of world's oceans (and potentially the ones to be constructed in the near future). This is mainly because the map does not contain large bathymetric gaps, areas where depth values are computed using interpolation, causing large uncertainties. Therefore, DEMs of 100 m, 200 m and 400 m resolution have been generated such that the TAN performance is evaluated with respect to DEMs of comparable quality to those obtained via world ocean models. The driving factor for selecting these particular resolutions is the GEBCO Seabed 2030 project, which will map 98.3% of the oceans with similar resolutions.

The generated lower resolution DEMs are not simply smoothed counterparts of the initial 50 m DEM. Instead, for each new grid cell of the specified resolution, all but one underlying grid cell of 50 m size are randomly discarded. Then, the depth of the randomly selected remaining cell is used to compose the depth value of the larger grid cell. This sub-sampling process aims to take into account the sparse bathymetric data in global bathymetric maps, where depth values are typically constrained by little or no bathymetric data (e.g. by applying interpolation). For example, approximately 82% of the grid cells in the GEBCO_2014 digital bathymetric model of the world's oceans have no actual depth measurements (Weatherall et al., 2015).

Figure 4 shows an example of the constructed DEMs. M42 and M44 of the DynOPO deployments took place in areas sharing the same bathymetric grid map, whereas M41 has its own grid. Figure 5 presents

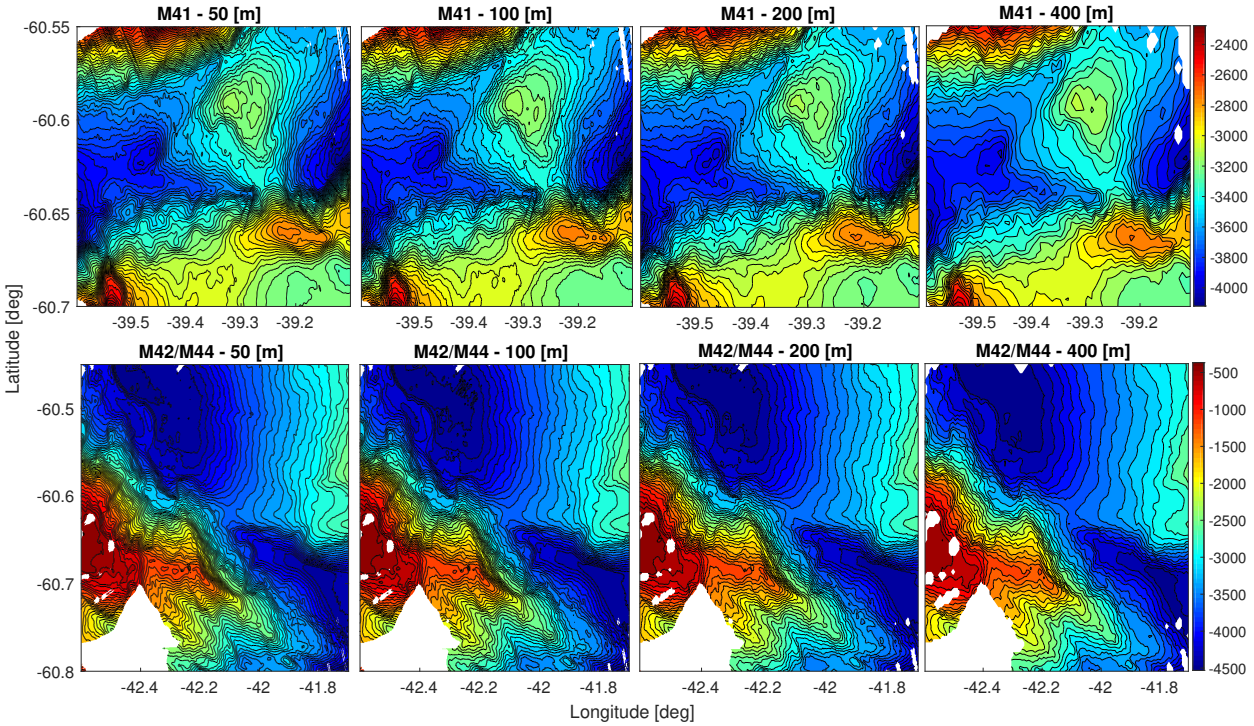

Figure 4: Contour plots of DEMs for M41, M42 and M44 with varying resolution and quality.

the magnitude of the introduced errors via the sub-sampling process. The accuracy degradation is shown through the absolute vertical deviation of the created DEMs from the original 50 m resolution map. The comparison for each grid cell is performed after re-scaling the newly generated maps back to 50 m resolution via bi-linear interpolation. The bi-linear interpolation is selected as it is used by the TAN algorithm to compute depth values in inter-grid locations, see Section 4.2.

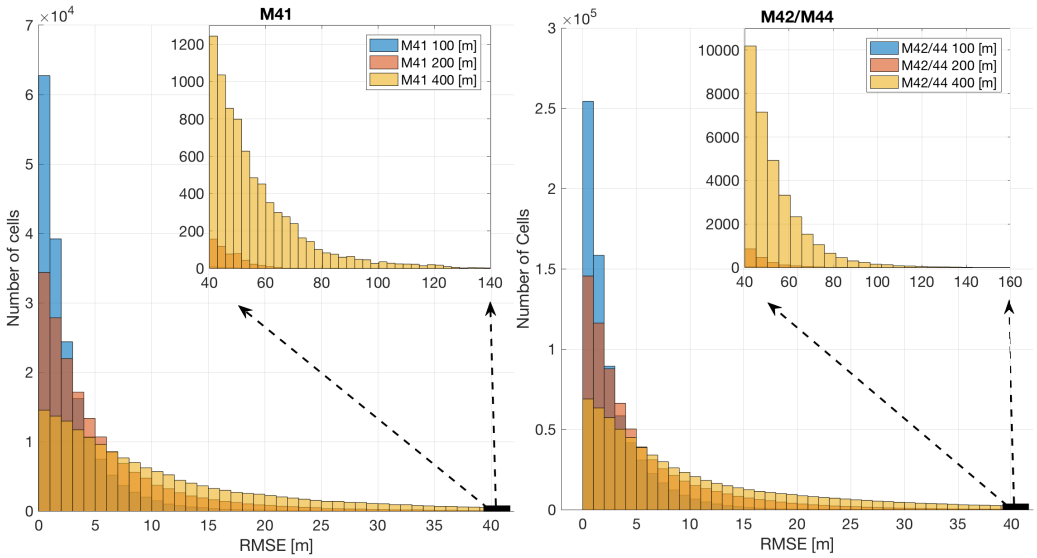

Figure 5: Calculating the grid cell vertical deviation, expressed in RMSE, of DEMs of 100 m, 200 m and 400 m resolution from the original 50 m resolution map. To perform a direct terrain value comparison, low-resolution DEMs are re-scaled to 50 m resolution via bi-linear interpolation. M42 and M44 took place in regions coved by the same bathymetric map.

5.2 RBPF Likelihood Function

For the DynOPO deployments, two Teledyne RDI 300 kHz DVLs were used (one facing downwards and another facing upwards). To avoid interference between them, the sensors were configured into 0.5 Hz pinging mode with pings offset by one second. However, close inspection has shown that it is possible for the strongest echo from one transducer to be received by the sidelobes of another transducer of the same module. Therefore, to avoid feeding the TAN algorithm with dubious range measurements, the interference is managed by a computationally simple and effective approach as detailed in Salavasidis (2019). Each valid beam-range $\{r_{k,i},\ i = 1, ..., N_r\}$, provided by the $\{N_r = 4\}$-beam downward-facing DVL, which successfully passes the interference test is projected into the 3-D space. Valid range measurements are then used to form individual bathymetric observations via Equation (6), which adds the vertical component of the projected vector to the computed operating depth, d_k.

The uncertainty for each valid i-th bathymetric observation at time k is modeled by $\mathcal{N}(0, \sigma^2_{\omega_{k,i}})$, which is the distribution that aims to capture errors in the range measurements, the hydrostatic pressure and the pressure-to-depth conversion. Given the statistical independence between the DVL beam-range error and the error in the vehicle's derived depth, the observation uncertainty $\sigma_{\omega_{k,i}}$ can be computed using:

$$\begin{aligned} \sigma^2_{\omega_{k,i}} &= \sigma^2_s(r_{k,i}) + \sigma^2_d(d_k) = (K_1 \cdot r_{k,i})^2 + (K_2 \cdot d_k)^2 \\ &= (3.3 \times 10^{-3} \cdot r_{k,i})^2 + (3.3 \times 10^{-4} \cdot d_k)^2, \end{aligned} \tag{27}$$

where $\sigma_s(r_{k,i})$ refers to the standard deviation of the i-th beam and $\sigma_d(d_k)$ represents the uncertainty of

the derived depth, both varying proportionally to the measured quantity (Salavasidis, 2019). The constants K_1 and K_2 in Equation (27) are defined empirically using data collected by these particular sensors from many past experiments. The maximum expected depth error using the Sea-Bird SBE 52-MP CTD probe is approximately 0.1% of d_k ($K_1 = 3.3 \times 10^{-3}$ using the 3-σ rule of thumb), while the maximum range error for the 300 kHz Teledyne RDI DVL corresponds to approximately 1% of the measured beam-range ($K_2 = 3.3 \times 10^{-4}$ using the 3-σ rule of thumb). When operating extremely close to the ocean bottom and σ_s almost vanishes, a beam-range invariant term should be added in the above uncertainty model (Salavasidis, 2019).

In the case of bathymetric map uncertainty, the models described in Section 3.1.2 are used. Errors of *Type 1* are considered to be sufficiently represented by $\sigma_h(\mathbf{x}_{k,i})$ and Equation (2), which quantifies the uncertainty due to the inaccurate ship-based sonar measurements. For errors of *Type 2* and *Type 3*, which have mostly been introduced by the gridding algorithm and the sub-sampling process described in Section 5.1, it is assumed that the spatially-invariant error model $\mathcal{N}(0, \sigma_g^2)$ is sufficient. Furthermore, assuming statistical independence between the ship-based sonar errors and the errors introduced by the gridding algorithm, the DEM's uncertainty can be computed as:

$$\sigma_m^2(\mathbf{x}_{k,i}) = \sigma_h^2(\mathbf{x}_{k,i}) + \sigma_g^2, \tag{28}$$

where $\sigma_m(\mathbf{x}_{k,i})$ is the standard deviation (in metres) assigned to the depth value of the grid map at the location of interest $\mathbf{x}_{k,i}$ (i-th DVL beam hitting the ocean bottom, see Equation (5)).

The noise intensity σ_g must be selected to be sufficiently high in order to avoid underestimating errors in coarse DEMs. Despite the degraded performance that TAN may experience in well-mapped regions, it is more important to treat a disputed terrain feature more conservatively and avoid extracting localisation information that may eventually lead to filter divergence. For a conservative TAN setup, $\sigma_g = 50$ m is assumed sufficient for the original DEM. To take into account the errors introduced by the grid sub-sampling, $\sigma_g = 100$ m is used for the 100 m resolution DEM. Due to the apparent error increase in DEMs of 200 m and 400 m resolution (see Figure 5), σ_g is increased to 150 m. Note that although these intensities may overestimate the actual map error, they enable a conservative filter convergence in position and water currents.

Given the statistical independence between the map errors and errors in bathymetric observations, σ_{z_k} can be computed by Equation (10), thus enabling the use of the likelihood function given by Equation (11).

5.3 RBPF Reset Capability Setup

Using TAN in areas with strong time-varying water currents and where only coarse bathymetric maps are available can be very challenging. This can become further complicated when the vehicle operates along bathymetric contour lines (as happened during the DynOPO experiments, e.g. see Figures 3), which results in a partially observable state (the speed of the along-track water currents is not observable). To deal with such a situation, the RBPF algorithm uses the reset capability of Section 4.8. Besides recovering from the most common causes for the filter's divergence (e.g. becoming stuck in a bathymetric local minimum), the reset capability is particularly important because it allows recovery from a prolonged time drift caused by the unobservable state parameters while operating along bathymetric contour lines. The re-initialisation function effectively exploits the achieved particle distribution, which is likely of an ellipsoid shape with the major axis along the bathymetric contour line, and stretches this spread further by artificially increasing the filter's search area along the uncertain/critical direction.

The smoothing factors α_s and α_f (see Section 4.8) have been set to 0.005 and 0.05, respectively. Moderate modification of these parameters is not expected to significantly affect the performance of TAN. In the case of β, this value depends on the quality of the bathymetric map. For the most reliable DEM (original 50 m resolution map), the most conservative reset scheme is used to avoid frequent resets ($\beta = 0.85$). As the quality of the map decreases and the likelihood function flattens, resulting in higher insensitivity to

bathymetric mismatches, β must increase. In this work, the maximum value of β used is 0.95 for the DEM of 400 m resolution, as the filter is mostly likely to become stuck in local minima or deviate when the vehicle operates along bathymetric contour lines. For the 100 m and 200 m resolution DEMs, a middle value is used ($\beta = 0.9$). In terms of the re-initialisation itself, the particle spatial distribution is approximated by a normal density with first and second moments defined by the filter's point estimates. When a reset is required, the covariance of the normal particle distribution is scaled by a factor of five. Moreover, to avoid consecutive re-initialisations, a reset is allowed only if the RBPF algorithm has executed a minimum of 100 observation updates from the previous reset and hence providing an elapsed time between resets so that the filter can potentially re-convergence to a solution. Note that no considerable effort has been directed towards tuning the reset parameters. Instead, they have been defined following the simple rationale of avoiding very frequent resets, whilst allowing re-initialisations when massive drops in the summation of the particle weights occur. Moreover, it is possible to develop more efficient re-initialisation strategies by taking into account, for example, the shape of the posterior density function for determining whether a reset is required and where particles need to be re-spread. However, to maintain computational simplicity, the described reset strategy is used and alternative options will be investigated in future research.

5.4 RBPF Parameters

To enable TAN, the parameters of the RBPF algorithm needs to be specified (e.g. the number of particles N and the intensity of the process noise). It is known that an increased number of particles allows the filter to more accurately represent the posterior density function (Arulampalam et al., 2002). However, this comes at the expense of increased computational power requirements and hence an extremely large number of particles would be infeasible. The computer installed on ALR6000 can handle over 20000 particles in real-time. However, a number of particles beyond a certain value would not result in a substantial performance increase, considering the dimension of the estimation problem (Claus and Bachmayer, 2015; Salavasidis et al., 2018). To assure filter robustness against false convergence, $N = 10000$ is selected. This relatively large number of particles allows the filter: a) to effectively cover the state space when the filter is highly uncertain (e.g. while initialising the filter after a long descent); b) to achieve fast convergence to the vehicle's true position; and c) to maintain this convergence throughout the entire mission length. It is important to highlight that the filter is assumed to be convergent to a solution when the posterior probability density function is unimodal (or nearly unimodal) and the filter's uncertainty satisfies a pre-selected uncertainty threshold (Meduna, 2011).

The appropriate selection of the process noise covariance matrix ($Q_k^{\boldsymbol{x}}$ and $Q_k^{\mathbf{c}}$) is of high importance for the RBPF algorithm. The intensity of the process noise must be carefully selected to take into account unmodeled dynamics and to provide sufficient, but not too large, particle spread. Avoiding an unnecessarily high intensity for the process noise is important for mainly two reasons: a) the number of particles may be insufficient to cover a large search area; and b) potential terrain similarities within the search area may lead the filter to become trapped in local minima. To distinguish the vehicle's operation phase between descending, hence drifting due to the systematic errors introduced by the unknown water currents, with the phase where range measurements are available and TAN is enabled, two process noise intensities are used: a) high intensity, $\boldsymbol{\eta}_k^{\boldsymbol{x}} \sim \mathcal{N}(0, Q_k^{\boldsymbol{x},H})$, while descending; and b) low-intensity, $\boldsymbol{\eta}_k^{\boldsymbol{x}} \sim \mathcal{N}(0, Q_k^{\boldsymbol{x},L})$, when TAN is active and estimates the water currents. This distinction allows Equation (24) to use a process noise of a sufficient intensity, $Q_k^{\boldsymbol{x},H}$, such that the time propagation of the vehicle's position, $\{\hat{\boldsymbol{x}}_k, P_k^{\hat{x}}\}$, does not underestimate the long-term effect of the water currents when descending. Then, once TAN is enabled, the noise intensity drops as the RBPF algorithm deals explicitly with water currents[4].

Table 2 specifies the process noise intensity for the two operation phases at each DynOPO mission. $Q_k^{\boldsymbol{x},H}$ depends on the time required for the vehicle to approach the ocean bottom and enable TAN. M41 began in an area with approximately 3400 m water depth. Assuming on average a drift of 0.15-0.2 m/s caused by water currents (which can be obtained from spatio-temporal historical data of the area) and descent

[4]Recall that despite the bottom-tracking regime, only propeller-based speed estimates are used in order to imitate mid-water column operations.

duration of 3.5 hours (including the time required for the compass calibration during descent, see Salavasidis et al. (2019)), a noise intensity of 36 m^2 per direction is considered sufficient to capture the localisation error (assuming $\Delta t = 1$ sec). Since the water depth at the beginning of M42 and M44 is shallower (approximately

Table 2: Process noise covariance matrix for each DynOPO mission

	M41	**M42/M44**	**Units**
$Q_k^{\boldsymbol{x},H}$	$\Delta t^2 \cdot diag(36, 36)$	$\Delta t^2 \cdot diag(16, 16)$	m^2
$Q_k^{\boldsymbol{x},L}$	$\Delta t^2 \cdot diag(0.25, 0.25)$	$\Delta t^2 \cdot diag(0.25, 0.25)$	m^2
$Q_k^{\boldsymbol{c}}$	$\Delta t^2 \cdot diag(10^{-6}, 10^{-6})$	$\Delta t^2 \cdot diag(10^{-6}, 10^{-6})$	m^2/s^2

500 m and 1150 m, respectively), the required time to reach the ocean bottom is approximately 1 hour and 1.5 hours, respectively (including the time for compass calibration). The systematic error introduced by water currents in M42 and M44 can therefore be captured with an intensity of 16 m^2 per direction. Once the vehicle approaches the ocean bottom and bathymetric observations become available, TAN is triggered and the process noise intensity reduces to 0.25 m^2 per direction, which is assumed sufficient to represent modeling and motion sensor errors.

It can be inferred from Figure 3 that M41 is governed by relatively weak water currents, compared to M42 and M44. Prior knowledge about the speed and the variability of the water currents can significantly benefit the filter design. When strong water currents are not to be expected, the dimension of the navigation filter can be reduced by deactivating the KF (Salavasidis et al., 2019). Such a simplification can be important since fewer filter parameters need to be considered. For example, failing to estimate water currents may lead the RBPF into estimation divergence. Prior analysis has shown that the primary reason for the RBPF to fail is when water current estimates become unrealistic. However, water currents cannot always be predicted *a priori*, particularly during long-range missions where they may substantially vary with space and time. When an AUV is operating under unexpectedly strong water currents for long periods of time, the 2-D filter will almost certainly fail. As a result, the full RBPF with a conservative setup for $Q_k^{\boldsymbol{c}}$ is preferred (see Table 2). In addition, enabling the filter with the reset capability can be important to avoid unrealistic estimates of the water currents.

In terms of the initial position uncertainty, it is assumed that the GPS error is normally distributed around the GPS measurement with 5 m standard deviation in the north and east directions. The KF is initialised assuming that water currents are adequately represented by a zero-mean normal distribution and 0.1 m/s standard deviation per direction.

6 TAN using Water-Relative Velocity: Results

Given the experimental environment and the RBPF setup described previously, this section evaluates the TAN performance with respect to a bathymetric map of varying resolution whilst operating in areas with strong and time-varying water currents. A set of ten independent runs for each of the three DynOPO missions are performed to numerically assess: a) the filter accuracy in estimating the state vector, in comparison to the DR position, the vehicle's baseline position (USBL measurements) and the baseline speed of the water currents[5]; b) the filter precision, by reporting the estimation repeatability among the ten independent runs[6]; and c) estimation uncertainty, by comparing the localisation error with the respective 3-σ confidence bounds. However, note that lacking constant access to the vehicle's ground truth position complicates the evaluation

[5]It have been stated that the TAN algorithm estimates the speed of the water currents but, in reality, it estimates the velocity bias in the north and east directions. However, it can be assumed that biases in velocity are in the majority introduced by the ocean currents, given the successful real-time implementation of sensor calibration and error correction techniques (Salavasidis et al., 2019). The speed of the water currents is calculated by subtracting the propeller-derived speed in the north and east directions from the corresponding BR speed provided by the downward-facing DVL (thus possible only when BR velocity is available).

[6]Filter results are considered repeatable when all runs produce nearly identical point estimates for each k time-step.

process. In fact, USBL measurements are not evenly distributed in time (e.g. see Figure 3) and a large number of the measurements are mainly concentrated in locations where TAN is initialised (beginning of missions) or the vehicle performs sharp manoeuvres (e.g. in M42). As a consequence, the average TAN error appears to be larger (e.g. because of the use of low-accuracy tilt sensors which causes poor beam-range projection) compared to what it is in reality.

The discussion given here is mainly concentrated around M44, which is the longest mission and is considered the most challenging for implementing TAN, because of both the presence of strong time-varying water currents and the vehicle traversing areas with low topographic information (along-track). The interested reader is referred to Salavasidis (2019) where the complete description of the results can be found. Table 3 introduces notation frequently used in the following subsection. A detailed description of the evaluation metrics used can be found in Salavasidis et al. (2019).

Table 3: Notations

k^*	Time of the last USBL measurement while TAN enabled
$RMSE^{10}$	$RMSE$ averaged over time and 10 runs (with enabled TAN)
$RMSE^{10}_{k^*}$	$RMSE$ averaged over 10 runs at $k = k^*$
TAN 50	TAN using 50 m resolution DEM
TAN 100	TAN using 100 m resolution DEM
TAN 200	TAN using 200 m resolution DEM
TAN 400	TAN using 400 m resolution DEM

6.1 Mission 41, 42 & 44

The performance of the TAN algorithm after executing ten independent runs for each mission and each resolution of the bathymetric map is summarised in Table 4. Also, Figure 6 shows the estimated path of the vehicle whilst using speed estimates derived from the rotation speed of the propeller of the vehicle, as obtained by the TAN algorithm (averaged over all mission runs) and DR estimation (DR*).

Table 4: TAN results for M41, M42 and M44.

		TAN 50	**TAN 100**	**TAN 200**	**TAN 400**	**DR**	**DR***	**Units**
M41	$RMSE^{10}$	0.24	0.33	0.38	0.89	1.2	2.6	km
	$RMSE^{10}_{k^*}$	0.2	0.39	0.57	0.65	1.1	2.9	km
	$RMSE^{10}_{k^*}$ w.r.t. DT	0.39	0.77	1.12	1.25	2.1	5.7	%
M42	$RMSE^{10}$	0.29	0.3	0.48	1.43	0.4	13	km
	$RMSE^{10}_{k^*}$	0.14	0.16	0.2	1.6	0.3	33	km
	$RMSE^{10}_{k^*}$ w.r.t. DT	0.12	0.14	0.17	1.6	0.3	28.4	%
M44	$RMSE^{10}$	0.31	0.3	0.3	0.78	0.4	14.5	km
	$RMSE^{10}_{k^*}$	0.07	0.18	0.23	0.24	0.4	40	km
	$RMSE^{10}_{k^*}$ w.r.t. DT	0.04	0.1	0.13	0.13	0.2	21.9	%

w.r.t.: with respect to.

Despite the progressive decrease of the map resolution via sub-sampling, no filter divergence occurred during the ten independent runs for each mission. In particular during M41, the averaged RMSE over the mission length (after filter initialisation and before ascending) and over the ten runs, $RMSE^{10}$, is 0.24 km for TAN 50, 0.33 km for TAN 200, 0.38 km for TAN 200 and 0.89 km for TAN 400, whereas the corresponding DR error is 1.2 km using BR velocity and 2.6 km when deriving speed through the propeller RPM. When considering the average localisation error with respect to the last available USBL measurement before ascending, $RMSE^{10}_{k^*}$, TAN results vary, subject to the resolution of the employed bathymetric map, from 0.2 km to 0.65 km

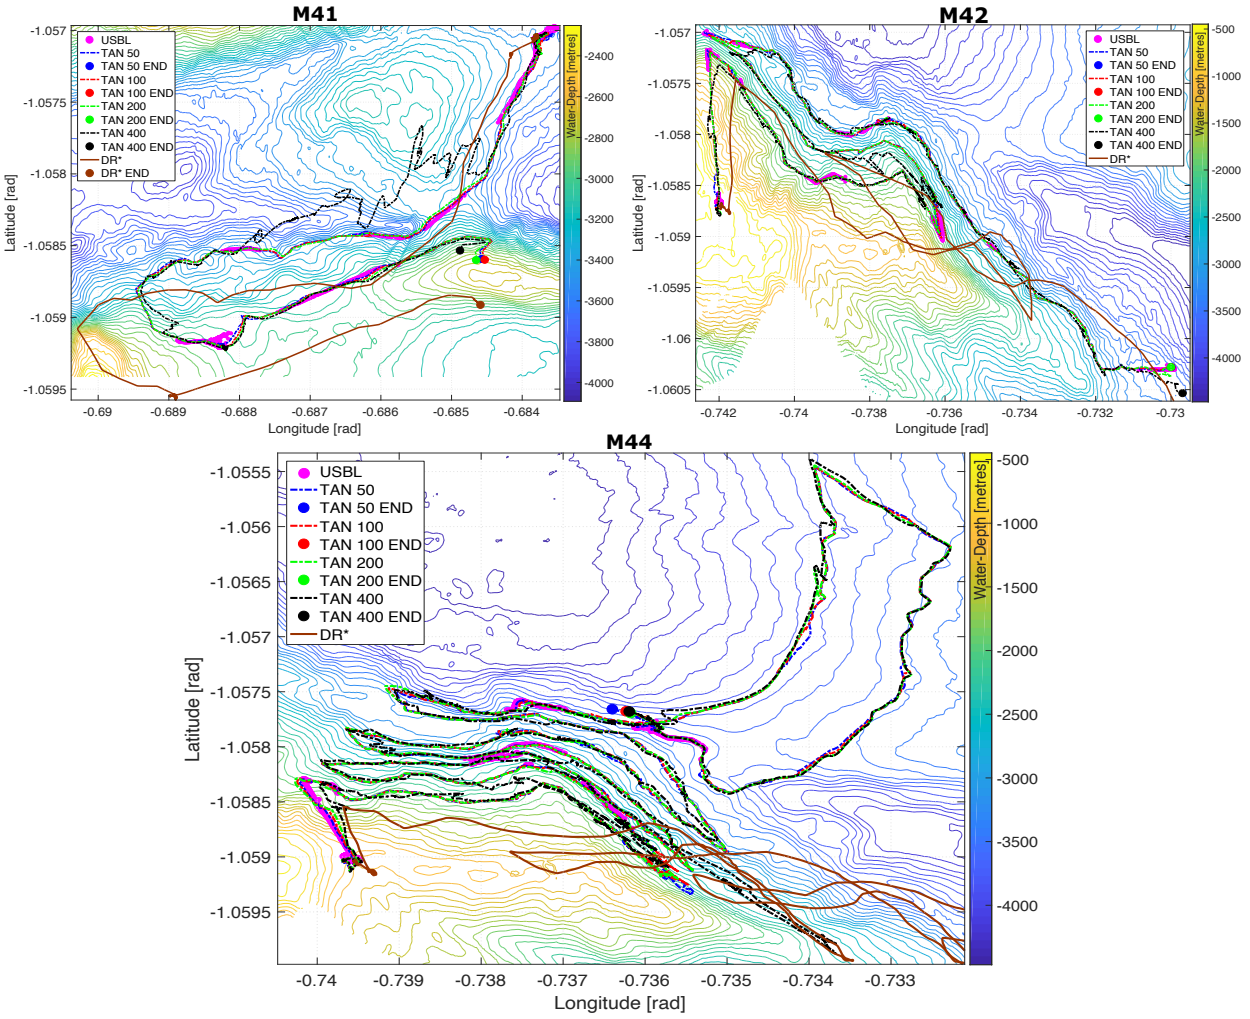

Figure 6: The path of the AUV, averaged over the ten independent runs, as estimated using DR* and TAN for each DynOPO mission while employing DEMs of 50, 100, 200 and 400 resolution. The vehicle's baseline position (USBL measurements) is shown in magenta. 4248, 2529 and 4077 USBL measurements are available covering approximately 3.92%, 1.44% and 1.48% of the length of M41, M42 and M44, respectively. For the visualisation, the 50 m resolution DEM is used. Note that DR* position estimates are not fully visualised due to the corresponding large drift in position.

(see Table 4). The corresponding DR error is 1.1 km and 2.9 km, respectively, with and without using BR velocity.

In M42, TAN estimates appear to be convergent on a localisation solution that agrees with ship-based USBL to within 0.29 km on average for TAN 50, 0.3 km for TAN 100, 0.48 km for TAN 200 and 1.43 km for TAN 400. Using DR navigation, the average error is 0.4 km given BR velocity and becomes 13 km by utilising speed estimates derived from the propeller RPM. For the average error $RMSE^{10}_{k^*}$ at time k^*, this varies from 0.14 km to 1.6 km, dependent on the resolution of the employed bathymetric map (see Table 4). The corresponding DR error is 0.3 km and 28.4 km with and without using BR velocity, respectively. Notice also that TAN 400 estimates appear to deviate, compared to TAN 50-200 experiments, at the end of the mission. It appears that the filter reset capability was triggered in all ten independent runs at the end of the mission. While in the initialisation phase, it is common for the filter to approximate a highly multi-modal posterior density and provide low-accuracy position estimates.

In M44, TAN estimates appear to be convergent on a localisation solution that agrees with ship-based USBL to within 0.31 km on average for TAN 50, 0.3 km for TAN 100, 0.3 km for TAN 200 and 0.78 km for TAN 400. Using DR navigation, the average error is 0.4 km given BR velocity and becomes 14.5 km by utilising speed estimates derived from the propeller RPM. In terms of the localisation error with respect to the last available USBL measurement before ascending, $RMSE^{10}_{k^*}$, this varies from 0.07 km to 0.25 km, depending on the resolution of the employed bathymetric map (see Table 4). With and without using BR velocity the corresponding DR error is 0.4 km and 40 km, respectively.

The expected inverse relationship between the map coarseness and TAN accuracy is strongly present in all three missions when decreasing the map resolution from 200 m to 400 m. In contrast, when utilising maps of 50-200 m resolution this relationship appears to weaken and the TAN algorithm re-produces estimates with similar accuracy, in terms of $RMSE^{10}$. One interpretation may be that the terrain variability along the vehicle's path is still sufficient for the algorithm to extract localising information. Another explanation can be that the noise in the USBL measurements is what actually drives the $RMSE^{10}$ performance measure, making the effect of the map degradation non-identifiable. However, the most likely reason is that the limited number of USBL measurements does not allow firm conclusions to be drawn regarding exactly how the map resolution affects the TAN performance. Conversely, TAN becomes clearly less efficient when employing the 400 m resolution map (e.g. see M41 in Figure 6). The sub-sampling process appears to introduce large map errors (see Figure 5), something that occasionally makes the filter approximate a highly multi-modal posterior probability density function due to the presence of multiple local minima. Despite this, one of the modes appears to always represent the vehicle's true position, allowing the filter to eventually convergence into a relatively narrow unimodal posterior density function.

To further analyse the navigation performance during M44, Figure 7 shows the estimation error in the north and east directions with the respective 3-σ confidence bounds for each considered resolution of the bathymetric map (the evaluation is possible only when USBL measurements are available). At the beginning of the mission, the vehicle is demanded to descend to a depth of approximately 1150 m. While in the descent phase and out of bottom-tracking range, TAN is inactive and position estimates are computed using DR techniques via Equation (24).

Due to unknown water currents, the navigation error appears to grow over 0.5 km at approximately 3 km of DT (just before TAN is triggered). Once bathymetric observations become available, the accumulated navigation error drops below 50 m regardless of the resolution of the employed bathymetric map. From this point onwards, the position uncertainty is visualised as if the particle distribution was bi-variate normal uncorrelated in each direction. Although this assumption may be crude and occasionally invalid (when the approximated state posterior density is highly multimodal), hence making the filter appear overconfident, Figure 7 shows that the navigation error always falls within the uncertainty bounds (i.e. the navigation filter never underestimates the localisation error). The maximum navigation error appears to be approximately 1.5 km for TAN 50, 1.4 km for TAN 100, 0.6 km for TAN 200 and 3 km for TAN 400. Given that the only implementation difference between these experiments is the map uncertainty σ_m, which is used in the likelihood function, this observation indicates that the errors affecting the original 50 m resolution map and the sub-sampled 100 m resolution map are likely to be underestimated.

Figure 7 also reveals the inverse relationship between map resolution and TAN efficiency, with the estimation uncertainty increasing on average when the map resolution is decreased. This can be attributed, first, to the reduced navigational information contained within low-resolution DEMs and, second, to the likelihood function which flattens by reducing the map resolution. Focusing in particular on TAN 400, position estimates exhibit large uncertainty at approximately 70 km of DT. Close inspection revealed that multiple re-initialisations occurred, causing this increase in the uncertainty of the filter (larger particle spread). Multiple factors occurred simultaneously causing these resets. There were: a) poor matching between the map and the on-board bathymetric observations because of the presence of multiple local minima; b) prolonged time intervals where bathymetric observations were lacking, as the DVL occasionally went out of the bottom-tracking regime due to steep downward slopes; and c) sharp manoeuvring. In fact, the vehicle occasionally was required to perform sharp manoeuvres to reach the demanded waypoints (turns of almost

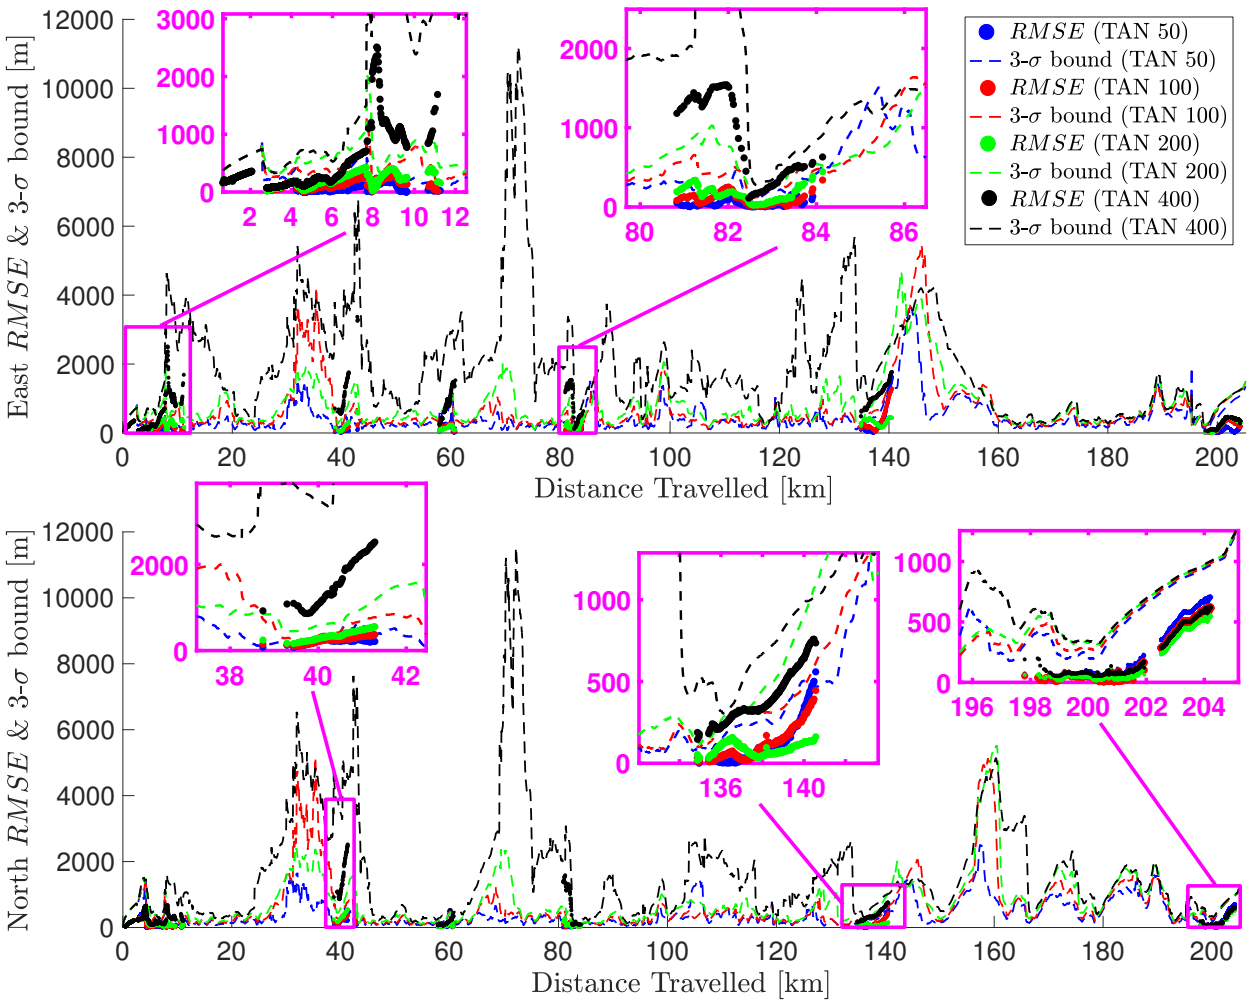

Figure 7: TAN estimation error over DT, averaged over 10 runs, with the corresponding 3-σ bound in the north and east directions, while employing four DEMs of varying resolution for M44.

180°, see Figure 6 and M44) and/or to maintain the demanded altitude of 90 m while operating in areas with rough terrain. All these manoeuvres have influenced the quality of the bathymetric observations, as the employed tilt sensors are unable to accurately measure the vehicle's pitch and roll angles (e.g. when actively controlling the vehicle's altitude or guiding the vehicle between waypoints), thus resulting in a poor beam-range projection. Under these circumstances, the posterior probability density function becomes highly multi-modal and the filter unable to provide high-quality position estimates (e.g. point estimates fall between the approximated modes).

Despite these challenges, the RBPF algorithm is able to eventually converge to the true mode of the posterior density and, most importantly, to never underestimate the localisation error (see Figure 7). To further support this statement, Figure 8 shows the minimum distance between particles and USBL measurements for the most challenging scenario of utilising the 400 m resolution map. It additionally provides the percentage of the particles that fall within the grid cell where USBL measurements are co-located. According to these plots, there are always particles near the true position of the vehicle.

The inverse relationship between map resolution and TAN effectiveness can also be shown by evaluating the filter's accuracy in estimating the speed of the water currents. Figure 9 compares water current estimates computed by the RBPF algorithm with the baseline water currents, calculated by subtracting the speed

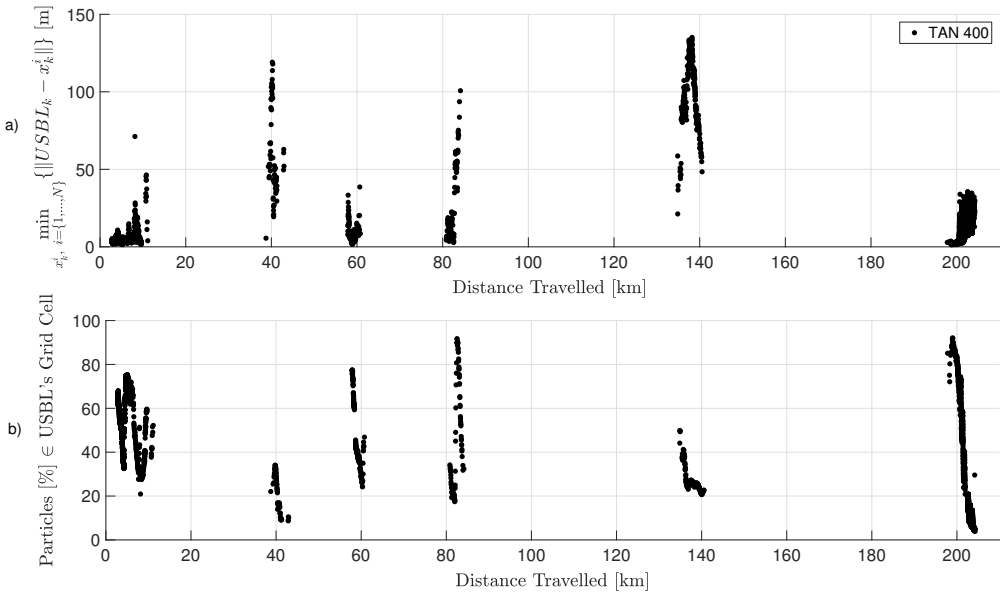

Figure 8: M44 and TAN 400 - a) distance between the closest particle to USBL measurements; and b) percentage of particles fallen in the grid cell where USBL measurements are co-located.

computed from propeller RPM from the BR velocity (when available). As the mapping function between propeller RPM and forward speed can be inaccurate, e.g. when the vehicle accelerates/decelerates, the reference water currents appear to be noisy. However, information can still be inferred from the noisy dataset regarding the speed and the speed variability of the water currents, even though smoothing techniques can be applied to refine these reference values.

As Figure 9 shows, the RBPF algorithm is able to follow the overall pattern in the baseline water currents, regardless of the resolution of the employed map, though sometimes it responded slowly to rapid variations in the speed of the water currents due to the modeling assumption that ocean currents vary slowly with time. The average deviation of the RBPF estimates from the baseline water currents is 0.12 m/s for TAN 50, 0.12 m/s for TAN 100, 0.13 m/s for TAN 200 and 0.16 m/s for TAN 400, whereas baseline water currents cause on average a drift of 0.21 m/s.

To demonstrate the ability of the filter to cope with gaps in bathymetric observations during M44, Figure 10.a shows the elapsed time between consecutive bathymetric observations. As the DVL pings the seabed at 0.5 Hz, the minimum elapsed time between observations is 2 seconds, whereas the maximum delay between observations appears to be 1146 seconds at approximately 7.5 km of DT. This large gap causes a localisation drift visible in Figure 7.

Figure 10.b shows the effectiveness of the DVL in obtaining range measurements from the moment of approaching and obtaining echoes from the seafloor until the begin of the ascent phase. From all attempts to ping the seabed at 0.5 Hz with all four transducers, the DVL fails to obtain at least a single echo from the seabed 26% of the time (e.g. because of potentially steep down terrain morphology or DVL drop-outs). A four-beam return is achieved in 15% of the attempts, whereas one-beam, two-beam and three-beam range data are obtained in 22%, 29% and 8% of the attempts, respectively. However, recall that due to the use of the interference detection algorithm (Salavasidis, 2019), it is likely that valid beam-ranges are occasionally rejected (as false positives) and thus the actual percentage of obtaining multiple echoes simultaneously may be underestimated.

Regarding the filter re-initialisation capability, it appears that resets occurred even when relying on the original (and most accurate) 50 m resolution bathymetric map. Close examination revealed that these resets

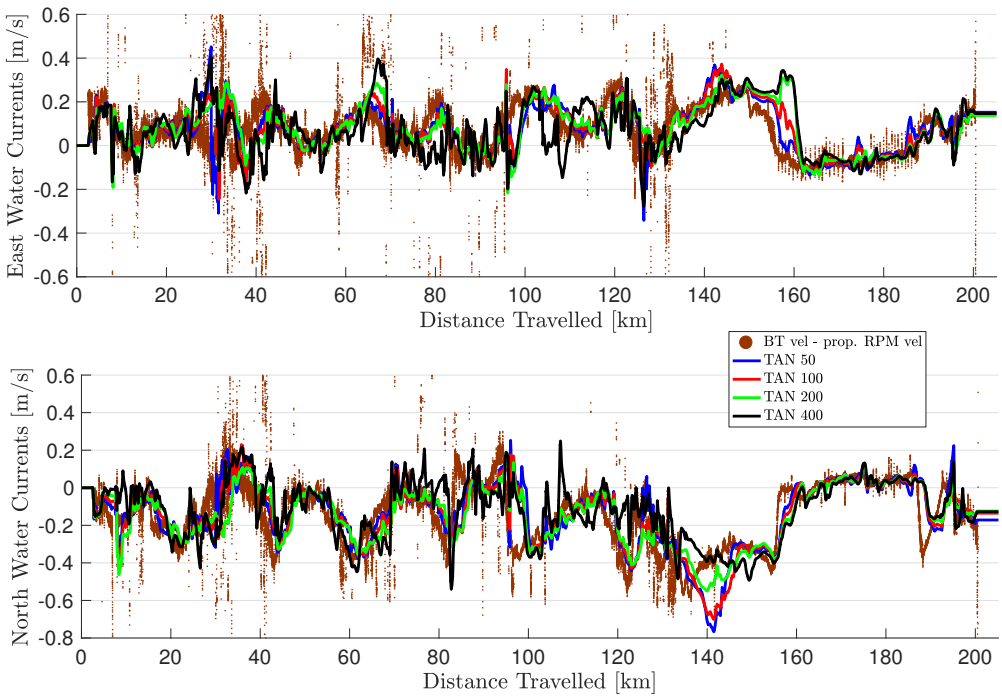

Figure 9: RBPF water current estimates compared to baseline water currents over DT while varying the resolution of the employed bathymetric map during M44. Baseline water currents are computed by subtracting the vehicle's velocity derived from the rotation speed of the vehicle's propeller from the DVL's bottom-relative velocity.

were triggered at locations where the vehicle performed sharp manoeuvres, and thus obtained low accuracy bathymetric observations due to a poor beam-range projection, and/or traversing areas with rapid change in the speed of the water currents. Hence, some of these resets are of high importance for the filter to escape from local minima, particularly at approximately 30-35 km of DT. It was also observed that the reset frequency increases by decreasing the quality of the map. For the most challenging setup (utilising the 400 m resolution map), multiple resets were triggered, allowing the filter to cope with the map sparsity, poor beam-projection whilst manoeuvring, and the rapidly changing water currents. Experiments performed without utilising the reset capability showed that the filter can be overconfident, leading into divergence in two experiments out of a total of ten independent runs.

To demonstrate the increase in robustness achieved using the filter re-initialisation capability, Figure 11 visualises the filter's behaviour via eight snapshots during M44 while employing the 400 m resolution map (worst case scenario). Due to the lack of a constant ground truth position, which complicates the demonstration of the effectiveness of the reset, TAN position estimates obtained while using BR velocity are used as the vehicle's baseline position (Salavasidis et al., 2019). While the RBPF appears to be able to accurately track the position of the vehicle (snapshot A), it suddenly converges to a local minimum (snapshot B) due to potential errors in the map and bathymetric observations. From this moment onwards, the RBPF would not be able to re-converge to the vehicle's true position without a reset mechanism. Instead, the filter is now able to detect the divergence and the particles are re-initialised (snapshot C). However, it also appears that the particle spread is not sufficiently large for the filter to completely recover and it eventually becomes trapped in a second local minimum (snapshot D). This problem arises when a false convergence is not detected at an early stage, and detection in this particular example required approximately an hour. However, the poor performance of the filter is detected again and a new re-initialisation is performed (snapshot E). Following the second reset, the posterior probability density function becomes highly multi-modal (snapshot F) and the filter remains incapable of providing good quality position estimates. A third reset is therefore triggered and performed (snapshot G) which results in a high filter uncertainty, despite that the actual estimation

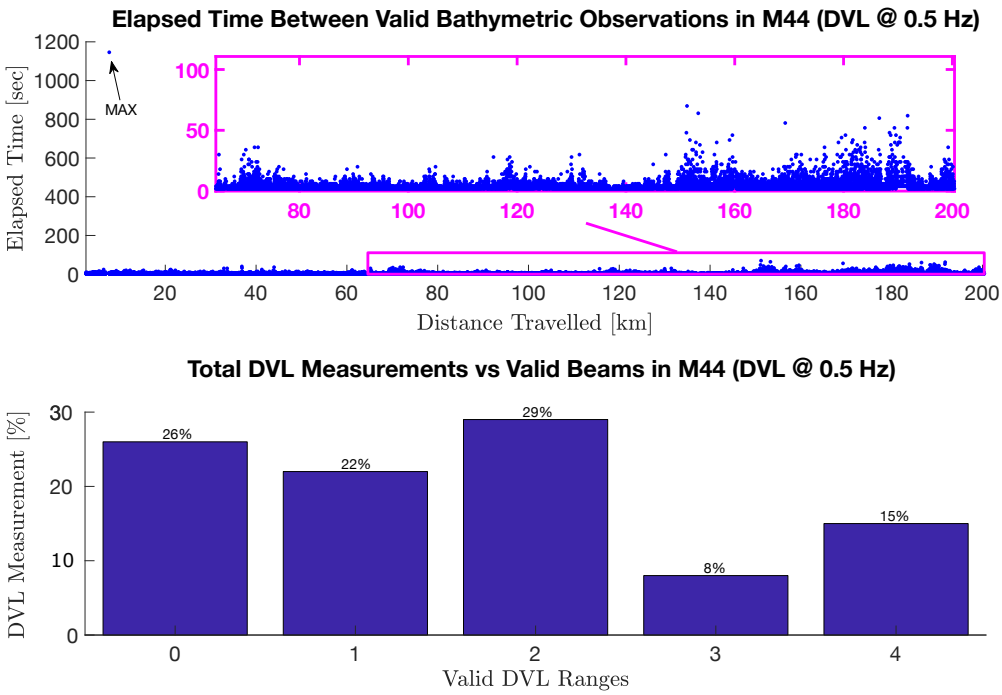

Figure 10: M44 - a) time elapsed between bathymetric observations over the DT; and b) the number of valid range measurements per attempt to ping the seabed, expressed as a percentage of the total DVL attempts to ping the seabed (once in the bottom-tracking regime and before ascending).

error is significantly reduced. Although not visualised, the filter does not converge to a unimodal probability density function because the vehicle is demanded to perform sharp manoeuvres that cause large errors when projecting the DVL's range measurements into the 3-D space. As a result, the filter tracks multiple hypotheses of the true vehicle's position for approximately three hours until it eventually converges to the single valid solution (snapshot H).

7 Conclusion

The increasing demand for conducting research activities on larger temporal and spatial scales in remote environments has led to the development of long-range AUVs. Platforms such as the ALR vehicles can facilitate the attempt of previously impossible missions, such as persistent monitoring and continuous multi-month data-collection mid-water column within ice-covered oceans (which host a diverse biological community and oceanographic features of high scientific interest). However, at present, weak water-relative AUV navigation challenges the feasibility of undertaking multi-month missions in GPS-denied environments.

In response to this limitation, this work has developed a TAN technique that relies on low-resolution bathymetric maps to limit the localisation error that an AUV experiences when operating in areas with strong water currents. To comply with the low-energy budget of long-range AUVs, the resulting TAN algorithm uses basic low-power sensors and is computationally feasible in real-time on processing boards with low computational power.

The performance of the navigation system has been evaluated using field data collected during three deep-water multi-day operations of ALR6000 in the Southern Ocean. Navigation challenges associated with these deployments, together with the unprecedented mission length ($\approx$ 200 km range in 77 hours), comprise a unique test case for evaluating TAN in realistic and demanding conditions. To imitate operations mid-water

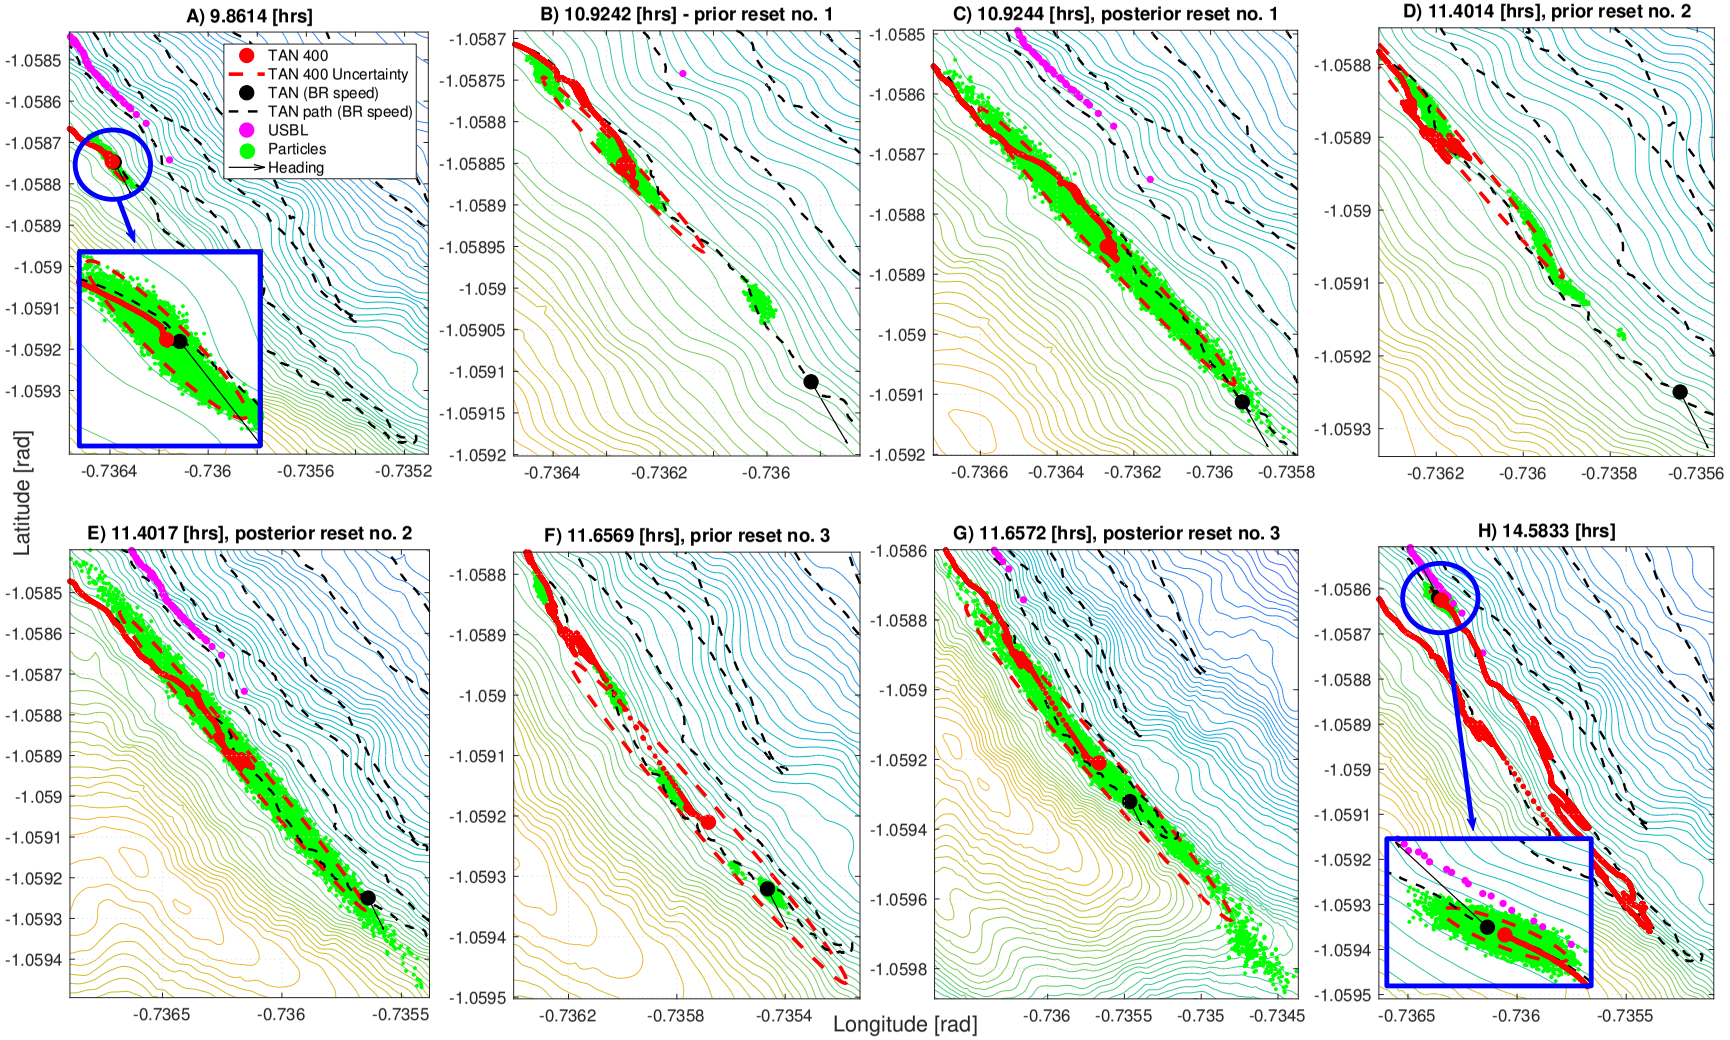

Figure 11: Visualisation of the reset capability. Eight snapshots capture a challenging instance of performing TAN during M44 and while using a 400 m resolution map. Three re-initialisations are performed within a five-hour interval in order for the filter to recover from local minima and to converge into a unimodal posterior probability density function from a highly multi-modal one.

column in the deep ocean, the available bottom-tracking velocity has been replaced by water-relative velocity estimates (derived from the rotation speed of the vehicle's propeller). To withstand the drift caused by water currents, the TAN uses the RBPF algorithm to estimate the 2-D speed of the water currents, in addition to the vehicle's 2-D position.

In operation, the RBPF algorithm relies on a small number of low-power navigation sensors, requiring no additional ranging sensors as the downward-facing DVL is treated as a four-beam sonar. Each available measurement from the DVL forms an individual bathymetric observation, after projection into the 3-D space. Moreover, it has been observed that DVL beams of the same device may interfere with each other, hence these beams are detected and rejected.

Given that the world's oceans are typically mapped in square grids of hundreds of metres, TAN was further evaluated relative to the resolution of the bathymetric map used. In particular, DEMs of 100 m, 200 m and 400 m resolution were generated by removing bathymetric data from the original 50 m resolution DEM of the area. To counter coarse bathymetric maps, the filter robustness was enhanced by integrating the RBPF algorithm with a re-initialisation capability, which allows the filter to escape from local minima and false convergence.

Multiple experiments were performed for each of the three missions and for each bathymetric map. The performance of the RBPF algorithm has been examined in terms of the filter's accuracy in estimating the vehicle's position together with the speed of the water currents. The DR navigation experienced rapid error growth (exceeding 40 km in a mission of only three-days duration), but the TAN was sufficiently robust, when integrated with a reset capability, to deal with both coarse bathymetric maps and strong time-varying water currents, despite traversing areas with low topographic information (along-track).

No divergence occurred throughout multiple independent experiments. Instead, the RBPF algorithm was able to maintain a bounded error with high estimation repeatability achieved between all runs for each employed bathymetric map. Furthermore, an inverse relationship between the map's resolution and the accuracy of TAN has been established, in both position and water current estimates. The localisation solution provided by TAN agreed with the ship-based USBL to within a few hundreds of metres (up to 1.5 kilometres) on average, depending on the coarseness of the employed DEM. Although bathymetric observations could be unavailable for long periods of time, the filter never exhibited overconfident behaviour, with the error estimates to be always within the calculated 3-σ confidence bounds.

7.1 Future Research Directions

While this work provides steps towards solving the general long-range navigation problem, it has also identified the following topics for consideration in future research:

1. *Improving the Robustness of the Navigation Filter*

   A future goal is to improve the RBPF robustness against false convergence. This is critically important for operations in the 12% of the world's oceans seasonally covered by ice. To enhance the robustness of the filter developed in this work, a re-initialisation framework has been developed. However, resetting an estimation process can be risky, especially when performed frequently. As an example where a reset leads to a failure, consider that the state posterior distribution is highly multimodal and particles are spread over a large region. If a reset is suddenly triggered, it is possible that the filter will be re-initialised in an area that misses one or more hypothetical solutions previously available (one or more modes of the approximated state distribution). The current re-initialisation scheme essentially fails since it assumes that the particle spread can be approximated by a unimodal Gaussian distribution.

   Alternative options must be investigated where the actual modes of the posterior density are preserved. An approach that can retain the modes of the posterior distribution is the injection of new

particles instead of resetting the filter. However, neither of these options can actually provide a general solution if, for example, false convergence is not detected at an early stage[7]. In fact, detecting a filter's suspicious behaviour is an equally difficult task, especially when the shape of the likelihood function varies with space.

One possible method to tackle the divergence problem in a broader sense is to borrow techniques from Interval Analysis (IA) (Jaulin et al., 2001) to represent the uncertainty in the vehicle's position. Using such techniques simultaneously with RBPF, good quality RBPF estimates can be used to update the results provided by IA. In turn, the computed position uncertainty based on IA will provide support to the RBPF when suspicious behaviour is detected and actions are required.

2. *Enabling TAN with Global Bathymetric Maps*

   One of the key contributions of this work was to demonstrate that TAN has the potential to navigate AUVs using bathymetric maps of similar resolution to those that the oceanographic community aims to map the world's oceans (Mayer et al., 2018). A future target is to enable TAN by using global bathymetric models. As a result, preparatory steps will be required which focus on quantifying the errors in global terrain models. Hence, a valuable extension to the current work would be the development of an appropriate probabilistic framework that takes into account terrain dependencies whilst computing the bathymetric map uncertainty.

3. *Closing the Navigation-Control Loop*

   A key outstanding challenge to obtaining a complete TAN system for real-time applications is to close the control-navigation loop. Sample-based estimation techniques often compute a multi-modal state posterior density function. Given this distribution, position estimates calculated by the mean of the state distribution can be highly biased and jumpy, causing discontinuities in the estimated AUV path (e.g. see M41 in Figure 6). It is important for the control algorithms onboard an AUV to eliminate poor and discontinuous estimates as they can lead to non-optimal control performance or, in certain circumstances, even to control instability. Therefore, a mechanism to discount inaccurate position estimates needs to be in place so that the overall trajectory is sufficiently smooth (relevant work can be found in Meduna 2011 and Ånonsen et al. 2013).

4. Multi-information Based Navigation

   Future work will also focus on developing a fully integrated geophysical-based navigation solution where other geophysical quantities are included as navigation references. Although the Earth's gravity is currently only used in simulation for AUVs (Pasnani and Seto, 2018), magnetic-based navigation has already been tested in the physical domain (Teixeira, 2007). A combined magnetic- and terrain-based navigation could increase the filter robustness by exploiting the strengths and compensating for the weaknesses of each component. Recent studies have also shown that environmental information such as temperature and salinity can potentially be used as navigation aids for long-range missions (Munafò et al., 2019).

Acknowledgement

This work was supported by the ROBOCADEMY (FP7 Marie Curie Programme ITN Grant Agreement Number 608096) and the NERC Oceanids programme. EPA was supported by NERC grants NE/K012843/1 (DynOPO) and NE/N018095/1 (ORCHESTRA). The authors would also like to thank the officers, crew and scientific party of the RRS James Clark Ross cruise JR16005 (DynOPO) for assistance in acquiring the data used here.

[7]Unless injecting particles or resetting the filter over the entire state space, which can be impractical for long-range AUV applications.

References

Abrahamsen, P. E. (2020). Gridded bathymetric compilation of selected areas within the orkney passage, scotia sea from multibeam echosounder data collected by multiple vessels (1989 — 2017) [data set]. In *UK Polar Data Centre*. Natural Environment Research Council, UK Research and Innovation. doi: 10.5285/1221BBBD-843E-46D6-B96C-9C559AECFB20.

Ånonsen, K. B., Hagen, O. K., Hegrenæs, Ø., and Hagen, P. E. (2013). The HUGIN AUV terrain navigation module. In *Oceans-San Diego, 2013*, pages 1–8. IEEE.

Anonsen, K. B. and Hallingstad, O. (2006). Terrain Aided Underwater Navigation using Point Mass and Particle Filters. In *Proceedings of the IEEE/ION Position Location and Navigation Symposium*, pages 1027–1035.

Arulampalam, M. S., Maskell, S., Gordon, N., and Clapp, T. (2002). A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on signal processing*, 50(2):174–188.

Bachmayer, R., Leonard, N. E., et al. (2004). Underwater Gliders: Recent Developments and Future Applications. In *Proceedings of the 2004 International Symposium on Underwater Technology (IEEE Cat. No. 04EX869)*, pages 195–200. IEEE. doi: 10.1109/UT.2004.1405540.

Bergem, O. (1993). Bathymetric Navigation of Autonomous Underwater Vehicles using a Multibeam Sonar and a Kalman Filter with Relative Measurement Covariance Matrices. *Dr. Scient thesis, University of Trondheim, Norway*.

British Antarctic Survey (2017). Dynamics of the Orkney Passage (DynOPO). https://www.bas.ac.uk/project/dynamics-of-the-orkney-passage-outflow/. Accessed: 28-January-2020.

Carpenter, J., Clifford, P., and Fearnhead, P. (1999). Improved particle filter for nonlinear problems. *IEE Proceedings-Radar, Sonar and Navigation*, 146(1):2–7. doi: 10.1049/ip-rsn:19990255.

Claus, B. and Bachmayer, R. (2015). Terrain-aided Navigation for an Underwater Glider. *Journal of Field Robotics*, 32(7):935–951.

Desmet, P. (1997). Effects of Interpolation Errors on the Analysis of DEMs. *Earth Surface Processes and Landforms: The Journal of the British Geomorphological Group*, 22(6):563–580.

Di Massa, D. E. (1997). *Terrain-Relative Navigation for Autonomous Underwater Vehicles*. PhD thesis, Massachusetts Institute of Technology.

Di Massa, D. E. and Stewart, W. (1997). Terrain-Relative Navigation for Autonomous Underwater Vehicles. In *Oceans' 97. MTS/IEEE Conference Proceedings*, volume 1, pages 541–546. IEEE. doi: 10.1109/OCEANS.1997.634423.

Donovan, G. T. (2011). Development and testing of a real-time terrain navigation method for AUVs. In *OCEANS 2011*, pages 1–9. IEEE.

Doucet, A., Godsill, S., and Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and computing*, 10(3):197–208.

Fofonoff, N. P. and Millard, R. C. (1983). Algorithms for the computation of fundamental properties of seawater. *UNESCO Technical Papers in Marine Science 44*.

Furlong, M. E., Paxton, D., Stevenson, P., Pebody, M., McPhail, S. D., and Perrett, J. (2012). Autosub Long Range: A long range deep diving AUV for ocean monitoring. In *Autonomous Underwater Vehicles (AUV), 2012 IEEE/OES*, pages 1–7. IEEE.

Google Earth (2017). Orkney Passage, Antarctic Peninsula, 59° 13′ 20″S, 49° 57′ 35″W, 3,504.29 km. https://earth.google.com/web/@-58.24792076,-54.57702361,-3401.62319381a,3512896.3445127d,35y,0h,0t,0r. Accessed: 10-August-2017.

Hare, R., Eakins, B., and Amante, C. (2011). Modelling Bathymetric Uncertainty. *The International Hydrographic Review*, (6).

Hayes, D., Boyd, T., and Patterson, M. (2007). Sensors and Instrument Requirements for Autonomous Underwater Vehicles. *Collaborative Autosub Science in Extreme Environments*, page 39.

Hobson, B. W., Kieft, B., et al. (2018). An Autonomous Vehicle Based Open Ocean Lagrangian Observatory. In *2018 IEEE/OES Autonomous Underwater Vehicles (AUV)*. IEEE. doi: 10.1109/AUV.2018.8729719.

International Hydrographic Organization (2008). IHO Standards for Hydrographic Surveys, 5th edition, Monaco. Special Publication No. 44, (S-44). http://hdl.handle.net/11329/388.

Jakuba, M. V., Breier, J. A., Gómez-Ibáñez, D., Tradd, K., and Saito, M. (2018). Clio: An Autonomous Vertical Sampling Vehicle for Global Ocean Biogeochemical Mapping. In *2018 IEEE/OES Autonomous Underwater Vehicles (AUV)*. IEEE. doi: 10.1109/AUV.2018.8729797.

Jaulin, L., Kieffer, M., Didrit, O., and Walter, E. (2001). Interval Analysis. In *Applied Interval Analysis*, pages 11–43. Springer.

Jensen, F. B., Kuperman, W. A., Porter, M. B., and Schmidt, H. (2011). *Computational ocean acoustics*. Springer Science & Business Media.

Li, T., Bolic, M., and Djuric, P. M. (2015). Resampling Methods for Particle Filtering: Classification, Implementation, and Strategies. *IEEE Signal Processing Magazine*, 32(3):70–86.

Li, Z. (1992). Variation of the Accuracy of Digital Terrain Models with Sampling Interval. *The Photogrammetric Record*, 14(79):113–128.

Mandt, M., Gade, K., and Jalving, B. (2001). Integrating DGPS-USBL position measurements with inertial navigation in the HUGIN 3000 AUV. In *Proceedings of the 8th Saint Petersburg International Conference on Integrated Navigation Systems, Saint Petersburg, Russia*, pages 28–30.

Marks, K. and Smith, W. (2008). An uncertainty model for deep ocean single beam and multibeam echo sounder data. *Marine Geophysical Researches*, 29(4):239–250.

Mayer, L., Jakobsson, M., Allen, G., Dorschel, B., Falconer, R., Ferrini, V., Lamarche, G., Snaith, H., and Weatherall, P. (2018). The nippon foundation—gebco seabed 2030 project: The quest to see the world's oceans completely mapped by 2030. *Geosciences*, 8(2):63.

McPhail, S. D. and Pebody, M. (1998). Navigation and control of an autonomous underwater vehicle using a distributed, networked, control architecture. *Underwater Technology*, 23(1):19–30.

Medagoda, L., Williams, S. B., Pizarro, O., Kinsey, J. C., and Jakuba, M. V. (2016). Mid-water current aided localization for autonomous underwater vehicles. *Autonomous Robots*, 40(7):1207–1227. doi: 10.1007/s10514-016-9547-3.

Meduna, D., Rock, S. M., and McEwen, R. (2009). AUV Terrain Relative Navigation Using Coarse Maps. In *Proceedings of the 2009 Unmanned Untethered Submersible Technology Conference*, volume 9.

Meduna, D. K. (2011). *Terrain Relative Navigation for Sensor-Limited Systems with Application to Underwater Vehicles*. PhD thesis, Stanford University.

Meduna, D. K., Rock, S. M., and McEwen, R. (2008). Low-Cost Terrain Relative Navigation for Long-Range AUVs. In *OCEANS 2008*, pages 1–7. IEEE. doi: 10.1109/OCEANS.2008.5152043.

Meduna, D. K., Rock, S. M., and McEwen, R. S. (2010). Closed-Loop Terrain Relative Navigation for AUVs with Non-Inertial Grade Navigation Sensors. In *2010 IEEE/OES Autonomous Underwater Vehicles*, pages 1–8. IEEE. doi: 10.1109/AUV.2010.5779659.

Melo, J. and Matos, A. (2017). Survey on advances on terrain based navigation for autonomous underwater vehicles. *Ocean Engineering*, 139:250–264. doi: 10.1016/j.oceaneng.2017.04.047.

Monckton, C. G. (1994). An Investigation into the Spatial Structure of Error in Digital Elevation Data. *Innovations in GIS*, 1:201–211.

Munafò, A., Fanelli, F., Salavasidis, G., Storto, A., and Oddo, P. (2019). Navigation of AUVs based on Ocean Fields Variability. In *OCEANS 2019-Marseille*, pages 1–9. IEEE. doi: 10.1109/OCEANSE.2019.8867309.

Naveira Garabato, A. C. and Forryan, A. (2017). RRS James Clark Ross Cruise JR16005: Dynamics of the Orkney Passage Outflow (DynOPO) cruise report. University of Southampton, National Oceanography Centre, Southampton, U.K. https://www.bodc.ac.uk/resources/inventories/cruise_inventory/reports/jr16005.pdf.

Naveira Garabato, A. C., Frajka-Williams, E. E., Spingys, C. P., Legg, S., Polzin, K. L., Forryan, A., Abrahamsen, E. P., Buckingham, C. E., Griffies, S. M., McPhail, S. D., et al. (2019). Rapid mixing and exchange of deep-ocean waters in an abyssal boundary current. *Proceedings of the National Academy of Sciences*, 116(27):13233–13238. doi: 10.1073/pnas.1904087116.

Nordlund, P.-J. (2002). *Sequential Monte Carlo Filters and Integrated Navigation*. PhD thesis, Department of Electrical Engineering, Linköpings University.

Nygren, I. and Jansson, M. (2004). Terrain navigation for underwater vehicles using the correlator method. *IEEE Journal of Oceanic Engineering*, 29(3):906–915. doi: 10.1109/JOE.2004.833222.

Pasnani, P. and Seto, M. L. (2018). Terrain-Based Localization and Mapping for Autonomous Underwater Vehicles using Particle Filters with Marine Gravity Anomalies. *IFAC-PapersOnLine*, 51(29):354–359.

Paull, L., Saeedi, S., Seto, M., and Li, H. (2014). AUV navigation and localization: A review. *Oceanic Engineering, IEEE Journal of*, 39(1):131–149. doi: 10.1109/JOE.2013.2278891.

Phillips, A., Haroutunian, M., Man, S., Murphy, A., Boyd, S., Blake, J., and Griffiths, G. (2012). *Nature in engineering for monitoring the oceans: comparison of the energetic costs of marine animals and AUVs*, pages 373–405. 77. The Institution of Engineering and Technology.

Phillips, A., Haroutunian, M., Murphy, A. J., Boyd, S., Blake, J., and Griffiths, G. (2017). Understanding the power requirements of autonomous underwater systems, Part I: An analytical model for optimum swimming speeds and cost of transport. *Ocean Engineering*, 133:271–279.

Phillips, A. B., Salavasidis, G., Kingsland, M., Harris, C., Pebody, M., Roper, D., Templeton, R., McPhail, S., Prampart, T., Wood, T., Taylor, R., and Jones, T. (2018). Autonomous Surface/Subsurface Survey System Field Trials. In *2018 IEEE/OES Autonomous Underwater Vehicles (AUV)*. IEEE. doi: 10.1109/AUV.2018.8729740.

Rees, W. G. (2000). The Accuracy of Digital Elevation Models Interpolated to Higher Resolutions. *International journal of remote sensing*, 21(1):7–20. doi: 10.1080/014311600210957.

Roper, D. T., Phillips, A. B., Harris, C. A., Salavasidis, G., Pebody, M., Templeton, R., Amma, S. V. S., Smart, M., and McPhail, S. (2017). Autosub Long Range 1500: An Ultra-Endurance AUV with 6000 Km Range. In *OCEANS 2017-Aberdeen*, pages 1–5. IEEE. doi: 10.1109/OCEANSE.2017.8084928.

Salavasidis, G. (2019). *Terrain-Aided Navigation for Long-Range AUVs Operating in Uncertain Environments*. PhD thesis, University of Southampton.

Salavasidis, G., Harris, C., McPhail, S., Phillips, A. B., and Rogers, E. (2016). Terrain-Aided Navigation for long range AUV operations at Arctic latitudes. In *Autonomous Underwater Vehicles (AUV), 2016 IEEE/OES*, pages 115–123. IEEE. doi: 10.1109/AUV.2016.7778658.

Salavasidis, G., Munafò, A., Fenucci, D., Harris, C. A., Prampart, T., Templeton, R., Smart, M., Roper, D. T., Pebody, M., McPhail, S. D., Rogers, E., and Phillips, A. B. (2020). *Ultra-Endurance AUVs: Energy Requirements and Terrain-Aided Navigation*. The Institution of Engineering and Technology.

Salavasidis, G., Munafò, A., Harris, C. A., McPhail, S. D., Rogers, E., and Phillips, A. B. (2018). Towards Arctic AUV Navigation. *IFAC-PapersOnLine*, 51(29):287–292. doi: 10.1016/j.ifacol.2018.09.517.

Salavasidis, G., Munafò, A., Harris, C. A., Prampart, T., Templeton, R., Smart, M., Roper, D. T., Pebody, M., McPhail, S. D., Rogers, E., and Phillips, A. B. (2019). Terrain–Aided Navigation for Long–Endurance and Deep–Rated Autonomous Underwater Vehicles. *Journal of Field Robotics*, 36(2):447–474. doi: 10.1002/rob.2183.

Schon, T., Gustafsson, F., and Nordlund, P.-J. (2005). Marginalized Particle Filters for Mixed Linear/Nonlinear State-Space Models. *IEEE Transactions on signal processing*, 53(7):2279–2289.

Siouris, G. M. (2004). *Missile guidance and control systems*. Springer Science & Business Media.

Stammer, D., Ray, R. D., Andersen, O. B., Arbic, B. K., Bosch, W., Carrère, L., Cheng, Y., Chinn, D. S., Dushaw, B. D., Egbert, G. D., et al. (2014). Accuracy Assessment of Global Barotropic Ocean Tide Models. *Reviews of Geophysics*, 52(3):243–282. doi: 10.1002/2014RG000450.

Teixeira, F. C. (2007). *Terrain-Aided Navigation and Geophysical Navigation of Autonomous Underwater Vehicles*. PhD thesis, Dynamical Systems and Ocean Robotics Lab, Lisbon.

Teixeira, F. C., Quintas, J., Maurya, P., and Pascoal, A. (2017). Robust particle filter formulations with application to terrain-aided navigation. *International Journal of Adaptive Control and Signal Processing*, 31(4):608–651. doi: 10.1002/acs.2692.

Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic Robotics*. MIT press.

Vickery, K. (1998). Acoustic positioning systems. A practical overview of current systems. In *Autonomous Underwater Vehicles, 1998. AUV'98. Proceedings of the 1998 Workshop on*, pages 5–17. IEEE.

Visbeck, M. (2002). Deep velocity profiling using lowered acoustic doppler current profilers: Bottom track and inverse solutions. *Journal of Atmospheric and Oceanic Technology*, 19(5):794–807.

Weatherall, P., Marks, K. M., Jakobsson, M., et al. (2015). A New Digital Bathymetric Model of the World's Oceans. *Earth and Space Science*, 2(8):331–345. doi: 10.1002/2015EA000107.

Wechsler, S. (2007). Uncertainties associated with digital elevation models for hydrologic applications: a review. *Hydrology and Earth System Sciences*, 11(4):1481–1500. doi: 10.5194/hess-11-1481-2007.

Yoerger, D. R., Curran, M., Fujii, J., et al. (2018). Mesobot: An Autonomous Underwater Vehicle for Tracking and Sampling Midwater Targets. In *2018 IEEE/OES Autonomous Underwater Vehicles (AUV)*. IEEE. doi: 10.1109/AUV.2018.8729822.