AUV Swarms for Monitoring Rapidly Evolving Ocean Phenomena

Thomas Lowndes

Supervisors: Alexander B. Phillips, Catherine A. Harris, Eric Rogers, Bing Chu, Ekaterina Popova

January 27, 2020

Abstract

Autonomous Underwater Vehicles (AUVs) have been commonplace in ocean science for years, taking the place of labour intensive and expensive ship operations. However, certain ocean phenomena, particularly those that change rapidly and require persistent monitoring over a large area, present challenges for AUVs, typically slow and energy restricted vehicles. Phenomena such as Harmful Algal Blooms (HABs) and oil spills fall into this category and have the potential to cause millions of dollars worth of damage to marine ecosystems, fisheries, human health and even tourism in coastal regions.

The emergence of a new class of micro AUVs allows much larger numbers of vehicles to be deployed than previously possible, making an AUV swarm (a large group of simple, small and low-cost robots) a feasible possibility. Due to the large number of vehicles, a swarm can resolve the temporal and spatial scales required to produce a cohesive data set of rapidly changing ocean phenomena. However, implementing an AUV swarm gives rise to the typical challenges associated with AUV operation as well as a new challenge, networking vehicles in a scalable manner. Communication between AUVs is a significant challenge, satellite and WiFi methods only work on the surface and while acoustic communication provides a solution to underwater communication, it suffers from low bandwidth, high latency and high packet loss. This work focusses on subjecting established swarming algorithms to the communication constraints typical of AUV operation, evaluating whether the swarm is able to achieve an acceptable level of coordination.

In order to evaluate the performance of the swarm, an individual AUV model is developed based on the ecoSUB, a micro AUV suitable for swarm deployments. This model is then validated against vehicle deployments, highlighting factors unaccounted for in the model. In order to simulate a swarm, satellite and acoustic

communication are added to the simulation, allowing vehicles to communicate and communication constraints to be imposed.

Firstly, an emergent flocking algorithm is evaluated, aiming to unify the vehicles as one group travelling in the same direction at the same speed. Swarms of 5 to 20 vehicles are tasked with following the flocking algorithm when subject to varying acoustic communication constraints. The results show a good level of flocking coordination across most swarm sizes and communication constraints, with the swarm consistently performing better than a single vehicle when completing a hill climb of a sensory function.

Secondly, a coverage control algorithm is evaluated, aiming to control the distribution of vehicles across an area, important when data is required at specific spatial scales in order to be most informative. Vehicle-in-loop experiments introduce a real vehicle to the simulated swarm and show the disturbances of vehicle drift and sensor noise to have little effect on the performance of the swarm. In simulation, the scalability of the algorithm is evaluated for increasing swarm sizes and communication constraints. Operating on satellite communication alone, the swarm is unable to achieve a stable distribution. Performance improves considerably with the introduction of acoustic communication, resulting in a stable distribution and a well coordinated swarm.

This work concludes that emergent flocking and coverage control are both suitable distributed control mechanisms for an AUV swarm, given acoustic communication is possible. Performance decrease with increasing swarm size and communication constraints is observed in performance metrics across both behaviours. However the performance decrease is not significant enough within the practically feasible ranges considered to state a limit of scalability.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Autonomous Underwater Vehicles (AUVs) have been commonplace in ocean science for years, taking the place of labour intensive and expensive ship operations. However, certain ocean phenomena, particularly those that change rapidly and require persistent monitoring over a large area, present challenges for AUVs which are typically slow and energy-restricted vehicles. The key philosophy of swarm robotics is to use large numbers of small and simple robots to perform the same task as one, large and expensive robot. This provides a solution to monitoring rapidly changing ocean phenomena.

However, a major barrier to the deployment of an AUV Swarm has been the risk of operations. AUVs have traditionally been: large vehicles requiring crane deployment from a ship; costing in the range of hundreds of thousands to millions of dollars and unreliable vehicles at risk of failure (Viquez et al. [2016], Brito et al. [2010]). The risk of traditional AUV operations has meant the deployment of an AUV swarm has been logistically and financially infeasible as well as a huge risk to future science . Now, a new class of small, simple and extremely low-cost AUVs, referred to as Micro AUVs, has ignited the prospect of an ocean monitoring swarm.

This thesis investigates the impact of underwater communication constraints on typical swarming algorithms, a key consideration when implementing an ocean-

going AUV swarm. The following section expands on the motivation for this work, taking a closer look at specific ocean phenomena and expanding on the benefits and challenges of an AUV swarm. Moving on, the aim and objectives of this work are presented, followed by publications, contributions and finally the structure of the thesis.

1.1 Motivation

Oil spills and Harmful Algal Blooms (HABs) are two examples of extremely transient ocean events which pose a huge risk to both human and marine life. While these events may seem distant and unrelated they share a number of key impacts. Both HABs and oil spills can: decimate fish, marine mammal and bird populations; impact human health and lead to a huge decline in tourism in affected regions. As a result, the cost of one of these events due to damage caused and response efforts can rise into the millions of dollars (Lee et al. [2018] , Sanseverino et al. [2016]).

When a HAB manifests in a coastal region it can result in mass fish deaths, deaths of marine mammals and birds, it can cause toxins to enter the food web and impact human health and even impact tourism in the area as blooms often discolour coastal waters and generate noxious foams. Very little is understood about the physical and biological processes leading to the formation of a HAB. Figure 1.1 shows two example HABs, from Paracas Bay, Peru and Monterey Bay, California. One issue is that the term 'Harmful Algal Bloom' is an umbrella term, and the phenomena included are highly diverse with a wide range of biological and physical initiation mechanisms [Cullen, 2007]. The other issue, and the one examined in this work, is that the spatial and temporal scales of a HAB mean rapid, extensive and sustained sampling is required to produce a cohesive data set [Babin et al., 2005]. As HABs require monitoring in both the temporal and spatial scales, a single ship or AUV is not capable of monitoring a HAB and

capturing the complete evolution over space and time to produce a cohesive data set [Griffiths, 2007].

Ryan et al. [2005] present an example study in which satellite, aircraft, ship and mooring based sensing were used to achieve the spatio-temporal resolution required to understand the evolution of a HAB. However, this study required extensive pre-existing infrastructure which is not rapidly deployable, nor financially feasible for every possible sampling location, as HABs occur in nearly every coastal region on the planet and can initiate, evolve and die off in a matter of weeks hence a system which is rapidly deployable is key to future HAB monitoring [Ryan et al., 2005] [Babin et al., 2005].

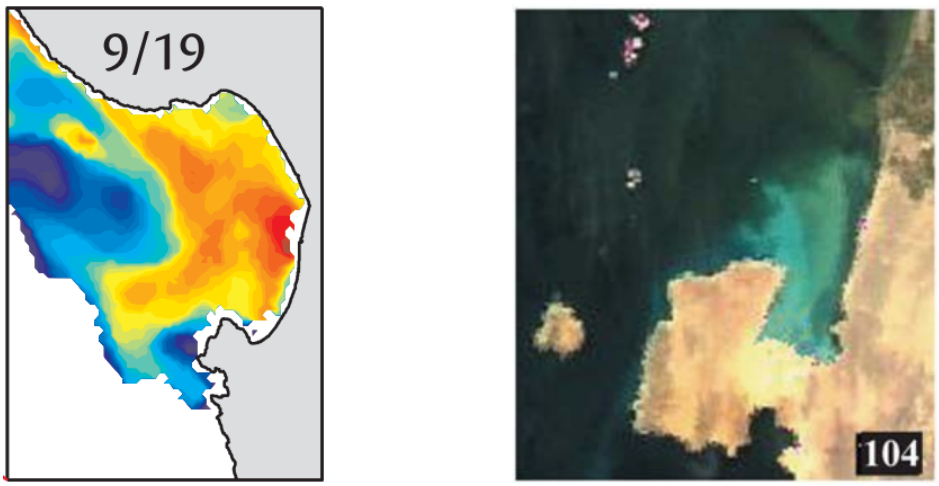

(a) Monterey Bay Algal Bloom (b) Paracas Bay Algal Bloom

Figure 1.1: Satellite imagery of Harmful Algal Blooms in a) Monterey Bay, California [Ryan et al., 2005] and b) Paracas Bay, Peru [Babin et al., 2005]

Moving on to oil spills, Figure 1.2 shows images of the slicks from the Exxon Valdez grounding and the Ixtoc-1 well blow-out. The impacts of an oil spill on the environment are very similar to the impacts of a HAB as they create a toxic environment for marine life and the oil eventually arrives on coastlines impacting tourism. While monitoring the extent and spread of the oil is of interest, this is a near impossible task as oil from both the Exxon Valdez and Ixtoc-1 spills was detected hundreds of miles from the initial source, making it extremely difficult to monitor the entire extent of the spill [NOAA, 2014]. Another topic of

(a) Grounded Exxon Valdez and surface slick

(b) Ixtoc-1 and surface slick

Figure 1.2: Examples of Oil Spills

interest surrounding oil spills is the effectiveness of response efforts, particularly the use of chemical dispersants, to reduce environmental impact. Dispersants are used to break up the oil slick so it can be absorbed into the water column, preventing oil reaching the shoreline and reducing the number of sea bird and marine mammal deaths. However, the use of dispersants is a trade-off between protecting the coastline and protecting the water column. For example, while a shoreline mangrove is protected through the use of dispersants, absorption of oil into the water column will create a toxic environment for a coral reef. Dispersants used in the Exxon Valdez spill were deemed ineffective and showed no significant benefit after spraying of 170,000 litres. The same dispersant was also used in the Deepwater Horizon spill, 21 years later, in much larger quantities [NOAA, 2014]. Several studies have concluded dispersants are in fact successful in mitigating coastal impacts. However, there is insufficient understanding of the actual concentrations and behaviour of the dispersed oil from field experiments, limiting the certainty of dispersants effect on the environment [National Research Council, 2005].

HABs and oil spills both have devastating effects on the environment and cost millions in mitigation and response efforts. Evolving on a timescale of weeks and undergoing rapid evolution due to internal processes and external disturbances mean a single sensing platform cannot produce an cohesive data set, capturing

the full evolution of an oil spill or HAB, hence another approach is needed. It should be noted that HABs and oil spills have been presented as examples of transient ocean phenomena, these are both ultimately chemical plumes and hence any chemical plume, for example those from hydro-thermal vents, are also covered by this work.

As shown in Figure 1.1, satellite sensing can be used to obtain the extent of a HAB. However, satellite observations suffer from susceptibility to cloud cover, provide little data below the surface of the ocean and offer low resolution in both spatial and temporal scales [Das et al., 2010]. As discussed previously, Ryan et al. [2005] achieve a cohesive data set through a wide array of sensing technology but heavily depended on pre-existing infrastructure i.e. ship deployed moorings which require anchoring to the seabed.

AUVs have been used to sample both Oil Spills [Ryan et al., 2013] and HABs [Robbins et al., 2006] and provide a vital third dimension to the survey by sampling at desired depths below the surface. However, both Willcox et al. [2001] and Griffiths [2007] discuss the difficulty in surveying rapidly evolving ocean phenomena using just a single AUV. The challenges of limited speed and endurance of a single AUV impose serious constraints on the dimensions and design of a survey, impacting the spatial and temporal resolutions achievable and hence the quality of the data.

As no approach has been shown to provide the temporal and spatial monitoring required while also being financially and logistically feasible to rapidly deploy in any sampling location, this work presents a new approach using an AUV swarm.

Griffiths [2007] discussed the prospect of using multiple low-cost AUVs to mitigate the challenges of an individual vehicle surveying a HAB. Recent developments in technology have led to the emergence of a new class of small, simple and low-cost AUVs known as micro AUVs making the prospect of an AUV swarm a reality. The field of swarm robotics concerns the use of large numbers of small,

(a) Autosub Long Range Deployment

(b) ecoSUB deployment

Figure 1.3: Deployment of (a) a traditional ship-based AUV, Autosub Long Range, compared to (b) an ecoSUB micro AUV from an Autonomous Surface Vessel (ASV).

simple robots to achieve tasks, boasting robustness to vehicle failure through distributed control, ease of deployment and reduced financial risk. Brooks and Flynn [1989] present one of the earliest arguments for swarm robotics, arguing complex systems ultimately cost more, take longer to develop and fail more catastrophically and hence a swarm of small, simple robots reduces time between missions and the 'launch mass' while increasing reliability in terms of overall mission success. Micro AUVs are ideal candidates for an AUV swarm, as their small size reduces the logistics of deployment, Figure 1.3 shows the reduced logistics of deplying a micro AUV, which can be deployed from other autonomous platforms, compared to a conventional ship-based AUV. Micro AUVs also cost a fraction of a normal AUV due to their simplicity and size and means an entire swarm may cost less than a single, sophisticated AUV. Deploying a swarm of AUVs inherently provides monitoring in both the spatial and temporal domains and hence is well suited to monitoring extremely transient ocean events such as oil spills and HABs. A coverage control algorithm complements this aspect of a swarm and allows the distribution of vehicles throughout the area to be controlled, enabling adaptation to the sensed distribution and improving data quality aiming to produce a cohesive data set.

There are two key challenges specific to implementing swarming algorithms in an AUV swarm. Firstly, the challenge of operating in the dynamic ocean envi-

ronment, common to all AUVs, creates significant uncertainty in vehicle localisation. This is especially relevant to a swarm AUV as they typically use low-cost, simple navigation sensors and rely solely on dead-reckoning for submerged navigation. Secondly, creating cooperative networks of AUVs proves challenging due to severely limited inter-vehicle communication. While submerged, AUVs can communicate using acoustic communication. However, this suffers from high latency, extremely low bandwidth and severe packet loss [Walls and Eustice, 2014]. In addition, acoustic communication requires a Medium Access Control (MAC) protocol as all swarm members utilise the same medium, the water column, to transmit their data. Satellite communication is common to most, if not all AUVs for the purpose of operation. However, the attenuation of electromagnetic signals in water restrict satellite communication to the surface. While able to transmit large quantities of data in comparison to acoustic communication, satellite communication is limited to small data packets to increase robustness and avoid connection issues in long data transmissions.

This work focusses on evaluating the impact of the severe communication limitations between swarm members on a swarms ability to perform various algorithms. These algorithms include emergent flocking, presented in Lowndes et al. [2017], and coverage control, presented in this thesis. Communication is seen as the key challenge for the implementation of an AUV swarm as navigation can be solved on an individual vehicle level, through techniques such as terrain aided navigation [Salavasidis et al., 2019], or on a swarm level through acoustic techniques [Webster et al., 2013]. In addition, the near-surface characteristic of the example ocean events means regular surfacing to correct for navigational error through GPS position is a feasible possibility, bounding error acquired through dead reckoning [Das et al., 2010].

1.2 Aim and Objectives

This thesis aims to evaluate potential approaches to sampling extremely transient ocean events, such as oil spills and harmful algal blooms, using an AUV swarm subject to communication constraints typical of AUV operation through:

- Identifying current swarm algorithms from literature which could be suitable for monitoring ocean features
- Developing a model of an AUV swarm, including vehicle dynamics and communication channels typical of AUV operation
- Validating individual vehicle dynamics through physical AUV deployments
- Using the AUV swarm model to evaluate the scalability of identified swarming algorithms to:
  - Increasing swarm size
  - Increasing communication constraints

1.3 Publications Arising from Thesis

Thomas S. Lowndes, Alexander B. Phillips, Catherine A. Harris, Eric Rogers and Ekaterina Popova, 2017. Evaluating the Capabilities of a Flight-Style Swarm AUV to Perform Emergent and Adaptive Behaviours. *Towards Autonomous Robotic Systems (TAROS) 2017*, proceedings. pp. 237 - 246, Guildford, UK, July 19-21 2017.

1.4 Main Contributions

The contributions of this thesis are as follows:

- Implementation and analysis of a coverage control algorithm on an ecoSUB AUV in vehicle-in-loop experiments
- Consideration of AUV communication limitations, the extent of which had not previously been covered by literature, and the effect on the performance of an AUV swarm acting under the control of emergent flocking and coverage control algorithms.

1.5 Thesis Structure

Following on from the introduction, the problem outline chapter gives an overview of why no current solutions to monitoring rapidly changing ocean phenomena exist and the challenges which arise in using AUVs and implementing an AUV swarm. After outlining the problem, a comprehensive literature review of the field is presented, covering current AUV survey techniques; the extent to which AUV swarms are currently being developed and used; both AUV and non-AUV approaches to monitoring the environment, focussing on coverage control and persistent monitoring; and finally methods of overcoming communication constraints in swarms of robots.

The following chapters detail how individual vehicles are modelled on the ecoSUB AUV and how communication channels typical of AUV operation are modelled in simulation. Chapter 4 introduced the AUV Model as well as Satellite and communication and how the limitations of these communication channels are accounted for in simulation. Chapter 5 presents validation of vehicle dynamics against real deployment data collected by the ecoSUB AUV in the Shetland Isles.

The thesis moves on to examine typical swarming approaches to environmental monitoring including emergent flocking, hill climb and coverage control algorithms, presented in Chapters 6 & 7. The studies evaluate the performance of swarms of increasing size and communication constraints, aiming to evaluate whether a sufficient level of coordination is achievable. The Coverage Control chapter includes vehicle-in-loop experiments, introducing drift and sensor noise into the swarm through the introduction of a real vehicle interacting with multiple simulated vehicles.

Finally, conclusions from the work presented in this thesis are summarised in Chapter 8 and future work discussed in Chapter 9.

Chapter 2

Problem Outline

This chapter outlines the problem and explain why an AUV swarm is a suitable solution to monitoring rapidly changing ocean events. Firstly, the chapter will discuss the spatial and temporal scales of such events, outlining key parameters where possible. However, as it is so difficult to monitor the evolution of rapidly changing phenomena in their entirety, information is sparse. Secondly, the chapter will discuss the challenges faced in general AUV operation and the challenges that arise in implementing a collaborative AUV swarm.

2.1 Scales of Rapidly Evolving Ocean Phenomena

Figure 2.1, from Dickey [2003], presents the spatial and temporal scales at which various ocean processes must be observed in order for the data to provide useful information. For example, this figure tells us if a number of measurements were taken 1km apart every week, if a plankton bloom had occurred, useful information on the bloom would be gained. Whereas if measurements were taken only at the start of each year, the data would not provide any useful information on the bloom as the temporal scale is not sufficiently fine.

For this work, plankton blooms, or harmful algal blooms, have already been

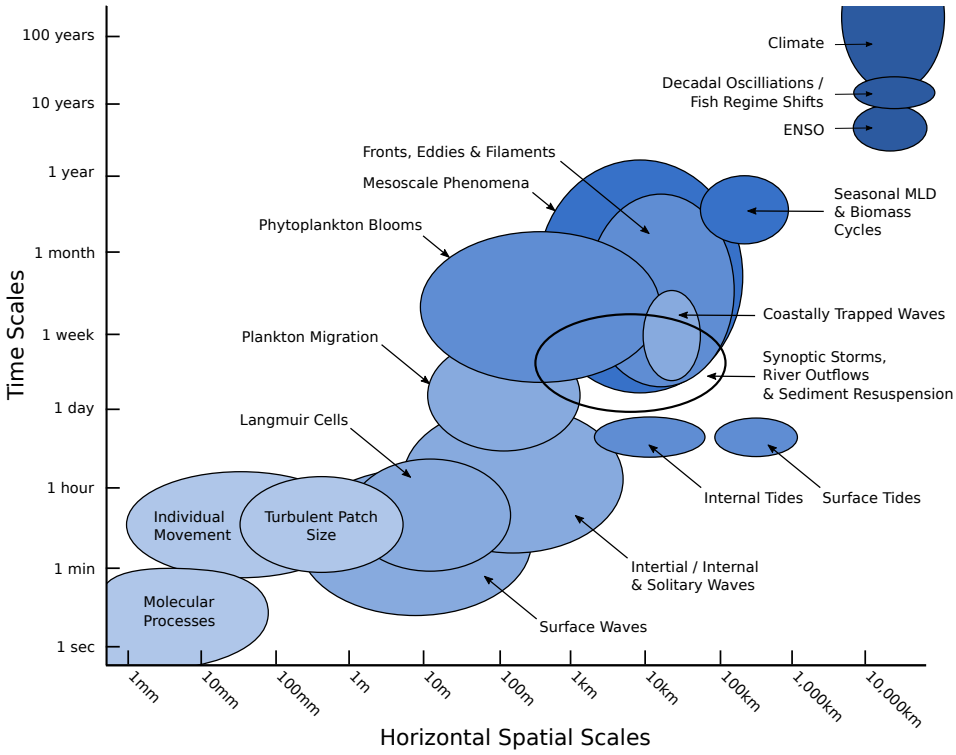

Figure 2.1: Spatial and Temporal Scales of Ocean Phenomena. Reproduced here from Dickey [2003]

mentioned as one of the motivating phenomena. From Figure 2.1, plankton blooms must be monitored on a time scale from daily to monthly and on a spatial scale from 10m to 10km to gain useful information. This is the temporal range described as 'rapidly changing'. Table 2.1 presents two case study HABs, one in Monterey Bay, CA and another in Paracas Bay Peru. These case studies echo the spatial and temporal scales discussed previously, occurring over large areas and evolving on a scale of days to weeks. It is the combination of this temporal scale and the spatial scale which make these types of phenomena difficult to monitor for a single AUV.

Table 2.1: Example HABs and their spatial and temporal scales.

	Paracas Bay, Peru	Monterey Bay, CA
Spatial (Surface), km^2	56	800
Temporal, days	12	11
Depth, m	20	20
Reference	Babin et al. [2005]	Ryan et al. [2005]

Following on from HABs, oil spills have also been discussed as potential applications for an AUV swarm. Table 2.2 presents 3 case study oil spills, a tanker grounding and 2 oil well blow-outs. The temporal scale of the Exxon Valdez spill is similar to a HAB, while the well blow-outs from Ixtoc-1 and the Macondo well lasted much longer due to the huge quantities of source oil meaning the spill only stopped when the well was capped. Spatially, the Exxon Valdez spill is again similar to a HAB. Of the two well blow-outs, the Macondo spill is actually closer to the spatial scale of a HAB. This was due to a large amount of oil being diverted into a subsurface plume at approximately 1120m below the surface. The subsurface plume formed due to the extreme source depth of the well at 1500m. In the case of Ixtoc-1 however, the shallow water meant all oil rose to the surface, causing a much larger surface slick.

Table 2.2: Example oil spill events and their spatial and temporal scales.

	Exxon Valdez	Ixtoc-1	Macondo
Type	Tanker	Well blowout	Well blowout
Spatial, km^2	1.92	42 - 300	12.6
Temporal, days	10 - 30	290	150
Depth, m	0	50	1500
References	NOAA [2014]	Jernelöv and Lindén [1981]	Camilli et al. [2010]

In summary, phenomena such as HABs and oil spills change rapidly and occur over a large area. It is the combination of the temporal and spatial scales of these phenomena which presents a problem for current AUV monitoring approaches. Similarities can be drawn between the scales of HABs and oil spills, the two motivating phenomena for this work. However, examining Figure 2.1 more closely and incorporating the spatial and temporal scales of oil spills, Figure 2.2 shows there are several other phenomena of scientific interest which fall in this 'rapidly changing' category.

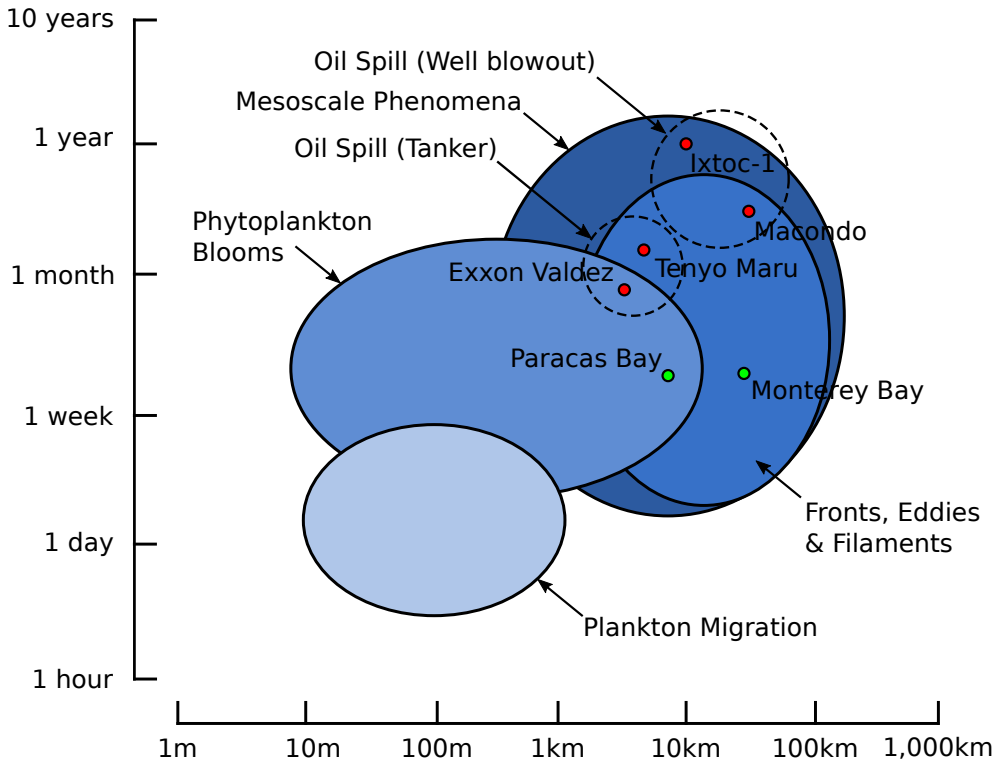

Figure 2.2: Rapidly Evolving Ocean Phenomena including case study oil spills. [Dickey, 2003]

2.2 Survey Error - Capturing the Snapshot

Griffiths [2007] presents a method of assessing the suitability of different marine autonomous vehicles (MAV) to the particular challenge of surveying a HAB. Griffiths [2007] uses a survey error metric, ϵ, derived by Willcox et al. [2001] which defines the quality of data collected by surveys of varying length and time scales. This error metric can be used to argue why surveying a HAB presents such a problem for current MAV and why an AUV swarm can be the solution to this. The length scale is representative of the distance between passes in a lawnmower survey, as shown in Figure 2.3, hence a smaller length scale is representative of less distance between passes and hence a finer survey of the area. The time scale is representative of the total time taken for the survey, hence a low time scale is representative of a vehicle completing a survey very quickly and a high time scale is representative of a vehicle taking a long time to complete a survey. In addition, the survey error metric takes into account the length and time scales

of the dominant process, the HAB. Due to this, it is not as trivial as doing a extremely long survey (high time scale) with very short distances between each pass of the lawn mower survey (low length scale) as the feature will change over the course of the survey and measurements taken at the start and end of the survey may not be able to be evaluated as one feature due to the change.

Figure 2.3: Example of a lawnmower survey and length scale, λ_s

The ideal survey of the HAB would take an instantaneous snapshot of the event at a single point in time at a very fine spatial resolution. Unfortunately this is not achievable as a vehicle can only sample one location at a time. If two vehicles performing the exact same lawnmower survey of a HAB, one at a low time scale and one at a high time scale, the first vehicles 'image' of the HAB will be very close to the ideal snapshot as the HAB has not changed significantly throughout the course of the survey. The second vehicles 'image' of the HAB will will gradually move further from the ideal snapshot as the HAB changes during the course of the survey. Due to the high time scale, the second vehicle will not produce a coherent snapshot of the event whereas the first vehicle will. Examining the length scale metric, a low length scale is preferable as this is closer to the ideal snapshot. If the survey has a very high length scale, there is larger distances between the passes of the lawnmower survey and hence the HAB may change significantly between the passes and this would not be captured by the survey.

The resultant is a survey error metric which ranges from 0, representing a complete, coherent snapshot of the HAB, to 1, representing a survey which ob-

tains little useful information. The resultant of the survey error metric using the scales of a HAB is shown in Figure 2.4. Details of the survey error metric and how it is calculated can be found in Griffiths [2007] and Willcox et al. [2001].

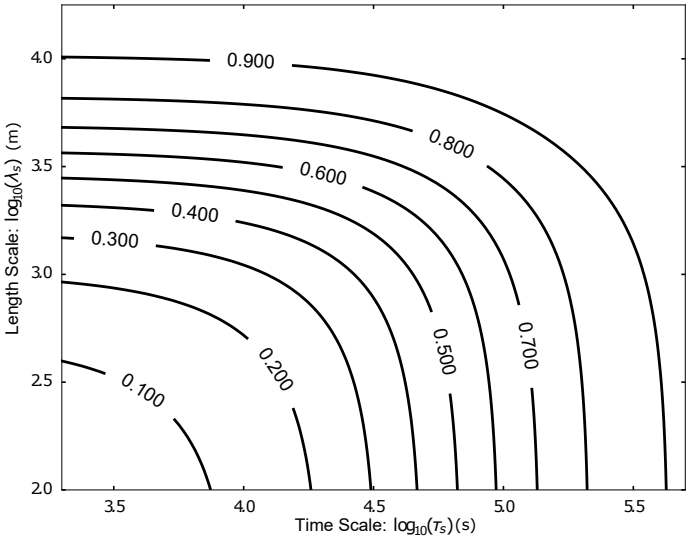

Figure 2.4: Griffiths Survey Error Metric given the time and length scales of a survey. The value of the error metric is shown by the contours.

Next, given a survey area of $A = 30km^2$, Griffiths [2007] simulate the resolvable time and length scales given different vehicle speeds using Equation 2.1. The values used for forward speed, V, are shown in Table 2.3. The AUV Swarm has been added to those presented by Griffiths [2007] to allow comparison for this work.

Table 2.3: Vehicles used in Survey Error Metric simulation.

Vehicle	V (m/s)	Number of Vehicles
Glider	0.30	1
Glider Fleet	0.30	5
AUV Swarm	0.75	10
AUV	2.00	1
ASV	6.00	1

$$\lambda_s = \frac{\frac{2A}{\tau_s V}}{\left(1 - \frac{2\sqrt{A}}{\tau_s V}\right)} \tag{2.1}$$

The result of Equation 2.1 is super imposed on to the survey error metric and shown in Figure 2.5. The line for each vehicle represents the minimum resolvable

scales, hence the vehicle can resolve scales and achieve error metrics to the right of these lines by going slower, but not to the left as the resolvable scales are based on the maximum speed.

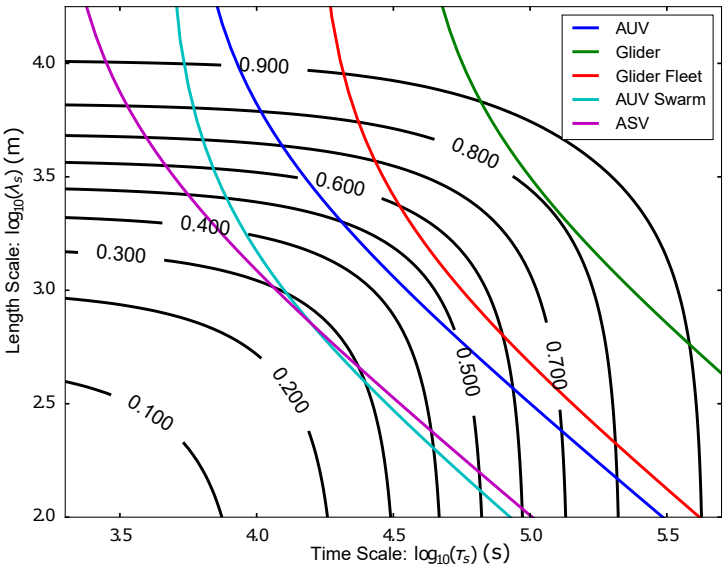

Figure 2.5: Griffiths Survey Error Metric given the time and length scales of a survey. The value of the error metric is shown by the contours. The survey error metric achievable by each vehicle in Table 2.3 are shown by the coloured lines.

As can be seen in Figure 2.5, the glider is able to achieve a minimum survey error metric $\epsilon \approx 0.8$, capturing little useful information about the environment. The AUV is able to move considerably faster than the glider and hence achieves a minimum value of $\epsilon \approx 0.46$. Interestingly, if the survey area is divided equally among a fleet of five gliders, they achieve a minimum value of $\epsilon \approx 0.55$. This is much closer to the value achieved by an AUV with the extended endurance, months compared to days, that comes with a glider. The lack of consideration of endurance is a particular drawback of this survey error metric. The maximum endurance of some commerical AUVs is as little as 5 hours whereas, as stated previously, a glider can stay at sea for months, clearly impacting their ability to survey an environment in its entirety. A fleet of five gliders will also have a cost similar to that of a single AUV. However one downside not captured by this simulation is the navigational error of the glider being considerably higher than an AUV. The final vehicle presented by Griffiths [2007] is an Autonomous Surface

Vessel (ASV) travelling at 6m/s. The ASV is able to achieve a much lower survey error metric at $\epsilon \approx .28$ but is obviously limited to surface operation, gaining little insight below the surface, a dimension not captured by the simulation. A speed of 6m/s is also unlikely to optimise endurance and hence the length of time the ASV can maintain a presence surveying the environment is brought into question.

For this work, the AUV swarm is added to the survey error metric plot from Griffiths [2007]. A swarm of 10 AUVs, dividing the survey area equally, is able to achieve a minimum value of $\epsilon \approx .28$, performance similar to the ASV. However the AUV swarm is able to dive, capturing three dimensions in its survey as oppose to two. This is a significant benefit of the AUV swarm over the ASV. Of course, the AUV swarm is not the only vehicle considered which is able to survey underwater, gliders and single AUVs are also capable of this. However, the cost of an AUV swarm will be similar to that of a single glider and a fraction of the cost of a sophisticated AUV hence there is a clear benefit to the AUV swarm over other underwater vehicles.

To summarise, the survey error metric presented by Griffiths [2007] demonstrates how the data collected in a survey of a HAB can be improved significantly through the use of an AUV swarm over a single AUV. Other considerations such as cost also demonstrate the advantages of an AUV swarm over an ASV or glider fleet.

2.3 Challenges of AUV Operation

In this section, the challenges of general AUV operation and those arising from the implementation of an AUV swarm will be presented.

2.3.1 Navigation

On the surface, AUVs have access to GPS and hence can obtain their position. However, once submerged the attenuation of electromagnetic waves in water

means GPS is no longer available [Paull et al., 2014]. Navigation while submerged typically relies on dead reckoning with periodic surfacing to attain a GPS fix and correct for drift error (Figure 2.6). Dead reckoning can be improved through the use of very accurate inertial navigation sensors as well as a Doppler Velocity Log (DVL), or Acoustic Doppler Current Profiler (ADCP), to sense velocity relative to the sea floor, both large, expensive and power demanding sensors. A DVL still does not solve all problems as the sensor must have bottom lock, meaning the sea floor is within range of the sensor's acoustic ping. Without a DVL, an AUV has no way of measuring its velocity relative to the sea floor, it can only infer it from the motor rpm or use a flow meter to measure velocity relative to the water column. Figure 2.7 gives indicative values for the accuracy of navigation along with cost and power consumption for various AUVs. From Figure 2.7, in order to reduce the error in navigation, both cost and power consumption increase. This is contrary to an AUV swarm as the swarm members must be small, simple and low-cost hence AUVs suitable for a swarm, for example the AUV second from the bottom in Figure 2.7, are limited in their navigational accuracy.

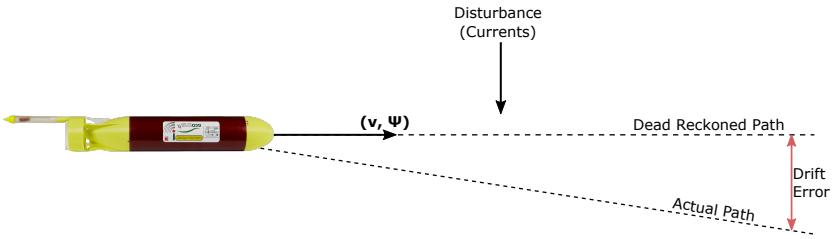

Figure 2.6: Dead Reckoning navigation.

Acoustic navigation offers a very accurate alternative to dead-reckoning through Long Base Line (LBL) and Ultra Short Base Line (USBL) technologies. LBL systems require a number of fixed acoustic transponders at known locations, often located on the seabed. The AUV is able to interrogate the LBL network and receives a response from each transponder. From the time of flight, the AUV is able to calculate the distance to each node and hence locate itself within the network. The significant disadvantage of the LBL approach, apart from only being able to locate within range of the network, is the cost of deploying and

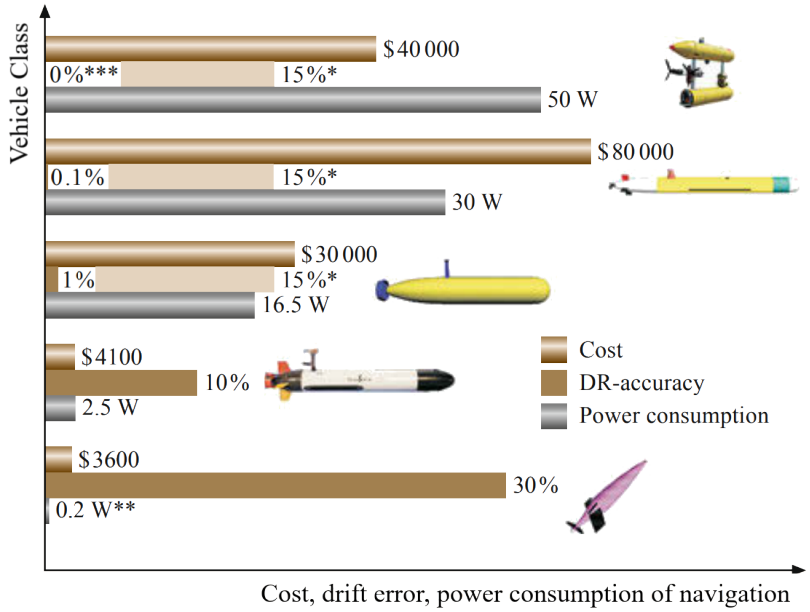

Figure 2.7: Navigation accuracy as percentage of distance travelled, power consumption and cost of various AUV sensor suites (bottom to top): (1) Glider with compass and altitude sensor; (2) low-cost AUV with compass, altitude sensor and flow meter; (3) medium range AUV with Inertial Navigation Systems (INS), DVL and LBL; (4) high-end AUV with Fibre Optic Gyroscope based INS, DVL and LBL; (5) special-task AUV with INS, DVL and Simultaneous Localisation and Mapping (SLAM). * Drift in mid water-column when DVL cannot obtain bottom or surface lock, ** Assuming 10% Duty cycle during which the navigation sensors are powered, *** Assuming that the vehicle was close enough to the sea floor throughout entire mission to take pictures and revisit places. Figure taken from Leonard and Bahr [2016]

maintaining the transponder network on the seabed. USBL differs from LBL as it uses several transponders mounted on the same unit, one carried by a ship and the other by the AUV. The AUV is able to attain both range and bearing information by interrogating the ship based transponder and hence estimate its location.

As one of the benefits of an AUV swarm is that it does not require a ship, as the individual vehicles are small enough to be launched from virtually any platform, traditional USBL or LBL acoustic positioning do not make sense for an AUV swarm. However, collaborative acoustic localisation approaches are being developed which allow swarm members to attain better positioning estimates by interrogating the acoustic network carried by the swarm. This can be likened

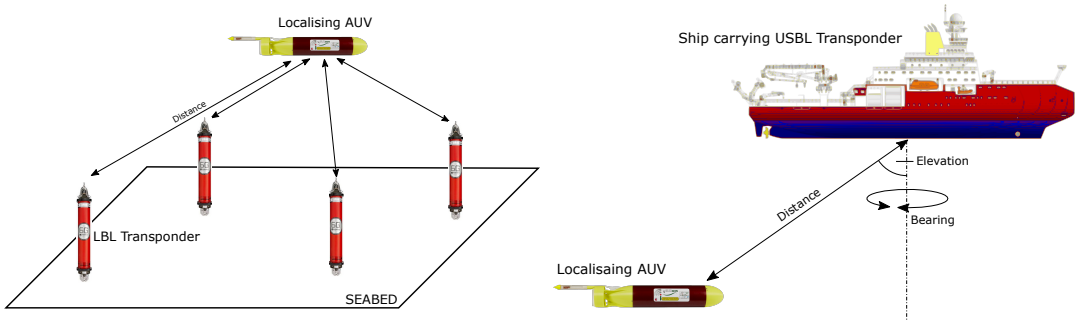

(a) LBL

(b) USBL

Figure 2.8: The most common forms of acoustic localisation: (a) Long Baseline (LBL) and (b) Ultrashort Baseline (USBL)

to an LBL system where each swarm member has a transponder and is able to interrogate the network to attain a new location estimate, an example of this is presented by Webster et al. [2013]. Additionally, methods exist which share information across the network which can in turn be used for localisation. Liu et al. [2014] introduces an approach where the most recent AUV to surface, and hence attain a GPS fix, communicates this to the rest of the swarm, reducing the navigational error. Alternatively, Tan et al. [2016] share bathymetric measurements across the swarm in order to perform terrain aided navigation, using measurements of water depths against pre-existing maps to determine a location estimate.

In summary, GPS positioning is available on the surface and bounds error as long as an AUV is able to surface. Submerged navigation primarily relies on dead reckoning. However, it is difficult to improve the accuracy of dead reckoning without large and powerful sensors. Finally, new cooperative acoustic techniques can be used by swarm members to attain better location estimates and reduce navigational error, given each swarm member is equipped with an acoustic modem.

2.3.2 Communication

Communication is vital for a swarm of AUVs to be able to work collaboratively, whether for the sharing of sensor measurements for intelligent monitoring or

simply sharing location to achieve the desired formation. Figure 2.9 presents the data transmission rates of various communication technologies found in AUVs. This section discusses each of these technologies and their suitability for inter-vehicle communication in a swarm.

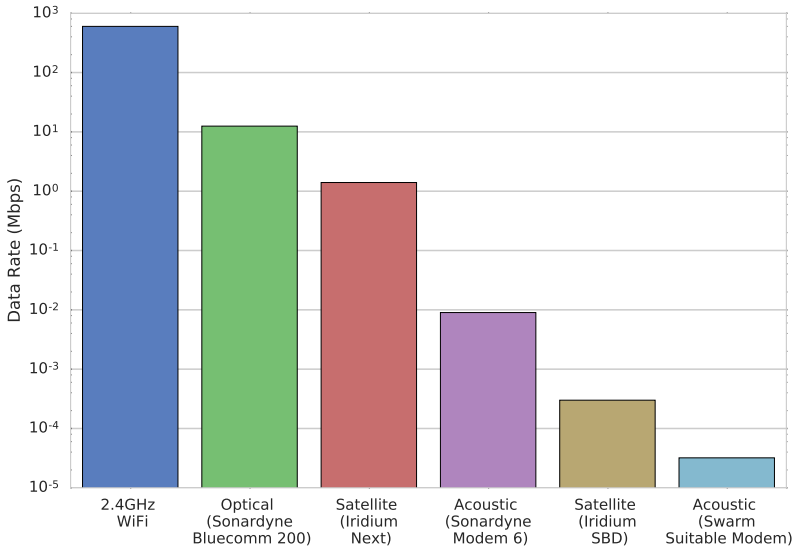

Figure 2.9: Data transmission rates of various communication technologies found in AUVs.

All AUVs are typically fitted with WiFi communication in order to allow the operator to send new commands to the vehicle and to retrieve data while it is on the surface. Considering inter-vehicle communication, the short range of WiFi requires the vehicles to be close together and on the surface in order to communicate. This could prove difficult in the vast, dynamic environment that is the ocean. The AUVs could travel to a mutually agreed location in order to communicate. However this would reduce the productivity of the AUVs during deployment and use up extremely valuable resources. The surface requirement and short range of WiFi deems it unsuitable for inter-vehicle communication.

Optical technologies are available for underwater communication and achieve the data rate closest to WiFi. However, optical communication has a short range of only 150m, performance is variable between day and night and dependant on the amount of particulate matter in the water-column.

AUVs operating in the ocean for long periods of time are typically equipped

with satellite communication technologies allowing for remote piloting from a land base. Satellite communication is only available on the surface and occurs in a client-server topology. For an example of how this would operate in a swarm, an AUV could send a message to a server via satellite containing its own location and any other information. The server could then respond with the last known locations of the other AUVs in the swarm. Iridium Short Burst Data (SBD) is often used in AUVs as the short bursts of communication avoid connection issues as waves washing over the antenna can affect communication. The size of an SBD message is modem dependant. A modem suitable for a micro AUV has a limit of 340 bytes per message and takes 20 seconds to send a message with a perfect view of the sky however in larger AUVs SBD messages can be up to 2 kilobytes. If the view of the sky is restricted, sending a message can take several minutes. Iridium Next provides a higher data transmission rate of 1.4 Mbps but is not yet available.

Finally, acoustic technologies are the predominant form of underwater communication, sending acoustic signals through the water column similar to how whales and dolphins communicate. The Sonardyne Modem 6 can achieve a data rate of 9 Kbps at a range of 1km. Unfortunately, considering the small size and low energy capacity of a swarm AUV, this modem is far too large. A swarm suitable modem, one small enough to fit on the vehicle and with low power consumption, can achieve a raw bit rate of 40 bps with a similar range to the Modem 6 [Morozs et al., 2018]. Acoustic communication is notoriously unreliable, suffering high levels of packet loss from phenomena such as beam bending and multipath effects.

In summary, WiFi and optical communication technologies achieve high data rates but are limited by the short range of communication. Satellite communication occurs between a client and a server in 340 byte messages and is only available on the surface. Acoustic communication is the predominant technology for underwater communication but suffers from low data transmission rates.

2.4 Summary

This chapter has shown the temporal and spatial scales at which rapidly evolving ocean phenomena occur. HABs and oil spills fall into this category along with a myriad of other phenomena. It is the combination of the temporal and spatial scales which make rapidly changing phenomena difficult to monitor. The survey error metric presented by Griffiths [2007] demonstrates how the data collected in a survey of a HAB can be improved significantly through the use of an AUV swarm over a single AUV or even a fleet of gliders.

Independent navigation is still a significant challenge in AUV operation. GPS positioning is available on the surface and provides a ground truth as long as an AUV is able to surface. Submerged navigation primarily relies on dead reckoning. However, it is difficult to improve the accuracy of dead reckoning without large and powerful sensors. Finally, new cooperative acoustic techniques can be used by swarm members to obtain better location estimates and reduce navigational error, given each swarm member is equipped with an acoustic modem.

Another challenge in implementing an AUV swarm is enabling the members to communicate in order to collaborate. WiFi and optical communication technologies achieve high data rates but are too short range to be useful. Satellite communication occurs between a client and a server in short burst messages and is only available on the surface. Acoustic communication is the predominant technology for underwater communication but suffers from low data transmission rates.

Chapter 3

Literature Review

This literature review begins by discussing the state of the art in AUV operation and examples of current AUV survey techniques to gain an understanding of how AUV operations are currently conducted and establish the level of autonomy used in the field. Secondly, the review looks at the current literature on the control and operation of AUV swarms with the focus of the field being on emergent flocking algorithms. Moving on, the review examines current approaches to monitoring the environment by autonomous platforms. For this section the focus broadens beyond techniques soley developed for AUVs monitoring the ocean to different environments with the view that these could be adapted for ocean monitoring. Through the review of AUV swarm literature and literature on monitoring the environment, the approaches which will be evaluated in varying swarm sizes and subject to communication constraints can be identified. Next, the review collates literature concerning how developers of multi robot teams and robot swarms have overcome communication constraints. This forms one of the major challenges to the implementation of an AUV swarm and the effect of communication constraints on coordination is a key area of investigation in this thesis. Finally, based on the review of literature, areas for investigation are recommended.

3.1 State of the Art in AUV Operation

Allowing online Adaptive Mission Planning (AMP) allows an AUV to make decisions based on knowledge of its own health and the environment. For example, if more energy is used in a section of a mission than expected, the AUV can decide to reduce the size of the survey area to allow for safe recovery. AMP is a common focus of AUV research. However, AMP is still rarely implemented in AUV operation. Brito et al. [2012] states the unpredictability of AMP is a key factor preventing the transition from research to commercial AUV operation. More recently, Brito et al. [2018] find the paucity of demonstration trials, which could be seen to correlate with predictability, is preventing the transition. The logistics, financial risk and risk to future science are all barriers to more trials being carried out in order to validate AMP approaches. Due to the lack of adoption of AMP, the state of the art approach to sampling using an AUV is to use predefined paths such as lawnmower surveys, continuous transects or virtual mooring approaches. These predefined paths minimise the complexity of the mission and reduce the risk of vehicle loss, while providing operators with confidence through obvious progression towards the goal and predictability in survey times [Pebody, 2007].

3.1.1 Examples of Current AUV Survey Techniques

Yoerger et al. [2007] present a common approach to AUV surveys using nested lawnmower surveys. An initial low resolution lawnmower survey, covering a large area is used to look for initial signs of hydrothermal activity. Once signs are detected, progressively higher resolution surveys can be performed to home in on hydrothermal vent sites. The common approach is for the AUV to surface between surveys in order for the operator to review the data, and plan the next survey. Figure 3.1 presents the survey paths from the high coverage, low resolution survey, Phase 1, and the progressively higher resolution surveys, Phases 2 and 3. Another example of nested surveys is presented by Wagner et al. [2013]

for cold-seep habitat mapping.

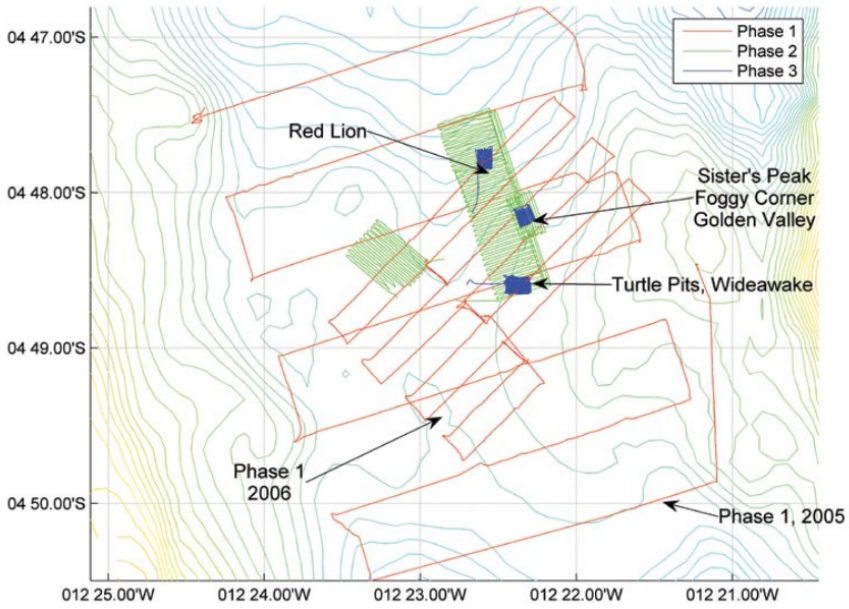

Figure 3.1: Progressively higher resolution lawnmower surveys performed by the Autonomous Benthic Explorer AUV in search of hydrothermal vent sites [Yoerger et al., 2007].

Robbins et al. [2006] present results from a REMUS AUV sampling a HAB. MODIS satellite imagery was used to inform two REMUS surveys, a 6km transect at depths of 1.5m and 4.5m intended to intersect frontal gradients along the bloom edge and a box transect performed at depths of 1.5, 3 and 8m. These surveys are shown in Figure 3.2.

3.1.2 Multi-Vehicle Deployments

While AMP has not yet transitioned from research to commercial deployments, in order to improve ocean surveys multi-vehicle deployments are increasingly being used. The series of MASSMO missions coordinated by the National Oceanography Centre, UK demonstrate the value of several vehicles operating in the same area [National Oceanography Centre, 2018]. However, there is a large step from multiple vehicles operating in the same area to the vehicles collaboratively working towards a common goal through autonomous sharing of data. This collaborative approach is still a topic of research and faces the same challenges as

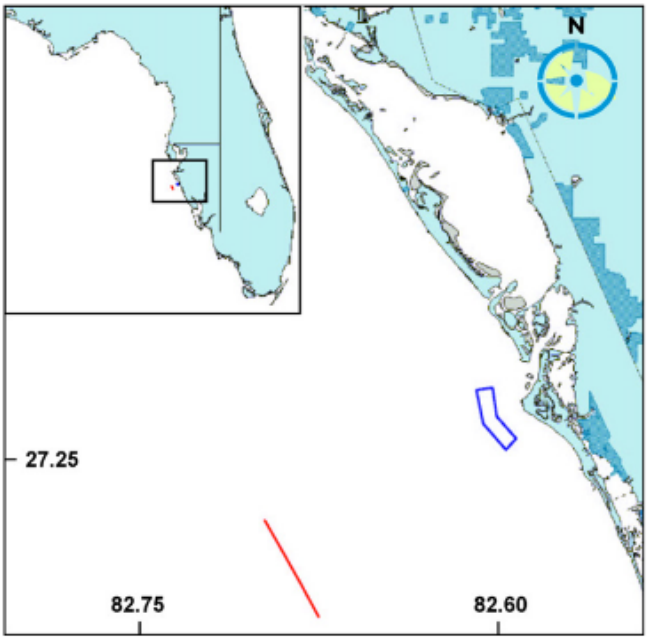

Figure 3.2: The transect (red) and box (blue) surveys performed by the REMUS AUV to sample a HAB along the west Florida shelf, January 2005, [Robbins et al., 2006]

.

AMP. These challenges are amplified as a vehicle may now make decisions based on both its own data and data received from other vehicles.

3.1.3 Summary

Simple lawnmower or transect surveys are still commonplace in commercial AUV operation and unpredictability and lack of proving trials are key obstacles, preventing the transition of AMP from research to commercial deployments. Multi-vehicle deployments are becoming more common, however the vehicles are piloted individually and do not share data for autonomous decision making. The implementation of an AUV swarm has potential to break down the barriers to adaptive mission planning and collaborative AUV missions as the logistics and financial risk are significantly reduced.

3.2 Current AUV Swarm Literature

In AUV research, there is a large focus on multi-vehicle approaches while the field of AUV swarms is still in its infancy. A multi-vehicle approach uses a small, possibly heterogeneous, team of sophisticated AUVs while a swarm consists of large numbers of simple, homogeneous vehicles. A strict definition of a swarm taken from Dorigo and Sahin [2004] is as follows:

- The swarm should be almost entirely homogeneous due to the high redundancy required
- Each swarm member should have only local sensing and communication capabilities
- The swarm should be scalable and not restricted to a maximum number of members.
- The implementation of a swarm must lead to increased task performance

A common approach to coordination of an AUV swarm is formation control or formation flying. This is likely due to its popularity in the multi-AUV field and application to AUV tasks such as bathymetric mapping, where tight control of relative vehicle locations is required to ensure individual vehicle sensor footprints align and provide an overlap. Formation flying could have applications in environmental monitoring. However, in order to achieve the spatio-temporal resolution required, a divide and conquer approach is seen as superior [Nieto-Granda et al., 2014]. See the following for formation control between AUVs: Amory et al. [2014], Millán et al. [2014], Lobo Pereira et al. [2015], Kemna et al. [2015].

Isokeit et al. [2017] presents a cooperative swarm behaviour between two MONSUN Swarm AUVs for insitu environmental sampling. The presented approach aims to divide sampling locations between a swarm of AUVs by first sorting the points into clusters with minimal inter-point distances then minimising mission time and balancing energy levels of individual vehicles. The results show the coordinated, two vehicle swarm reduces mission time by 43% compared to a single vehicle. The authors claim the approach is highly scalable. However

coordination relies on one vehicle to act as the coordinator to gather data and assign tasks based on it's knowledge of the swarm, implying a limit to scalability.

3.2.1 Emergent Behaviours

Within the field there is a specific focus on emergent flocking, likely due to its popularity in general swarm robotics. Read et al. [2014] and McGookin and McColgan [2017] investigate flocking performance in a simulated swarm of AUVs with varying communication constraints. Read et al. [2014] assume lossless, instantaneous communication between 11 CoCoRo AUVs. The authors investigate the impact on flocking performance when communication is firstly restricted to an AUVs local swarm, i.e. all vehicles within a defined distance, and then further to immediate neighbours. The study finds communication to be vital to the successful coordination of the swarm and restricting an individual AUVs knowledge of the swarm actually improves performance even in the presence of navigation error. McGookin and McColgan [2017] examines flocking in a swarm of 12, biomimetic AUVs, varying Time Divison Multiple Access (TDMA) frame lengths from 0.5 to 24 seconds. The study finds the algorithm is not able to successfully coordinate the swarm of AUVs at any of the evaluated frame lengths.

3.2.2 Summary

A swarm is defined as a group of small, simple and homogeneous AUVs. Formation control is a common approach to coordination of AUVs. However, this is more suited to seabed mapping tasks as it does not provide the coverage required for environmental monitoring. An approach is presented for environmental monitoring in swarm AUVs but uses a centralised coordinating vehicle and requires known sampling locations. Finally, emergent flocking is shown to be a poor approach to controlling an AUV swarm as the communication delay, even when using acoustic communication, means the swarm is not able to successfully coordinate.

3.3 Monitoring the Environment

Several adaptive approaches exist for controlling a robot swarm to achieve a sampling methodology. A key approach is to track the boundary of a chemical plume and studies often cite oil spills and HABs as potential applications [Petillo et al., 2012][Clark and Fierro, 2007][Susca et al., 2008]. While tracking a boundary could be useful, no information on the chemical distribution, the variation with depth or the underlying physical processes leading to the event can be collected by sampling where the plume is not present. In addition, the boundary of a plume is likely to be highly complex in the dynamic ocean environment and the likelihood of the entire swarm losing the tracked feature is much higher if all the members are positioned at the boundary. Other approaches to environmental monitoring include coverage control, aiming to distribute vehicles throughout an environment, and persistent monitoring, driving vehicles to continuously sweep an environment. These approaches are better suited to monitoring oil spills and HABs as they have the potential to capture the evolution of the feature in its entirety, including the maximum concentration, the chemical distribution and the boundary.

3.3.1 Current AUV Approaches

This section details current approaches to monitoring extremely transient ocean events such as oil spills and HABs, focussing on approaches developed specifically for AUV operation.

AUV Operations in Response to the Macondo Well Blowout

Ryan et al. [2013] and Camilli et al. [2010] detail the AUV operations carried out in the Gulf of Mexico following the Macondo well blowout in 2010. The fast survey speed and deep diving capabilities of the AUVs were described as essential for mapping and sampling target features near the wellhead. Ryan

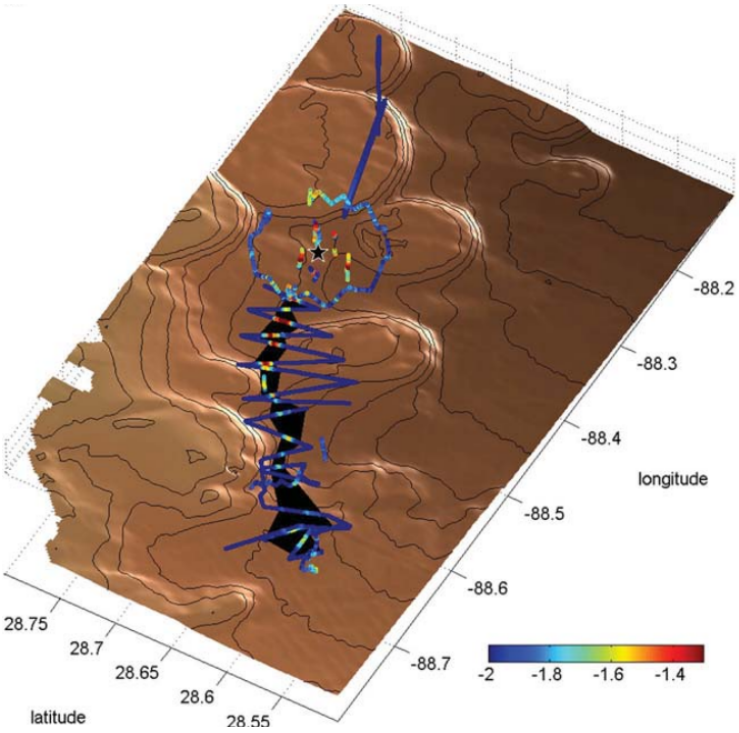

Figure 3.3: Path of the Sentry AUVs second survey in response to the 2010 Macondo well blowout. Relative methane signal intensity is shown on the colour bar. [Camilli et al., 2010]

et al. [2013] details the operations of the Dorado AUV, targeting the maximum signal localized by ship hydrocasts. The AUV performed vertical section surveys perpendicular to the vector between the well head and survey location. A total of 9 parallel, vertical sections were performed between 900 and 1200m below the surface. Camilli et al. [2010] details the operations of the Sentry AUV, deployed nearly 2 months after the initial well blowout. The AUV conducted three surveys, at depths in excess of 1000m for a total of 64 hours covering a distance of 2350km. The first survey was a east north-east radial projection from the source, conducted to rule out oil spreading to the North. The second survey dove to a constant depth of 1120m in a zig-zag path southwest of the source, this survey is shown in Figure 3.3. Finally, the third survey extended to the southwest where hydrocarbon readings diminished at 27km from the source. On this survey the AUV used dynamic re-tasking to track lines at differing depths to determine a hydrocarbon maxima at 1160m.

Oil Spill Patch Tracking

Kukulya et al. [2016] present a patch tracking algorithm for a single, long range AUV with the aim of mapping oil spills. The approach is shown in Figure 3.4 and aims to continuously transect the plume, revisiting the centre as frequently as possible using an algorithm originally designed for HAB tracking and presented by Godin et al. [2011]. While the approach is reasonably straight forward, one clear problem is the inability to perceive multiple maxima in the chemical distribution. The algorithm also does not appear robust to irregular shaped features as it relies on the AUV to detect when it has completed a transect of the plume purely through sensor measurements.

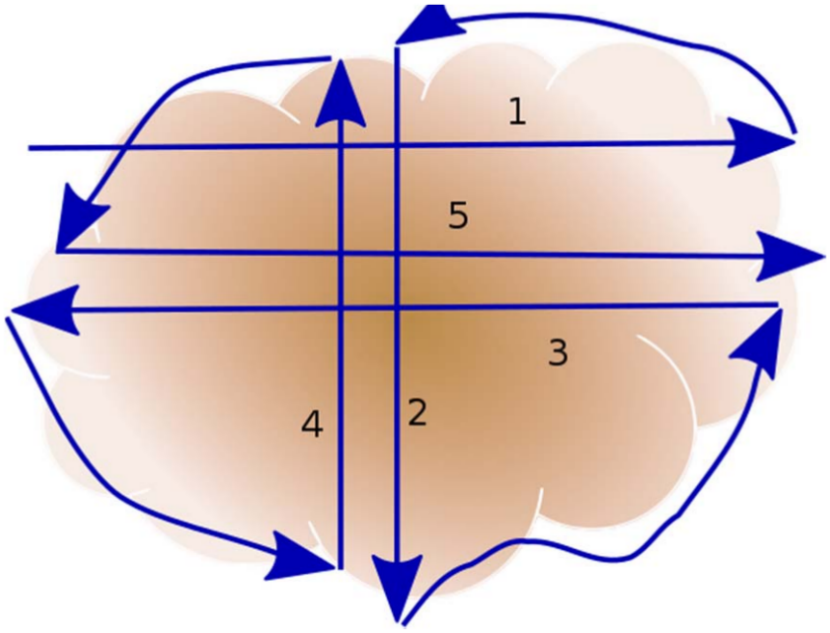

Figure 3.4: Example of the process of transecting a patch presented by Kukulya et al. [2016]

Fox et al. [2012] introduce plan-based policy learning to the patch tracking problem. The authors highlight one of the major problems in patch tracking is not the tracking, but relocating the patch once lost. The study examines using previous experience of tracking similar patches to aid in tracking the patch and reacting to losing the patch entirely. The approach shows improvement over a static policy and demonstrates robustness to different patch shapes.

Predicting Bloom Trajectory

Das et al. [2010] utilise data from a range of sources to predict the trajectory of a HAB patch and in turn inform the survey path of an AUV. The authors discuss the various difficulties in sampling a HAB including the large spatial characteristics, the complex variability in coastal waters, poor understanding of the underlying ecology driving the bloom, delay in obtaining satellite data and the spatial sparsity of mooring data. Predicting the trajectory of the bloom over months requires better ecological understanding, however to predict the trajectory over a week, the advective effect of ocean currents on the bloom can be used. The study shows the approach is promising for tracking large blooms and planning simple lawn-mower style surveys with multi-vehicle control stated as future work. Smith et al. [2010] also present an approach to track plankton blooms based on ocean model predictions, implementing the approach using two underwater gliders. Similar to Das et al. [2010], the approach takes remote sensed satellite data and uses an ocean model to predict the short term movement of the bloom e.g. 16 hours. This is then used to plan a sampling survey for the two gliders, one focussing on the centroid of the plume and the other on the boundary. One problem outlined by Smith et al. [2010] is that upon separation of the bloom into multiple features, only one of these can be continually tracked. The study reports good coverage of the plume centroid, but poor performance in tracking the boundary in a deployment spanning 3 days. The poor boundary tracking is due to the ROMS ocean model advecting the plume in the opposite direction to currents observed by HF radar, used as a baseline. This is possibly due to the low magnitude of the currents showing no one, dominant, driving force behind the currents.

Occupancy Grids

Jakuba and Yoerger [2008] use occupancy grids (OG) to locate all plume sources in environments where the number of sources is unknown and confidently con-

clude no other sources exist. The algorithm uses both plume detection and plume absence to inform the binary state of the OG cells and measures water current velocity, allowing plume propagation to be tracked back to the source. The AUV measures its position within an acoustic transponder net, allowing plume measurements to be recorded against location, vital for the construction of a consistent OG map. The paper constructs the OG maps in post processing, however states online computation would allow autonomous design of surveys. Saigol [2011] extends the work of Jakuba and Yoerger [2008] to online planning for hydrothermal vent prospecting using an AUV.

Summary

AUV monitoring in response to the Macondo well blowout in 2010 shows commercial AUV operation is still limited to using basic survey patterns including transects and lawnmower surveys. Approaches to patch tracking have been developed. However, approaches implemented in AUVs are extremely simple, cannot comprehend multiple features and is susceptible to irregular shaped features. Predicting bloom trajectory is shown to be a useful tool in planning the trajectory of AUVs and occupancy grids are a useful way to represent the distribution of information in an environment for ocean monitoring.

3.3.2 Coverage Control

Coverage control aims to drive individual swarm members to a desired distribution in order to cover an area. Examples in literature use wheeled robots and show distributions defined by inter-vehicle separation, a Gaussian function and adaptive distributions based on sensor measurements. Approaches utilise artificial potential fields and Voronoi partitions to achieve coverage. An example of the initial positions, sensory function and final positions of a swarm controlled by a coverage control algorithm is shown in Figure 3.5. Coverage control has not yet been used in an AUV swarm and no literature was found implementing

coverage control in an AUV swarm.

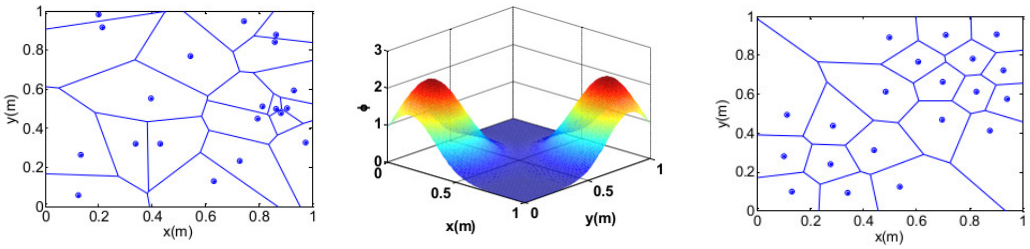

(a) Initial Positions (b) Sensory Function (c) Final Positions

Figure 3.5: Example of a swarm following an adaptive coverage control algorithm in a bimodal Gaussian sensory environment, presented by Dirafzoon et al. [2010]

Artificial Potential Approaches

Howard et al. [2002] present a mobile sensor network deployment algorithm, synonymous to coverage control, using potential fields. In the approach, virtual forces are used to repel swarm members away from each other and from obstacles, forcing them to spread out across an area. A viscous friction force opposing the virtual force then allows the swarm to achieve an equilibrium state. The authors claim, as each member is able to sense range and bearing to neighbours, no communication is required and the approach is therefore highly scalable. However, the final distribution of vehicles in terms of inter-vehicle separation is difficult to predict as it depends on external phenomena, such as obstacles in the environment, and internal factors, such as the relative weights given to the virtual forces. Poduri and Sukhatme [2004] introduce a different approach, using artificial potentials to maximise coverage but also to maintain a minimum number of neighbours for each node in the network. Instead of the viscous friction force used by Howard et al. [2002], Poduri and Sukhatme [2004] use an attraction force to impose the minimum neighbour constraint and therefore limit the expansion of the swarm and achieve a stable state.

Voronoi Partition Approaches

Another approach to the coverage control problem is to use Voronoi partitions, also known as Voronoi tessellations, Schwager and Rus [2010] argue approaches to coverage control that utilise Voronoi partitioning are best suited to distributed actuation and sensing tasks. Decentralised coverage control in a mobile sensor network using Voronoi partitioning is introduced by Cortés et al. [2004]. The algorithm presented drives agents towards the centroid of their Voronoi cell, achieving a distribution defined through a gradient descent algorithm applied to a Gaussian density function. The authors argue the properties of a Voronoi diagram determine the average number of neighbours to be less than or equal to six, leading to a scalable network, and the final distribution of vehicles is directly controlled through the Gaussian function. The algorithm by Cortés et al. [2004] is adapted by Lindhé et al. [2005] to achieve flocking with obstacle avoidance, using a distribution defined by inter-vehicle separation. The paper emphasises the minimal communication required compared to other flocking behaviours. The minimal communication dependency when using Voronoi partitions not for coverage control but for task allocation between a group of AUVs is also emphasised by Kemna et al. [2017].

Finally, the state of the art in coverage control is being able to distribute robots in an unknown environment and have them configure to the optimal sensing distribution through in situ sensor measurements. Schwager et al. [2017] presents the nuclear disaster at Fukushima, Japan in March 2011 as an example application, stating the immediate deployment of a group of robots could provide invaluable information on radiation levels, focussing on areas of high interest. Dirafzoon et al. [2011] achieves this using neural networks, Schwager et al. [2007] uses a basis function tuned through learned parameters and Lynch et al. [2008] uses basis functions in conjunction with Kalman filtering. One problem encountered is gathering enough information to accurately represent the environment before moving to the optimal sensing locations. Schwager et al. [2008] build on

the work presented by Schwager et al. [2007], introducing an exploration factor. The approach uses a coverage force to attract the vehicle towards the optimal sensing position but also subjects the vehicles to an exploration force, forcing the vehicle to circle the optimal sensing location to gather more information on the environment. Schwager et al. [2017] takes a different approach and introduce an exploration phase, in which the robots divide the sampling region using a K-means clustering algorithm. Each robot determines a path to cover its sub-region before moving on to perform coverage. Results presented indicate the exploration phase bounds the learning error and creates a robust, repeatable solution to the coverage problem.

AUV Constraints in Coverage Control

AUVs are commonly flight-style, or non-holonomic, for efficiency and simplicity in control. Kwok and Martìnez [2010] argue non-trivial vehicle dynamics can invalidate the stability proofs of coverage control algorithms where omni-directional vehicles are assumed. Non-holonomic vehicles are considered by Lindhé et al. [2005] and Luna et al. [2010], the latter uses a non-linear steering control law to guarantee stability in adaptive coverage control. However, Luna et al. [2010] carry out simulation experiments in a $1m^2$ area, hence the non-holonomic nature appears to be a problem. However, in the open ocean, AUVs are likely to be hundreds of metres away from other swarm members, and hence the effect of not considering the non-holonomic nature of vehicles in the controller design is likely negligible. Visually, it can be considered as attempting to parallel park a car in a busy street compared to parking a car in an large car park with only one or two other vehicles.

One key observation is that none of the mentioned papers include considerations for a communication delay, a severe limitation for a swarm of AUVs. Lynch et al. [2008] particularly, require information to disperse throughout the network which will significantly increase the delay in bandwidth limited communications.

Communication latency is introduced into the coverage control problem by Bullo et al. [2012], Miah et al. [2015] and Patel et al. [2016]. Bullo et al. [2012] introduce 'gossip' communication defined as asynchronous, short range, pairwise and potentially unreliable communication where the robot pair that communicates can be selected deterministically or randomly at each time step. Approaches to Voronoi coverage control discussed thus far define regions based on each agents location, meaning each movement of an agent changes both its own region and neighbouring regions. However, the algorithm presented by Bullo et al. [2012] dictates how to update the regions, ignoring the agents position. As the agents position is no longer of concern, the agent is free to explore its region and move to meet a neighbor in order to 'gossip' and update its region. Moving back to classic Voronoi coverage approaches, Miah et al. [2015] present a stochastic, intermittent communication network with simulated packet loss but no consideration of range limitations. The algorithm uses full state estimators on each vehicle to bridge the communication gap to achieve non-uniform coverage. Asynchronous one-to-basestation communication is examined by Patel et al. [2016], choosing a random agent to communicate at each time interval. The approach uses overlapping territories to achieve a centroidal Voronoi partition and divide a confined sampling region between agents.

Summary

Literature on coverage control using both artificial potential fields and Voronoi partitions has been reviewed. There is significantly more literature on Voronoi coverage control, possibly as the final distribution of vehicles is provable and more controlled than use of artificial potentials. Work has been done to extend Voronoi coverage control to unknown environments, in which the distribution of sensory information must be learnt, and varying communication constraints. No literature was found implementing coverage control in an AUV swarm.

3.3.3 Persistent Monitoring

As opposed to the lawnmower survey in which the vehicle is able to complete the task, a persistent monitoring task is never complete. Persistent monitoring requires a vehicle to continually sweep the environment, ensuring repeated coverage of areas of interest. Figure 3.6 presents the example of a robot constantly patrolling a car park, looking for new arrivals and checking for cars over staying their booked time. In the example of an oil spill, a vehicle may be required to persistently survey between the source of the spill and nearby marine protected areas to monitor the progress of oil towards protected ecosystems. The two key challenges in persistent monitoring appear to be planning the path through the area to optimise information metrics and controlling the speed of the vehicle along the path, in turn controlling sampling resolution.

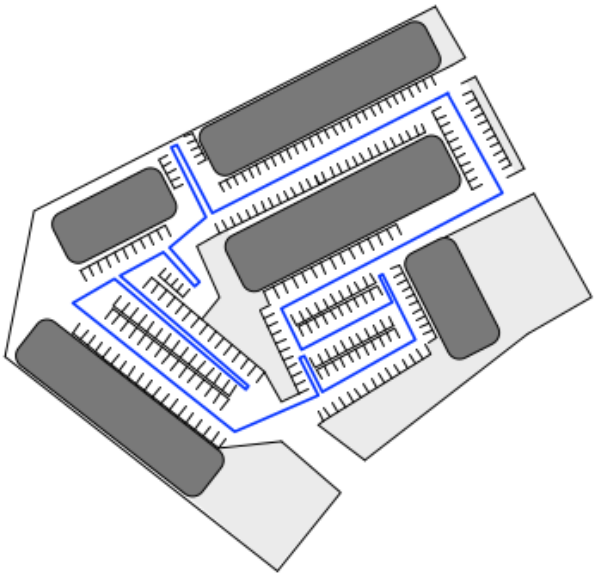

Figure 3.6: Example of a robot persistently monitoring a car park to detect any new arrivals or over stays, [Asghar and Smith, 2014]

.

Smith et al. [2011a] presents a persistent monitoring approach for an underwater glider, optimising a cost function to plan a path which maximises the information value while minimising deviations from the path by avoiding areas of high current. Through glider deployments in Monterey Bay, California, the al-

gorithm was shown to cover more distance, in less time and at a denser sampling resolution within predefined areas of high interest. Smith et al. [2011b] presents a solution to the persistent monitoring problem using a speed controller along a predefined path to minimise an accumulation function. The accumulation function represents uncertainty in the environment, growing with a 'production rate' in areas out of range of the sensor footprint of the robot and decreasing in areas observed by the robot. An optimal speed controller is determined to control the vehicles speed along the predetermined path and keep the accumulation function as low as possible. Smith et al. [2012] extends this method to multiple robots, however concludes stating unknown and time-varying production rates as future work, requiring online learning and calculation of the speed controllers.

Lan and Schwager [2016] present an approach to plan periodic trajectories for robots to estimate a spatio-temporal field in a dynamic environment. Two path planning algorithms are used to minimise the largest eigenvalue of the error covariance matrix of the Kalman filter. The eigenvalue represents where the error in the maintained Kalman filter model is most likely to change and hence by minimising the largest eigenvalue, the likelihood of the error in the model increasing can be minimised. The approach is demonstrated using ocean temperature data from the Caribbean Sea. Garg and Ayanian [2014] also present a solution to the persistent monitoring problem in a spatio-temporal environment evolving in a stochastic manner. It is not possible to model real world stochastic phenomena, Garg and Ayanian [2014] focus on modeling information metrics such as entropy, which represent the informativeness of sensor readings, to guide a small team of robots through the environment. The use of Gaussian Processes in the study allows for non-parametric modeling which suits the assumption that no prior knowledge of the phenomena is available to the robots.

3.4 Overcoming Communication Constraints

The substantial communication delay is a problem unique to underwater operation. This section presents how communication delays have been overcome in areas outside of those discussed so far.

As mentioned previously, Miah et al. [2015] use full state estimators to bridge the communication gap in an intermittent, stochastic network. Modeling the state of other swarm members is a common way to approach the problem. Lobo Pereira et al. [2015] present an approach to formation keeping using a fast low-level controller and a slow high-level controller. If there is packet loss or delay in the network, the low-level controller is able to estimate the states of other AUVs until communication is restored using Model Predictive Control (MPC). The controller is shown to be able to withstand temporary communication dropouts and recover formation. Lobo Pereira et al. [2015] state non-linear modelling as future work. Yamchi and Esfanjani [2017] presents a similar approach using MPC to account for communication delays in networked robots. However, the authors state the system becomes unstable at 180ms delay in the communication network, significantly less than the magnitude of delays in acoustic communication.

Millán et al. [2014] present an approach using H_2/H_∞ control in combination with a feedforward controller. The feedforward controller makes formation control in the presence of disturbances i.e. communication delays and packet loss, possible. The controller assumes all communications are affected by an average delay and, even though this is a significant assumption, still improves controller performance.

Finally, Hollinger and Singh [2010] use receding horizon planning to ensure periodic connectivity in a multi-robot team. This allows the robots to search, survey and cover an environment, alleviating the constraint of requiring constant connectivity for success but ensuring periodic communication.

3.5 Recommended Areas for Investigation

This section details some of the areas for more research from the reviewed literature. Through identifying the areas for more research, the contributions and novelty of this thesis can be justified.

Within the literature, there is a large focus on emergent flocking for AUV swarms. However, the field has not yet considered the application of coverage control algorithms, arguably more useful for ocean science, to an AUV swarm and the communication constraints associated with AUV operation. Different communication constraints are considered in general coverage control, however these do not cover the extensive limitations of AUV communication and hence this is a an area for more research crucial to the implementation of an AUV swarm.

As coverage control has not yet been applied to AUV swarms, we look to the wider field of general swarm robotics. Often in coverage control, only two dimensional environments are considered. However the third dimension, depth below the ocean surface, is important to monitoring ocean features as both chlorophyll blooms and oil spills can have subsurface maxima. Identification of subsurface maxima for ocean sampling is shown by Zhang et al. [2010]. However, transferring this approach to a distributed algorithm for an AUV swarm is not trivial, hence coverage control with consideration of the chemical distribution in the third dimension is an area for more research which would improve monitoring performance.

The effect of a dynamic environment is also not considered in general coverage control literature. A dynamic environment such as the ocean can have an effect on both the feature being monitored and the navigation of vehicles. This is a key consideration for monitoring ocean events which is still unanswered.

Exploration versus exploitation is a common debate in optimisation, how much a domain should be explored for solutions before exploiting the knowledge to find an optimum solution. Exploration has been considered in general

coverage control literature. However the approach is either to explore the space immediately adjacent to the sampling location or to cover the entire environment before performing the coverage control algorithm. In the open ocean, exploring immediately adjacent to the sampling location is unlikely to provide any new information and exploring the entire environment is potentially a huge task. In addition, a common approach in optimisation literature is to initially perform a random walk. However this does not increase predictability for operators nor ensure the safety of the vehicle. An efficient approach to exploration, possibly based on the spatial characteristics of the feature being monitored, for coverage control in the ocean environment is an area for more research.

Persistent monitoring is a partial solution to the exploration problem as the vehicle continually sweeps the environment, looking for changes. However, there is still an optimisation factor as the aim is for a vehicle to spend more time sampling where the environment changes more rapidly and planning a path through the environment is still required. Persistent monitoring approaches often require predefined paths, knowledge of the environment or machine learning approaches to plan a path through the environment.

Chapter 4

The Model

This chapter details the modelling of individual vehicle dynamics and behaviour. In addition, the assumptions made for acoustic and satellite communication, the two channels by which the swarm can communicate, are presented to allow the simulated vehicles to exchange data in a manner similar to a real AUV swarm. Through modeling the swarm, potential swarming algorithms can be subject to underwater communication constraints. Through running simulations, the ability of the swarm to achieve an acceptable level of coordination between swarm members is determined and the scalability of the algorithm to larger swarm sizes and more severe communication constraints can be evaluated.

In order to validate the model, data on how the vehicle behaves in reality is required. A deployment to the Shetland Isles in 2017, location shown in Figure 4.1, provides some of the first data concerning how the ecoSUB behaves in open water. During the deployment, the ecoSUB completed a series of missions imitating a pipline survey. The data from these missions captures not only vehicle dynamics but also drift error from disturbances in the form of currents and tides. Comparing this data against a simulated vehicle, given the same starting position and waypoints, allows the parameters used in the model to be validated and provides insight into the magnitude of drift error from the dead reckoning navigation.

Figure 4.1: Location of ecoSUB trials in the Shetland Isles.

The next sections present: the development of the model; the method of validation, which details the mission given to the real vehicle and the parameters to be compared; the comparison between the real and simulated vehicle dynamics; and the concluding remarks.

4.1 The AUV Model

4.1.1 Vehicle Dynamics

Figure 4.2: The ecoSUBμ, a small, low cost AUV suitable for use in a swarm configuration.

The individual vehicles in the simulated swarm are modelled on the ecoSUBμ (Figure 4.2), a small, low cost AUV suitable for swarm deployments (Phillips et al. [2017]). The simulation of each individual swarm member follows Algorithm 1.

Algorithm 1 outlines the simulation of a single swarm member. All actions

are encompassed in a simulation time loop. Each time step represents 0.5s and the simulation time i.e. the time simulated not the time the simulation is executing on the computer, is defined by the user. At each timestep, the swarm member starts by evaluating (a) whether it is on the surface and (b) whether is has communicated via satellite since being on the surface (Line 2). If the vehicle is on the surface and hasn't yet communicated via satellite, the vehicle exchanges data with the simulated base station and enters a time delay typical of satellite communication (Lines 3 - 5). Next, the vehicle checks if the current timestep coincides with its designated frame in the TDMA loop (Line 6). If it does, the vehicle broadcasts its state to the rest of the swarm via the acoustic channel (Line 7). Moving on, if the vehicle is communicating via satellite it must be stationary on the surface. Line 9 checks the vehicle isn't in the satellite communication delay and hence performs the required behaviour. If the swarm is flocking, the vehicle sets its speed and heading demand according to the flocking algorithm (Line 10 & 11). If the vehicle is executing the Voronoi Coverage Control algorithm, it sets the waypoint in the waypoint following behaviour according to the coverage control algorithm. Line 14 checks whether the vehicle is (a) within the acceptance radius of the waypoint, or (b) exceeding the time out given to reach the waypoint. If either of these is true, the vehicle returns to the surface to communicate via satellite and acquire new information. Once the vehicle has checked communication channels and set any demands according to the desired behaviour, the vehicle dynamics are modelled. Line 20 sets the yaw rate based on the error between the yaw of the vehicle and the yaw demand. For example, if the yaw of the vehicle is extremely different to the yaw demand, it will apply the maximum yaw rate. Finally, Line 21 calculates the vehicles position based on dead reckoning taking into account the previous location, heading, pitch and speed of the vehicle to calculate the new position. This navigation is perfect due to the lack of simulated disturbances.

The ecoSUBμ's flight style, or non-holonomic, nature is modeled in simulation

and hence the vehicle requires forward motion to make changes in yaw and depth. In addition to the flight style nature, the key parameters used in the ecoSUBμ model are shown in Table 4.1.

Table 4.1: Characteristics of ecoSUBμ vehicle used in modeling

Parameter	Value
Max Speed	1.0m/s
Pitch Range	-40°/+40°
Operating Speed	0.7m/s
Max Yaw Rate	11.5°/s
Max Pitch Rate	10.0°/s

The Equation 4.1 simulates the vehicle dynamics. In Equation 4.1, z is the depth of the vehicle below the surface of the ocean, v is vehicle speed, Δt is the time elapsed since the last calculation, ψ is yaw, θ is pitch and Φ is 1 degree of latitude in kilometres (111.111km).

$$\begin{bmatrix} \varphi_{t+\Delta t} \\ \lambda_{t+\Delta t} \\ z_{t+\Delta t} \end{bmatrix} = \begin{bmatrix} \varphi_t \\ \lambda_t \\ z_t \end{bmatrix} + v \begin{bmatrix} \cos\theta \cos\psi \\ \cos\theta \sin\psi \\ \sin\theta \end{bmatrix} \Delta t H \tag{4.1}$$

$$H = \begin{bmatrix} 1/\Phi \\ 1/\Phi \cos\varphi_t \\ 1 \end{bmatrix} \tag{4.2}$$

Buoyancy forces nor gravity are accounted for in simulation. When ballasting an AUV, the aim is to make the AUV as neutrally buoyant as possible i.e. the buoyancy is balanced by the weight and there is no resultant force. This is the aim as if the AUV is neutrally buoyant it requires minimal energy to dive to the required depths below the surface of the ocean.

4.1.2 Navigation

The vehicle is assumed to have perfect navigation using dead reckoning and following Equation 4.1. Perfect navigation is a significant assumption to make

in the dynamic ocean environment. However, the near-surface characteristics of potential applications e.g. HABs and surface oil slicks, mean regular surfacing to correct for drift error is a feasible possibility, bounding navigation error through GPS positioning. In future, it may be possible to implement acoustic localisation techniques using the information communicated by each swarm member, reducing navigation error without increasing the load on the acoustic channel.

4.1.3 Waypoint Following Behaviour

The simulated vehicle uses a simple, line of sight waypoint following behaviour. Given a latitude, longitude and depth, with acceptance ranges on these values, the pitch and yaw demand can be set. The behaviour is limited by a time-out at which the AUV returns to the surface to reflect the safety precautions taken in real operation.

4.2 Modeling Communication

In order to simulate acoustic and satellite communication between swarm members the following assumptions were made.

4.2.1 Satellite Communication

- Is only possible on the surface.
- Follows a one-to-basestation topology.
- The basestation is able to store received data.
- The vehicle experiences a delay, T_s, on the surface typical of the data exchange process.

An appropriate value of T_s for the ecoSUBμ vehicle has been determined from field data as three minutes. The frequency of satellite communication is controlled by the time-out to reach a given waypoint. Time outs of five and ten minutes are used in simulation, as in reality the vehicle will need to surface to correct for drift

error within this time frame. The time out defines the minimum frequency of satellite communication as if an AUV reaches the waypoint before the time-out has passed, it will surface to communicate and re-task.

4.2.2 Acoustic Communication

- Is not possible during satellite communication.
- Follows a Time Division Multiple Access (TDMA) protocol with a defined frame length, T_f.
- Has limited range of 2km.
- Is lossless.

When an AUV communicates via the acoustic channel, any AUV within range receives the communicated data. Typically in acoustic communication, the sender will send a message and wait for a response from the recipient for verification. As acoustic messages can become corrupt while travelling through the water column, responses will often include the information sent to the recipient for verification leading to more reliable communication. However, in this work we assume one-way broadcasting which, while not as reliable as two way, hugely improves the scalability of the approach when the acoustic channel is in high demand, shown in Figure 4.3. In addition, neighbours will move in and out of acoustic communication range, therefore monitoring which neighbours are in range to send messages to, and hence expect responses from, is not possible. Communication is also assumed to be lossless, this allows the effect of the latency in acoustic communication to be solely evaluated before introducing further hindrances.

The TDMA Medium Access Control (MAC) protocol used in this study allocates each swarm member a sequential time slot, or frame, in which to broadcast its state information. The frame length for TDMA, T_f, is determined by a transmission time, based on the bandwidth of the acoustic modem, B, and the estimated message size, S, and a guard time, based on the propagation delay of the acoustic signal through the water column, U_p, and the maximum range of the

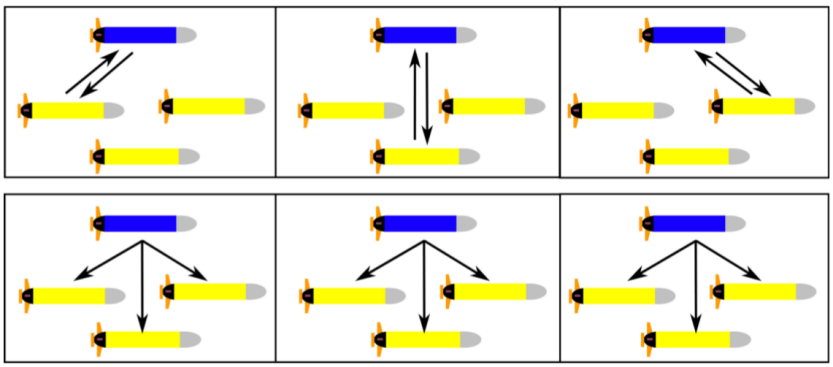

Figure 4.3: An AUV communicating with swarm members in a two way communication topology (Upper) and one way broadcasting (Lower)

modem, R. The equation for T_f is shown by Equation 4.3, the rounding of the guard time is for simplicity of simulation.

$$T_f = \frac{B}{S} + ceil\left(\frac{R}{U_p}\right) \tag{4.3}$$

Using $B = 4\,bytes/s$, $S = 8\,bytes$, $U_p = 1500m/s$ and $R = 2km$, $T_f = 4s$. $T_f = 6s, 8s, 10s$ are also used in this work as in future, larger messages containing information such as sensor measurements may need to be communicated or more infrequent communication could be used to save energy.

4.3 Method of Validation

This section details the missions passed to the ecoSUB vehicle in deployment which will be replicated in simulation. In addition, the key parameters which will be examined in the comparison of the real vehicle and the simulated vehicle will be discussed.

4.3.1 The Mission

The Shetland deployment aimed to demonstrate the potential of the ecoSUB AUV to survey sub-sea pipelines in the Shetland Isles. Data used in this comparison is taken from a mission performed in preparation for the actual trials. The mission consisted of navigating to four waypoints at varying depths, in a zig-zag

Algorithm 1 Algorithm to simulate individual swarm members. Let t be elapsed time, n be the size of the swarm, i be the unique identifier of the AUV $i \in \mathbb{Z}$, $p_i(\varphi, \lambda, z)$ be the location (Latitude, Longitude, Depth) of AUV i, P be the positions of swarm members as known by the vehicle, V be the velocities of swarm members as known by the vehicle, $w(\varphi, \lambda, z)$ be the location of the waypoint for the vehicle to move to, T_f be the frame length in which each vehicle communicates, R_{acc} be the acceptance radius on waypoints and Ω defines whether the vehicle has communicated via satellite.

while $t < t_{Run}$ **do**
  if $p_{i,z} = 0$ AND Ω = False **then**
    Exchange data with basestation
    Enter satellite delay, T_s seconds (Sec. 4.2.1)
    $\Omega \leftarrow$ True
  else if $|(t - floor(t/T_f n)T_f n)/T_f| = i$ **then**
    Send Acoustic Message (Sec. 4.2.2)
  end if
  if NOT in satellite delay **then**
    if Flocking = True **then**
      $\psi_{demand}, v_{demand} \leftarrow Flocking(P, V)$
    else if Coverage Control = True **then**
      $w \leftarrow$ Voronoi Coverage(p_i,P) (Alg.2, pg. 114)
      if $\|p_i - w\| < R_{acc}$ or $t_{timeout} < t$ **then**
        $w \leftarrow (p_{i,\varphi}, p_{i,\lambda}, 0)$
        $\Omega \leftarrow$ False
      end if
    end if
  end if
  Model vehicle dynamics (Tab 4.1)
  Localise via dead reckoning (Eq. 4.1)
end while

pattern imitating crossing a sub-sea pipeline. The waypoints used in the mission are detailed in Table 4.2.

The waypoints are subject to an acceptance radius of 5m within which the vehicle determines the waypoint has been successfully reached. In addition to the waypoints detailed in Table 4.2, the ecoSUB vehicle started at a location of 60.391747°, -1.362773°. After the final waypoint is reached the AUV will halt and rise to the surface through its buoyancy.

Table 4.2: Waypoints used in mission

Waypoint	Latitude (°)	Longitude (°)	Depth (m)
1	60.39210	-1.36119	0
2	60.39076	-1.36100	3
3	60.39133	-1.35854	5
4	60.38999	-1.35835	8

4.3.2 Model Parameters

Table 4.1 in the previous section (Page 62) details the values given to parameters which define the vehicle dynamics in the model. The parameters include: maximum speed, pitch range, operating speed, maximum yaw rate and maximum pitch rate. The degrees of freedom are labelled on the diagram of an AUV in Figure 4.4.

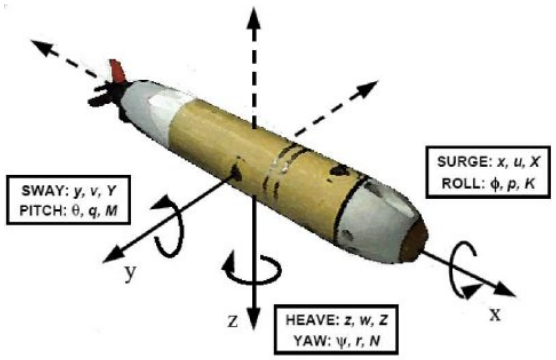

Figure 4.4: Degrees of Freedom of an AUV [Carlos et al., 2008]

The ecoSUB currently cannot measure its own speed, therefore infers the speed from the motor demand given by the operator. As this is then used in the dead reckoning navigation and there is no other source of location data when the vehicle is submerged, there is no way to obtain an actual value of submerged speed from the data. However, the ecoSUB moves to the first waypoint on the surface, hence navigates via GPS. From this initial portion of the mission a speed estimate can be obtained.

The pitch and yaw of the vehicle are measured by the inertial measurement unit on the ecoSUB and hence data from the entire mission can be compared directly to the simulation. In addition, depth is inferred from hydrostatic pressure

measured by a sensor onboard the vehicle, hence this can also be compared directly.

Finally, when the ecoSUB surfaces at the end of the mission, it acquires a GPS position and, through comparing its last position estimate from the dead reckoning navigation to this GPS position, an estimate of drift error can be obtained, often expressed as percentage of distance moved.

4.4 Comparison against Simulation

4.4.1 Vehicle Paths

Figure 4.5 shows the dead reckoned path of the real vehicle and the path travelled by the simulated vehicle. It can be seen, for the section up until waypoint 1 the paths of the real and simulated vehicle align well, this is because the real vehicle is on the surface for this section, navigating via GPS. As mentioned previously, from this section a velocity estimate of 0.6m/s was obtained and was then used in simulation. However, the speed of the vehicle may change when the entire body is submerged, increasing drag, and when it must use some speed to maintain depth.

Through waypoints 2, 3 and 4 the vehicle paths do not align so well due to the real vehicle now navigating by dead reckoning, hence the path shown here is only the estimated path and the actual path is unknown. The GPS fix attained by the vehicle upon reaching waypoint 4 and surfacing is shown by the 'x' and provides an estimate of the dead reckoning navigational error. The error is calculated to be 14% of the distance travelled, not too dissimilar to the value of 10% stated by Leonard and Bahr [2016] for a similar class of AUV.

4.4.2 Dive Profile

Moving on to compare the dive profiles of the real and simulated vehicles. Figure 4.6 presents the depth of the vehicle against time elapsed. After the initial surface

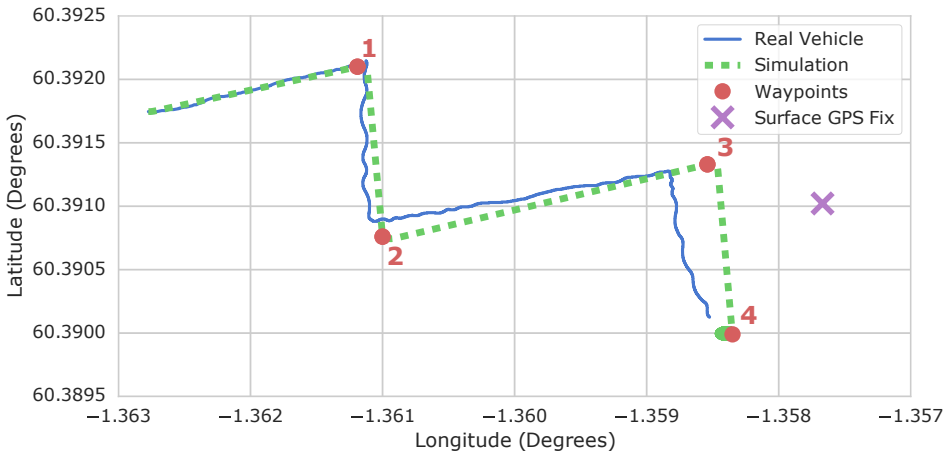

Figure 4.5: Overview of the dead reckoned path by the ecoSUB and navigated by the simulated vehicle. The vehicles move from west to east and the GPS position obtained upon surfacing is also shown.

period, the two profiles differ for the same reason as the vehicle paths, the switch from GPS to dead reckoning navigation. In addition, the speed of the real vehicle is likely not constant as it moves through the water, whereas the speed of the simulated vehicle is presenting another reason for the difference. However, the rate at which the two vehicles change their depth match well, with the simulated vehicle diving slightly faster then the real vehicle.

For the portion of the dive at 5m the real vehicle can be seen to maintain its depth whereas the simulated vehicle continues diving at a slow rate. This is due to a difference in the waypoint following behaviour between the two vehicles. The real vehicle controls its depth independently of its position in latitude and longitude. Conversely, the simulated vehicle determines its pitch angle based on the distance to the waypoint and the waypoints depth, calculating a sufficient pitch angle to achieve the required depth upon reaching the waypoint.

4.4.3 Yaw Behaviour

Figure 4.7 presents the yaw, or compass heading, of the real and simulated vehicles throughout the mission. The yaw of the real vehicle can be seen to fluctuate largely. While the demand is not shown by the data, this is indicative of a con-

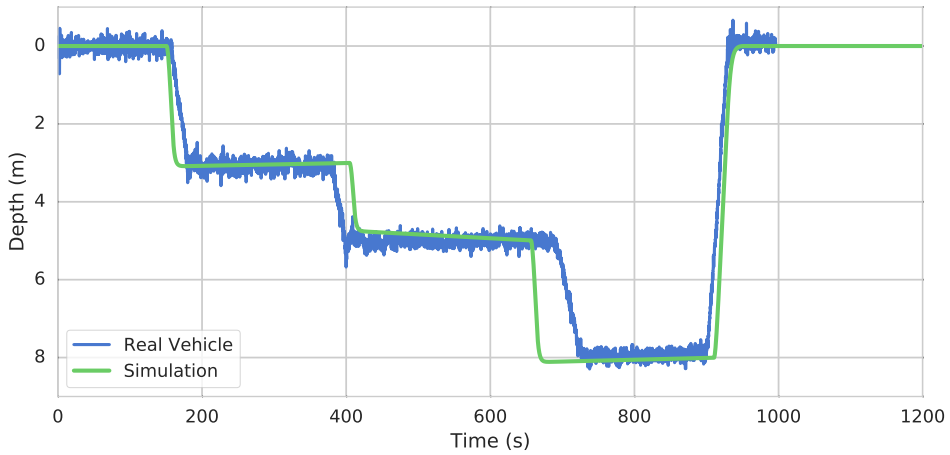

Figure 4.6: Dive profiles of the real and simulated vehicles throughout the Shetland deployment

troller tuning problem. Regardless, the two plots match well, with the real vehicle achieving a slightly higher yaw rate indicated by the gradient of the line between the sections of 'constant' yaw.

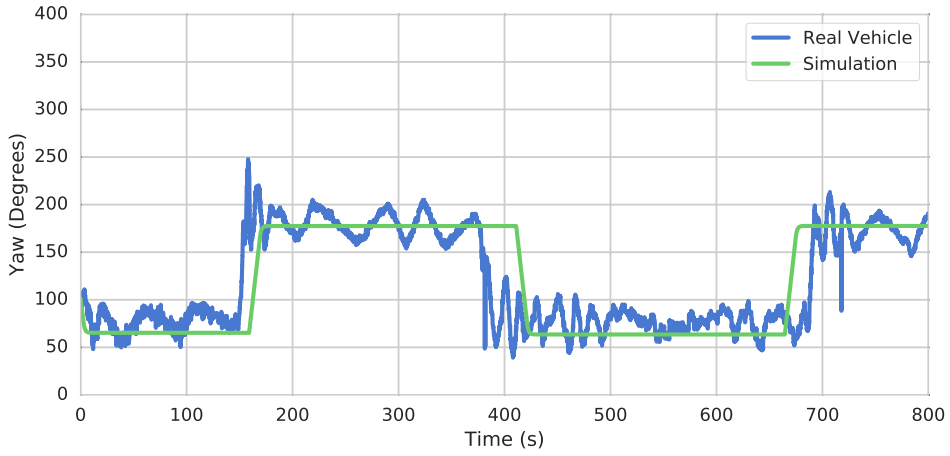

Figure 4.7: Yaw against Time for both the real and simulated vehicles

4.4.4 Pitch Behaviour

The final behaviour to be compared is the pitch, presented in Figure 4.8. The maximum pitch of 40° correlates well with the maximum pitch achieved by the vehicle, bearing in mind sensor noise may account for some peaks in the real vehicle data. The largest difference between the two plots is that the real vehicle

maintains a pitch of approximately 20° between depth changes. The vehicle adopts this pitch angle to counteract its buoyancy and maintain depth. However, the buoyancy forces are not accounted for in simulation, hence why a positive pitch is required at the end of the mission for the vehicle to drive itself to the surface. Finally, it is difficult to determine a pitch rate due to sensor noise present in the data, hence from observation the rate at which the vehicle is able to adjust its pitch seems to match well between the real and simulated vehicles, Figure 4.9 shows the period of pitch change from 150s to 200s. Using the period from 150s to 160s, crude pitch rate estimations of $6.4°/s$ and $6.8°/s$ can be made for the simulated and real vehicle respectively.

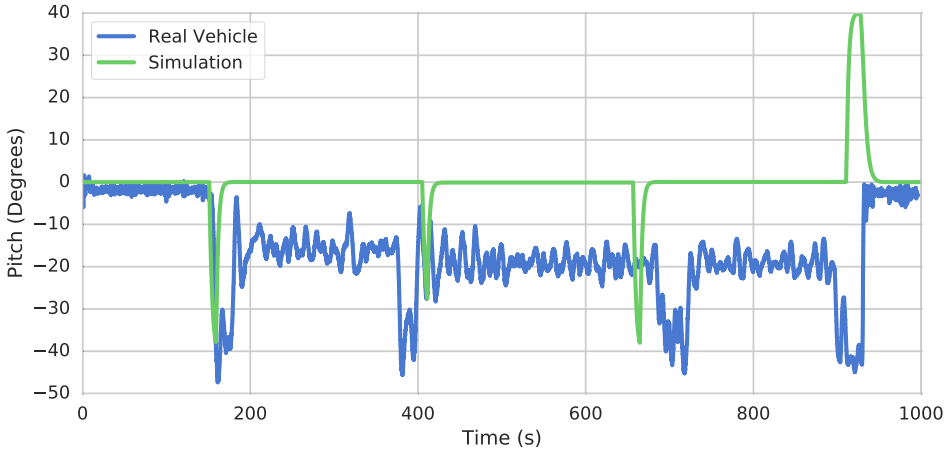

Figure 4.8: Pitch against Time for both the real and simulated vehicles

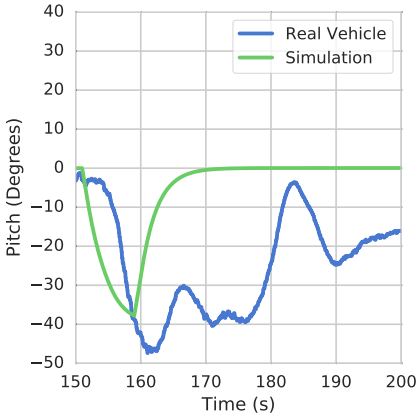

Figure 4.9: Pitch against Time for both the real and simulated vehicles from 150s to 200s

4.5 Summary

The individual AUVs in simulation are modelled on the ecoSUBμ, a small, low-cost micro AUV. The vehicle dynamics are modelled by a pitch rate, a yaw rate and a maximum speed. Navigation is performed via dead reckoning and is assumed perfect. While underwater navigation is far from perfect, the potential applications of the AUV swarm concern near surface ocean phenomena such as oil spills and harmful algal blooms, making frequent surfacing to correct for navigation drift error a feasible possibility. In addition, acoustic localisation techniques are present in the literature enabling an AUV swarm to reduce navigation error without the need to surface. The swarm is able to communicate via satellite and acoustic communication with the relevant delays. Satellite communication is available on the surface and hence frequency of communication is controlled by AUVs successfully reaching waypoints or reaching the defined time-out for the waypoint following behaviour. Acoustic communication is implemented as a one way broadcast to all AUVs in range of the sender and a TDMA MAC protocol is introduced as all AUVs must share a single transmission medium. Finally, this model will be used to simulate an AUV swarm of increasing swarm size under increasingly severe communication constraints. Through simulation, the level of coordination between swarm members can be evaluated, allowing the suitability of swarming algorithms to constraints typical of AUV operation to be determined.

The data collected from a deployment in the Shetland Isles is compared against simulation to highlight differences unaccounted for in simulation. The mission tasked the ecoSUB AUV with moving between 4 waypoints at varying depths before rising to the surface.

The vehicle paths showed how the dead reckoning of the real vehicle, the differences in waypoint following behaviour and the likely inconsistent speed of the real vehicle causes the paths of the real and simulated vehicles to diverge once the first waypoint is reached. Initially the real ecoSUB vehicle moved across

the surface, navigating via GPS and allowing a speed estimate of 0.6m/s to be obtained. In addition, the GPS position obtained on surfacing allows a dead reckoning error of 14% of distance travelled to be calculated. The dive profile and yaw behaviour of the vehicles correlate well between the real and simulated vehicles. Finally the pitch behaviour of the real vehicle differs considerably to the simulated due to having to overcome buoyancy unaccounted for in the simulation in order to reduce complexity.

The results could be brought closer to the real vehicle through the introduction of factors unaccounted for by simulation: drift error, inconsistent speed and buoyancy. However, the aim of the model is to evaluate the effects of underwater communication constraints on swarming algorithms and swarm coordination, hence introducing drift error, for example, may shroud the impact of communication delay and hinder the interpretation of results by introducing more variables.

Chapter 5

Emergent Flocking

5.1 Introduction

This study examines the effect of acoustic communication constraints on the ability of an AUV swarm to coordinate swarm members using an emergent flocking approach and perform environmental monitoring through a hill climb algorithm. The study aims to evaluate the scalability of the approaches, meaning performance of the swarm must remain consistent, or increase, with a) more severe communication constraints and b) increases in swarm size.

In order for multiple vehicles to collaborate on a task they need to remain within communication range while underway, hence a control algorithm is required to maintain an aggregated swarm and avoid collisions between vehicles. This is typically achieved through formation control, in which the positions of vehicles relative to one another are strictly controlled. An example application for formation control would be a multi-vehicle team surveying the seabed. Strict formation control is required in this example to ensure the correct swath overlap between sonar beams is achieved. However, formation control often requires a centralised, leader-follower topology, removing the benefits gained from a distributed architecture and introducing a single point of failure. Flocking provides a distributed and easily implemented solution to this problem, aiming to maintain

a cohesive swarm, travelling in a unanimous direction at low relative velocities.

When controlled by the flocking algorithm the vehicles are controlled relative to one another. However, the vehicles need a method of controlling their absolute position in space, in turn deciding where sensor measurements are taken. Through communication of sensor information, each member of the swarm can make decisions based on wider knowledge of the environment. In this study, this knowledge is used to perform a hill climb optimisation algorithm. Formally, a hill climb is a local search algorithm which starts with an initial solution and attempts to reach a better solution by making incremental changes. In this study, the algorithm attempts to reach a higher sensor value by making changes to the heading, speed and hence position of the vehicle. This is synonymous with several situations in ocean science including tracking a chemical plume from hydrothermal vents or locating the source of an oil spill.

The following section presents the method used, including details of the flocking and hill climb algorithms, the communication constraints introduced between swarm members and the set up of the simulation experiments. Following on, the flocking and hill climb results are presented and the scalability of the approaches is discussed against increasing communication constraints and swarm size. An earlier version of this work was accepted to, and presented at, TAROS 2017 and later published in Springer Lecture Notes in Computer Science. The study was revisited due to the development of low power, low cost and long range acoustic modems suitable for a swarm AUV. These new 'nano' acoustic modems increase the communication range and give a clear data transmission rate, bringing simulation results closer to what is achievable in reality [Morozs et al., 2018].

Thomas S. Lowndes, Alexander B. Phillips, Catherine A. Harris, Eric Rogers and Ekaterina Popova, 2017. Evaluating the Capabilities of a Flight-Style Swarm AUV to Perform Emergent and Adaptive Behaviours. *Towards Autonomous Robotic Systems (TAROS) 2017*, proceedings. pp. 237 - 246, Guildford, UK, July 19-21 2017.

5.2 Method

This section details the specific mechanics of the flocking and hill climb algorithms along with the metrics used to evaluate the performance of the swarm. In addition, the communication constraints imposed on the swarm are presented, followed by the simulation configuration.

5.2.1 Flocking Algorithm

Reynolds [1987] presents a distributed behavioural model aiming to create animations of flocking behaviours performed by flocks of birds in nature. This approach is commonly referred to as Reynolds or Boids flocking. Each vehicle, or boid, evaluates three rules based on its own state and the states of its local neighbours, local defined as within communication range. The three rules are depicted in Figure 5.1 and are as follows:

- **Separation** to avoid crowding local flockmates
- **Cohesion** to move towards the flock centre
- **Alignment** to match the velocity and heading of local flockmates

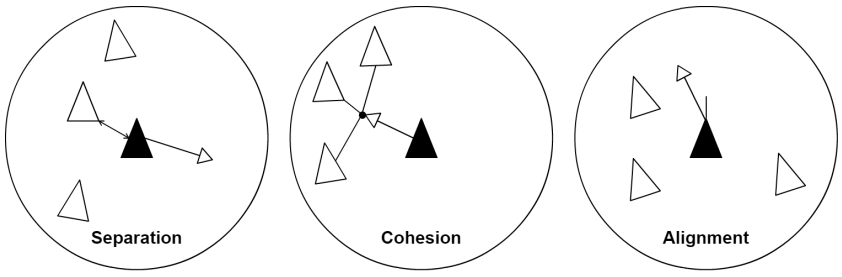

Figure 5.1: The three rules of Reynolds' flocking

These rules give three acceleration vectors which are then subjected to a weighted average to determine the heading and speed demand of the flock member. This algorithm was altered slightly for application in this study due to the issues discussed herein. In the boids algorithm, the separation vector is calculated considering the position of the nearest neighbour, while the cohesion vector

is calculated using the centre of the flock. Using two different points of reference for these rules results in no defined optimum state with regards to the distribution of vehicles. No optimum state for the distribution of vehicles means it is difficult to prove the vehicles are demonstrating the desired behaviour. In computer simulation prior to physical deployments, this leads to difficulties in measuring performance and difficulties in proving a stable behaviour. During physical deployments the operators will likely only receive location updates from vehicles, hence with no defined optimum distribution the operators have no way of guaranteeing the vehicles are behaving correctly and working towards the goal. If operators cannot guarantee the vehicles are performing the desired behaviour, they are unlikely to risk the AUVs themselves, ship and operator time and the responsibility to collect data for scientists by continuing with the deployment. Secondly, the weighted average in Reynold's flocking requires tuning in order to create a balanced system of rules. The different reference points for the separation and cohesion rules mean tuning the weighted average is difficult and even when tuned it cannot guarantee collisions will be avoided. For these reasons, the number of rules was reduced to two, as follows:

- **Separation** maintains a desired distance to the closest swarm member
- **Alignment** to match the velocity and heading of local flockmates

In the new rules, the separation and cohesion rules from Reynolds [1987] have been merged into one and the alignment rule is largely kept the same. The single separation rule now only has a single point of reference, the nearest neighbour, and aims to maintain a user defined distance to this neighbour, resulting in an defined optimum state. Each rule results in a vector, $\mathbf{v}$, each of which consists of a direction, $\mathbf{v}_\theta$ in the form of a compass bearing, and a magnitude, $\mathbf{v}_r$, in the form of a velocity demand. r is used to denote the magnitude to maintain convention with polar vectors. The choice to use a velocity rather than an acceleration, as in Reynolds [1987], is due to the acceleration of an AUV typically being controlled

by a low-level controller whereas the velocity can be controlled by the operator.

Separation Vector

The separation vector, **s**, drives the vehicles to a desired distance separation, avoiding collisions while keeping the vehicles within communication range of at least one other vehicle. To evaluate the separation vector, the swarm member evaluates the distance to its nearest neighbour and generates a vector either towards or away from the neighbour depending on whether the distance is more or less than the desired separation set by the operator. A graphical representation of the method for determining the separation vector is shown in Figure 5.2.

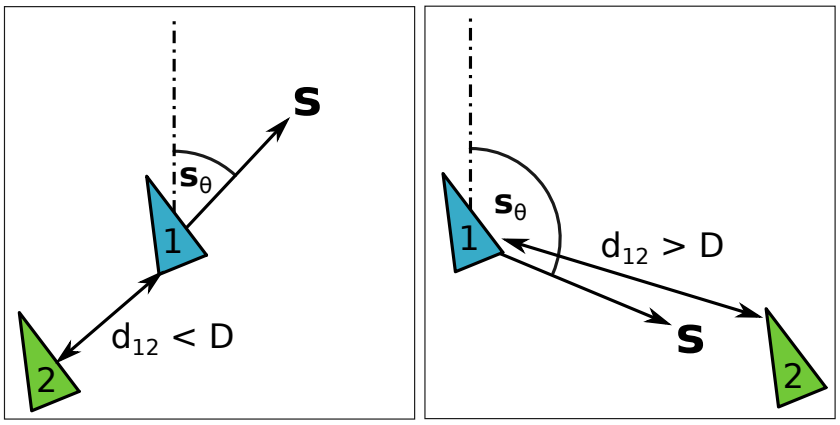

Figure 5.2: Determining the direction of the separation vector for distance less than (left) and greater than (right) the desired separation.

The magnitude of the separation vector, s_r is determined by Equation 5.1 where $\mathbf{c}_i$ is the position of the swarm member, $\mathbf{c}_j$ is the position of the nearest neighbour, D is the desired separation set by the operator and V_{max} is the maximum velocity of the vehicle. Under this equation, when the closest neighbour is D meters away, the magnitude of the vector decreases to zero leaving only the alignment vector acting on the vehicle.

$$\mathbf{s}_{i_r} = \left(\frac{|D - |\mathbf{c}_i - \mathbf{c}_j||}{D} \right) V_{max} \tag{5.1}$$

This method, having removed the consideration of the centre of the flock in the cohesion rule, does have the potential to push vehicles further apart if they

are not each others nearest neighbour. However, once the vehicles reach the desired separation to their nearest neighbour the effect of the separation vector falls to zero leaving the alignment vector in control. The alignment vector then keeps the vehicles at the same distances with their neighbours and hence does contribute to maintaining a cohesive swarm. The ability of the flock to maintain a cohesive swarm under the new rules is evaluated in the discussion section and the metrics used are discussed later in this section.

Alignment Vector

The alignment vector, $\mathbf{a}$, is determined in very much the same way as in Reynolds flocking. The direction of the vector, $\mathbf{a}_\theta$ is determined for a swarm member by evaluating the average of headings across the swarm, as shown in Figure 5.3. $\mathbf{a}_\theta$ is evaluated by Equation 5.2, where ψ_j is the heading of vehicle j, $\boldsymbol{\psi}_j$ is the unit vector heading of vehicle j, N is the number of vehicles in the swarm and i is the evaluating swarm members unique identifier. The magnitude is set at a constant 0.5m/s, causing the vehicles to slow to a more efficient, unanimous speed once the magnitude of the separation vector reduces to 0.

$$\mathbf{a}_{i_\theta} = tan^{-1}\left(\frac{\boldsymbol{\Psi}_2}{\boldsymbol{\Psi}_1}\right) \tag{5.2}$$

Where

$$\boldsymbol{\Psi} = \sum_{j=0}^{j=N, j\neq i} \boldsymbol{\psi}_j, \quad \boldsymbol{\psi}_j = \begin{bmatrix} cos(\psi_j) \\ sin(\psi_j) \end{bmatrix} \tag{5.3}$$

Once the separation and alignment vectors have been evaluated, the average of these vectors is taken resulting in a yaw demand and a velocity demand to be passed to the vehicles control systems.

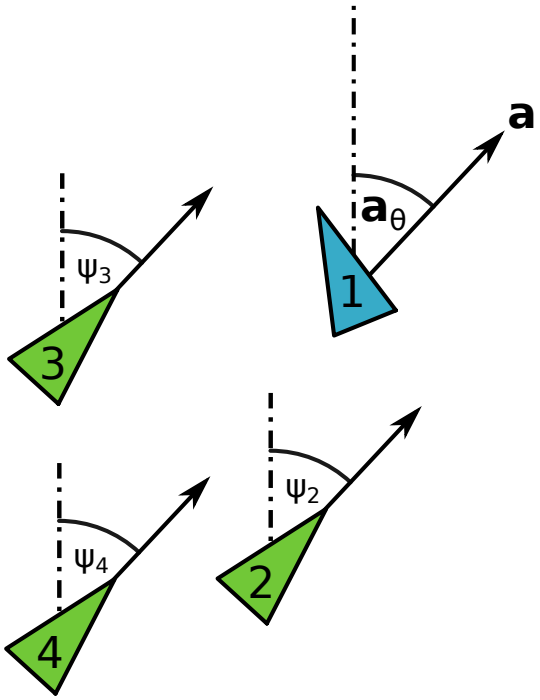

Figure 5.3: Determining alignment vector direction, $\mathbf{a}_\theta$ for vehicle i=1.

Performance Metrics

It is difficult to quantify the performance of a flocking behaviour as there is no defined optimum configuration for the vehicles as in strict formation control. Read et al. [2014] use several metrics for assessing the organisation of a flock. These metrics include polarisation, a measure of alignment of the vehicles and angular momentum, a measure of rotation about the flocks centre. These metrics are taken from Couzin et al. [2002] and the equations for polarisation and angular momentum, p_{group} and m_{group} respectively, are shown by Equation 5.4 and Equation 5.5 respectively. In these equations, N is the number of swarm members, $\mathbf{c}_{group}$ is the centre of the swarm, $\mathbf{c}_i$ is the position vector for swarm member i and $\mathbf{v}_i$ is the unit direction vectors of swarm member i. Polarisation and angular momentum are calculated at every time step. Polarisation ranges from 0 to 1, 1 indicating all headings are aligned and all vehicles are travelling at the same velocity. Angular momentum also ranges from 0 to 1, with a low angular momentum being desirable. However, a high polarisation and high angular momentum can also be indicative of a well organised flock while a low polarisation tends to

indicate a lack of coordination.

$$p_{group}(t) = \frac{1}{N}\left|\sum_{i=1}^{N}\mathbf{v}_i(t)\right| \tag{5.4}$$

$$m_{group}(t) = \frac{1}{N}\left|\sum_{i=1}^{N}\mathbf{r}_{ic}(t)\mathbf{v}_i(t)\right| \tag{5.5}$$

where

$$\mathbf{r}_{ic} = \mathbf{c}_i - \mathbf{c}_{group}$$

and

$$\mathbf{c}_{group}(t) = \frac{1}{N}\sum_{i=1}^{N}\mathbf{c}_i(t)$$

The separation between a swarm member and its nearest neighbour is also evaluated for the swarm as the flocking algorithm should drive the swarm members to a desired separation. For each swarm member, the distance to its nearest neighbour is evaluated, then the average is evaluated across the swarm at each time step.

Polarisation, angular momentum and separation to nearest neighbour will be the 3 main metrics which will be used to evaluate the performance of the flock across all communication constraints and swarm sizes.

As mentioned previously, the algorithm does have the potential to push swarm members further apart if they are not nearest neighbours. In order to ensure the vehicles maintain a cohesive swarm, this study evaluates the rate of change of the mean distance to the centre of the flock across the swarm. If this metric is positive, it shows the swarm is spreading out and vehicles are moving further from the flocks centre. If this value is equal to zero, it shows the vehicles are moving in a stable distribution about the flocks centre. This metric will indicate whether the alignment vector is enough to maintain a cohesive swarm. In order to calculate the metric, the centre of the flock is determined by taking an average

of the swarm member's positions. The distance from each swarm member to the flocks centre is then calculated and a mean is taken for that timestep. This is repeated until a mean distance to flock centre is calculated for each timestep. The rate of change of the mean distance to the flocks centre is evaluated over six, 200s windows, starting at 1800s in order to allow the swarm to reach a stable distribution before evaluating.

Summary

An adaptation of the Reynolds flocking algorithm [Reynolds, 1987] is used in this study. An average of two vectors, controlling separation between swarm members and alignment with other swarm members, gives a yaw and velocity demand which controls the vehicles position in order to achieve the flocking behaviour. The flocking performance is measured by: polarisation, a measure of alignment between vehicles; angular momentum, a measure of rotation about the swarms centre; and separation between swarm members to ensure collisions are being avoided and communication maintained. To ensure the flock is maintaining a stable distribution, the rate of change of the mean distance to the flocks centre will be evaluated.

5.2.2 Hill Climb Algorithm

A hill climb algorithm is implemented alongside the flocking algorithm to allow the vehicles to follow the Gaussian sensor function and identify the location of the highest sensor value or maxima of the function. The Gaussian sensor function imitates an undisturbed chemical plume in the ocean, with a maximum concentration of the chemical at the source location and dissipating to a point where the chemical is undetectable at the edges. Examples of such chemical plumes would include the Exxon-Valdez and Ixtoc-1 oil spills or a hydrothermal vent in the ocean floor as these all take the form of a chemical plume originating at a single source and dissipating into the ocean at a boundary. Figure 5.4 shows

a photograph of the Ixtoc-1 oil spill showing a a source location of the leaking oil and clear edge to the surface slick. Figure 5.5 shows a swarm of 5 AUVs performing the hill climb algorithm to successfully locate the highest sensor value of the sensory function.

Figure 5.4: Ixtoc-1 oil spill

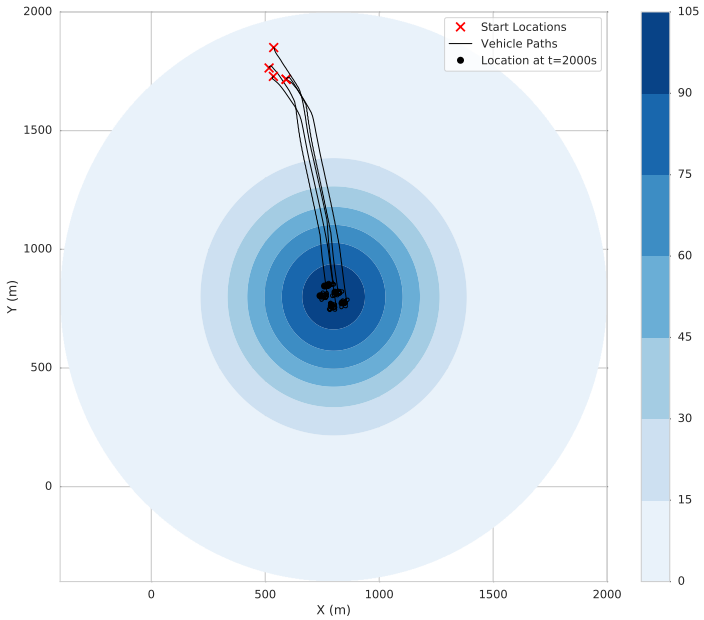

Figure 5.5: An example of the hill climb algorithm for a swarm of 5 vehicles, the sensory function and individual vehicle paths are shown.

It is assumed the vehicles are fitted with an appropriate sensor for the sensory function and that the sensor can only make point measurements at the vehicles current position. In order to implement the hill climb, the vehicles sensor mea-

surement must be communicated to other swarm members.

The hill climb approach is implemented by introducing an additional vector, **h**, in addition to the alignment and separation vectors required by the flocking algorithm. Through the sharing of sensor measurements, a swarm member can evaluate the direction of increasing sensor values across the swarm and hence use this to move in a direction which will increase its own sensor measurement. This is shown in Figure 5.6 where vehicle 1 is at the lowest sensor value and vehicle 4 is at the highest sensor value, hence in this example $\mathbf{h}_\theta$ is determined as the yaw angle of vehicle 4 relative to vehicle 1. The magnitude of the hill climb vector is set at a constant 0.5m/s, equal to the magnitude of the alignment vector. It is important to find the balance between the flocking and hill climb algorithms, to ensure neither is detrimental to the other and neither is dominant. To ensure a balance, the flocking performance metrics are still evaluated when the swarm is performing the hill climb algorithm.

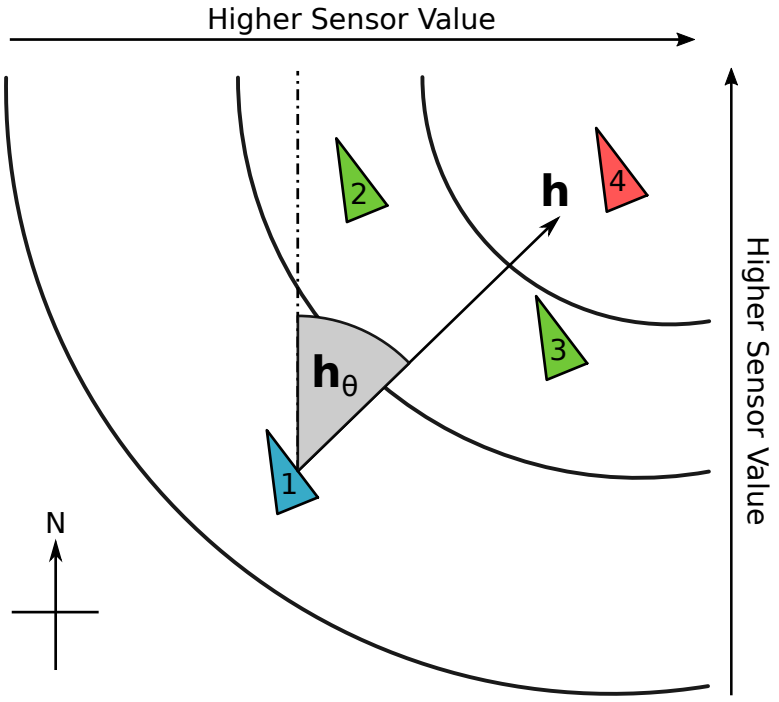

Figure 5.6: Determination of hill climb vector by a swarm member using knowledge of local sensor measurements.

To evaluate whether the swarm offers an improvement in performance over a single vehicle, experiments are conducted to measure the performance of a

single vehicle given the same task of locating the maxima of the sensory function. A single vehicle does not have the wider knowledge of the environment made possible by the swarm and hence must use a different method to determine the hill climb vector. A method was developed using 2 rules presented in order of priority:

1. If the current sensor value is less than 99% of the sensed maximum move towards location of sensed maximum
2. If the current sensor measurement is less than the previous sensor measurement, randomly generate a new yaw demand.

This approach allows the AUV to explore the sensed distribution by randomly generating a new heading demand if the sensed values are not increasing and hence the AUV is not making progress towards the goal. This is provided by rule 2. However, to avoid the AUV losing the feature entirely and to ensure the AUV continues to make progress, if the sensed value falls below 99% of the maximum sensed value, the yaw demand is set in the direction of the location at which the maximum sensor value was measured. This is shown by rule 1 and bounds the exploration of the sensor function by ensuring the AUV is pulled back towards the location where it achieved the highest sensor value.

To measure the performance of the hill climb algorithm, the time taken to identify the sensory maxima is used. Time is always a factor in AUV operations due to the limited energy capacity of onboard batteries. If the swarm cannot locate the maxima in less time than the single vehicle, the swarm provides no improvement to task performance. If the swarm is able to locate the maxima in less time, this is an increase in task performance and means the swarm has more energy remaining than the single vehicle. However, this raises the question of how much of an improvement justifies the use of a swarm over a single vehicle. For example, if a single AUV uses 75% of its energy to locate the maxima and a swarm of five vehicles use 50% each, energy spent by each vehicle in the swarm is less than by the single vehicle and the swarm can continue monitoring for twice as long as the single vehicle, 50% remaining versus 25% remaining. However, the

total energy used by all swarm members is much greater than by the single vehicle implying less efficient operation as the swarm have essentially used 250% of a single vehicles energy capacity, compared to 75% used by the single vehicle. As any improvement in performance achieved by the swarm could result in successful location of the maxima in a situation where the single vehicle is unsuccessful, this study will accept any increase in performance to satisfy this criteria.

The time taken to locate the maxima is determined in post processing. Firstly, the time at which the maximum value sensed by the swarm exceeds 90% of the maxima is determined, denoted as t_{90}. The mean value sensed by the swarm is then taken across the remaining run time. If the mean value is greater than 80% of the maxima, this is deemed a success and t_{90} is used as the time to reach the maxima. This is shown in Figure 5.7.

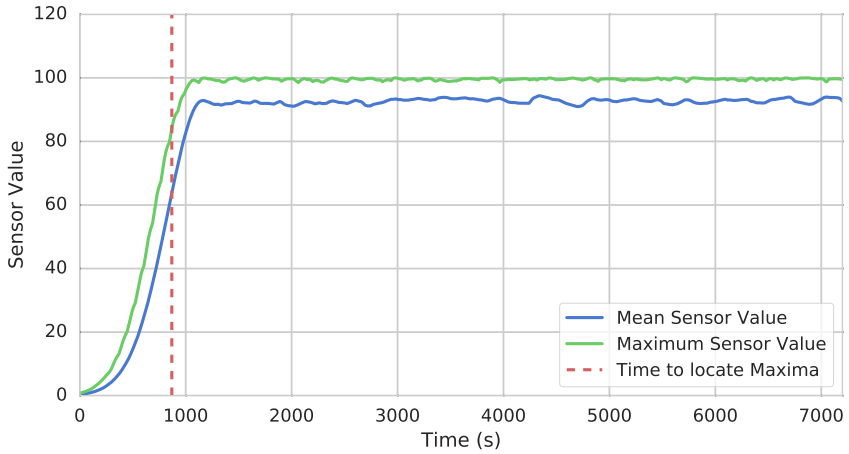

Figure 5.7: Exemplar plot of a single experiment demonstrating how the mean and maximum sensor values across the swarm vary throughout the hill climb behaviour. t_{90} is indicated by the red dotted line.

In summary, the aim of a hill climb algorithm is to drive the vehicles to a maximum sensor value. Due to the sharing of information, a swarm member is able to evaluate the direction of the sensory function across the swarm and hence include this as a vector in the averaging of the flocking vectors. A single vehicle can only take point measurements at its location, hence 2 rules are used to determine the yaw demand of the vehicle, one emphasising exploration while the other bounds the exploration and pulls the vehicle back to the highest sensed

value. Finally, hill climb performance for both the swarm and the single vehicle is measured in the time taken to determine the location of the maxima.

5.2.3 Communication Assumptions

Underwater communication is one of the main challenges associated with AUV operation and one of the biggest barriers to implementing an AUV swarm. This work introduces the range and data rate constraints of an acoustic modem to the flocking and hill climb algorithms. The range of the acoustic modem used for this study is 1km and the communication is assumed lossless.

As all communication travels through the same medium, a Time Division Multiple Access (TDMA) medium access control protocol is used. This means each AUV is sequentially allocated a time slot, or frame, to broadcast their state to other swarm members. One-way acoustic broadcasting is assumed and increases the scalability of the algorithm. However the allocation of sequential time slots could still pose a limiting factor. Table 5.1 details the four frame lengths used and maximum message size which can be transmitted within these frames, details of how the frame lengths are determined can be found in Section 4.2.2 on Page 64.

Table 5.1: Communication Contraints - Frame lengths used for the TDMA protocol and maximum message size (S_{max})

T_f (s)	T_g (s)	T_t (s)	$S_{max}(bytes)$
4.0	1.0	3.0	12
6.0	1.0	5.0	20
8.0	1.0	7.0	28
10.0	1.0	9.0	36

For the flocking and hill climb algorithms to function, a vehicles position, its heading and its sensor value must be communicated. If the vehicle position is represented as latitude and longitude, which is typical for AUV operation, this will require 16 bytes for full accuracy. In addition, heading will range from 0 to 360° hence 2 bytes are required. Finally, the number of bytes reuqired to transmit

the sensor value can vary depending on sensor specification, in this study 2 bytes are assigned. This gives a total of 20 bytes meaning the required information can be sent within $T_f = 6s$. A frame length of four seconds with $S_{max} = 12$ bytes is possible if the position of a vehicle can be represented as x and y in meters from a mutually known origin. Whereas $T_f = 8s$, $10s$ allow for larger sensor data or extra information to be sent.

5.2.4 Simulation Configuration

Swarms of 5, 10, 15 and 20 vehicles are first simulated, under the control of the flocking algorithm alone and subject to increasing communication constraints.

In order to determine the range of swarm sizes used in the study, the method of simulating marine autonomous vehicle surveys of a HAB presented by Griffiths [2007] is used. This method was previously discussed in Section 2.2 (Page 28). A swarm of 5 vehicles achieves a lower survey error metric than a single AUV at all length and time scales. As a key requirement for the swarm is to provide improved task performance, this sets the lower bound of swarm sizes used. Furthermore, a swarm of 20 vehicles achieves a minimum survey error metric of approximately 0.2, a considerable improvement compared to a single AUV at 0.46 and a glider fleet at 0.55. In addition, Griffiths [2007] states in order to resolve the time and length scales required of rapidly evolving phenomena, a 20km by 20km lawnmower survey with 1km between tracks would be a suitable survey for a single AUV. By dividing this into 20 1km survey tracks significant benefits can be gained from the use of a 20 vehicle swarm. While increasing the swarm size would reduce the error metric further, the cost of a 20 vehicle swarm would be approximately equal to the cost of a single glider and hence this forms the upper limit to the swarm sizes used.

Moving on, the hill climb algorithm is introduced and simulations at each swarm size and communication constraint are repeated. The flocking simulation is run for 1 hour simulated time and the hill climb simulations are run for 2 hours

simulated time, both at a resolution of 1 second.

A 5 vehicle swarm represents the smallest group of vehicles classified as a 'swarm' while a 20 vehicle swarm represents a swarm of financial cost approximately equal to a typical AUV such as a Slocum glider. It is also important to consider the logistics of operation. As swarm AUVs are much smaller than a typical AUV, they are considerably easier to launch and this can be done from a ship, small boat or even other autonomous vehicles. However, coordinating the launch, piloting and recovery of 20 vehicles is a considerable task hence the maximum swarm size.

The swarm is subject to a range of communication constraints. These include ideal communication, meaning all AUVs have complete knowledge of the heading, position and sensor value of all swarm members at every time step, and acoustic communication following a TDMA protocol for each of the frame lengths shown in Table 5.1. Ideal communication is used to set a bench mark whereas the range and data rate of the acoustic communication has been chosen to be typical of AUV operation.

The initial starting locations of the swarm are randomly seeded within a 150m by 150m area for each repeat and 100 repeat experiments for each swarm size and communication constraint are performed. Paired seeds are used to permit fair comparison, hence the starting positions for each repeat experiment are the same across all swarm sizes and communication constraints. A swarm of 10 vehicles requires 5 more starting locations than a swarm of 5 vehicles, hence the first 5 vehicles of the 10 vehicle swarm start at the same locations as the 5 vehicle swarm, then the remaining 5 vehicles are distributed according to the randomly seeded positions. For the hill climb algorithm, the centre of the starting area is 1km from the maxima of the sensory function, imitating a swarm being deployed after initial detection of a chemical plume.

For the hill climb algorithm, the sensory function is a radial Gaussian function which imitates a two dimensional chemical plume. The function dissipates to zero

at a distance of 1.2km from the maxima. The hill climb algorithm is performed using the pure Gaussian function, then again with sensor noise applied to the function. The sensory function, with and without noise, is shown in Figure 5.8 against distance from the maxima and the envionment in space is shown in Figure 5.9.

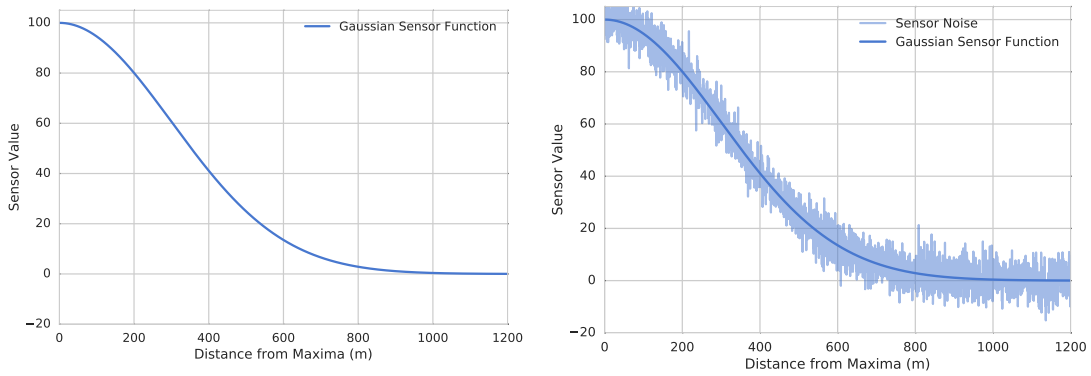

(a) Sensor Function

(b) Sensor Function with added arbitrary noise

Figure 5.8: Sensor value against distance from maxima for Gaussian sensor function used in hill climb simulation to imitate chemical plume

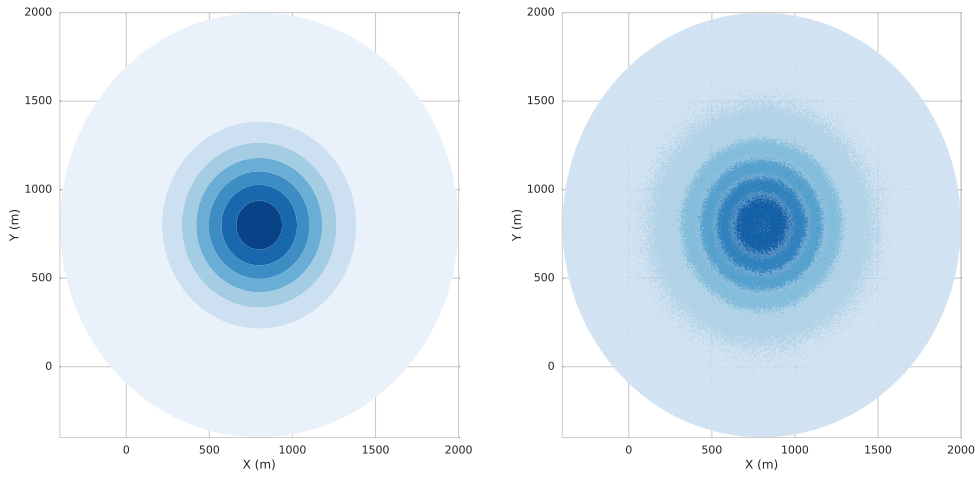

(a) Sensor Function

(b) Sensor Function with added Noise

Figure 5.9: Gaussian sensor function over XY space used in hill climb simulation to imitate chemical plume. Colour bars not shown as sensor values are arbitrary.

The vehicles are released within a 150m square starting area, the centre of which is 1000m from the maxima. These initial conditions reflect the vehicles being launched from a support vessel after an initial detection of the edge of

a plume. All starting locations of vehicles across the 100 repeats are shown in Figure 5.10.

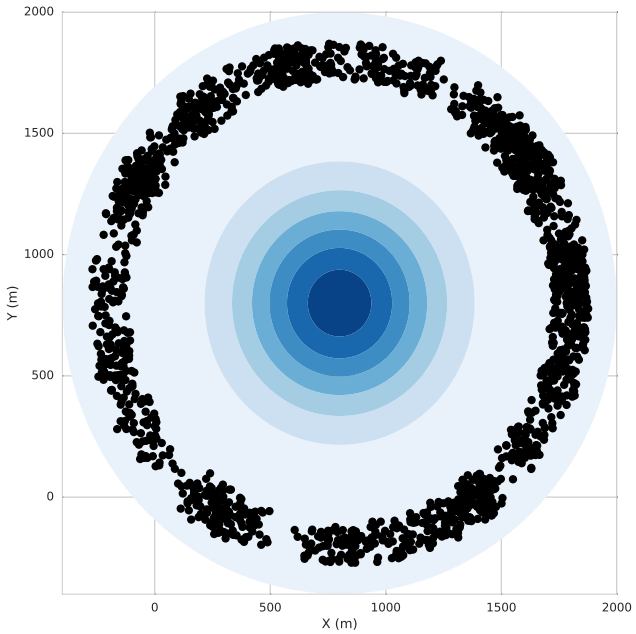

Figure 5.10: All possible starting locations for an AUV (black markers) across the 100 repeat experiments plotted on top of the sensory function.

In summary, the results from these experiments will show whether the swarm is able to flock under increasingly severe communication constraints, both from increasing TDMA frame length and swarm size. The hill climb experiments will demonstrate whether there is a benefit to using a swarm over a single vehicle and whether the benefit increases with swarm size. A successful result would demonstrate scalability, meaning there is negligible decrease in performance with increasing communication constraints and swarm size in both flocking and the hill climb. An increase in task performance with swarm size indicates further success.

5.3 Results & Discussion

The following section presents the results from the flocking and hill climb experiments. Throughout the results we are looking for scalability with both increasing swarm size and communication constraints. For flocking performance, scalability

with swarm size is indicated by consistent performance in the flocking performance metrics. However, when examining the hill climb results there is little reason to deploy a swarm of 20 vehicles when a swarm of 10 vehicles will perform just as well, especially when considering the increased cost and risk in deploying a greater number of vehicles. Hence, an increase in performance with swarm size is desirable. Conversely, when evaluating scalability with increasing communication constraints, consistent performance is acceptable in both the flocking and hill climb results.

5.3.1 Flocking Performance

Examining the flocking performance of the swarm, first the results from a single experiment will be examined and discussed in detail. This experiment simulated a swarm of 5 vehicles for $T_f = 4s$. Figure 5.11 presents the 5 vehicle paths as they come to an equilibrium state from their starting positions, shown by the red hashed line, and proceed to flock in a north-easterly direction. This is a good example of the coordination expected between vehicles under the flocking algorithm.

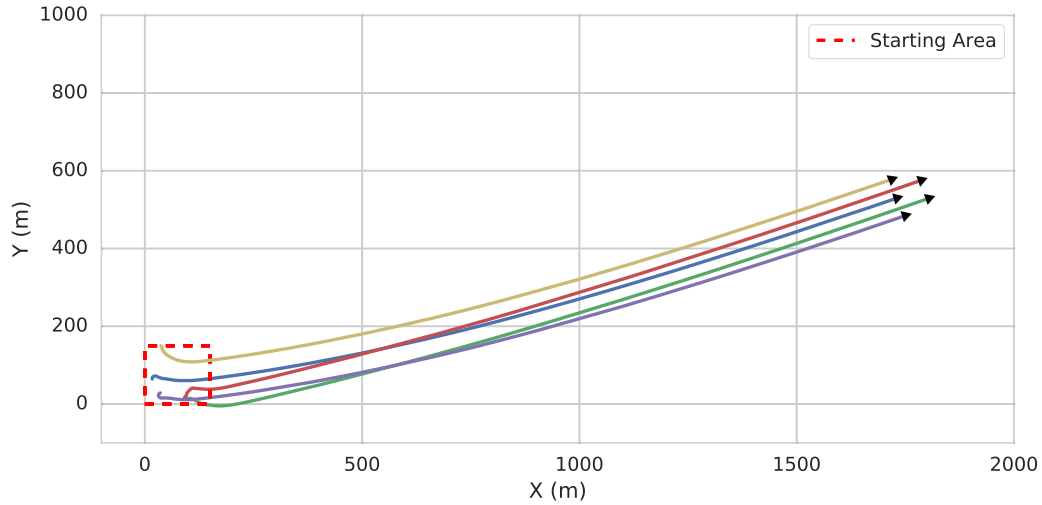

Figure 5.11: Paths and final locations of 5 AUVs controlled by the flocking algorithm, $T_f = 4s$

Figure 5.12 presents the three performance metrics by which the organisation

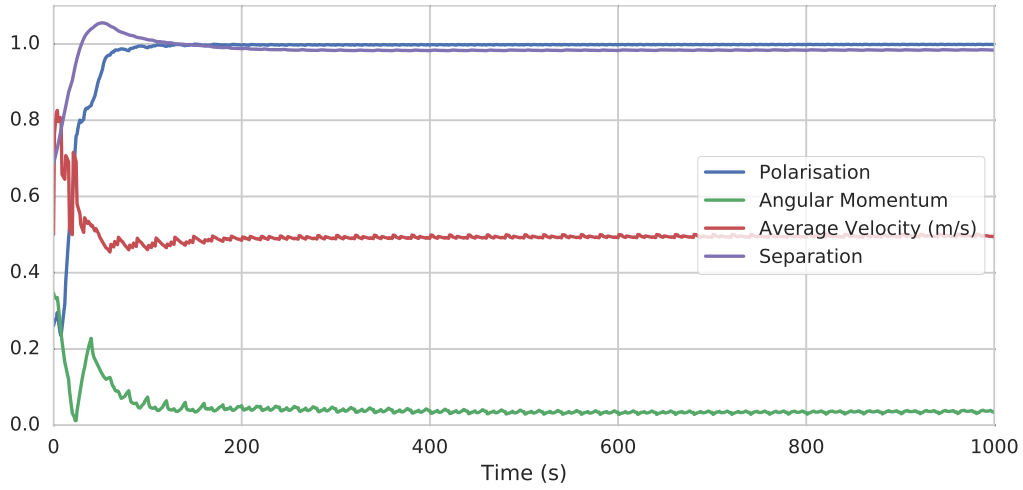

Figure 5.12: The three flocking performance metrics and angular momentum plotted against time for a simulation of 5 vehicles controlled by the flocking algorithm, $T_f = 4s$.

of the swarm is measured along with the average velocity across the swarm. It can be seen that polarisation tends to 1 and angular momentum tends to a minimum value close to zero, this is the desired result. In addition, the separation in this plot is normalised by the desired separation, hence as this plot tends to one the vehicles are moving to the desired separation. Finally as the separation tends to one, resulting in the corresponding vector magnitude reducing to 0 as the vehicles approach the desired separation, the average velocity tends to 0.5m/s, the value set as the magnitude of the alignment vector. This plot indicates the rules are having their intended effect, resulting in the desired flocking behaviour.

In order to ensure the flock is able to achieve a stable distribution under the adapted rules, the rate of change of mean distance to the centre of the flock is evaluated. This is shown in Figure 5.13 for the ideal communication case across all studied swarm sizes. From Figure 5.13 it can be seen that the metric is less than 0.002m/s across all swarm sizes in all experiments, with the median rate of change for each swarm size falling to approximately 0.0005m/s and below. Figure 5.14 shows example plots of the mean distance to the flock centre against time for 3 single experiments. These plots indicate the type behaviour which is represented by this metric, with a rate of change of 0.001m/s indicating a rela-

tively consistent mean distance to the flock centre with some small fluctuations. Considering the algorithm only controls the a vehicles position relative to it's nearest neighbour, this demonstrates the adapted rules are able to form stable distributions of vehicles.

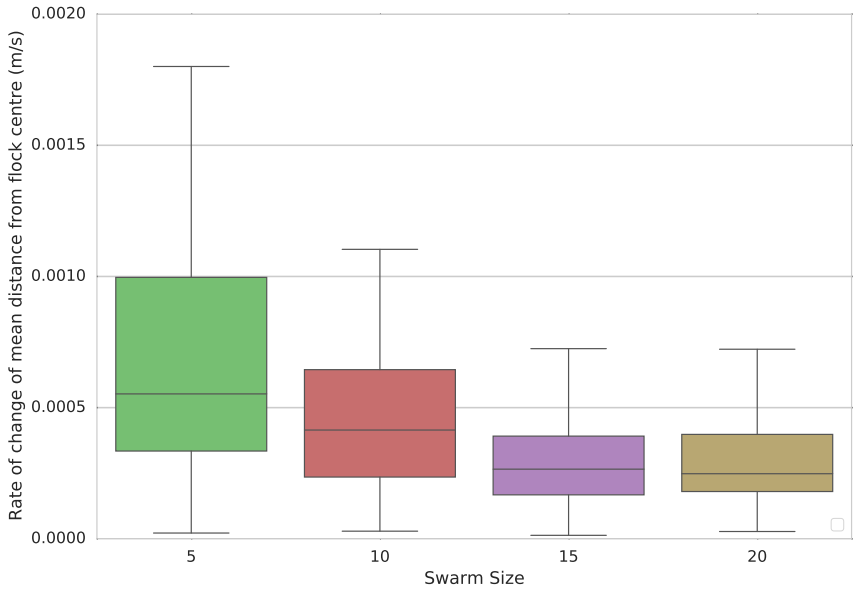

Figure 5.13: The rate of change in mean distance to the centre of the flock for the ideal communication case and swarm sizes of 5 to 20 vehicles. 100 repeat experiments were performed at each swarm size.

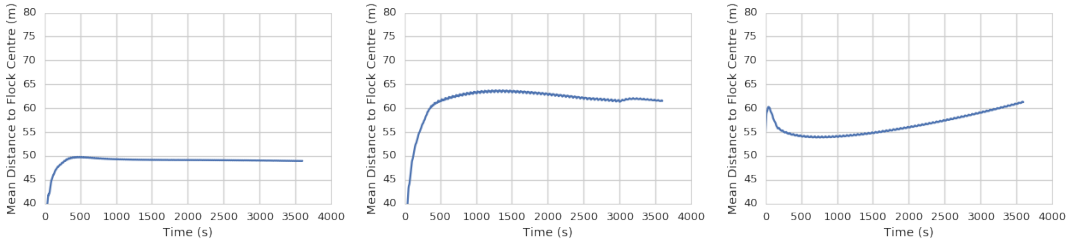

(a) 0.0001m/s (b) 0.001m/s (c) 0.003m/s

Figure 5.14: 3 Example plots of Mean distance to flock centre against time for 3 single experiments and their respective rate of change from 1800s onwards.

Moving on to examine the effect of increasing swarm size and TDMA frame length on the performance of the swarm, Figure 5.15 presents the three main flocking performance metrics on separate plots, comparing a swarm of 5 vehicles to a swarm of 20 vehicles for $T_f = 4s$. It can be seen that the 20 vehicle swarm takes longer to settle to the desired values, high polarisation, low angular

momentum and separation of 50m, than the 5 vehicle swarm. This can only be expected as, using the TDMA protocol, the time to complete a single communication cycle is increased from 20 seconds in the 5 vehicle swarm to 80 seconds in the 20 vehicle swarm. An increase in the time to form a well organised flock with an increase in swarm size doesn't inherently support the conclusion that this is a scalable solution. However, a 20 vehicle swarm forming a well organised flocking swarm in approximately 10 minutes would still be acceptable if the swarm was at sea, performing monitoring tasks, for several hours.

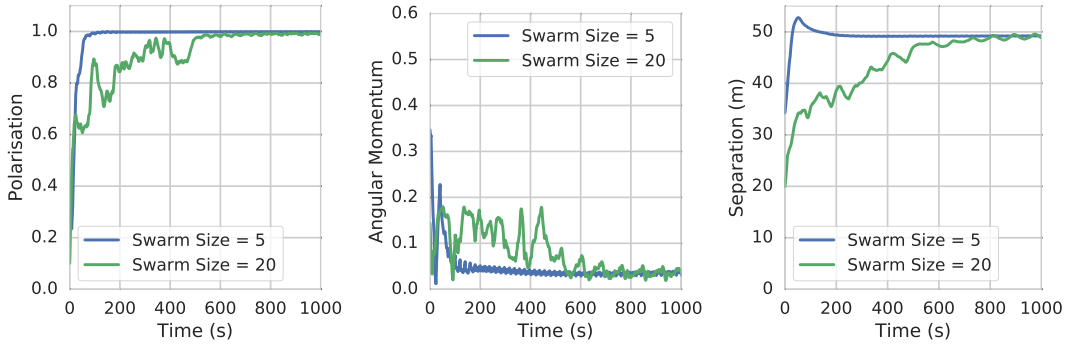

(a) Polarisation (b) Angular Momentum (c) Separation

Figure 5.15: Flocking metrics for $T_f = 4s$ for swarm sizes of 5 and 20 vehicles.

Considering a swarm of 5 vehicles and varying TDMA frame length, Figure 5.16 presents a similar decrease in performance for increase in TDMA frame length. However, rather than simply taking longer to settle to the optimum values, increasing T_f appears to introduce a steady state error. This is especially clear in the angular momentum plot (Figure 5.16b) as the result for $T_f = 10s$ settles to a constant value closer to 0.1 whereas the result for $T_f = 4s$ settles closer to 0. Again, a decrease in performance can only be expected with increasing communication constraints, but given the performance shown in polarisation and separation plots, the performance of the swarm subject to $T_f = 10s$ can still be described as well organised.

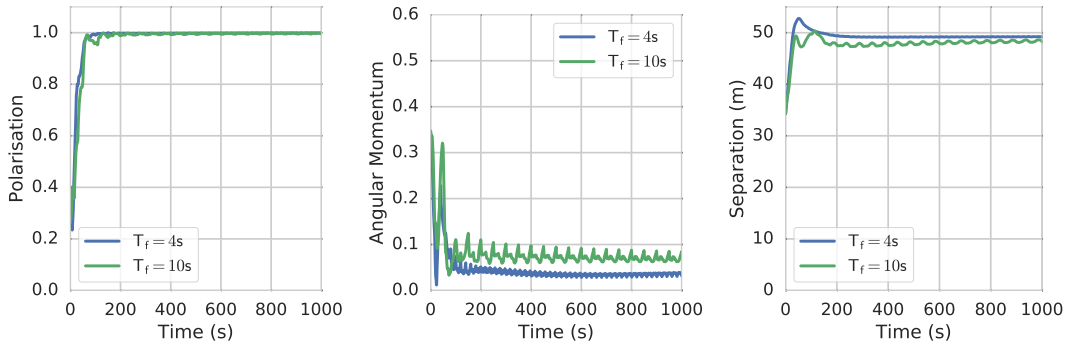

(a) Polarisation (b) Angular Momentum (c) Separation

Figure 5.16: Flocking metrics for swarm size of 5 vehicles increasing T_f from 4s to 10s.

Figure 5.17 presents the polarisation across 100 repeat experiments at each swarm size and communication constraint. It can be seen that the polarisation of the swarm decreases with increases in swarm size and TDMA frame length. However, even for a swarm of 20 vehicles and $T_f = 10s$, the polarisation remains above 0.85. Combining this with the results presented in Figure 5.18, the swarm achieves a low angular momentum of less than 0.1 across the majority of the experiments, indicating an organised flock at all swarm sizes and communication constraints. This implies the flocking algorithm is scalable with swarm size from 5 to 20 vehicles, but there is doubt in the scalability beyond this as the performance degrades further. Interestingly, the angular momentum of the swarm decreases with swarm size for ideal communication and frame lengths of 4 and 6 seconds while still demonstrating some decrease for longer frame lengths of 8 and 10 seconds. In addition, a decrease in the variability in angular momentum can be seen across all swarm sizes, indicated by the shrinking whiskers on the box plot, showing a more consistent result with increases in swarm size.

Examining a less extreme case, it was determined previously, with current acoustic communication technology, a $T_f = 6s$ is needed to send the information required by the flocking and hill climb algorithms. Across all swarm sizes the polarisation remains above 0.95 for $T_f = 6s$ indicating good flocking performance and a well organised flock. In addition, the median angular momentum for $T_f =$

6s (Figure 5.18) remains close to 0.05 across all swarm sizes supporting this conclusion.

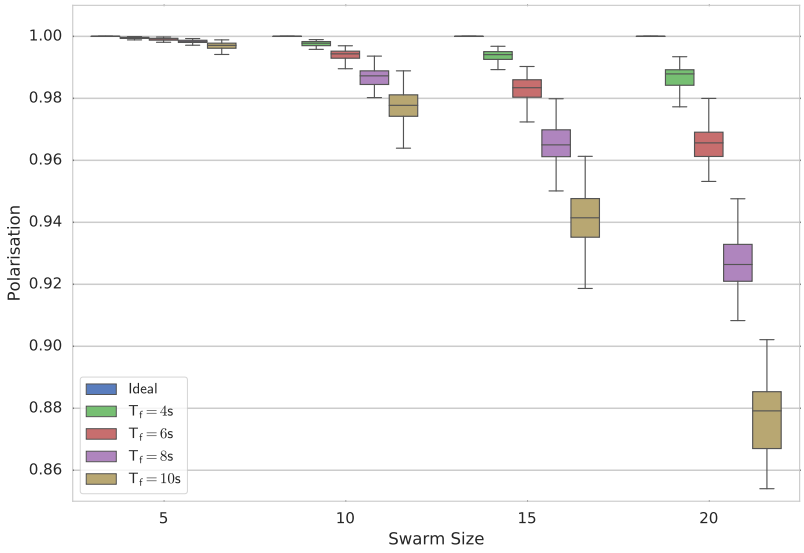

Figure 5.17: Polarisation over the 100 repeat experiments for swarm sizes from 5 to 20 vehicles and T_f from 4 to 10 seconds. The polarisation of each repeat is taken as the median across the entire run time of 1 simulated hour. The whiskers show the extents of the data, the upper and lower quartile are represented by the box and the median is shown as the line within the box.

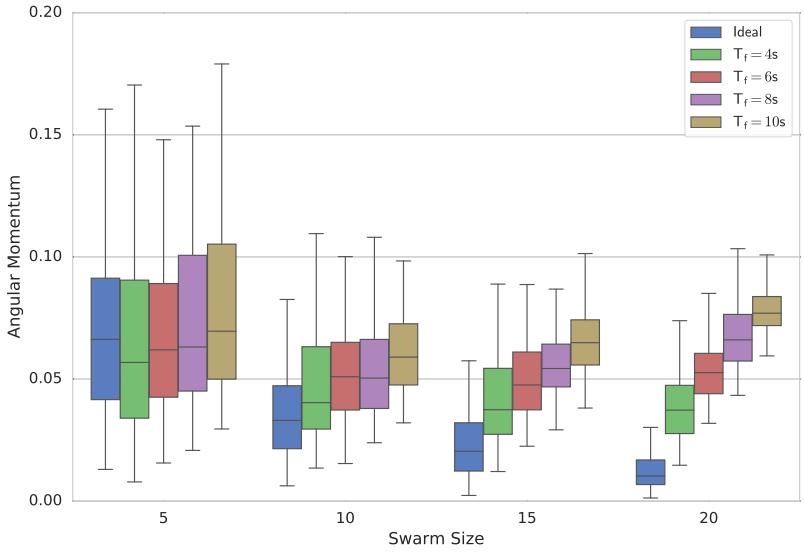

Figure 5.18: Angular momentum over the 100 repeat experiments for swarm sizes from 5 to 20 vehicles and T_f from 4 to 10 seconds. The angular momentum of each repeat is taken as the median across the entire run time of 1 simulated hour.

For the final evaluation of flocking performance, Figure 5.19 presents the average separation across the 100 repeat experiments for each swarm size and com-

munication constraint. This is clearly where increasing swarm size and TDMA frame length have the most detrimental impact on flocking performance. For swarms of 15 and 20 vehicles, errors of 10% or more can be seen for the majority of communication constraints, with the median error increasing beyond this in the most extreme communication cases. However, for the required frame length of six seconds discussed previously the median remains within 10% of the desired separation of 50m, error which can be attributed to the flight-style nature of the vehicles.

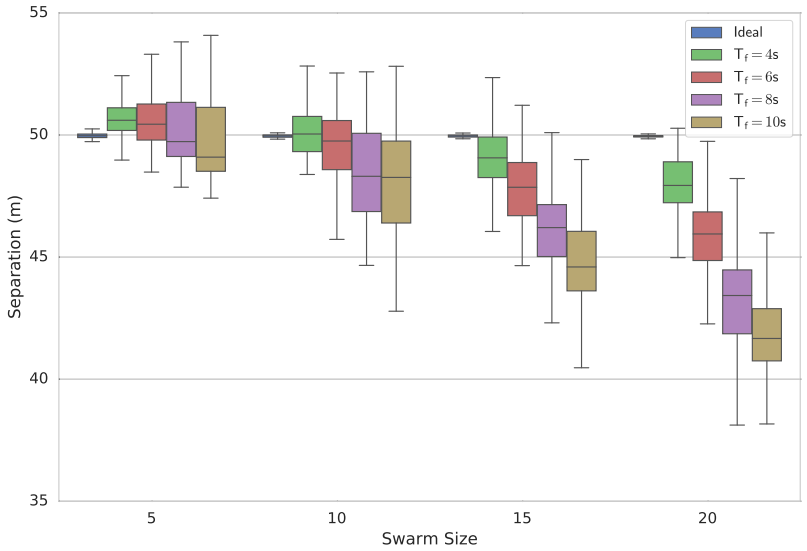

Figure 5.19: Separation over the 100 repeat experiments for swarm sizes from 5 to 20 vehicles and T_f from 4 to 10 seconds. The separation of each repeat is taken as the median across the entire run time of 1 simulated hour.

In summary, the flocking results presented indicate a well organised flock at all swarm sizes and communication constraints and hence demonstrate the scalability of this approach to coordination of the swarm. Examining the absolute results, flocking performance does decrease with increasing swarm size and communication constraints, indicating a limit to scalability will be reached eventually, only to be expected when the vehicles are communicating less frequently. The swarm separation is where growing communication constraints seem to have the largest impact, with the error from desired separation of 50m growing beyond 10% in the more severe communication constraints for larger swarm sizes.

5.3.2 Hill Climb Performance

The following sections present the results of the hill climb experiments. As with the flocking experiments, 100 repeats were performed for each swarm size (5,10,15,20) for each communication constraint (Ideal, 4s, 6s, 8s, 10s). Firstly the results are presented for a hill climb of the Gaussian sensory function, then the performance relative to a single vehicle is examined to evaluate the increase in task performance and finally results are presented for a hill climb of the Gaussian function with added sensor noise in an effort to reduce the gap between simulation and reality.

Gaussian Sensory Function

The first set of results concern the hill climb algorithm applied to the Gaussian sensor function shown in Figure 5.8b. Figure 5.20 presents the success rate of the swarm when tasked with locating the maxima of the sensor function using the hill climb algorithm. The success rate of the swarm is shown to increase with swarm size, surpassing 97% in all cases for swarms of 10, 15 and 20 vehicles. This demonstrates a highly scalable approach to larger swarm sizes.

A swarm of 5 vehicles achieves the lowest success rates, demonstrating a clear decrease in success with increasing communication constraints. The success rate falls to its lowest at $T_f = 10s$ with a value of 88%. Firstly, this is due to the poor directional perception in small swarm sizes. As the direction of the hill climb vector, $\mathbf{h}_\theta$, is determined by the heading from the minimum to the maximum sensor measurement known by a vehicle, the outcome is limited by the possible combinations of measurements. In a swarm of 5 vehicles, $\mathbf{h}_\theta$ is limited to one of 20 possible directions. The increase in possible directions is shown by Figure 5.21. This does not inherently lead to better performance as the benefit is dependent on the distribution of vehicles. For example, if all the vehicles are exactly distributed along a line, there will only be 2 possible directions for $\mathbf{h}_\theta$. Secondly, the increase in T_f means the swarm are making decisions based

on increasingly uncertain information, affecting both the hill climb and flocking algorithms. It is a combination of these two factors that lead to the declining success rate in this case.

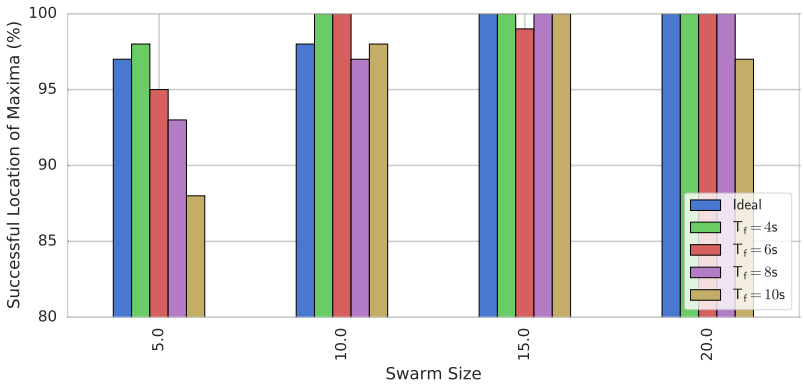

Figure 5.20: Success rate of swarm locating the maxima for each swarm size and communication constraint over 100 repeats.

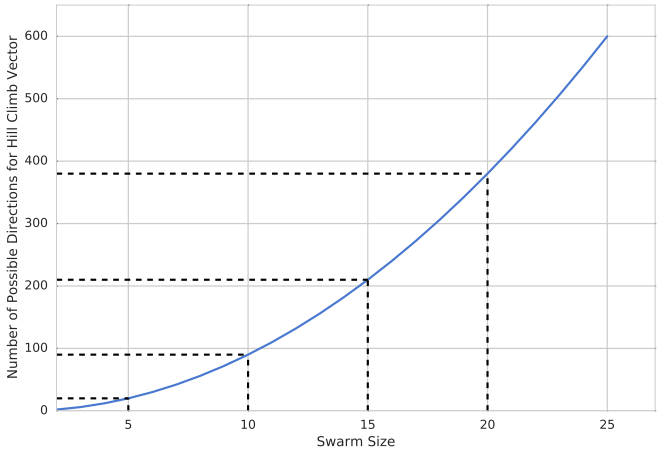

Figure 5.21: Number of possible directions for the hill climb vector against swarm size. Values for 5, 10, 15 and 20 vehicle swarms are indicated.

Figure 5.22 presents the time to locate the maxima for all swarm sizes and communication constraints. It can be seen across all swarm sizes that an increase in T_f leads to an increase in the time taken to locate the maxima and hence a decrease in performance. This is expected as the AUVs are acting on increasingly out-of-date information which affects both the flocking and hill climb algorithms. Examining increasing swarm size, the performance of the swarm becomes more consistent with swarm size, indicated by the size of the whiskers shrinking from

significantly from five to ten vehicles. Consistency in performance increases the predictability of the hill climb algorithm, an important factor found to be one of the largest barriers to adoption of adaptive algorithms in AUV operations [Brito et al., 2012].

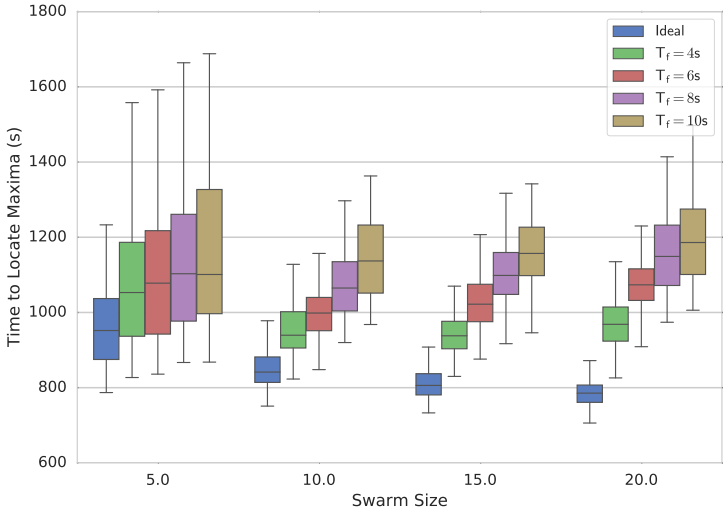

Figure 5.22: Time to locate the maxima of the Gaussian sensor function for each swarm size and communication constraint. The whiskers show the extents of the data, the box shows the upper and lower quartiles and the line within the box shows the median.

For ease of comparison, Figure 5.23 presents the median time to locate the maxima for each swarm size and communication constraint. Examining the ideal case, the time to locate the maxima decreases with increasing swarm size, which is an extremely positive result for the swarm indicating scalability to larger swarm sizes. However, beyond the ideal case this trend becomes less clear. For T_f = 4s, 6s and 8s, an increase in performance is seen from 5 to 10 vehicles, but then this performance decreases again when increased to 15 and 20 vehicles. This result appears to be a trade-off between the positive impact of increased knowledge of the environment, leading to increased directional perception in the hill climb algorithm, and the negative impact of increasing swarm size leading to longer communication cycles and hence more uncertainty in the knowledge of each AUV. These results indicate a limit in scalability to larger swarm sizes.

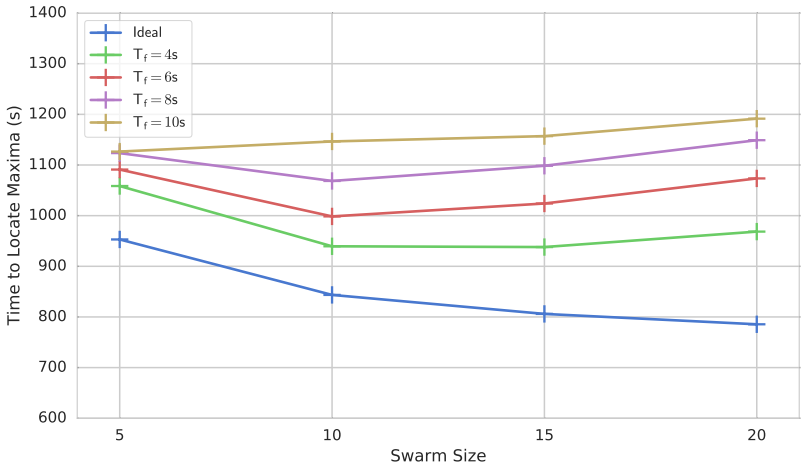

Figure 5.23: Median time to locate the maxima of the Gaussian sensor function for each swarm size and communication constraint. Full results in Figure 5.22

Finally in this section, the flocking performance metrics are examined for the hill climb experiments. Figure 5.24 presents the three metrics. A decrease in flocking performance is evident in both polarisation and angular momentum due to the introduction of the hill climb vector. For example, in Figure 5.24a five of the cases result in a polarisation of less than 0.9, compared to one in the pure flocking analysis (Figure 5.17, pg. 97).This is expected as the introduction of the hill climb vector means dilution of the flocking vectors when the vector average is calculated. Separation seems largely unaffected by the introduction of the hill climb vector. This is due to the magnitude of the separation vector being variable, hence it is able to overcome the effect of the additional vector in the averaging process and drive the vehicles to the desired separation. The performance generally remains satisfactory, achieving a high polarisation and a low angular momentum with separation remaining consistent with pure flocking results (Figure 5.19, pg. 98).

In summary, increasing TDMA frame length generally has a negative effect on the performance of the swarm performing the hill climb algorithm and hence results in poor scalability. A longer frame length results in a greater uncertainty in the information held by each swarm member and hence poorer decision making. Increasing swarm size increases the information available to each AUV leading to

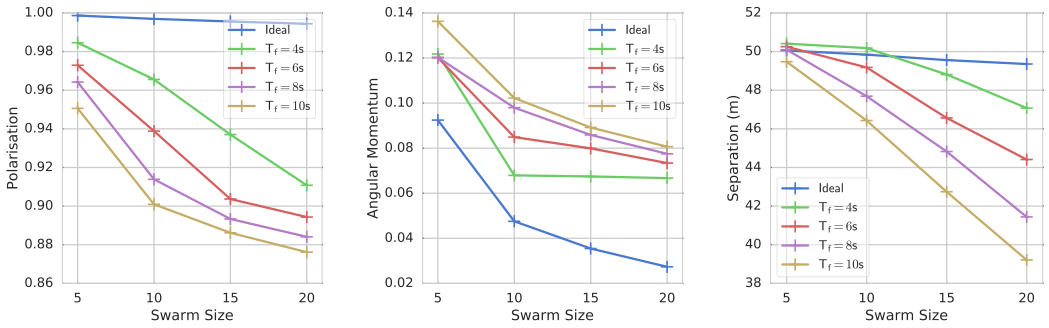

(a) Polarisation (b) Angular Momentum (c) Separation

Figure 5.24: Median Flocking metrics throughout the hill climb experiments for all swarm sizes and communication constraints.

better decision making and increases in success rate, consistency and hill climb performance of the swarm. However, this improvement is bounded by the time to complete one communication cycle which increases with swarm size, limiting scalability. Finally polarisation and angular momentum are negatively affected by the introduction of the hill climb vector, with separation remaining consistent with pure flocking performance.

Performance Relative to a Single Vehicle

In this section, the hill climb performance of the swarm is measured against that of a single vehicle to evaluate whether the swarm offers an increase in task performance. Figure 5.25 presents the hill climb results shown in Figure 5.22 alongside the results for the single vehicle. As the median time for a single vehicle to locate the maxima of the sensory function is greater than that of the swarm, regardless of swarm size or communication constraint, there is an increase in performance offered by the swarm. The improvement indicates the benefit gained from sharing sensor information outweighs the effect of coordinating swarm members.

The performance of the swarm relative to a single vehicle is shown in Figure 5.26. A value of 1 indicates performance equal to that of a single vehicle, with values greater than 1 indicating better performance than a single vehicle and less than 1 indicating worse performance. As mentioned previously, the swarm

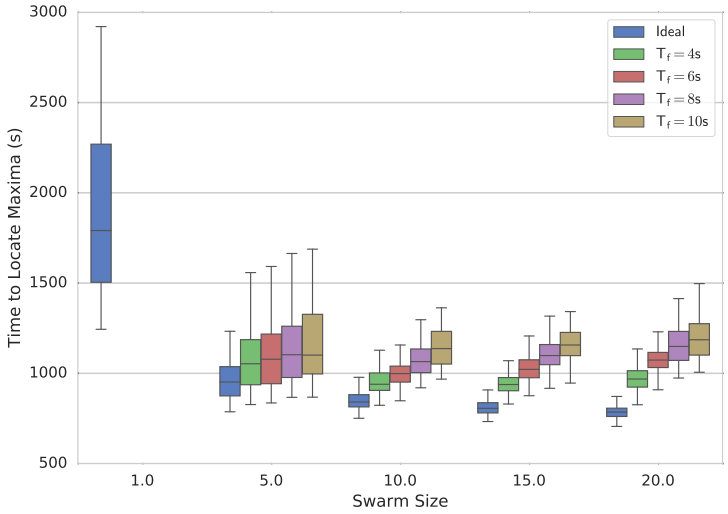

Figure 5.25: Time to locate the maxima of the Gaussian sensor function for each swarm size, including a single vehicle, and communication constraint. The whiskers show the extents of the data, the box shows the upper and lower quartiles and the line within the box shows the median.

provides an increase in performance across all swarm sizes and communication constraints. The relative performance follows a similar trend as previously, with an increase in performance as the swarm size increases from 5 to 10 vehicles, then a decrease in performance at larger swarm sizes. Comparing the trend in the ideal results with the constrained communication results, the trend turns from increasing with swarm sizes to increasing then decreasing, indicating the increasing time to complete a communication cycle introduces the limit to scalability.

Comparing the relative performance to the absolute performance results (Figure 5.23), for $T_f = 10s$ the absolute results show a consistent increase in the time to locate the maxima with swarm size and hence a decrease in performance. However, in Figure 5.26 we see a trend more similar to other communication constraints, with a slight increase in performance from 5 to 10 vehicles and remaining relatively constant from 10 to 15 vehicles. While increasing the swarm size for $T_f = 10$s doesn't offer significant increase in task performance as the relative performance hovers around 1.5 through all swarm sizes, using a larger swarm may be warranted by the increase in success rate from 88% for a swarm

of 5 vehicles to 100% for a swarm of 15 vehicles, shown in Figure 5.20.

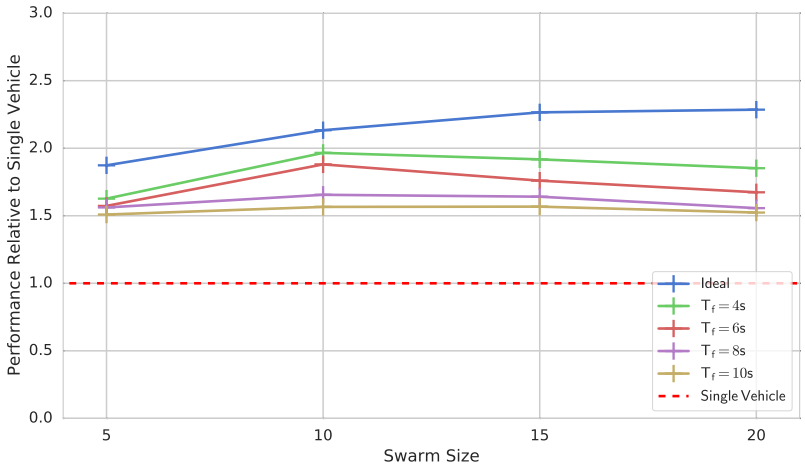

Figure 5.26: Relative performance of the swarm compared to a single vehicle. A value of 1.0 indicates equal performance to a single vehicle.

In summary, the swarm provides an increase in task performance over a single vehicle. However, this increase does not scale with swarm size as the time to complete one communication cycle increases. While relative performance of the swarm is shown to remain constant with increase in swarm size, larger swarm sizes offer no improvement in relative performance. However, an increase in swarm size may be warranted to increase other performance metrics, such as success rate.

Gaussian Sensory Function with Sensor Noise

For the final set of results, the swarm was tasked with performing the hill climb, using the Gaussian sensor function as the environment, while subject to sensory noise. Subject to sensory noise, the single vehicle was not successful in locating the maxima in any of the 100 repeat experiments, making little progress from the starting position in each case. Figure 5.27 presents the success rate of the swarm subject to sensor noise. There is an evident decrease in the success rate compared to the noiseless results (Figure 5.20). However, the majority of results maintain a success rate of 60% or greater. The trend of a decrease in success rate with increase in communication constraint is now evident across more of the swarm sizes, rather than just the 5 vehicle swarm as in the noiseless results.

Originally this was attributed to the poor decision making based on growing uncertainty in the information held by each vehicle. This is now emphasised as there is an uncertainty introduced at the point of measurement due to the sensor noise. Examining scalability with swarm size, there is negligible change in the success rate beyond a swarm of 10 vehicles.

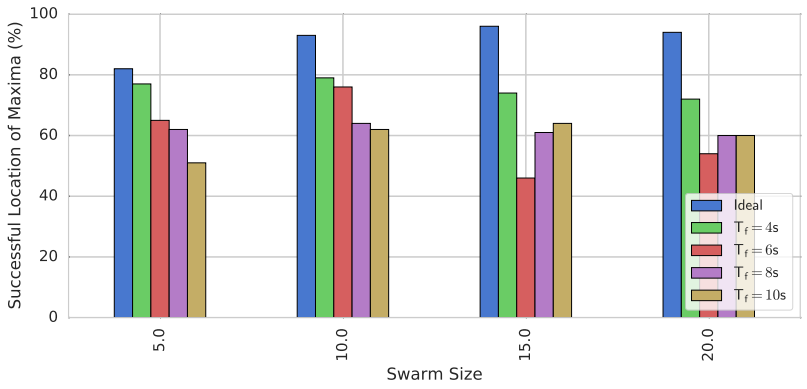

Figure 5.27: Success rate of the swarm performing the hill climb algorithm subject to random sensor noise.

The time to locate the maxima for each swarm size and communication constraint is shown in Figure 5.28. Comparing this to the noiseless results, again there is an evident increase in the time to locate the maxima and hence a decrease in performance. The common trend of decreasing performance with increasing communication constraint is still prevalent in the results. The previously observed relationship of increased consistency in performance with swarm size is not evident enough in this plot to come to the same conclusion.

Examining the results in Figure 5.28 independently, there is a decrease in median time to locate the maxima as the swarm size increases, this is more clearly shown in Figure 5.29. The decrease in time and hence increase in performance is thought to be due to the increased knowledge of the environment leading to an averaging of sensor noise and hence a performance improvement in the noisy environment.

In summary, while there is a performance decrease from the noiseless results, the simple fact the swarm is able to perform the hill climb in an environment

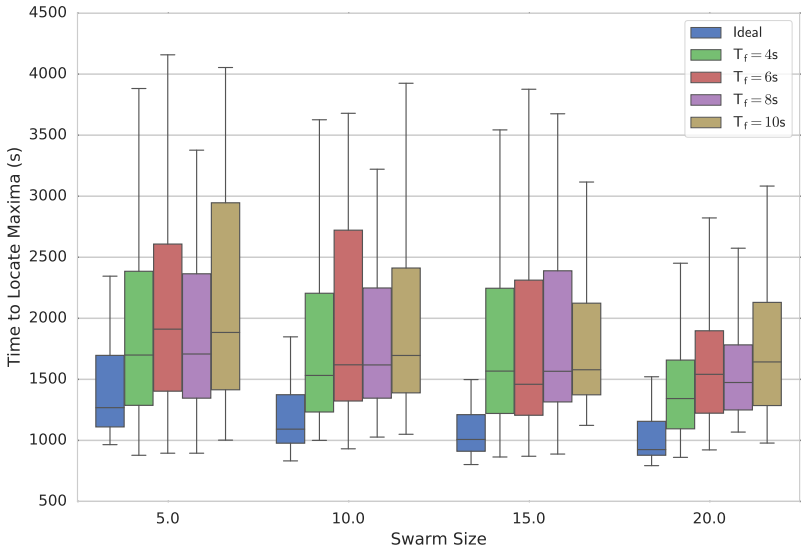

Figure 5.28: Hill climb performance of the swarm for all swarm sizes and communication constraints subject to random sensor noise.

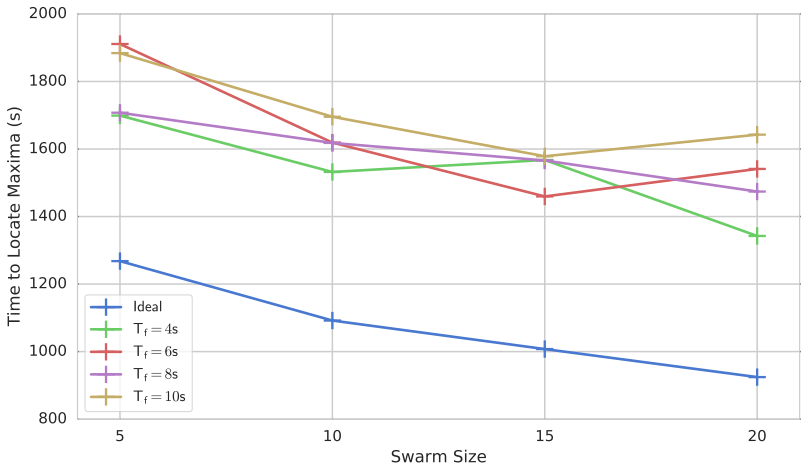

Figure 5.29: Hill climb performance of the swarm for all swarm sizes and communication constraints subject to random sensor noise.

where the single vehicle fails entirely is an extremely positive result in favour of the use of a swarm.

5.4 Summary

Examples of unrealistic communication assumptions including instantaneous, simultaneous and lossless communication in simulation of AUV swarms can be found in Amory et al. [2014] and Read et al. [2014]. Due to the extensive limi-

tations of underwater communication, this study subjects swarm behaviours to more realistic communication constraints and evaluates the ability of the swarm to perform emergent flocking and hill climb algorithms. The emergent flocking algorithm is adapted from Reynolds flocking to allow the separation between swarm members to be controlled. The hill climb algorithm introduces an additional yaw demand into the flocking algorithm, driving the vehicle in the direction of increasing sensor measurements.

The swarms of 5 to 20 vehicles are shown to coordinate and form a well organised flock when controlled by the emergent flocking algorithm and subject to increasing communication constraints. While performance does degrade with increasing swarm size and communication constraints, the impact is not significant enough to conclude emergent flocking is not a suitable method of coordination. In practice, while it is possible, it is unlikely to deploy a swarm of more than 20 vehicles which are required to flock as a single group hence, while ideally the approach would be scalable to any size swarm, these results indicate emergent flocking is a scalable solution to coordination of an AUV swarm.

When performing the hill climb algorithm, increasing communication constraints has a negative impact on performance. However, increasing the swarm size from 5 to 10 vehicles is shown to increase success rate, consistency in performance and reduce the time to reach the maxima of the sensory function. Unfortunately, this improvement is bounded by the length of a communication cycle, evident as performance starts to worsen when the swarm size is increased to 15 and then 20 vehicles. When compared to the performance of a single vehicle, the swarm offers a consistent increase in relative performance to a single vehicle. While performance does not exclusively increase with swarm size, larger swarm sizes are shown to offer higher success rates in locating the maxima, still offering a benefit. The final study examined the hill climb performance subject to random sensor noise introduced into the environment. The single vehicle was unable to locate the maxima in any of the 100 repeat experiments while the

swarm achieved success rates of 60% in most cases, peaking at 79%. These results demonstrate the sharing of sensor information across a swarm in uncertain, noisy environments such as the ocean can achieve tasks impossible by a single vehicle.

The findings in this study indicate that emergent flocking, accompanied by a hill climb algorithm, can coordinate a swarm of AUVs to form a well organised flock and provides a scalable increase in task performance over the use of a single vehicle for the purposes of ocean monitoring.

Chapter 6

Coverage Control

6.1 Introduction

This study implements the Voronoi coverage control algorithm presented in Lindhé et al. [2005] in an AUV swarm. Coverage control can be a useful tool in monitoring the ocean, as events such as HABs and oil spills often occur over a large spatial area. Furthermore, simultaneous monitoring across a large area, at a given spatial resolution, is a monitoring solution only rivalled by remote sensing, which can only observe the top metres of the ocean, or a mesh of sensor moorings, which is expensive to deploy and maintain and simply relies on being in the right place at the right time. Henceforth, an AUV swarm controlled by a coverage control algorithm provides an unrivalled mobile three dimensional monitoring solution for rapidly evolving ocean phenomena.

The aim of the study is to evaluate whether the coverage control algorithm is suitable for an AUV swarm by subjecting it to communication limitations typical of AUV operation. The algorithm was chosen due to the predictability of the final distribution of vehicles, shown to be a key consideration for AUV operators [Brito et al., 2012], and minimal communication dependency often emphasised in Voronoi approaches.

Firstly the swarm simulation is compared against Vehicle-in-Loop (VIL) ex-

periments. The VIL experiments introduce a real ecoSUB AUV to the simulated swarm, introducing disturbances such as drift due to currents and GPS positioning error. All vehicles in the swarm, real or simulated, are controlled by the coverage control algorithm and subject to the same communication constraints. Following on, a scalability study is performed in simulation and the performance of the swarm at increasing swarm sizes and subject to various communication constraints is evaluated. Communication constraints are chosen to represent a spectrum of AUV communication, from frequent peer to peer communication via acoustic communication to infrequent, one-to-basestation, client server satellite communication.

In coverage control literature, Miah et al. [2015] and Patel et al. [2016] investigate communication constraints which can be likened to those in AUV operation. However, these works do not extend to the extreme communication latencies experienced in AUV communication. This work evaluates the performance of the swarm under the coverage control algorithm subject to a spectrum of communication constraints representative of AUV operation. This work also includes vehicle in loop experiments, using a real ecoSUB AUV interacting with several simulated vehicles to validate the algorithm for practical implementation in ocean monitoring.

6.2 Method

The general methodology for this programme of research is to use real vehicle deployments to validate simulation before using the simulation to evaluate the scalability of certain algorithms or approaches to increasing swarm sizes, communication constraints and survey areas. This section focuses on the methodology used to evaluate whether the communication constraints of AUV operation, especially in larger swarms, are too limiting for an AUV swarm to perform a coverage control algorithm. The coverage control algorithm and metrics for measuring

performance are presented. Following on, the approach the vehicle-in-loop experiments and the method used are presented. The vehicle-in-loop experiments involve a single real ecoSUB vehicle interacting with a simulated swarm and are an initial step towards full swarm validation. Finally, the approach to simulating a swarm of AUVs is presented with a method for the scalability study carried out.

6.2.1 Coverage Control

Coverage control aims to drive individual swarm members to a desired distribution in order for the vehicles to sample the environment across an area. In this work, the algorithm presented by Lindhé et al. [2005] is used to control the inter-vehicle separation using Voronoi partitioning, removing the flocking driven by a navigation function also presented by Lindhé et al. [2005].

Constructing a Voronoi partition requires a set of points in a plane which can be divided into regions. Each region in a Voronoi partition represents the subset of the plane closest to the point within the region than any other. An example of a Voronoi partition can be seen in Figure 6.1a. Furthermore, the convex hull of the points, shown in Figure 6.1b, is the smallest convex set which contains all the points and defines which of the points have unbounded Voronoi regions which extend to infinity.

Algorithm The coverage control algorithm, presented by Lindhé et al. [2005] and reproduced in Algorithm 2, drives the swarm to a target inter-vehicle separation. The algorithm relies on a swarm member determining the centroid of its Voronoi region. However, it is possible for a members region to extend to infinity and hence it would be impossible to determine the centroid, the case for vehicles in the convex hull. To resolve this, if the member is in the convex hull and hence has an unbounded region, the algorithm defines mirror neighbours according to Equation 6.1 taking into account the desired inter-vehicle separa-

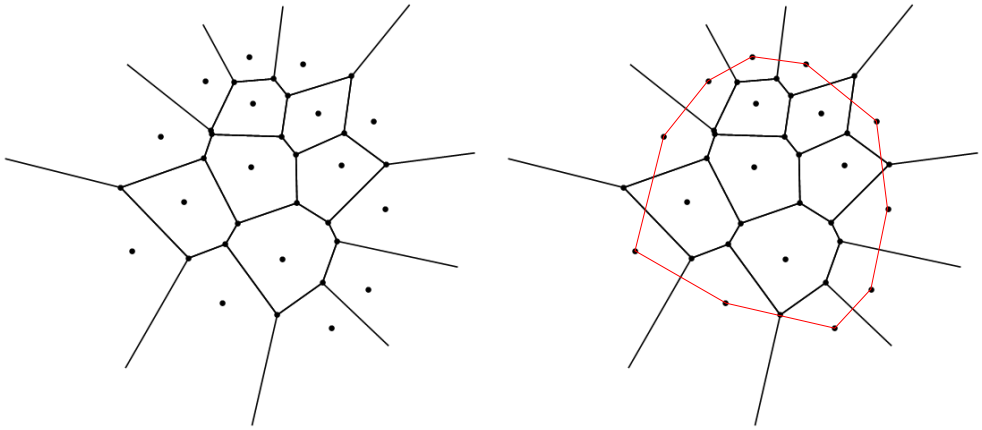

(a) Voronoi partition (b) Convex Hull

Figure 6.1: Example of Voronoi Partition and Convex Hull

tion. These neighbours are then included in the Voronoi partition and bound a members region, allowing the centroid to be determined. Figure 6.2 displays the steps of the algorithm.

As mentioned previously, the navigation function considered in Lindhé et al. [2005] is removed for this work. This means when the centroid is calculated, instead of performing two integrals to determine the centre of mass according to the navigation function and the centroid of the region, only the latter needs to be determined reducing computation demand.

$$\hat{p}_j = p - d\frac{p_j - p}{\|p_j - p\|} \tag{6.1}$$

Algorithm 2 Coverage control algorithm from Lindhé et al. [2005], aiming to control the inter-vehicle separation. Let P be the positions of swarm members as known by the vehicle, p_i the position of the evaluating swarm member, d the desired inter-vehicle separation and n the swarm size.

function VORONOI COVERAGE(p_i, P)
  if p_i is not inside the convex hull of P **then**
    Determine neighbours of p_i, $P_i \subset P$
    for $p_j \in P_i$ **do**
      Create mirror neighbour, $\hat{p}_j$ (Eq. 6.1)
    end for
    $\hat{P}_j = \bigcup_{j=1}^{n} \hat{p}_j$
  end if
  $\hat{P}_i = P_i \cup \hat{P}_j \cup p_i$
  Determine Voronoi region, V, based on $\hat{P}_i$
  Determine v, the region enclosing p_i
  C_v = centroid of v
  return C_v as waypoint
end function

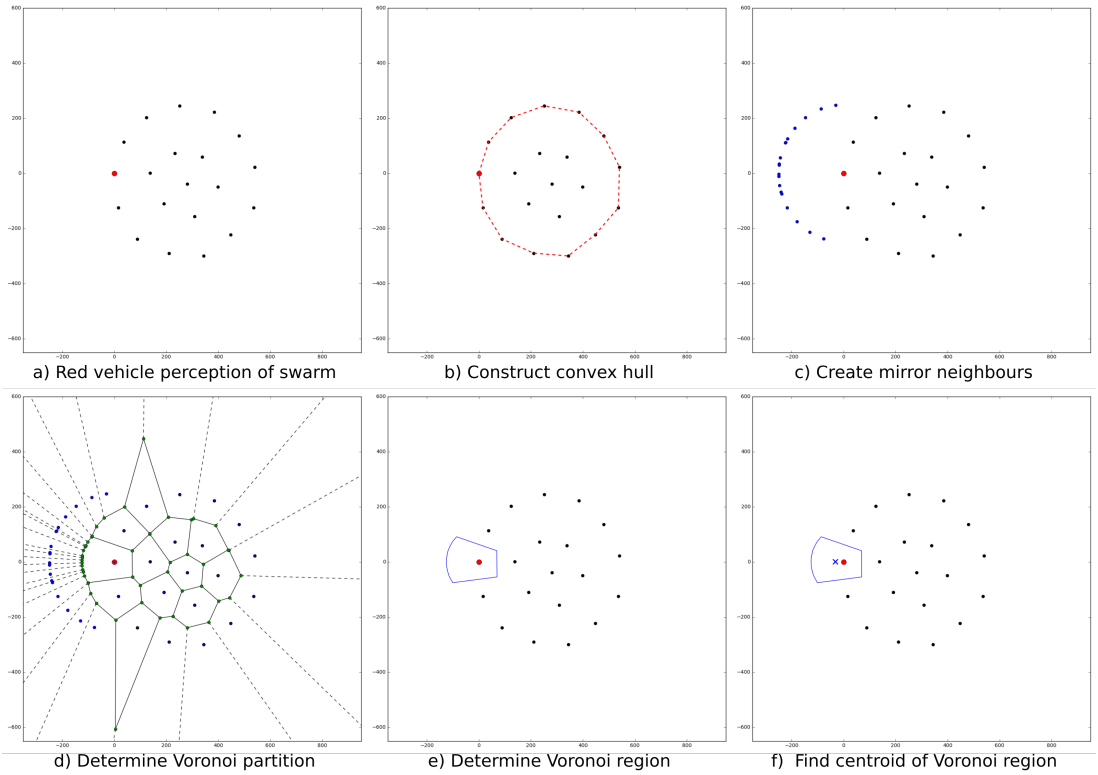

Figure 6.2: Steps of the coverage control algorithm. Black markers represent swarm members, blue are mirror neighbours and the red, larger marker indicates the member evaluating its next waypoint.

6.2.2 Measuring Swarm Performance

This section presents the metrics by which the scalability of the algorithm is measured. All performance metrics are calculated based on the inter-vehicle separation as this is what is being driven by the coverage control algorithm. The inter-vehicle separation is calculated as the median distance to a vehicles immediate neighbours then a median is taken across the swarm, resulting in a value for each timestep. Figure 6.3 shows this and an immediate neighbour can be defined as a vehicle whose Voronoi region shares an edge with the region of the evaluating vehicle.

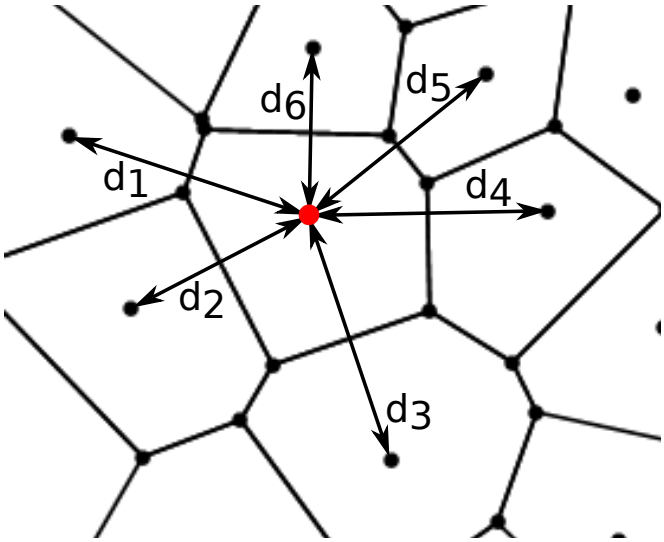

Figure 6.3: Calculation of inter-vehicle separation. The inter-vehicle separation is taken as the median of values $d_j, j = 1, 2, 3, .., 6$.

Criteria are presented by Dorigo and Sahin [2004] to define what constitutes a robotic swarm. They are as follows:

- The swarm should be almost entirely homogeneous due to the high redundancy required
- Each swarm member should have local sensing and communication capabilities
- The swarm should be scalable and not restricted to a maximum number of members.
- The implementation of a swarm must lead to increased task performance

The first two criteria are met by the characteristics of the vehicle. The last two, scalability and the increase in task performance, can be used to assess the performance of an AUV swarm. Scalability is of particular interest in AUV swarms as the communication limitations increase with swarm size when operating under a TDMA protocol. On the other hand, the increase in task performance is less important, as implementing an AUV swarm provides additional benefits including robustness to vehicle failure, reduced financial risk due to the simplicity and thus low cost of individual members, and ease of deployment.

Under the coverage control algorithm, the swarm is driven to an inter-vehicle separation defined by the user. Assuming convergence is reached within 90% of the run time, a final inter-vehicle separation is determined by averaging over the last 10% of the run i.e. approximately 100 seconds in vehicle-in-loop experiments and 1000 seconds in simulation. The time taken for the swarm to initially reach the desired inter-vehicle separation, the 'rise time' or t_{Rise}, is used as a performance metric however this does not give any indication of whether the swarm reaches a stable, equilibrium state. In addition to the rise time, the standard deviation of the inter-vehicle separation between t_{Rise} and the run time, the 'steady state standard deviation' or σ_{ss}, is used to represent the 'steadiness' of the system. Finally, to indicate how accurately the algorithm is able to control the inter-vehicle separation, the error between the final and desired inter-vehicle separation is evaluated, this is referred to as the steady state error and denoted by e_{ss}. These performance metrics allow the scalability of the algorithm to be evaluated when subject to increasing swarm sizes and varying communication constraints. Figure 6.4 illustrates these metrics on a typical plot of inter-vehicle separation against time for a swarm controlled by the coverage control algorithm. Figure 6.5 shows a comparison between two results demonstrating the difference between a low value and high value of the steady state standard deviation.

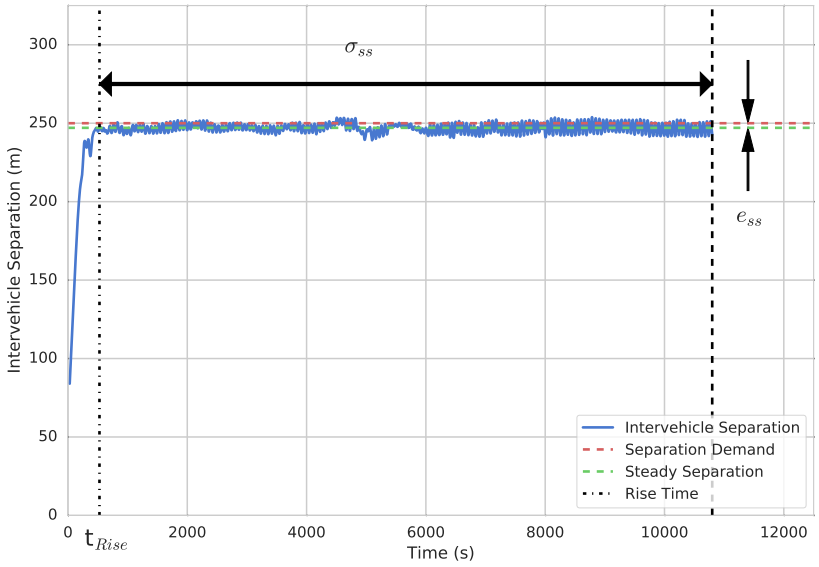

Figure 6.4: The three performance metrics shown on a typical plot of inter-vehicle separation against time for a swarm controlled by the coverage control algorithm.

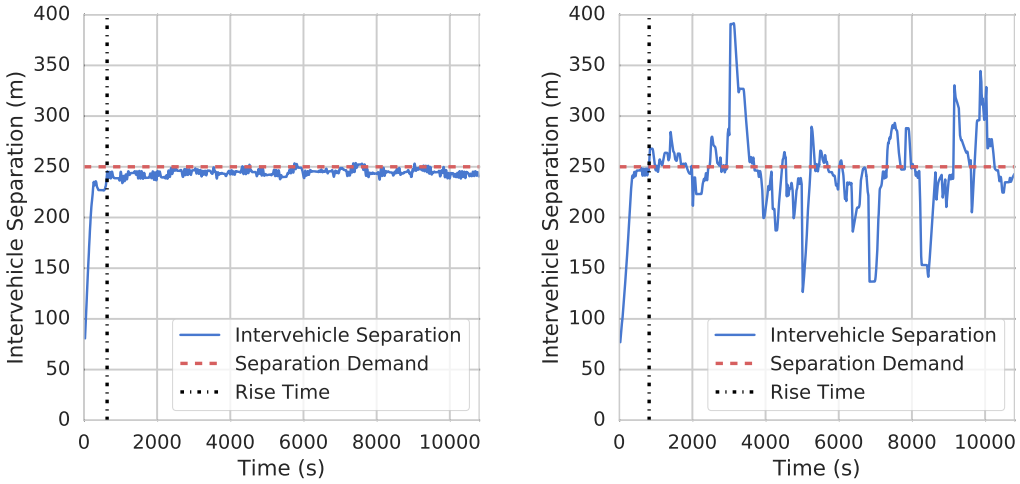

(a) (b)

Figure 6.5: Example distributions demonstrating the steady state standard deviation metric. (a) A very stable distribution achieving $\sigma_{ss} = 3.18$. (b) a less stable distribution with large fluctuations and thus a large standard deviation $\sigma_{ss} = 41.54$.

6.2.3 Vehicle in Loop Experiments

To validate the assumptions made in simulation, experiments with an ecoSUB vehicle are performed at Testwood Lake, Totton, UK. Figure 6.6 shows the ecoSUB vehicle undergoing humidity, pressure and battery checks at the lake before deployment.

The communication between real and simulated swarm members occurs over WiFi when the member either reaches its waypoint or reaches the time out (60 seconds). The WiFi communication reflect satellite communication in topology, meaning it follows a client-server topology. However, the frequency of communication, at least once a minute, is closer to the communication frequency of acoustic communication. Hence, while the communication over WiFi is not representative of either satellite or acoustic, the same communication constraints can be applied in simulation for comparison.

Figure 6.6: The ecoSUB vehicle being prepared for deployment at Testwood Lake, Totton, UK.

Integration with the ecoSUB AUV

At this point, complete integration with the vehicle was not a necessity, hence the coverage control algorithm was run on a laptop connected to the vehicle over WiFi. This also allowed monitoring of the vehicles health including battery life and hummidity, temperature and pressure inside the pressure vessel to ensure the vehicle did not have a leak. The laptop was connected to the Robotic Operating System (ROS) network onboard the ecoSUB and hence had access to all data available to the onboard software over the WiFi connection. Figure 6.7 shows a block diagram of the integration and the following sections explain the jobs of

the individual blocks.

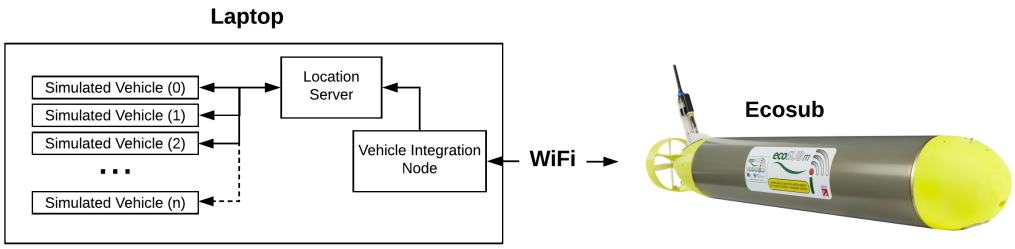

Figure 6.7: Block diagram showing the integration of the algorithm with the real vehicle and allowing real and simulated vehicles to exchange data.

Vehicle Integration Node The vehicle integration node monitors the state of the ecoSUB vehicle and sends new waypoints when needed. Through monitoring the state of the onboard state machine, the node can detect when the vehicle returns to its 'idle' state. Once an idle vehicle is detected, the node reads the vehicles latitude and longitude from the onboard data. This data is then sent to the location server, which provides the last known positions of other swarm members in response. Now the node has the most recent data available, the coverage control algorithm can be performed to determine a new waypoint. At this point, distance and bearing to the new waypoint are presented to the operator for validation, giving the operator chance to intervene. The new waypoint is then written into a mission file in the ecoSUB specific format and sent to the ecoSUB for execution.

Location Server The location server controls the flow of data between vehicles. It is programmed as a ROS service which means nodes must send a request and the service sends a response in return. For the ecoSUB vehicle, the vehicle integration node sends the vehicles latitude and longitude as a request to the location server. The location server receives this data, updates its knowledge of the vehicles location and responds with the last-known positions of all vehicles in the swarm.

Simulated Vehicle These nodes simulate other vehicles in the swarm. The simulation is the same as in Chapter 4 with some slight changes for integration with the location server. When the vehicles reach the given waypoint or exceed the given time out, they send a request to the location server containing their location, and receive the last known positions of the swarm. Using the location server for both real to simulated and simulated to simulated vehicle communication allows communication to be strictly controlled. When the vehicles communicate with the server they are subject to a delay, this imitates the time taken to safely retask the real ecoSUB.

Vehicle in Loop Experiments to Validate Simulation

Response to Static Nodes In order to verify the implementation of the coverage control algorithm and gain some measure of the error to be expected in the vehicle in loop experiments, the position of the real vehicle is passed to the coverage control algorithm along with static nodes. The static nodes are arranged in a equilateral triangle with side length 50m, the desired separation also passed to the coverage control algorithm. As there is only one moving vehicle, the ecoSUB, it's final position can be predicted. Figure 6.8 shows the experimental concept. Here the vehicle is subject to a 1 minute time out to reach the 5m acceptence radius on the waypoint returned by the coverage control algorithm.

Response to Simulated Vehicles Moving towards the deployment of a real swarm, the static nodes used previously are now replaced with simulated vehicles which follow the model outlined in Section Chapter 4. This experiment aims to quantify the introduction of a real vehicle as it represents disturbances unaccounted for in simulation. These include navigational error, variable speed and disturbances due to currents.

The ecoSUB and three simulated vehicles form the swarm and are controlled by the coverage control algorithm. The ecoSUB is initially sent to a waypoint in the centre of the lake, then starting positions for the simulated vehicles are

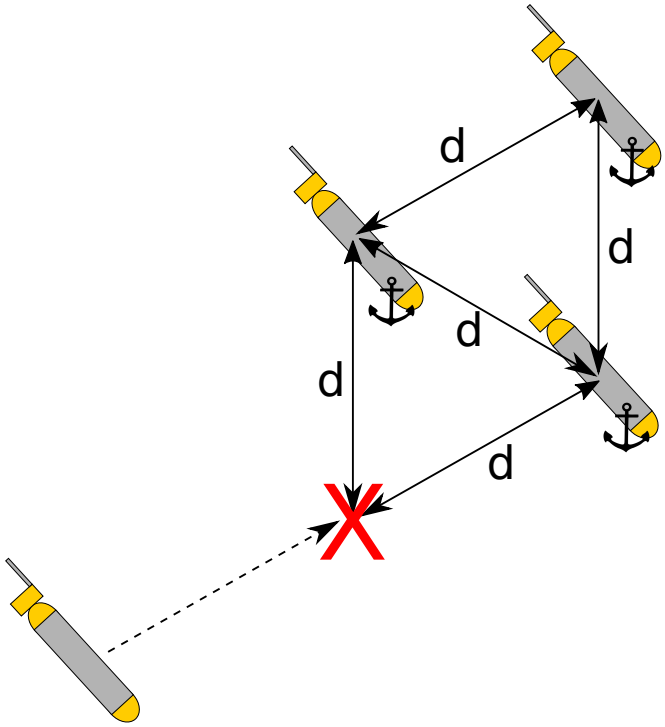

Figure 6.8: Experimental concept with static nodes. Fixed nodes are indicated by the anchor symbol, d is the desired inter-vehicle separation and the final location of the real ecoSUB is shown by the red X.

generated relative to its position at 10m separation. The starting positions are generated as follows; one 10m to the north of the real ecoSUB, one 10m away at 60° to the real ecoSUB and one 10m north of this position. This generates the initial starting distribution depicted in Figure 6.9 and allows the approximate movements of the real vehicle to be predicted.

The generation of the simulated vehicles relies on receiving a new position from the ecoSUB and this occurs only when the start is triggered by the operator, hence all AUVs begin moving towards the initial waypoint at the same time. Once the waypoint is reached or the timeout for the behaviour exceeded, all vehicles communicate with the base-station to re-task.

Communication between all vehicles occurs in a one-to-basestation fashion, following the assumptions for satellite communication in Section 4.2.1. Four repeat experiments are performed and the inter-vehicle separation is varied to introduce a disturbance to the stability of the distribution. Table 6.1 shows the values used for the parameters throughout the four repeat experiments.

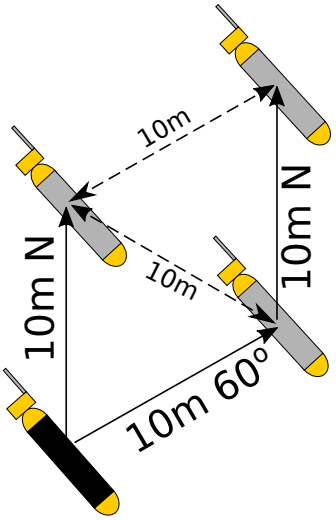

Figure 6.9: Starting positions of the simulated vehicles relative to the position of the real vehicle (Black Filled). Calculations are shown as solid lines and dashed lines represent resultant distances.

Table 6.1: Parameters used in experiments with a real vehicle for validation of simulation

Parameter	Value
Inter-vehicle Separation	100m
Time-out	60s
Acceptance Radius	3m
Surface delay (Simulated)	20s

6.2.4 Simulation Experiments to Evaluate Scalability

Simulations are performed in order to evaluate the scalability of the algorithm to large swarm sizes, hence longer TDMA communication cycles, and varying communication constraints. Voronoi coverage control in general boasts minimal dependency on communication. However the delays experienced between communication events in AUV operation are uniquely large and hence not considered in literature. This is the first step to verifying if the coverage control algorithm is a suitable candidate for environmental monitoring using an AUV swarm.

Simulations are performed with swarm sizes of 5, 10, 15 and 20 vehicles. While the philosophy of swarm robotics is often that a swarm should be scalable to any number of robots, the communication constraints of AUV operation impose a limit on this scalability, mainly due to bandwidth limits in both acoustic and satellite communication. The largest swarm size, 20 vehicles, represents a

deployment of vehicles significantly larger than what is typical even in large fleet deployments [National Oceanography Centre, 2018]. While 20 vehicles is not a maximum swarm size, it is still a significant logistical challenge given the way AUVs are currently operated and hence the limit in this work.

Each swarm size is subject to the following communication constraints, T_f denoting the framelength in the TDMA protocol and $t_{timeout}$ denoting the time an AUV has to reach a waypoint before surfacing to communicate:

- Ideal

  The vehicles have complete knowledge of the swarm at evey time step

- Operation on satellite and acoustic communication

  $T_f = 4s$ and $t_{timeout} = 10$ minutes

  $T_f = 6s$ and $t_{timeout} = 10$ minutes

- Operation on satellite communication only

  $t_{timeout} = 5$ minutes

  $t_{timeout} = 10$ minutes

Ideal communication removes the effect of communication and hence the result is representative of AUV vehicle dynamics and behaviour, acting as a benchmark against which to measure the impact of communication delay. At the other end of the spectrum, satellite only operation represents the vehicles communicating asynchronously, at least every 5 or 10 minutes. While acoustic communication is common in AUV operation between an AUV and an operator, vehicle to vehicle acoustic communication is less common. In operation, ocean disturbances may push vehicles beyond the range of the acoustic modem meaning without satellite communication the behaviour cannot be robust to the environment. In addition, vehicles are not always fitted with acoustic modems as standard and hence it limits the approach to assume acoustic communication is always available. These are all reasons for including operation on satellite communication only.

The initial locations of the vehicles are randomly seeded within a 150m by 150m starting area and 100 repeats are performed using paired random seeds to permit comparison between communication constraints at each swarm size.

Each simulation is run for 3 simulated hours. In operation, the inter-vehicle separation will be derived from the the spatial scale of the feature being monitored in order to provide the best data. To achieve both temporal and spatial data, sampling must be sustained at the desired resolution for a significant period of time. Hence, if the swarm is not able to achieve the desired sampling resolution within 3 hours, this represents a failure of the monitoring task. In addition, this does not provide confidence if moving to an adaptive approach where the vehicles are required to adapt quickly to changes in the sensed distribution.

6.3 Results & Discussion

Firstly, results from vehicle-in-loop experiments are presented using a single AUV interacting with static nodes and then multiple simulated vehicles controlled by the coverage control algorithm. These results are then compared to total simulation to evaluate the difference and validate the model of the ecoSUB AUV. Once validated, the simulation is used to evaluate the scalability of the coverage control algorithm when subject to communication constraints typical of AUV operation.

6.3.1 Vehicle-In-Loop Results and Simulation Validation

Response to Static Nodes

To first validate the implementation of the coverage control algorithm, the real vehicle responds to static nodes at locations which allow the Expected Final Location (EFL) of the real vehicle to be determined. Figure 6.10 shows an example run of this experiment. It can be seen that both the real vehicle and the simulation stop short of the EFL. This occurs as the distance to the waypoint output by the algorithm decreases as the inter-vehicle separation approaches the desired value. This means the waypoint output by the algorithm is now inside the acceptance radius for the waypoint following behaviour and hence the vehicle considers the behaviour a success. This could possibly be improved by reducing

the acceptance radius. However this causes 2 problems. First the acceptance radius in this experiment was 5m, already close to the error on GPS positioning. Second, the AUV may be unable to reach the waypoint causing it to circle around in an attempt to reach it due to the flight-style nature of the ecoSUB. At small scales this would have a significant impact on inter-vehicle separation. However would encourage persistence to reach the waypoint. A second observation is that the simulated AUV re-tasks more, indicated by plateaus in the plot. This is due to the simulated AUV reaching waypoints more successfully and could be due to differences in vehicle constraints or controller performance. However, at a later date issues with the ecoSUBs rudder became apparent which could have been present in this experiment. This means the vehicle could not control its heading as well as it should have been able to and hence, struggled to reach waypoints.

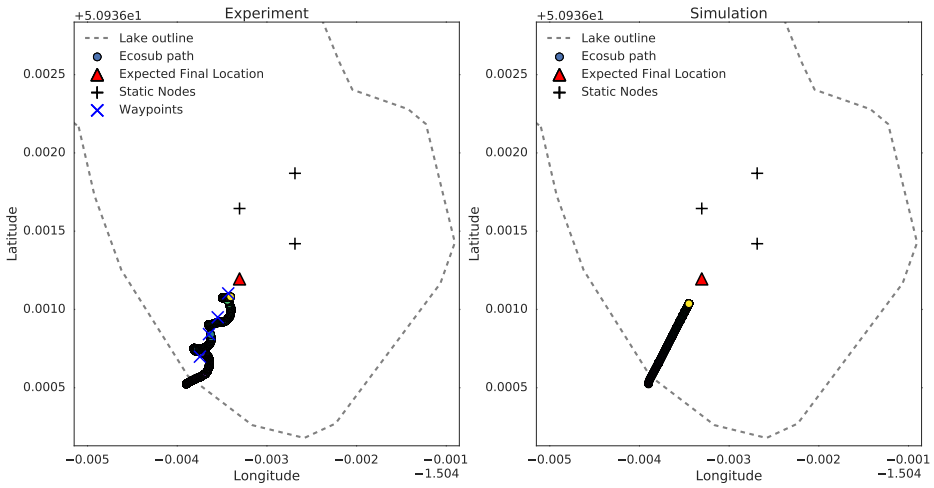

Figure 6.10: Example of the path taken by the experimental vehicle and a simulated vehicle in response to the static nodes under the coverage control algorithm.

Figure 6.11 shows the three repeat experiments performed. Similar behaviour is observed in each plot and the rate of progress towards the EFL and the final distance to EFL being similar to simulation in all three cases. The three plots also show the reduced rate of re-tasking for the real vehicle as it struggled to reach waypoints. Note, Figure 6.11b and 6.11c begin at a smaller distance to the EFL as WiFi connection issues meant the static nodes had to be moved closer to

the starting location of the real vehicle, and hence the WiFi antenna.

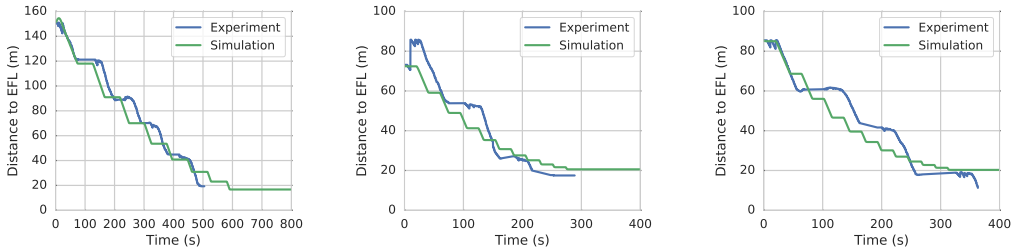

(a) (b) (c)

Figure 6.11: 3 repeat experiments showing vehicle response to static nodes under control of the coverage control algorithm.

In summary, the real vehicle responds to the static nodes with the desired behaviour when controlled by the coverage control algorithm, showing performance similar to that of a simulated vehicle under the same conditions.

Response to Simulated Vehicles

Here, the vehicle-in-loop experiments are compared to an entirely simulated swarm in which the vehicles operate under the same parameters and from the same starting locations. Figure 6.12 shows one of the four vehicle-in-loop experiments as a representative example and is compared to the equivalent simulation. First, the period of constant 100m inter-vehicle separation demand up to t=1500s will be examined. The simulation result rises to a steady state in 205s and remain at a constant steady-state error of 4.71%. This steady-state error exists because the distance to the waypoint output by the coverage control algorithm falls below the acceptance radius for a given waypoint, as seen in the response to static nodes, introducing error into the distribution. Due to the perfect navigation and lack of disturbances in simulation, the swarm remains in this steady state indefinitely. On the contrary, the vehicle-in-loop experiment achieves a similar rise time of 210s. However introduction of the real vehicle and environmental disturbances continually disrupt the state and the vehicles have to continually evaluate and adapt to maintain the inter-vehicle separation. As a result, the steady state error in the vehicle-in-loop results is lower at 0.55%, but the inter-vehicle sepa-

ration fluctuates more, indicated by the steady state standard deviation of 3.67 in comparison to 0.19 in the simulation. At t=1500s the separation distance is reduced, this is updated on the server, then communicated to the vehicles when they communicate with the basestation. This demonstrates the ability of the algorithm to move between stable states even with a real vehicle in the loop.

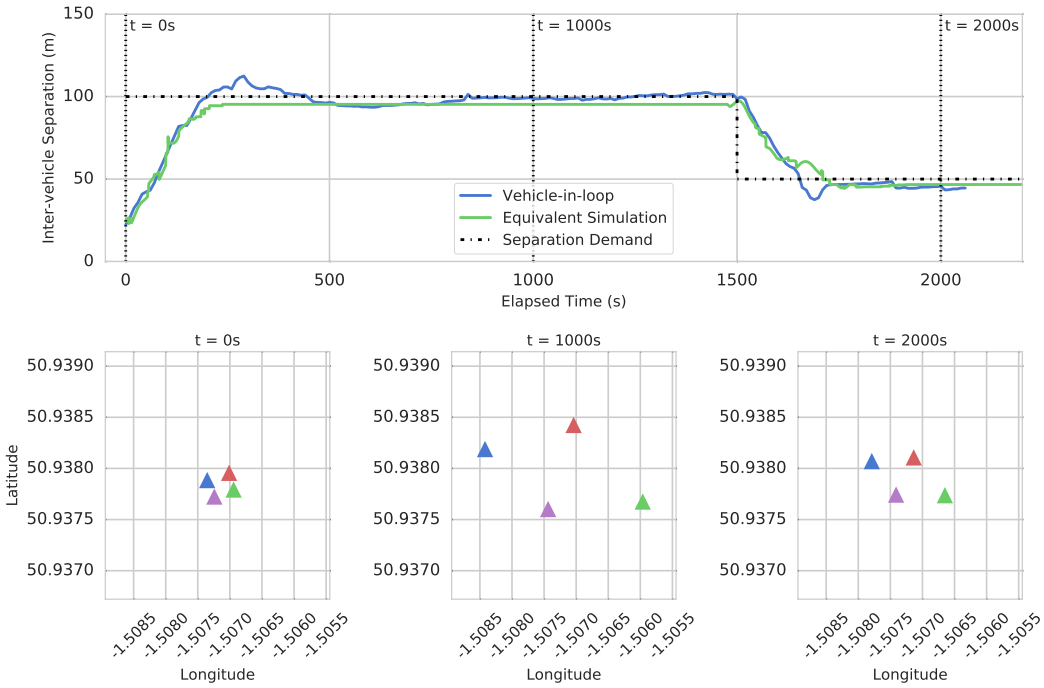

Figure 6.12: Top: Inter-vehicle separation against time for one of four vehicle-in-loop experiments and the equivalent simulation. Also shown is a change in separation demand at t=1500s. Bottom: 3 snapshots of the distribution of the vehicles in space at t = 0s showing the start positions, t = 1000s showing the distribution once the swarm has converged and t = 2000s when the inter-vehicle separation demand is reduced to 50m

Figure 6.13 shows inter-vehicle separation against time for the both vehicle-in-loop and simulation results, averaged across the four repeats. The rise time for the averaged simulation results is found to be 266s, this is higher than the vehicle-in-loop results which rise to a steady state in 210s. In Figure 6.13, the vehicle-in-loop experiments achieve a lower rise time due to overshoot, whereas the simulated experiments gradually rise to the steady state. The comparison between the averaged results for steady state error and standard deviation is very similar to those presented in Figure 6.12, with the vehicle-in-loop results

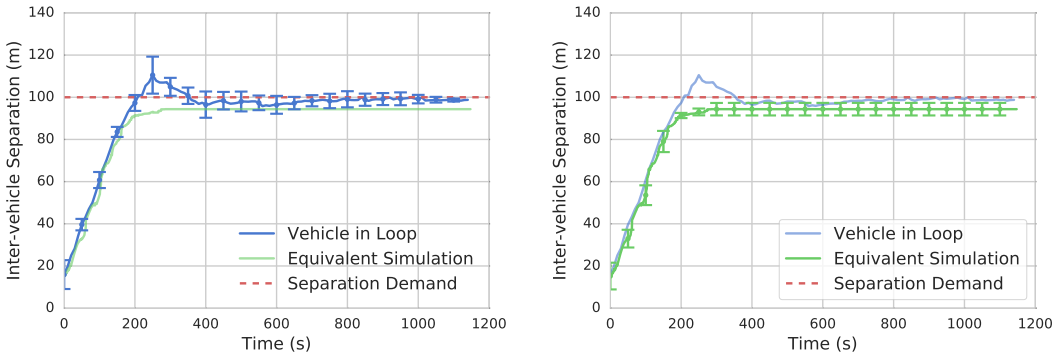

(a) (b)

Figure 6.13: Both plots show the inter-vehicle separation of the swarm against time for the vehicle-in-loop experiments and equivalent simulation results, both averaged over the 4 repeats with the standard deviation shown by the error bars. Plots separated for clarity.

achieving a lower error, 0.72% compared to 5.65% in simulation, but fluctuating more once the rise time has passed shown by the standard deviation of 2.82 compared to 0.08 in simulation.

In summary, the results indicate the introduction of a real vehicle, and hence current induced drift and positioning error, does not significantly negatively impact the performance of the swarm under the coverage control algorithm.

6.3.2 Simulation Results and Scalability Evaluation

As the introduction of a real vehicle and it's accompanying disturbances in the form of drift and sensor error does not hinder the ability of the swarm to perform the coverage control algorithm, the model is now used to evaluate swarms of increasing size subject to varying communication constraints. This section looks at swarms of 5, 10, 15 and 20 vehicles in total simulation and evaluates the scalability of the coverage control algorithm by examining the rise time, steady state standard deviation and steady state error.

Rise Time

Figure 6.14 shows the rise time for each swarm size for each of the five communication constraints. Examining the results of a swarm subject to the satellite only

constraints, both 5 and 10 minute time-outs show a clear increasing trend with swarm size indicating, whilst only using satellite communication, the algorithm is not scalable to larger swarm sizes. In addition, the range of rise times seen grows demonstrated by the increasing size of the whiskers with swarm size for the 5 and 10 minute time-outs, indicating unpredictability in performance. The satellite and acoustic communication constraints demonstrate a much more gradual increasing trend with swarm size, with the rise time for the $T_f = 6s$ constraint increasing at a rate of 51s per additional vehicle as oppose to 167s under satellite only constraints. Some increase in rise time with swarm size is expected. As the swarm size increases, so does the time to complete one communication cycle, actively reducing the rate at which a swarm member receives information from the same neighbour and hence slowing the rate of expansion. There is little performance difference between the two acoustic cases, demonstrating scalability to more infrequent communication within reason. More infrequent communication could be intentional to allow more data to be communicated or unintentionally through packet loss.

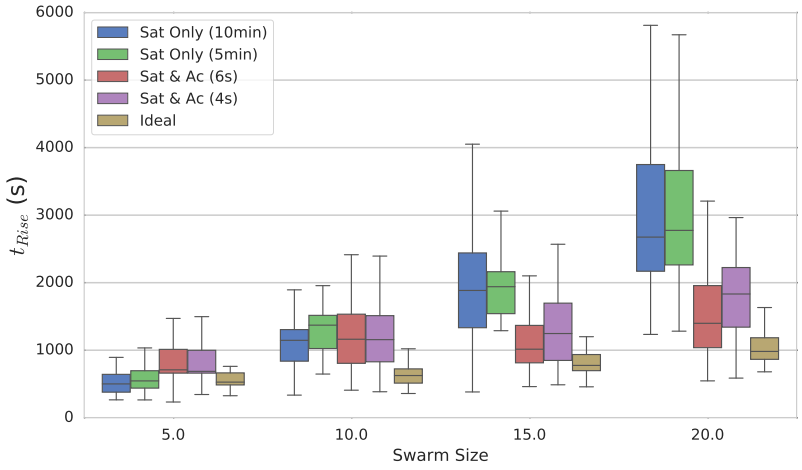

Figure 6.14: Mean rise time for each swarm size and each communication constraint across 100 repeats. Error bars show standard deviation.

Steady State Standard Deviation

Examining the steady state standard deviation, Figure 6.15 presents σ_{ss} for each swarm size given each communication constraint. This value is a measure of the fluctuation about the mean inter-vehicle separation from t_{Rise} to the run time. The satellite only communication constraints result in large values for σ_{ss}. Even though the fluctuation is large, implying poor stability, σ_{ss} does decrease with swarm size, implying increasing the swarm size leads to a more stable distribution. In addition, the size of the whiskers shrink, indicating more consistency in performance which is important for predictability. The satellite and acoustic communication constraints result in a very low σ_{ss} and a much more stable distribution of vehicles in general. A swarm of 5 vehicles achieves the best results, consistently achieving minimal values of σ_{ss}. From 5 to 10 vehicles there is an increase in σ_{ss}. However, the results from 10 vehicles onwards demonstrate little to no increase with swarm size, demonstrating scalability. σ_{ss} does increase when increasing T_f from 4 to 6 seconds, hence more infrequent acoustic communication leads to worse performance.

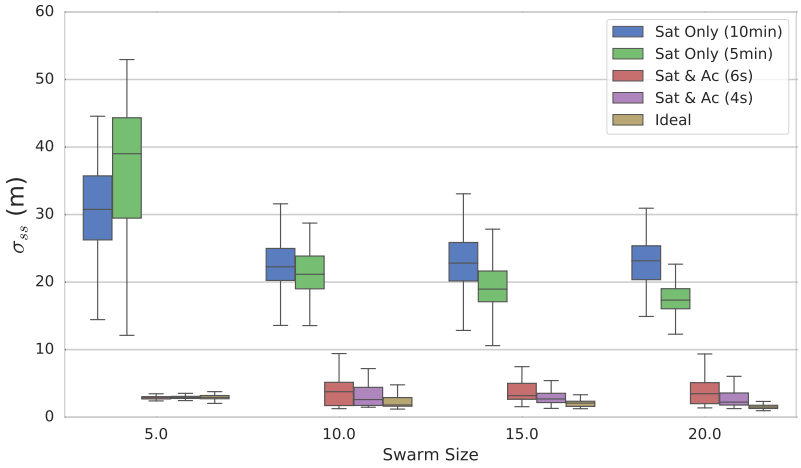

Figure 6.15: Steady state standard deviation for each swarm size and each communication constraint averaged over 100 repeats. Error bars show standard deviation.

Steady State Error

The steady state error, Figure 6.16, again shows poor performance for the satellite only cases, increasing with swarm size and surpassing a median error of 10% for a swarm of 20 vehicles. The range of values achieved across the 100 repeats for e_{ss} also increases demonstrating unpredictability in the performance. For the satellite and acoustic communication constraints, the results presented show median performance similar to the ideal communication case across all swarm sizes, demonstrating scalability. This shows the steady state error is less effected by the communication constraints and more so by the flight-style nature of the vehicles.

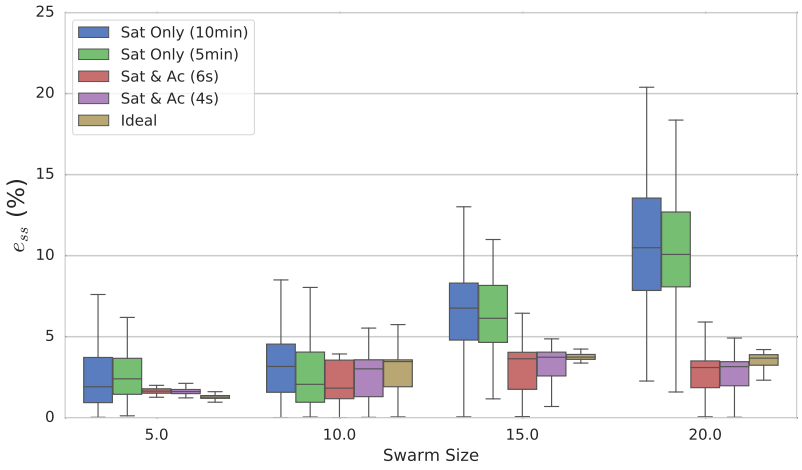

Figure 6.16: Mean steady state error as a percentage of desired separation (250m) for each swarm size and each communication constraint across 100 repeats. Error bars show standard deviation.

Summary

In summary, it is a common trend that, when operating under satellite communication alone, the swarm performs poorly, demonstrating an unstable distribution of vehicles unable to maintain the desired inter-vehicle separation. Conversely, through the introduction of acoustic communication performance is significantly improved. Operating on acoustic communication, the swarm is demonstrates a small increase in rise time with swarm size, however steady state deviation and

error remain relatively consistent with increasing swarm size, demonstrating scalability of the approach. Steady state standard deviation is the only performance metric to show a clear, albeit small, improvement with more frequent communication. This leads to the conclusion that the Voronoi approach is also scalable to more severe acoustic communication constraints with more investigation needed into the limit to this scalability.

6.4 Summary

The work included in this report presents the implementation of a coverage control algorithm, presented by Lindhé et al. [2005], for an AUV swarm. The first aim of this work is to validate simulation through experiments with a real ecoSUB vehicle interacting with a simulated swarm. The second aim is to use the validated simulation to evaluate the performance of increasing swarm sizes and varying communication constraints.

Validation experiments have been presented in which a real ecoSUB AUV interacts with static nodes and multiple simulated AUVs. The experiments using static nodes validated the implementation of the coverage control algorithm and showed similar performance in the vehicle in loop and simulation results. These results also highlighted the issue that, even though the vehicle was still beyond the acceptance radius of the expected final position, it was within the acceptance radius of the waypoint output by the algorithm, causing a significant steady state error at small inter-vehicle separation.

Through introducing a real vehicle, the negative impact of drift and GPS positioning error was shown to be small and the consistent disturbance of the swarm lead to a smaller steady state error compared to the simulated swarm. Through introducing more real vehicles into the swarm, the algorithm can be validated for use in an ocean going AUV swarm.

Finally, the simulation is used to evaluate swarms of increasing size subject

to communication constraints typical of AUV operation. When operating only using satellite communication, the swarm is unable to achieve a steady distribution of vehicles due to large uncertainty in an individual members knowledge of the swarm. However, introducing acoustic communication increases the communication frequency and the swarm is able to achieve a stable distribution and demonstrate scalability to large swarm sizes and consistent performance with more infrequent communication. The results presented indicate Voronoi coverage control to be an appropriate control mechanism for an AUV swarm when operating on acoustic communication and could provide an unrivalled monitoring solution for rapidly changing ocean phenomena.

Chapter 7

Conclusions

Through reviewing the relevant literature, it is found that a large amount of the research for multi-AUV control is focused on formation keeping, with very little research in the areas of coverage control, arguable more useful in ocean monitoring. While general coverage control literature considers communication constraints, they are not as severe and extensive as those found in AUV operation. Finally, no studies consider the 3D nature or dynamics of the ocean environment and none involve experimental validation of swarm control using AUVs.

An simulated AUV swarm is developed through modelling the ecoSUBμ, a swarm suitable AUV, and the communication channels between AUVs. The vehicle dynamics are modelled by a pitch rate, a yaw rate and a maximum speed. Navigation is performed via dead reckoning and is assumed perfect, justified by frequent surfacing to acquire a GPS position and acoustic localisation techniques both being feasible possibilities to reduce navigation error. The swarm is able to communicate via satellite and acoustic communication with the relevant delays. The model is used to simulate an AUV swarms of varying size and under increasingly severe communication constraints. Through simulation, the level of coordination between swarm members can be evaluated, allowing the suitability of swarming algorithms to constraints typical of AUV operation to be determined.

Data collected from a deployment in the Shetland Isles is compared against

AUV simulation to highlight differences unaccounted for in simulation and validate the model. Simulation could be brought closer to reality through the introduction of factors unaccounted for by simulation: drift error, inconsistent speed and buoyancy. However, the aim of the model is to evaluate the effects of underwater communication constraints on swarming algorithms and swarm coordination, hence introducing drift error, for example, may shroud the impact of communication delay and hinders the interpretation of results. The correlation of simulation and reality is deemed sufficient for the level of complexity and any further complexity would lead to lengthy computation times.

Examples of unrealistic communication assumptions including instantaneous, simultaneous and lossless communication in simulation of AUV swarms can be found in Amory et al. [2014] and Read et al. [2014]. Due to the extensive limitations of underwater communication, this wprk aims to subject swarm behaviours to more realistic communication constraints and evaluates the ability of the swarm to perform emergent flocking and coverage control algorithms.

Using an algorithm adapted from Reynold's flocking, a study on the suitability of emergent flocking as a control mechanism for an AUV swarm is presented in Lowndes et al. [2017], published during the course of this work. It is found that close proximity flocking is extremely dependent on communication, and performance degrades as the swarm, and hence communication overhead, grows. As the performance would only degrade further with the reduction of assumptions and the increase of swarm and domain size it is concluded close proximity flocking is unlikely to be an appropriate control mechanism for AUV swarm control. Work presented by McGookin and McColgan [2017] reinforces this conclusion across a wider range of acoustic communication constraints.

Due to new information on the capabilities of swarm suitable acoustic communciation modems, a new emergent flocking study is conducted with a refined algorithm allowing for the separation between AUVs to be controlled. With the new communication constraints, mainly increased communication range, swarms

of 5 to 20 vehicles are shown to coordinate and form a well organised flock when controlled by the emergent flocking algorithm and subject to increasing communication constraints. While performance does degrade with increasing swarm size and communication constraints, the range of swarm sizes and communication constraints are thought to be the limits of what is feasible in practical application, hence these results indicate emergent flocking is a scalable solution to coordination of an AUV swarm.

Implementing the hill climb algorithm to simulate environmental monitoring, increasing communication constraints has a negative impact on performance while increasing the swarm size is shown to increase success rate, consistency in performance and reduce the time to reach the maxima of the sensory function. Unfortunately, this improvement is bounded by the length of a communication cycle. When compared to the performance of a single vehicle, the swarm offers a consistent increase in relative performance to a single vehicle. While performance does not exclusively increase with swarm size, larger swarm sizes are shown to offer higher success rates in locating the maxima, still offering a benefit. The final part of this study demonstrated the single vehicle was unable to locate the maxima in a noisy environment while the swarm achieved success rates of 60% in most cases, peaking at 79%. These results demonstrate the sharing of sensor information across a swarm in uncertain, noisy environments such as the ocean can achieve tasks impossible by a single vehicle.

Once AUVs have flocked to a monitoring site, it is necessary for the spatial distribution of vehicles to correlate with the spatial scale of the dominant process being monitored, thus introducing coverage control. Using the algorithm presented by Lindhé et al. [2005], validation experiments have been presented in which a real ecoSUB AUV interacts with static nodes and multiple simulated AUVs. The experiments using static nodes validated the implementation of the coverage control algorithm. Through introducing a real vehicle, the negative impact of drift and GPS positioning error was shown to be small and the consistent

disturbance of the swarm lead to a smaller steady state error compared to the simulated swarm.

The simulation is used to evaluate swarms of increasing size subject to communication constraints typical of AUV operation. When operating only using satellite communication, the swarm is unable to achieve a steady distribution of vehicles due to large uncertainty in an individual members knowledge of the swarm. However, through introducing acoustic communication the swarm is able to achieve a stable distribution and demonstrate scalability to large swarm sizes and consistent performance with more infrequent communication. The results presented indicate Voronoi coverage control to be an appropriate control mechanism for an AUV swarm when operating on acoustic communication and could provide an unrivalled monitoring solution for rapidly changing ocean phenomena.

Thus to conclude, this work has evaluated the impact of communication constraints typical of AUV operation on emergent flocking and coverage control algorithms. Both algorithms can produce a well coordinated, well organised swarm given acoustic communication is available between swarm members. Due to the use of the TDMA protocol, both algorithms show signs of a limit to the scalability with increasing swarm size and/or communication constraints. However, performance remains within the levels of acceptability. The vehicle-in-loop experiments presented represent an initial step to validate a coverage control algorithm for an AUV swarm, an approach which has the potential to provide unrivalled monitoring of rapidly changing ocean features.

Chapter 8

Future Work

While the AUV swarm has been shown to be able to perform swarming algorithms to a sufficient level of coordination, there are several problems in assuming the swarm is constantly connected over acoustic communication. Firstly, acoustic communication is notoriously unreliable. The channel is subject to the directionality of the modem, beam bending as the signal travels through the water column, noise interference, reflections and shadowing to give some examples. These all contribute to the reliability of the channel and lead to packet loss. In addition, the presence of currents in the ocean which exceed the speed of an AUV is likely and hence may cause AUVs to expand beyond acoustic communication range. At present, the swarm shows extremely poor performance under satellite communication alone and hence to build a robust solution, the swarm should be able to operate in absence of acoustic communication.

There are several other benefits to being able to operate on satellite communication. Acoustic communication consumes extra power, a severely restricted resource in AUVs, and ultimately shortens the survey length which is undesirable when attempting to capture the evolution of an event which could span days to months. In addition, not all AUVs are fitted with acoustic modems as standard and even those which are can often only communicate with modems from the same manufacturer. Standardisation of underwater communication is a future

target and work such as the JANUS underwater communications standard is a step towards this Potter et al. [2014]. However adoption and cooperation by commercial modem manufacturers is the key to the solution. The assumption that all AUVs in the swarm are able to communicate acoustically severely limits the vehicles which can be included in the swarm, which is not conducive to forming a large swarm. Finally, acoustic communication also restricts the area covered by the swarm if the success of the algorithm depends on the vehicles being within acoustic communication range of their neighbours.

8.1 Develop Swarm Modelling

When operating on satellite communication alone, the timeout to reach a given waypoint determines the minimum frequency of communication. For long time outs, significant error could accumulate between the position of an AUV and its position as known by the server. In an effort to compensate for this error, it is proposed that the movement and behaviour of each AUV be modelled to provide an estimate of the AUVs current position in space. Using this modelling, when swarm members communicate with the server, the estimated locations of the swarm can be used to perform the coverage control algorithm.

To improve model accuracy, currents can be modelled based on external data including ocean circulation models, including the NEMO and FOAM models, for example. In addition, information from the swarm members, for example accumulated navigational error on surfacing, can be used to improve the model.

8.2 Development of Advection Model

As in Das et al. [2010] and Smith et al. [2010], advection models can be used to estimate the evolution of a feature over short time scales. This would allow the delay between data sets attained from satellite data to be bridged to allow for effective positioning of the swarm.

This has been done in the aforementioned work at varying degrees of success. Smith et al. [2010] employ the ROMS ocean model to advect a bloom and measures its performance against a bloom advected using HF radar measurements. Repeatedly the ROMS model advects the bloom in the opposite direction to the HF radar, leading to path planning issues. This is possibly due to the low current speeds and the absence of any dominant driving force on the blooms position.

This work would improve upon that of Smith et al. [2010] by including vehicle measurements in the modelling. Smith et al. [2010] uses only two gliders, hence there is little data to enhance the model. However, with 10 AUVs, significantly more data can be used to improve modelling. In addition, adpative coverage control approaches could be used to determine sampling locations based on the sensory distribution rather than geometry of the bloom extent

Through attaining satellite data from a historic HAB as an example, the advection model can be evaluated by advecting the data between data sets and hence comparing the advected model to the new incoming data to determine error.

8.3 Coverage Control in an AUV Swarm for Environmental Monitoring

Maintaining a model of the bloom on a server, using various sources of data, allows relatively simple coverage control to be performed as the environment is 'known'. While adaptive coverage control approaches which deal with unknown environments have been presented in literature, they are always reactive approaches and hence do not predict how the environment will evolve.

This approach uses the advection model to predict the evolution of the feature being monitored and a swarm model to predict the locations of AUVs between surfacing events to create a technology toolchain for tracking and monitoring ocean features. When an AUV surfaces to re-task, the maintained feature and

swarm models can be used to retask the vehicle accordingly.

The approach can be evaluated in a similar manner to the advection model, using historic HAB data. Ryan et al. [2005] produced a coherent data set using a array of sensing technologies therefore this data set could be used as a defining standard. If this can't be achieved, any HAB satellite data could be used. By applying the approach between satellite data sets, the final distribution of vehicles at the timestamp of the new remote sensing data could be used to evaluate the accuracy of the advected model and how well the swarm was able to monitor the feature and produce a coherent data set.

References

A. Amory, T. Tosik, and E. Maehle. A load balancing behavior for underwater robot swarms to increase mission time and fault tolerance. *Proceedings of the International Parallel and Distributed Processing Symposium, IPDPS*, pages 1306–1313, 2014.

A. B. Asghar and S. L. Smith. Robot monitoring for the detection and confirmation of stochastic events. In *53rd IEEE Conference on Decision and Control*, pages 408–413, Dec 2014.

M. Babin, J. J. Cullen, C. S. Roesler, P. L. Babin, G. J. Doucette, M. Kahru, M. R. Lewis, C. A. Scholin, M. E. Sieracki, and H. M. Sosik. New approaches and technologies for observing harmful algal blooms. *Oceanography*, 18, June 2005.

M. P. Brito, G. Griffiths, and P. Challenor. Risk analysis for autonomous underwater vehicle operations in extreme environments. *Risk Analysis*, 30(12): 1771–1788, 2010.

M. P. Brito, N. Bose, R. Lewis, P. Alexander, G. Griffiths, and J. Ferguson. The role of adaptive mission planning and control in persistent autonomous underwater vehicles presence. In *AUV 2012 - Southampton, UK*, 2012.

M. P. Brito, R. S. Lewis, N. Bose, and G. Griffiths. Adaptive autonomous underwater vehicles: An assessment of their effectiveness for oceanographic applications. *IEEE Transactions on Engineering Management*, 66(1):98–111, 2018.

R. Brooks and A. Flynn. Fast, cheap and out of control: A robot invasion of the solar system. *Journal of the British Interplanetary Society*, 42:478–485, 1989.

F. Bullo, R. Carli, and P. Frasca. Gossip coverage control for robotic networks: Dynamical systems on the space of partitions. *SIAM Journal on Control and Optimization*, 50:419 – 447, 2012.

R. Camilli, C. M. Reddy, D. R. Yoerger, B. A. S. V. Mooy, M. V. Jackuba, J. C. Kinsey, C. P. McIntyre, S. P. Sylva, and J. V. Maloney. Tracking hydrocarbon plume transport and biodegradation at deepwater horizon. *Science*, 330:201–204, 2010.

J. Carlos, J. C. Cutipa Luque, and D. Donha. Two dof robust controller for a six dof underwater autonomous vehicle. *ABCM Symposium Series in Mechatronics*, 12 2008.

J. Clark and R. Fierro. Mobile robotic sensors for perimeter detection and tracking. *ISA Transactions*, 46(1):3 – 13, 2007.

J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage Control for Mobile Sensing Networks. *IEEE Transactions on Robotics and Automation*, 20(2): 243–255, 2004.

I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1):1 – 11, 2002.

J. J. Cullen. Observation and Prediction of Harmful Algal Blooms, Real-time Coastal Observing Systems for Marine Ecosystem Dynamics and Harmful Algal Blooms: Theory, Instrumentation and Modelling. *Nature*, pages 1 – 41, 2007.

J. Das, K. Rajan, S. Frolov, F. Py, J. Ryan, D. A. Caron, and G. S. Sukhatme. Towards marine bloom trajectory prediction for auv mission planning. In *International Conference on Robotics and Automation - Anchorage, Alaska*, May 2010.

T. D. Dickey. Emerging ocean observations for interdisciplinary data assimilation systems. In *Journal of Marine Systems*, number 40 - 41, pages 5 – 48, 2003.

A. Dirafzoon, M. B. Menhaj, and A. Afshar. Voronoi Based Coverage Control for Nonholonomic Mobile Robots with Collision Avoidance. *IEEE International Conference on Control Applications*, pages 1755–1760, 2010.

A. Dirafzoon, S. Emrani, S. M. A. Salehizadeh, and M. B. Menhaj. Coverage control in unknown environments using neural networks. *Artificial Intelligence Review*, 38(3):237–255, 2011.

M. Dorigo and E. Sahin. Swarm Robotics - A Special Issue Editorial. *Autonomous Robots*, 17 (2-3):111 – 113, 2004.

M. Fox, D. Long, and D. Magazzeni. Plan-based policy-learning for autonomous feature tracking. In *International Conference on Automated Planning and Scheduling - Sao Paulo*, June 2012.

S. Garg and N. Ayanian. Persistent monitoring of stochastic spatio-temporal phenomena with a small team of robots. In *Robotics: Science and Systems*, 2014.

M. A. Godin, Y. Zhang, J. P. Ryan, T. T. Hoover, and J. G. Bellingham. Phytoplankton bloom patch center localization by the tethys autonomous underwater vehicle. In *OCEANS'11 MTS/IEEE KONA*, pages 1–6, Sept 2011.

G. Griffiths. Gliders and autonomous Underwater Vehicle Observing Systems. *Nature, Real-time Coastal Observing Systems for Marine Ecosystem Dynamics and Harmful Algal Blooms: Theory, Instrumentation and Modelling*, pages 495 – 525, 2007.

G. Hollinger and S. Singh. Multi-robot coordination with periodic connectivity. In *IEEE International Conference on Robotics and Automation 2010 - Alaska*, May 2010.

A. Howard, M. J. Matarić, and G. S. Sukhatme. Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem. In *Distributed Autonomous Robotic Systems 5*, pages 299–308, Tokyo, 2002.

C. Isokeit, B. Meyer, and E. Maehle. Cooperative swarm behaviour for in situ underwater environmental measurements. In *OCEANS 2017 - Aberdeen*, pages 1–6, June 2017.

M. V. Jakuba and D. R. Yoerger. Autonomous Search for Hydrothermal Vent Fields with Occupancy Grid Maps. *Proceedings of the Australasian Conference on Robotics and Automation*, 2008.

A. Jernelöv and O. Lindén. Ixtoc i: A case study of the world's largest oil spill. *Ambio*, 10(6):299–306, 1981.

S. Kemna, D. A. Caron, and G. S. Sukhatme. Constraint-induced formation switching for adaptive environmental sampling. In *OCEANS 2015 - Genova*, pages 1–7, May 2015.

S. Kemna, J. G. Rogers III, C. Nieto-Granda, S. Young, and G. S. Sukhatme. Multi-Robot Coordination through Dynamic Voronoi Partitioning for Information Adaptive Sampling in Communication Constrained Environments. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2124 – 2130, 2017.

A. L. Kukulya, J. G. Bellingham, J. W. Kaeli, C. M. Reddy, M. A. Godin, and R. N. Conmy. Development of a propeller driven long range autonomous underwater vehicle (lrauv) for under-ice mapping of oil spills and environmental hazards: An arctic domain center of awareness project (adac). In *2016 IEEE/OES Autonomous Underwater Vehicles (AUV)*, pages 95–100, Nov 2016.

A. Kwok and S. Martìnez. Unicycle coverage control via hybrid modeling. *IEEE Transactions on Automatic Control*, 55(2):528–532, February 2010.

X. Lan and M. Schwager. Rapidly-exploring random cycles: Application to estimate spatio-temporal fields. *IEEE Transactions on Robotics*, 32.5:1230 – 1244, 2016.

Y. G. Lee, X. Garza Gomez, and R. M. Lee. Ultimate costs of the disaster: Seven years after the deepwater horizon oil spill. *Journal of Corporate Accounting and Finance*, 29(1):69–79, 1 2018.

J. J. Leonard and A. Bahr. Autonomous underwater vehicle navigation. *Springer Handbook of Ocean Engineering*, pages 341–358, 2016.

M. Lindhé, F. Ögren, and K. H. Johansson. Flocking with obstacle avoidance: A new distributed coordination algorithm based on Voronoi partitions. *Proceedings - IEEE International Conference on Robotics and Automation*, pages 1785–1790, 2005.

J. Liu, Z. Wang, Z. Peng, J. Cui, and L. Fiondella. Suave: Swarm underwater autonomous vehicle localization. In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*, pages 64–72, April 2014.

F. Lobo Pereira, J. Borges de Sousa, R. Gomes, and P. Calado. A model predictive control approach to auvs motion coordination. *Coordination Control of Distributed Systems*, pages 9–18, 2015.

T. S. Lowndes, A. B. Phillips, C. A. Harris, E. Rogers, B. Chu, and E. Popova. Evaluating the Capabilities of a Flight-Style Swarm AUV to Perform Emergent and Adaptive Behaviours. In *Towards Autonomous Robotic Systems. TAROS 2017. Springer Lecture Notes in Computer Science*, Surrey, UK, 2017.

J. M. Luna, R. Fierro, C. Abdallah, and J. Wood. An adaptive coverage control algorithm for deployment of nonholonomic mobile sensors. In *49th IEEE Conference on Decision and Control (CDC)*, pages 1250–1256, Dec 2010.

K. M. Lynch, I. B. Schwartz, P. Yang, and R. A. Freeman. Decentralized Environmental Modeling by Mobile Sensor Networks. *IEEE Transactions on Robotics*, 24:710–723, 2008.

E. W. McGookin and J. McColgan. Effect of communication delays on the successful coordination of a group of biomimetic auvs. In *OCEANS 2017 - Aberdeen*, pages 1–10, June 2017.

S. Miah, B. Nguyen, A. Bourque, and D. Spinello. Nonuniform Coverage Control With Stochastic Intermittent Communication. In *IEEE Transactions on Automatic Control*, volume 60, pages 1981–1986, 2015.

P. Millán, L. Orihuela, I. Jurado, and F. R. Rubio. Formation control of autonomous underwater vehicles subject to communication delays. *IEEE Transactions on Control Systems Technology*, 22(2):770–777, March 2014.

N. Morozs, P. D. Mitchell, Y. Zakharov, R. Mourya, Y. R. Petillot, T. Gibney, M. Dragone, B. Sherlock, J. A. Neasham, C. C. Tsimenidis, M. E. Sayed, A. C. McConnell, S. Aracri, and A. A. Stokes. Robust tda-mac for practical underwater sensor network deployment: Lessons from usmart sea trials. In *Proceedings of the Thirteenth ACM International Conference on Underwater Networks & Systems*, WUWNet '18, pages 11:1–11:8, New York, NY, USA, 2018. ACM.

National Oceanography Centre. Massmo: Marine autonomous systems in support of marine operations. projects.noc.ac.uk/massmo/, 2018. Accessed: 27/04/2018.

National Research Council. *Oil Spill Dispersants: Efficacy and Effects*. The National Academies Press, Washington, DC, 2005.

C. Nieto-Granda, I. John G. Rogers, and H. I. Christensen. Coordination strategies for multi-robot exploration and mapping. *The International Journal of Robotics Research*, 33(4):519–533, 2014.

NOAA. Twenty-Five Years After the Exxon Valdez Oil Spill : NOAA's Scientific Support, Monitoring and Research. Technical report, National Oceanic and Atmospheric Administration, 2014.

R. Patel, P. Frasca, J. W. Durham, R. Carli, and F. Bullo. Dynamic Partitioning and Coverage Control with Asynchronous One-to-Base-Station Communication. *IEEE Transactions on Control of Network Systems*, 3(1):24–33, 2016.

L. Paull, S. Saeedi, M. Seto, and H. Li. Auv navigation and localization: A review. *IEEE Journal of Oceanic Engineering*, 39(1):131–149, Jan 2014.

M. Pebody. The contribution of scripted command sequences and low level control behaviours to autonomous underwater vehicle control systems and their impact on reliability and mission success. In *OCEANS 2007 - Europe*, pages 1–5, June 2007.

S. Petillo, H. Schmidt, and A. Balasuriya. Constructing a distributed auv network for underwater plume-tracking operations. *International Journal of Distributed Sensor Networks*, 8(1), 2012.

A. B. Phillips, N. Gold, N. Linton, C. A. Harris, E. Richards, R. Templeton, S. Thune, J. Sitbon, M. Muller, I. Vincent, and T. Sloane. Agile design of low-cost autonomous underwater vehicles. In *OCEANS 2017 - Aberdeen*, pages 1–7, June 2017.

S. Poduri and G. S. Sukhatme. Constrained coverage for mobile sensor networks. In *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, volume 1, pages 165–171 Vol.1, April 2004.

J. Potter, J. Alves, D. Green, G. Zappa, I. Nissen, and K. McCoy. The janus underwater communications standard. In *2014 Underwater Communications and Networking (UComms)*, pages 1–4, Sept 2014.

M. Read, C. Moslinger, T. Dipper, D. Kengyel, J. Hilder, R. Thenius, A. Tyrrell, J. Timmis, and T. Schmickl. Profiling Underwater Swarm Robotic Shoaling

Performance Using Simulation. In *Towards Autonomous Robotic Systems*, volume 8069, pages 103–113, 2014.

C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Comput. Graph.*, 21(4):25–34, Aug. 1987.

I. Robbins, G. Kirkpatrick, S. Blackwell, J. Hillier, C. Knight, and M. Moline. Improved monitoring of habs using autonomous underwater vehicles (auv). *Harmful Algae*, 5(6):749 – 761, 2006.

J. Ryan, H. Dierssen, R. Kudela, C. Scholin, K. Johnson, J. Sullivan, A. Fischer, E. Rienecker, P. McEnaney, and F. Chavez. Coastal Ocean Physics and Red Tides. *Oceanology*, 18:246–255, 2005.

J. P. Ryan, Y. Zhang, H. Thomas, E. V. Rienecker, R. K. Nelson, and S. R. Cummings. A high resolution survey of a deep hydrocarbon plume in the gulf of mexico during the 2010 macondo blowout. *Monitoring and Modeling the Deepwater Horizon Oil Spill: A Record Breaking Enterprise*, pages 63–75, 2013.

Z. Saigol. *Automated Planning for Hydrothermal Vent Prospecting Using AUVs*. PhD thesis, 2011.

G. Salavasidis, A. Munafò, C. A. Harris, T. Prampart, R. Templeton, M. Smart, D. T. Roper, M. Pebody, S. D. McPhail, E. Rogers, and A. B. Phillips. Terrain-aided navigation for long-endurance and deep-rated autonomous underwater vehicles. *Journal of Field Robotics*, 36(2):447–474, 2019.

I. Sanseverino, D. Conduto, L. Pozzoli, S. Dobricic, and T. Lettieri. Algal bloom and its economic impact. *JRC Technical Reports, EU Science Hub, European Commission*, Jul 2016.

M. Schwager and D. Rus. Unifying Geometric, Probabilistic, and Potential Field Approaches to Multi-Robot Deployment. *Robotics Research*, pages 21–38, 2010.

M. Schwager, J.-J. Slotine, and D. Rus. Decentralized, Adaptive Control for Coverage with Networked Robots. In *IEEE International Conference on Robotics and Automation, Rome, Italy*, pages 3289–3294, 2007.

M. Schwager, F. Bullo, D. Skelly, and D. Rus. A ladybug exploration strategy for distributed adaptive coverage control. In *2008 IEEE International Conference on Robotics and Automation*, pages 2346–2353, May 2008.

M. Schwager, M. P. Vitus, S. Powers, D. Rus, and C. J. Tomlin. Robust Adaptive Coverage Control for Robotic Sensor Networks. *IEEE Transactions on Control of Network Systems*, 4:462 – 476, 2017.

R. N. Smith, J. Das, Y. Chao, D. A. Caron, B. H. Jones, and G. S. Sukhatme. Cooperative multi-auv tracking of phytoplankton blooms based on ocean model predictions. In *OCEANS 2010 IEEE - Sydney*, pages 1–10, May 2010.

R. N. Smith, M. Schwager, S. L. Smith, B. H. Jones, D. Rus, and G. S. Sukhatme. Persistent ocean monitoring with underwater gliders: adapting sampling resolution. *Journal of Field Robotics*, 28(5):714–741, 2011a.

S. L. Smith, M. Schwager, and D. Rus. Persistent monitoring of changing environments using a robot with limited range sensing. In *2011 IEEE International Conference on Robotics and Automation*, pages 5448–5455, May 2011b.

S. L. Smith, M. Schwager, and D. Rus. Persistent robotic tasks: Monitoring and sweeping in changing environments. *IEEE Transactions on Robotics*, 28(2): 410–426, April 2012.

S. Susca, F. Bullo, and S. Martinez. Monitoring environmental boundaries with a robotic sensor network. *IEEE Transactions on Control Systems Technology*, 16(2):288–296, March 2008.

Y. T. Tan, M. Chitre, and F. S. Hover. Cooperative bathymetry-based localization using low-cost autonomous underwater vehicles. *Autonomous Robots*, 40 (7):1187–1205, 2016.

O. A. Viquez, E. M. Fischell, N. R. Rypkema, and H. Schmidt. Design of a general autonomy payload for low-cost auv r&d. In *2016 IEEE/OES Autonomous Underwater Vehicles (AUV)*, pages 151–155, Nov 2016.

J. K. Wagner, M. H. McEntee, L. L. Brothers, C. R. German, C. L. Kaiser, D. R. Yoerger, and C. L. Van Dover. Dynamic Partitioning and Coverage Control with Asynchronous One-to-Base-Station Communication. *Deep-Sea Research Part II: Topical Studies in Oceanography*, 92:183–188, 2013.

J. M. Walls and R. M. Eustice. An origin state method for communication constrained cooperative localization with robustness to packet loss. *International Journal of Robotics Research*, 33(9):1191–1208, 2014.

S. E. Webster, J. M. Walls, L. L. Whitcomb, and R. M. Eustice. Decentralized extended information filter for single-beacon cooperative acoustic navigation: Theory and experiments. *IEEE Transactions on Robotics*, 29(4):957–974, Aug 2013.

J. S. Willcox, J. G. Bellingham, Yanwu Zhang, and A. B. Baggeroer. Performance metrics for oceanographic surveys with autonomous underwater vehicles. *IEEE Journal of Oceanic Engineering*, 26(4):711–725, Oct 2001.

M. H. Yamchi and R. M. Esfanjani. Distributed predictive formation control of networked mobile robots subject to communication delay. *Robotics and Autonomous Systems*, 91:194 – 207, 2017.

D. Yoerger, A. M. Bradley, M. Jakuba, M. Tivey, C. German, T. Shank, and R. W. Embley. Mid-ocean ridge exploration with an autonomous underwater vehicle. *Oceanography*, 20, Dec 2007.

Y. Zhang, R. S. McEwen, J. P. Ryan, and J. G. Bellingham. Design and tests of an adaptive triggering method for capturing peak samples in a thin phytoplankton layer by an autonomous underwater vehicle. *IEEE Journal of Oceanic Engineering*, 35(4):785–796, Oct 2010.