

Exploring Maneuvering Strategies for Heterogeneous Cooperative Navigation in Underwater Environments

by
Megan C. Flynn

S.B., Massachusetts Institute of Technology (2021)
Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of
Master of Science in Mechanical Engineering
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2023

©2023 Megan C. Flynn. All rights reserved.

The author hereby grants to MIT a nonexclusive, worldwide,
irrevocable, royalty-free license to exercise any and all rights under
copyright, including to reproduce, preserve, distribute and publicly
display copies of the thesis, or release the thesis under an open-access
license.

Author
Department of Mechanical Engineering
May 19, 2023

Certified by
Douglas Hart
Professor of Mechanical Engineering
Thesis Supervisor

Certified by
John J. Leonard
Samuel C. Collins Professor of Mechanical and Ocean Engineering
Thesis Supervisor

Accepted by
Nicolas G. Hadjiconstantinou
Chair, Graduate Program Committee

Exploring Maneuvering Strategies for Heterogeneous Cooperative Navigation in Underwater Environments

by

Megan C. Flynn

Submitted to the Department of Mechanical Engineering
on May 19, 2023, in partial fulfillment of the
requirements for the degree of
Master of Science in Mechanical Engineering

Abstract

Due to the challenges of the underwater environment and limited communication methods, undersea navigation is difficult. Autonomous underwater vehicles (AUVs) experience unbounded localization errors when operating below the surface. Range measurements between vehicles can be utilized to improve localization estimates. We define a two-agent team composed of a leader and a follower, in which the former has better navigational capabilities than the latter. The follower attempts to navigate to a destination while the leader aids in the follower's localization by providing range measurements from varied locations. Planning the relative motion between agents is vital to ensuring that meaningful range measurements are provided to support an effective estimation of the follower's pose.

This work explores five different maneuvering strategies based on geometric and observability principles. After designing the strategies, we tested their impact on the localization quality of the team through extensive simulation results. To investigate the resilience of the strategies to environmental conditions, we altered the simulated ocean currents. For additional study we allowed the leader to operate at a higher speed to explore the relationship between energy use and estimation performance.

Ultimately, the best maneuvering strategy was found to be the circling strategy due to its superior performance; however, the circling strategy used the most energy, especially with larger radii. Mission priorities may affect the selection of a maneuvering strategy; the zigzag and covariance squish strategies are still viable options as they do not suffer great performance loss when compared to the circling strategy.

Thesis Supervisor: Douglas Hart
Title: Professor of Mechanical Engineering

Thesis Supervisor: John J. Leonard
Title: Samuel C. Collins Professor of Mechanical and Ocean Engineering

Acknowledgments

I have so many people to thank for their help in making this thesis possible.

To my co-advisors, Professors Douglas Hart and John Leonard, thank you for your guidance and for helping me to explore this research area. Our discussions of the expansions of this work helped to bring context and purpose to my efforts. I am grateful to have had two stellar co-advisors.

Thank you to Brendan O’Neill for your help to bridge the gaps in my knowledge, debug several issues, and relate this work back to reality. It has been a joy and a privilege to work with you.

Thank you to Michael Benjamin for answering my questions about MOOS-IvP and for generously loaning me a laptop for several months during my initial research efforts. I’d like to acknowledge the work of Shir Goldfinger, who spent a summer developing the MOOS-IvP applications for advanced current fields that are mentioned in this thesis. Thank you to my lab mates across the Hart and Leonard labs for your feedback and listening ears.

The biggest thank you to my family, especially my parents, for always cheering me on and supporting my educational journey. To my partner, Josh, thank you for reminding me of my strength and for your endless compassion and understanding.

Contents

1	Introduction	15
1.1	Motivation	15
1.2	Problem Statement	19
1.3	Thesis Overview	19
2	Background	21
2.1	Acoustic Communication	21
2.1.1	WHOI Micromodem-2	24
2.2	Range-Only Single Beacon Navigation	25
2.3	Maneuvering Strategies	27
2.4	Nonlinear Least Squares Estimation and Factor Graphs	29
2.5	Mission Oriented Operating Suite with Interval Programming (MOOS-IvP)	38
2.6	Energy Usage for Underwater Vehicles	41
2.7	Problem Formulation	43
3	Methods	45
3.1	MOOS-IvP Setup	45
3.1.1	Leader Vehicle	48
3.1.2	Follower Vehicle	48
3.2	Ranging Protocol	49
3.3	State Estimation with iSAM2	50
3.4	Simulating Currents	52

3.5	Maneuvering Strategy Development	52
3.5.1	Zigzag	52
3.5.2	Circling	56
3.5.3	Diverse Relative Bearings	58
3.5.4	Covariance Squish	60
3.5.5	Optimal Point	61
3.6	Summary	63
4	Data Analysis	65
4.1	Summary of Trials	65
4.2	Circling Standard	66
4.3	Zigzag Standard	67
4.4	Advanced Strategies Standard	72
4.5	Comparing Standard Results	77
4.5.1	Energy Usage	77
4.5.2	Performance	77
4.6	Increased Current Magnitude	80
4.7	Varied Current Direction	82
4.8	Increased Leader Speed	89
4.9	Other Strategy Considerations	95
5	Conclusions	97
5.1	Contributions	97
5.2	Future Work	98

List of Figures

2-1	A simple example of a state estimation problem.	31
2-2	A simple state estimation problem with measurements.	32
2-3	A factor graph representation for the simple state estimation problem	33
2-4	A visualization of Powell's dogleg algorithm	37
2-5	Each vehicle is called a community in MOOS-IvP and operates around its central database of variable-value pairs, the MOOSDB.	40
2-6	A representation of MOOS' shoreside-vehicle topology for communica- tion and information flow.	41
3-1	MOOS-IvP coordinate system for XY plane and angles.	46
3-2	Visualization of how maneuvering strategies are integrated in the leader community.	47
3-3	MOOS-IvP topology for our mission setup highlighting where new ap- plications and behaviors are hosted.	47
3-4	Geometric properties of the zigzag and circling strategies.	53
3-5	Sample trajectories using the zigzag strategy with different radii and trail distances.	54
3-6	Relative leader trajectories for example zigzag patterns.	54
3-7	Example leader and follower trajectories for circling strategy.	57
3-8	Relative leader trajectories for example circling patterns.	57
4-1	Follower estimation performance as a function of circle radius for the standard mission configuration.	67

4-2	Follower estimation performance as a function of cubed leader speed for the standard mission configuration.	68
4-3	Follower estimation performance as a function of zigzag radius for the standard mission configuration.	69
4-4	Performance metrics as a function of cubed leader speed for the standard mission configuration.	69
4-5	Cubed leader speed is shown for all combinations of radius and trail distance.	70
4-6	Performance metrics for the zigzag strategy with standard mission configuration, broken out by radius and trail distance.	70
4-7	A visualization of the maximum designed relative angle of zigzags for all radii and trail distances combinations.	71
4-8	All follower error metrics as a function of maximum designed relative angle for the zigzag strategy at the standard mission configuration.	71
4-9	Example leader and follower trajectories for the covariance squish strategy at standard conditions.	73
4-10	Example leader and follower trajectories for the diverse angles strategy with the standard mission configuration.	74
4-11	Example leader and follower trajectories for the optimal point strategy with the standard mission configuration.	75
4-12	Comparison of all performance metrics for the three advanced strategies at the standard mission configuration.	75
4-13	Relative leader positions at range measurements for all strategies with the standard mission configuration.	76
4-14	Energy usage across strategies for the standard mission configuration, broken down by radius when appropriate.	77
4-15	Performance metrics across strategies based on radius when applicable.	78
4-16	Performance across strategies based on energy usage for standard mission configuration.	79

4-17 Comparison of mean values of performance metrics for all strategies with the standard mission configuration.	79
4-18 Leader and follower trajectories for a 30 meter circling strategy with increasing current magnitudes.	81
4-19 Differences in performance metrics for all circling radii as a result of increased current magnitude.	82
4-20 Performance of zigzag strategy as a function of maximum designed relative angle for varying magnitudes of applied current.	83
4-21 Performance of all strategies with increasing current magnitude.	83
4-22 Energy usage for all strategies with different angles of current applied.	84
4-23 Mean values of performance metrics for all strategies with different angles of currents applied.	85
4-24 Performance metrics for circling as a function of radius with different angles of current applied.	86
4-25 Performance metrics for zigzags as a function of radius with different angles of current applied.	86
4-26 Zigzag performance by maximum designed relative angle for 240 degree current applied.	87
4-27 Zigzag performance by radius and zigzag trail distance for applied current at 330 degrees.	88
4-28 Distribution of leader ranging locations with increasing speed. Example trajectories for 20 and 70 meter radius circles.	90
4-29 Differences in performance metrics for circling strategy with increased speed for all radii.	91
4-30 Evolution of leader and follower trajectories for increased leader speed with advanced maneuvering strategies.	92
4-31 Energy usage metric for all strategies for different desired leader speeds of 1.5, 2.0, and 2.5 m/s	93
4-32 All strategies' performance metrics as a function of cubed leader speed.	94

4-33 All performance metrics for the five strategies with increased leader speed. 94

List of Tables

4.1	Parameters for standard mission configuration.	66
4.2	Summary of the different mission configurations that were run and will be compared throughout this chapter.	66
4.3	Comparing mean values for all performance metrics across all five strategies for the standard mission configuration.	80

Chapter 1

Introduction

This chapter reviews the motivation behind this thesis and outlines the subsequent chapters.

1.1 Motivation

An ever-present challenge for mobile robotics involves the act of localizing a robot in space, or simply, estimating its position and orientation in the world. The development of sensing equipment, predictive models, and new cooperative strategies all provide numerous avenues for tackling localization. Various facets of mobile robotics research face different issues due to the wide array of operating environments, mission design, and robotic team makeup. Marine robotics specifically faces environmental challenges due to the adverse conditions for communication and sensing underwater. Extensive research efforts have explored the concept of cooperative navigation, in which a team of agents works together to overcome these problems.

Challenges in Underwater Environments

The main challenges in underwater environments involve the communication methods available, the turbulent and unpredictable environment, and the lack of visibility.

Due to the rapid attenuation of electromagnetic waves in water, radio communications, satellite-based methods, and other communication schemes that rely on

electromagnetic waves are not feasible for underwater environments. The main alternative that is utilized is acoustic communication. Information can be encoded in acoustic signals and may require communication nodes to be outfitted with appropriate acoustic modems or receivers. Reflection at the seafloor and the water surface can cause multi-path issues for communication, and dispersion of the signals can result in dropped messages [25].

When completing a transit, an autonomous underwater vehicle (AUV) must face an ever-changing environment of aquatic currents, which impact its ability to estimate its pose or state. While there are direct methods for measuring water column velocities such as using an acoustic Doppler current profiler (ADCP) or acoustic Doppler velocimeter (ADVs), these hardware solutions come with high cost and require more complex integration [42]. When a vehicle is under the water's surface, it does not have access to the Global Positioning System (GPS) or a global navigation satellite system (GNSS) and will experience unbounded growth of the error of its estimated position; the rate of error accumulation will greatly depend on the accuracy of its sensor suite, but the presence of strong, turbulent, and varying water currents further increases the rate of error accumulation [48]. Determining the impact of currents and improving the state estimation of underwater vehicles without the addition of hardware is an established area of research [18],[27],[42].

Land-based cooperative navigation often relies on optical sensor fusion to create a map of the environment for improved localization and navigation. The field of simultaneous localization and mapping (SLAM) research has experienced substantial advances for both land-based and underwater applications [13],[26],[45]. However, the use of optical sensors is sometimes not feasible in underwater environments due to the presence of debris, bubbles, wildlife, vegetation, and lack of light. Sonar is better suited for imaging in the underwater environment but is still impacted by the challenges of the medium. Thus, SLAM and its techniques have a place in marine robotics, but the challenges of underwater environments limit the direct use of land-based cooperative navigation methods. Often instead of imaging measurements, relative range and bearing measurements between nodes are the main tool used to

enhance the abilities of navigational sensors [36].

Heterogeneous Teaming

For underwater missions that require surveying, environmental sampling, or have multiple objectives, it can be beneficial to use several vehicles that are equipped with a different suite of sensors. In previous works [1],[10],[14] the separate responsibilities of surveying and maintaining a good position estimate were delegated to two sets of vehicles. The vehicle focused on the latter was dubbed the Communication and Navigation Aid (CNA) and often operated at the surface to allow for a GPS connection. The CNA could then communicate its GPS-derived position to the underwater survey vehicle to allow for longer missions with less surfacing overall. Additionally, this concept can be expanded to multiple survey vehicles with one or a handful of CNAs, allowing for improved scalability and cost reduction for larger teams [40].

In its most basic form, a team of vehicles can be heterogeneous because of their different sensor suites. This can be beneficial in reducing the total team hardware cost or may be necessary based on available resources. One set of vehicles can be outfitted with sonar to map out points of interest within a survey area, while another set of vehicles can use more advanced sensors to investigate the marked areas of interest instead of traversing through the entire field [43].

A specific use case of heterogeneous vehicle teams is mine countermeasures (MCM), where potentially dangerous targets need to be located and neutralized. A team of AUVs can be used to detect, localize, and classify points of interest in a survey area for later identification and neutralization. In this type of environment, it is preferable to have several more expendable vehicles exploring the potentially dangerous areas and a single well-equipped CNA to aid in localization and communication protocol. Even if a human operator is required to make the final classification, they do not have to be involved in the collection or observation of the targets because the team of vehicles is well equipped to accurately report the locations of interest [37],[43].

An interesting expansion of heterogeneous teams is the involvement of humans, specifically divers, with autonomous systems in underwater missions. Human-autonomy

teaming (HAT) comes with unique challenges and potential. Human divers can perform more detailed and tactile tasks when compared with robotic systems. For example, divers can make decisions based on real-time conditions, past experiences, and protocol with more nuance than a vehicle. Dive teams can be used to search for objects underwater and normally proceed in teams, utilizing a compass and kick counts to keep track of progress; these methods for human underwater navigation can be coarse in nature and thus prone to high error rates. Dive tablets have been used in previous HAT systems to provide mission feedback, aid in communication, and allow for acoustic messaging [29],[31].

The main advantage of using an AUV with a diver is the sensor suite and computation that becomes available. A vehicle is able to perform more advanced state estimation than a diver and can carry the load of additional sensors. These sensors may include a Doppler velocity log (DVL), acoustic Doppler current profiler (ADCP), and an inertial navigation system (INS) that will all improve the state estimation capabilities of the vehicle. With the addition of acoustic communication between the entities, an improvement in the vehicle’s state estimation can enhance the estimate of the diver’s position as well. Integrating an autonomous vehicle into a team of divers will result in new safety considerations. Team members need to communicate consistently and maintain a safe separation to avoid collisions or other interactions.

Possible missions that HAT can thrive in include transit and search missions or area-wide searches. In the former case, there is an established position of interest that the HAT team must travel to and recover or manipulate an object of interest, for example, lost equipment. The second case is an expansion on the first, where there is a large search area of interest and possibly many locations with objects that require inspection. The vehicle can aid in diver localization while also scanning for points of interest. The diver has unique manipulation abilities and can focus on smaller areas identified by the AUV from an initial scan or from inspecting during a maneuver [38].

Previous work from O’Neill and Pelletier [37],[38] sparked interest in the potential for exploring HAT applications of cooperative navigation, and serve as the main inspiration for the development of a teaming architecture for this thesis.

1.2 Problem Statement

This thesis will explore how maneuvering strategies can be used to improve the quality of an underwater team’s cooperative navigation efforts. Our specific interest is in the localization of two marine vehicles where one vehicle has significantly lower quality navigational sensors than the other. Localization is carried out by combining range measurements between the vehicles with dead reckoning information. Previous works have highlighted zigzag and circling geometric patterns as favorable from general observability principles. There remains a gap in the literature for a data-driven comparison between these geometric strategies and other advanced computational maneuvering strategies. This thesis will focus on how the selection of a maneuvering strategy impacts the localization performance and energy usage for our two vehicle team. The insights gained from the findings in this thesis can be applied to future implementations of cooperative navigation with heterogeneous teams.

1.3 Thesis Overview

Chapter 2 will review key concepts that shape the theoretical basis of our work. This chapter will frame previous methods of range-based navigation efforts and maneuvering strategies. We will outline the key mathematical concepts of nonlinear least squares estimation and factor graphs for a better understanding of the state estimation problem. Then, we will briefly review the mission-oriented operating suite with interval programming (MOOS-IvP) as a platform for simulation and autonomy. Finally, we will discuss energy usage for AUVs.

Chapter 3 will review the software developed for the execution of this thesis in the MOOS-IvP simulation environment. We will describe the new MOOS-IvP modules created for this work and will discuss the design of the five maneuvering strategies for this thesis.

Chapter 4 will describe the trials and subsequent results of this work. We will analyze the performance of our maneuvering strategies using error and energy met-

rics. Results for several mission configurations will be compared to comment on the robustness of each strategy.

Chapter 5 will discuss the conclusions drawn from the results of previous chapters and the contributions of this thesis. This chapter then will outline suggestions for future work that can be extended from the efforts of this thesis.

Chapter 2

Background

This chapter will discuss several areas of previous research including acoustic communications, cooperative navigation, and state estimation. The works noted in this chapter serve as the theoretical foundation for the problem formulation of this thesis in Section 2.7.

2.1 Acoustic Communication

When communicating through an aquatic medium, electromagnetic waves that are normally used in radio channels will attenuate rapidly, rendering them an unfit method of underwater communication. Instead, acoustic waves can be used for communication through encoding and decoding messages with transducers and receivers. Acoustic methods of communication bring specific challenges but ultimately can be overcome with the help of new hardware developments. A byproduct of communicating with acoustic waves is the ability to measure distance through the signal's time of flight [36].

Even when using acoustic communication there is still a degree of attenuation, or weakening, of the signal in water. This can be described by an attenuation coefficient, which represents the reduction in decibels of the signal per kilometer traveled through

the medium. For acoustic waves in the ocean,

$$\alpha' \simeq 3.3 \times 10^{-3} + \frac{0.11f^2}{1+f^2} + \frac{44f^2}{4100+f^2} + 3.0 \times 10^{-4}f^2 \quad (2.1)$$

where α' is the attenuation coefficient in decibels per kilometer and f is the frequency of the acoustic signal in kilohertz. The actual value of the attenuation coefficient is subject to changes in temperature, salinity, and pH of the ocean, but Equation 2.1 is a good estimation for determining appropriate operating frequencies for ocean acoustic signals [20]. Due to the relationship between frequency and attenuation in Equation 2.1, normal operating frequencies for AUV communications remain below 30 kilohertz to keep the attenuation manageable for distances up to several hundreds of meters. A competing interest is that lower frequencies are not capable of representing as much data compared to higher frequencies. AUVs normally utilize acoustic modems that operate between 10 and 28 kilohertz, which balances these two priorities of underwater signal performance [37].

Other issues that are observed in acoustic communication include frequency spreading, Doppler shift, and multipath propagation. When the acoustic source and receiver have relative motion, there is a spreading and a shift of the signal's frequency that complicates the ability of the receiver to understand the original signal. The shift of the signal's frequency is known as the Doppler shift and depends on the relative speed of the two agents; a higher relative speed between the two will worsen the effects of the Doppler shift. Multipath propagation is a phenomenon observed when an acoustic signal travels from source to receiver along multiple paths. This occurs most commonly when the signal reflects or refracts off the water surface or seafloor. Multipath propagation can be especially prevalent in shallow water conditions with reflective seafloor. If there is a significant change in sound speed through the water, the acoustic waves can also bend by Snell's law, creating another possible path of transmission. All of these potential issues require robust acoustic receivers that can overcome any degradation or disruption of the signal; thoughtful encoding of messages can aid in mitigating these challenges of the aquatic medium [25].

In acoustic communication schemes, information is typically grouped into packets. Individual symbols are represented by several bits encoded through signal frequency or phase. Qurasishi et al. [41] utilized Orthogonal Frequency-Division Multiplexing (ODFM) to encode each symbol with up to four carrier frequencies. The appearance of each frequency in a pulse served as a bit equal to 1, allowing for sixteen unique symbols per pulse. To decode the signal, a Fast Fourier Transform (FFT) was carried out to extract the frequencies present in each pulse signal. In anticipation of possible communication errors, a repetition procedure was added such that each data bit was transmitted twice, with two different sets of carrier frequencies. This communication scheme was deployed in field trials with two autonomous surface vehicles (ASVs), such that one vehicle acted as a beacon for the other. Outlier rejection and filtering of erroneous positions were carried out, which resulted in good performance of the ranging messages even in a challenging environment of a shallow lake. Range measurements were fused with an extended Kalman filter (EKF) to improve the state estimation of the ASV.

Beyond the communication capabilities that acoustic methods have to offer, the time of flight of messages can be used to measure distances. With an accurate speed of sound estimate, the one way travel time (OWTT) or two way travel time (TWTT) can be used to determine the distance between the source and receiver. In order to have an accurate sound speed estimate, an AUV may be outfitted with a conductivity, temperature, and depth (CTD) sensor to properly assess the physical properties of the surrounding water [36]. For OWTT calculations, the two acoustic entities must have accurately synchronized clocks, as the OWTT is calculated as the absolute time difference. For TWTT, the internal clocks need not be synchronized. The first vehicle sends a ping, which the second vehicle receives and then replies back to the first vehicle. Typically for TWTT systems, the turnaround time for the second vehicle to reply to the initial ping is standardized. This way, the range distance can be calculated as

$$d_{range} = \frac{T_{total} - T_{turnaround}}{2} \times v_{sound} \quad (2.2)$$

where T_{total} is the total travel time, $T_{turnaround}$ is a known turnaround period, and v_{sound} is the speed of sound in water given the output of the CTD sensor or the best estimate of the conditions. While the TWTT method has more lenient hardware requirements, the derivation of range can suffer from inaccuracy due to the vehicles' relative motion. In an extreme case, a vehicle may send a ping and then receive a reply after it has moved a considerable distance, thereby invalidating the static assumptions of Equation 2.2. However, sound travels quickly enough through water to keep the time spent transmitting messages short, making this source of error for range measurements relatively small for reasonable relative vehicle velocities [36].

2.1.1 WHOI Micromodem-2

The Acoustic Communication Group at Wood's Hole Oceanographic Institute (WHOI) is at the forefront of developing new acoustic modems that tackle some of the insidious issues described previously. In 2009 the Micromodem-2 was released as an updated version based on previous work with the WHOI Micromodem-1; improvements to the power usage, bandwidth, computational ability, memory, and more were achieved while retaining backward compatibility and a similar form factor [15].

The timekeeping accuracy of the real time clocks improved tenfold from the Micromodem-1 to the Micromodem-2, further honing the ability to calculate range measurements with TWTT in a ping sequence. As displayed in Equation 2.2, the distance between acoustic modems is calculated from the total travel time, a known turnaround time, and the speed of sound. This turnaround time on the WHOI Micromodem-2 is more precise and consistent due to the improvement in clock accuracy and the ability to prioritize message queues from expanded memory availability [15]. Through use in field trials paired with GPS and GNSS as ground truth positions, the dynamic ranging accuracy's normal fit was characterized with a bias of -0.41 meters and a standard deviation of 2.9 meters [37]. Curcio et al. [10] confirmed that there is acceptable performance of the WHOI Micromodem-1 with up to 500 meters of separation between modems.

Overall, the WHOI Micromodem-2 has proven to be a well-performing acoustic

modem and has been used in the field for work from O’Neill et al. [38], Benjamin et al. [6], and Vaganay et al. [48]. For the purpose of this thesis, the WHOI Micromodem-2 will serve as the grounding inspiration for performance and ranging protocol to be implemented in simulations.

2.2 Range-Only Single Beacon Navigation

In order to improve upon a simple dead reckoning solution derived from a vehicle’s onboard proprioceptive sensors, range measurements between the vehicle and another modem-equipped entity can be incorporated through a variety of methods. Many previous works have used techniques under the umbrella of range-only single beacon (ROSB) navigation, where time of flight measurements are utilized between the main vehicle and another entity, often called the beacon [17],[21],[28],[34],[48]. While the configuration of a beacon or beacon vehicle has evolved over time, the concept of fusing range measurements with odometry measurements has remained at the core of cooperative navigation efforts.

Long baseline (LBL) is a traditional localization method that uses range measurements from beacons at set locations, often fixed to the seafloor. By incorporating ranges from three or more beacons, trilateration is used to resolve the position of a vehicle. Some drawbacks include the overhead cost of setting up the beacons, as the installation can be time intensive and expensive, and must ensure that the beacon positions are accurate. Additionally, the operable area for an AUV using an LBL setup is limited by the static nature of the beacons and must be predetermined. Extending the area involves more costly setup and reconfiguration. GPS intelligent buoys (GIBs) can reduce some of the initial cost as they need not be installed on the seabed [36]. Ultrashort baseline (USBL) utilizes an array of transceivers and can measure range from time of flight as well as bearing from the relative phase difference between the transceivers. These arrays can be affixed to the bottom of support ships for increased flexibility and mobility, but generally are less robust to environmental factors and require more costly hardware [30].

Some early ROSB systems employed a static surface beacon as a means for navigation. This was especially useful when performing a homing routine in which a vehicle attempts to maneuver to a beacon that is sending acoustic pings. Vaganay et al. [47] developed an algorithm and procedure for the homing problem that involves two stages. First, an AUV attempted to initialize its position by executing a circular path around the beacon and accepting several range measurements. As the measurements were accepted, the AUV could get an estimate of the beacon's position and the underwater currents in the surrounding area. Once the first phase was complete, an extended Kalman filter was initialized and further refined the estimate of the vehicle's and beacon's positions based on range and odometry measurements. The vehicle then maneuvered towards the beacon in a spiral motion to allow for reduced error accumulation based on its relative bearing to the beacon.

As cooperative navigation techniques have become more advanced, the use of AUVs or ASVs as mobile beacons has become more common. Vaganay et al. [48] developed the moving long baseline (MLBL) concept for cooperative navigation between a heterogeneous team of AUVs. The makeup of the team included two CNA vehicles that operated underwater, a search classify map (SCM) vehicle that was outfitted with sensors to detect objects on the seabed, and a re-acquire identify vehicle (RI) that possessed sensors to further identify objects on the seabed. With accurate clock synchronization, the CNAs were able to provide range measurements using OWTT. This allows for the communication architecture to be scalable to handle a larger team of vehicles for future efforts. Even though the position error of the CNAs grows unbounded without any surfacing for a GPS fix, the vehicles were outfitted with superior INS and DVL to keep their position estimate to an error of 0.1% of distance traveled. A fixed geometry was used for the vehicle team; all members traveled in the same direction, and the CNA vehicles lagged behind to create a horizontal baseline for the other two vehicles. The error growth rate of the estimated position for the SCM and RI vehicles was successfully limited by implementing the MLBL concept in both simulation and field trials [48].

ROSB techniques have evolved over time to become more flexible and show promise

for heterogeneous underwater teams when paired with advances in navigational and acoustic communication hardware. A large body of literature has emerged that focuses on how best to maneuver groups of AUVs to improve their state estimation and mission performance.

2.3 Maneuvering Strategies

Building upon the developments of MLBL, Fallon et al. [14] explored using an ASV as an aid for cooperative navigation for one or more AUVs. The ASV could also handle communications, making it a communication and navigation aid (CNA) for the AUVs. The WHOI Micromodem was used for round trip range measurements between vehicles; the CNA derived its position estimate from GPS while the AUVs performed filtering with dead reckoning sensor data and range measurements. Fallon explored the concept of observability for a ROSB problem, stating that the system is observable if the observability matrix is full rank. From the derivation of the observability matrix of their system, Fallon defined observability criteria,

$$\det(O) = \frac{-\dot{x}_A (y_A - y_C) + \dot{y}_A (x_A - x_C)}{d_{AC}^2} \neq 0 \quad (2.3)$$

where (x_A, y_A) denotes the AUV location, (x_C, y_C) denotes the CNA location, and d_{AC} is the distance between the vehicles. Equation 2.3 must be true at the time a range measurement is made for the system to be observable. Additionally, Fallon notes that the range and relative position must change between successive range measurements to eliminate ambiguity and the growth of error in the direction orthogonal to the range measurement. In order to satisfy these conditions two general maneuvering patterns were chosen for the CNA to execute. The first was a continuous encirclement of the AUV and the second was a 45-degree zigzag motion lagging behind the AUV. Three different filtering and smoothing techniques were compared: the extended Kalman filter (EKF), particle filter (PF), and nonlinear least squares (NLS). The NLS and PF outperformed the EKF, but the PF became computationally expensive with a lot

of particles. These efforts by Fallon highlighted simple maneuvering strategies based on mathematical criteria and motivated the use of NLS estimation techniques that will be discussed in Section 2.4.

Quenzer et al. [40] use the observability Gramian to inform control decisions and generate paths for a CNA supporting one or many AUVs. Essentially, an observable linear system has a nonsingular Gramian for all time, allowing for the original state to be uniquely recovered from information output from start to finish. The local estimation condition number is defined as a ratio of the largest to smallest eigenvalues of the observability Gramian, and was used as a heuristic for one control algorithm. The other algorithm to plan CNA headings used path inertia. In the end, CNA planned trajectories resulted in spiraling motions that required bounding when proceeding away from the AUVs. The two greedy algorithms for CNA navigation based on observability definitions offered improved localization performance for the team of AUVs. From their results, Quenzer stated that using observability metrics over information gain metrics in a greedy motion planning system can provide better localization performance for a cooperative navigation team. This further motivates the use of observability for planning successful CNA trajectories.

Bahr et al. [4] built upon previous work with trilateration [1],[2] to develop a new path planning algorithm for cooperative navigation. Trilateration was performed as a localization technique to fuse an EKF estimate with range measurements while incorporating the covariance of the ranging vehicle [3]. A team of CNAs worked to aid in the localization of a team of AUVs, all operating underwater. A set transmission schedule of messages allowed for each vehicle to acoustically communicate without interruption; at each transmission, estimated position and covariance information was shared and range measurements were made with OWTT. The CNAs planned their movement such that their next transmission created the largest reduction in the sum of the trace differences for the AUVs' covariance matrices. This required some knowledge of the sensors deployed on all vehicles and the filtering methods used, as the CNA must extrapolate the AUV covariance matrix forward in time. As anticipated by previous work [50] CNA relative positions perpendicular to the AUV's

direction of travel were highly favored to reduce uncertainty.

Papalia et al. [35] investigated path planning for a group of AUVs where a subset of vehicles act as anchor nodes and are considered to have accurately known trajectories. This problem can be represented by a graph where each vehicle is a node and each edge is a range measurement between two vehicles. The edges of the graph form the Fisher information matrix (FIM) to describe the underlying information geometry and subsequently the quality of the localization solution. Localizability constraints were defined based on A-optimal design and E-optimal design. The former seeks to maximize the negative trace of the inverse of the FIM and the latter seeks to maximize the smallest eigenvalue of the FIM. Thresholds were defined to create allowable configurations, and prioritized path planning was completed for the set of robots from start to finish locations, using reachable and connected criteria as well. The resulting localization performance of the path planning algorithm was similar to or better than comparable methods including A* search, rapidly exploring random tree (RRT), and a potential field. This work stresses the importance of the configuration of a group of AUVs for accurate cooperative navigation and shows improved performance through additional computation and FIM-based planning.

2.4 Nonlinear Least Squares Estimation and Factor Graphs

This section will review the foundational derivations for nonlinear least squares (NLS) estimation and how a factor graph representation is beneficial to solving a cooperative navigation problem. Using NLS is necessary due to the nonlinear nature of range measurements. NLS methods for cooperative navigation state estimation perform well when compared to a particle filter (PF) or extended Kalman filter (EKF) as shown by Fallon et al. [14].

The core of state estimation is in choosing a set of states based on measurement information and predictive measurement models. A starting point for understanding

state estimation is the maximum a posterior or MAP estimate, which is defined as

$$X^{MAP} = \underset{X}{\operatorname{argmax}} p(X|Z) \quad (2.4)$$

where X is the set of unknown state variables, Z is the set of measurements, and $p(X|Z)$ is the posterior density, representing the probability of a possible set of state variables given the set of measurements [13]. In plain terms, the estimation routine aims to find the most likely series of states that would best explain the given measurements. Bayes' theorem can help in breaking down Equation 2.4, and states that

$$p(X|Z) = \frac{p(Z|X)p(X)}{p(Z)} \quad (2.5)$$

where $p(Z|X)$ is now the conditional probability density of the measurements Z given the states X and $p(X)$ is the probability or prior over all the states. $p(Z)$ is the probability of the set of measurements, which serves as a normalization and can be ignored for the maximization problem because the measurements are given; thus $p(Z)$ will not be altered by varying X during optimization [13].

Consider a simple setup as shown in Figure 2-1 where there are three unknown poses and two landmarks that are interacted with during the mission. The poses are linked by odometry information shown as arrows and the measurements to landmarks are shown as dashed lines. Figure 2-2 expands this problem to more explicitly show and name the measurements that make up Z . Note there is an absolute pose measurement for the initial state in order to fix the solution in space. As we continue through the section we will refer back to this simple example.

We can break down Equation 2.4 by using the likelihood function. The likelihood of states X given measurements Z is defined as any function proportional to $p(Z|X)$:

$$l(X; Z) \propto p(Z|X). \quad (2.6)$$

It is important to understand that the likelihood is only a function of the states X , and not Z , as the measurements act as set parameters for the estimation problem.




Figure 2-1: A simple example of a state estimation problem. Three successive poses are connected through odometry measurements. Two landmarks provided range measurements at poses two and three. To solve this problem, we desire to know the three poses and the landmark positions based on the measurements collected. Measurements could include odometry, range, bearing, and an initial absolute pose measurement.

By conditioning on measurements, we end up with functions that are not inherently Gaussian [13]. Rewriting Equation 2.4 with Bayes' theorem from Equation 2.5 and the relationship in Equation 2.6, we come to

$$X^{MAP} = \arg \max_X l(X; Z)p(X). \quad (2.7)$$

This results in an optimization problem that involves the likelihood of possible states X given measurements Z , and the prior probability distributions of the states $p(X)$. In order to efficiently solve this problem, a factor graph representation can be used for its separation of states and measurements, handling of non-Gaussian likelihood functions, and inherent structure for efficient computation.

Factor graphs can represent a joint probability density as a product of factors, similar to a Bayes net. To represent the posterior $p(X|Z)$ for the simple state estimation




Figure 2-2: An expansion on the state estimation example in Figure 2-1 to include the measurements that make up Z . The first measurement z_1 is an absolute pose measurement while z_2 and z_3 are range measurements to landmarks.

problem from Figures 2-1 and 2-2 with a Bayes net, we get the factorization:

$$\begin{aligned}
 p(X|Z) &\propto p(x_1)p(x_2|x_1)p(x_3|x_2) & (2.8) \\
 &\times p(l_1)p(l_2) \\
 &\times l(x_1; z_1) \times l(x_2, l_1; z_2)l(x_3, l_2; z_3)
 \end{aligned}$$

where there are conditional densities creating a Markov chain on the three poses, prior densities for the landmarks, a conditional density for the absolute pose measurement z_1 , and conditional densities for the range measurements z_2 and z_3 .

A factor graph has the advantage of being able to represent any factored function $\phi(X)$ over a set of variables X , not just probability densities. An entire factor graph contains variables and factors, which act as nodes and are connected by edges [13]. To factorize a global function $\phi(X)$, it is a product of factors as

$$\phi(X) = \prod_i \phi_i(X_i) \tag{2.9}$$

such that the edges of the graph encode the independence relationships within the factor graph.




Figure 2-3: A factor graph representation for the state estimation problem shown in Figure 2-1 where each black circle is a factor in the graph. The odometry measurements and range measurements are now represented with factors between the variable nodes. Priors are added to the landmarks and initial pose, which also has an absolute pose measurement to ground the solution.

Any Bayes net can be transformed into a factor graph. Each factor is a function of its neighbor nodes. Figure 2-3 shows the connectivity of the simple factor graph example, where each black circle represents a factor and is neighbored by one or more nodes that represent the pose and landmark locations. Our example factor graph factorization becomes

$$\phi(l_1, l_2, z_1, x_2, x_3) = \phi_1(x_1)\phi_2(x_2, x_1)\phi_3(x_3, x_2) \quad (2.10)$$

$$\times \phi_4(l_1)\phi_5(l_2)$$

$$\times \phi_6(x_1) \times \phi_7(x_2, l_1)\phi_8(x_3, l_2) \quad (2.11)$$

The composition of this factorization is similar to Equation 2.8 and can be combined with Equation 2.9 to rewrite our original MAP estimate as

$$X^{MAP} = \arg \max_X \prod_i \phi_i(X_i) \quad (2.12)$$

where the main factors that will be contained in the factor graph are odometry factors and range factors.

To link successive poses together, the use of odometry factors requires a measurement function that represents a difference between poses. Since odometry measurements are used to measure the relative difference in poses, they consist of a pose inversion and composition between adjacent poses. Range factors can link landmarks with poses and have a measurement function that predicts the distance between the two possible coordinates, which can have their own covariance distributions. These measurement functions $h_i(X_i)$ serve as a predictive model for the expected measurement value for a state X_i . We can assume that the actual measurements will have additive Gaussian noise with some known covariance Σ_i . This results in the equation for each factor in our graph to be defined as

$$\phi(X_i) \propto \exp\left\{-\frac{1}{2}\|h_i(X_i) - z_i\|_{\Sigma_i}^2\right\} \quad (2.13)$$

where $h_i(X_i)$ is the appropriate measurement model and z_i is the sensor measurement value. $\|e\|_{\Sigma_i}^2 \triangleq e^T \Sigma_i^{-1} e$ is the squared Mahalanobis distance, which allows measurement between a point like measurement z_i and a distribution with covariance Σ_i such as a measurement function $h_i(X_i)$ [13].

Now that we have defined the factorization of the factor graph representation and the equation for our factors, we can arrive at the MAP estimate for the NLS problem,

$$X^{MAP} = \arg \min_X \sum_i \|h_i(X_i) - z_i\|_{\Sigma_i}^2. \quad (2.14)$$

To perform minimization on nonlinear measurement functions, a Taylor expansion is required to linearize the measurement function $h_i(X_i)$. This is defined as

$$h_i(X_i) = h_i(X_i^0 + \Delta_i) \approx h_i(X_i^0) + H_i \Delta_i \quad (2.15)$$

where X_i^0 is a given linearization point and Δ_i is the state update vector defined as the difference between X_i and X_i^0 . H_i is the measurement Jacobian defined as the

multivariate partial derivative of $h_i()$ at X_i^0 ,

$$H_i \triangleq \left. \frac{\partial h_i(X_i)}{\partial X_i} \right|_{X_i^0}. \quad (2.16)$$

Substituting the Taylor expansion of Equation 2.15 into the NLS problem in Equation 2.14, we can arrive at a linear least squares problem for solving the state update vector Δ ,

$$\Delta^* = \underset{\Delta}{\operatorname{argmin}} \sum_i \|h_i(X_i^0) + H_i \Delta_i - z_i\|_{\Sigma_i}^2 \quad (2.17)$$

$$= \underset{\Delta}{\operatorname{argmin}} \sum_i \|H_i \Delta_i - \{z_i - h_i(X_i^0)\}\|_{\Sigma_i}^2 \quad (2.18)$$

where $z_i - h_i(X_i^0)$ is grouped to explicitly represent the prediction error at the linearization point. Δ^* is the solution to the locally linearized problem. To further simplify this expression, we can transform the Mahalanobis norm into a Euclidean norm with a change of variables. Dellaert and Kaess [13] show that pre-multiplying by the matrix square root of the covariance matrix Σ allows for

$$\|e\|_{\Sigma}^2 = e^T \Sigma^{-1} e = (\Sigma^{-1/2} e)^T (\Sigma^{-1/2} e) = \|\Sigma^{-1/2} e\|_2^2. \quad (2.19)$$

Then, by defining $A_i = \Sigma_i^{-1/2} H_i$ and $b_i = \Sigma_i^{-1/2} (z_i - h_i(X_i^0))$ we can obtain a standard least-squares problem,

$$\Delta^* = \underset{\Delta}{\operatorname{argmin}} \|A\Delta - b\|_2^2 \quad (2.20)$$

where all Jacobian matrices and prediction errors are combined into a large, sparse matrix A and a long vector b respectively. The matrix A has a block structure that mimics the structure of the factor graph and allows for improved computation due to its sparsity.

Several nonlinear optimization methods can be used to solve the NLS problem. The summation in Equation 2.14 can serve as the cost function for solving the NLS

problem. Typically the procedure of a nonlinear optimization method is to start at an initial estimate X^0 and iteratively update the solution to reduce a cost function until there is convergence.

Gradient descent (GD), as the name implies, uses the estimated gradient of a cost function $c(X)$ to find the largest direction for cost function reduction. Upon each iteration, a step is taken in this direction as a form of an updated solution where

$$\Delta_{GD} = -\alpha \nabla c(X)|_{X=X^t}. \quad (2.21)$$

The step size α for gradient descent is an important parameter; if α is too small many iterations are required, if α is too big the minima can be stepped over. GD methods often have slow convergence near the minimum and can terminate at local minima.

Gauss-Newton (GN) methods use second order updates to provide faster convergence and take advantage of the structure of the NLS problem. For a second order update the Hessian is approximated by the square of the combined Jacobian matrix A . To get the GN update Δ_{GN} , the normal equations for least-squares solutions must be solved, using

$$A^T A \Delta_{GN} = A^T b \quad (2.22)$$

where A and b are the same quantities as defined in Equation 2.20. With a good initial estimate, close to quadratic convergence is possible; however if the objective function is not well fit to a quadratic function, divergence of the estimation may occur.

To blend the two optimization methods above, we can use Powell's dogleg (PDL) algorithm [39], which establishes a trust region for the update to the current solution. This trust region must be defined to recognize that our linear assumptions can only be extended so far, and to guard against divergence. For the PDL algorithm the gradient descent and Gauss-Newton update steps are both computed. The GD update is applied to the current estimate, and then the estimate is updated in the direction of the GN update; the total step size from original estimate to new estimate is bounded by the trust region. For visualization in two dimensions, Figure 2-4 shows the combination of update steps to the last estimate using the PDL algorithm.

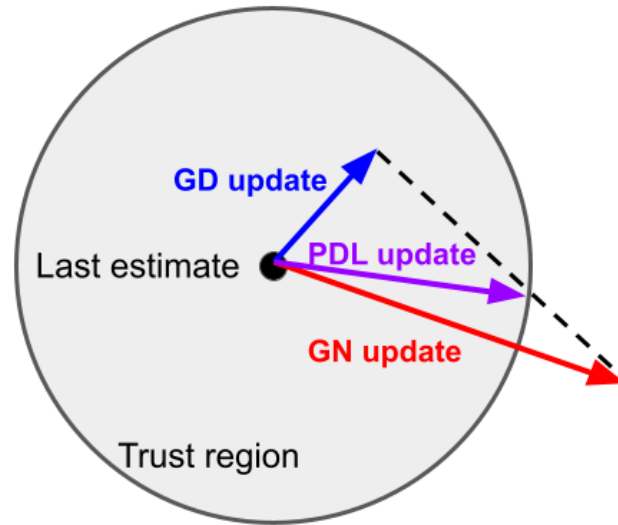


Figure 2-4: A visualization of Powell’s dogleg algorithm, which bounds the combination of a gradient descent (GD) and a Gauss-Newton (GN) update to a trust region in which the linear approximation is considered valid.

Kaess et al. created the incremental smoothing and mapping 2 (iSAM2) algorithm [22] to perform optimization for problems in which measurements are continually added and online updates are required. The original relationship between a factor graph and Bayes net is exploited such that only connected parts of the factor graph are affected by a new measurement. More specifically, a partial state update can take place, thereby reducing the computation for the update step significantly. Furthermore, linearization points can be reused between small updates and only are recalculated for states in which there is a significant change; again this occurs for a partial set of states instead of a full relinearization and is based on configurable thresholds. All of these aspects makes iSAM2 a desirable algorithm for online NLS state estimation. Kaess et al. [23] performed a comparison to other SLAM methods to show that iSAM2 is the most accurate and computationally efficient. Its use in works such as [38] and [37] to create the `ros_hat` repository motivates the use of iSAM2 in this thesis.

2.5 Mission Oriented Operating Suite with Interval Programming (MOOS-IvP)

The mission oriented operating suite with interval programming (MOOS-IvP), is an open-source middleware with modules for marine autonomy. Paul Newman first wrote MOOS-IvP in 2001 for the MIT Ocean Engineering and MIT Sea Grant programs to aid in their use of marine robotics [32]. The code base has been developed by engineers at both Oxford and MIT. It has lasted and evolved over time as a flexible architecture with new additions of modules such as swarm autonomy and advanced sensor drivers. The core set of MOOS-IvP capabilities is maintained by developers within MIT’s Laboratory for Autonomous Marine Sensing Systems (LAMSS) and Computer Science and Artificial Intelligence Lab (CSAIL), as well as the Oxford Robotics Institute (ORI) [7].

As a middleware, MOOS-IvP is intended to be run on a backseat autonomy computer that can dictate higher level decisions. The back seat computer communicates with a vehicle’s front seat computer, which in turn issues actuation commands and relays sensor data. In this way MOOS-IvP can be a platform-agnostic tool and is focused on solving problems of autonomy rather than improving the vehicle’s control loop. One of the strengths of MOOS-IvP is the use of interval programming for multi-objective optimization. Heading, speed, and depth decisions are made by the IvPHelm application, which solves multi-objective optimization at each time step of a mission. Users can configure several behaviors to run on the vehicle’s IvPHelm; each behavior outputs an objective function which is assigned a priority weighting that may change in time. Using interval programming, the IvPHelm selects the most optimal heading and speed while considering the benefits across all active behaviors. This is especially useful for a heterogeneous team where agents may need to consider their own priorities while balancing cooperative objectives [6],[9].

Interval programming (IvP) problems work with objective functions and define them as piecewise functions. Multi-objective optimization problems typically operate with several decision variables that comprise a multi-dimensional decision space. In

the IvP definition, we consider each decision variable to have a finite set of discrete values with upper and lower bounds. The IvP function is crafted out of interval pieces such that every point in the decision space is contained in no more than one piece. Every piece is defined on a set of boundary intervals for each decision variable. Any point within a piece can be evaluated on the IvP function from a linear interior function. For an IvP problem that spans n decision variables and incorporates k objective functions, we search for a set of values for our decision variables that maximize the quantity

$$w_1 f_1(x_1, \dots, x_n) + \dots + w_k f_k(x_1, \dots, x_n) \quad (2.23)$$

where f_i is the IvP piecewise function and w_i is the priority weight assigned to each function at the outset of the problem formulation [9]. To construct this piecewise representation it is necessary to make an approximation of an underlying nonlinear function. MOOS-IvP provides tools and guidance for creating these approximations in informed ways that allow for accurate and time-efficient computation. For MOOS-IvP instances of interval programming, the decision variables are normally limited to heading, speed, and depth. The objective functions are created from behaviors that come pre-designed or can be modeled by a user. The search methods that are utilized by MOOS-IvP to select the optimal heading and speed are efficient and ensure global optimality for finding the best solution from all pieces [8].

Within the MOOS-IvP operating environment, each vehicle is considered a community, which has a central variable-value database called the MOOSDB. For each community, several applications run during a mission. These applications publish and subscribe to a set of variables in the MOOSDB and carry out their own processes and utilities according to their design. This organization can be seen in Figure 2-5 where several applications are linked to the MOOSDB for an example community. In order to communicate with other vehicles, MOOS-IvP utilizes a shoreside-vehicle topology instead of direct inter-agent communication as seen in Figure 2-6. In essence, all communications between agents are filtered through the shoreside community, which often acts as a manager and/or simulator for realistic field conditions such as dropped

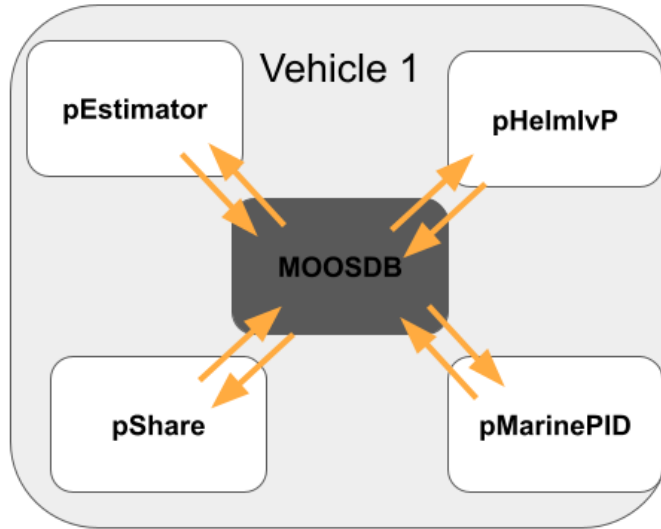


Figure 2-5: Each vehicle is called a community in MOOS-IvP and operates around its central database of variable-value pairs, the MOOSDB. Different applications run in the community and use a publish and subscribe system to update internal variables.

messages, blackout periods, or other communication restrictions [7]. MOOS-IvP has expanded its open-source codebase to include new modules as the middleware has become more popular. Many researchers have utilized its capabilities for advanced multi-vehicle simulations and field tests [6],[14],[34],[43].

One key example of utilizing MOOS-IvP’s capabilities for improved performance is the work of Benjamin et al. [6], where the IvP model was run on autonomous kayaks outfitted with acoustic modems. One kayak acted as a CNA with known position and remained stationary, while the other kayak balanced two mission objectives with two different MOOS-IvP behaviors. The first objective was to complete a survey of an area and was implemented as a set of desired waypoints. The second objective was to retain a suitable position estimation and was implemented as a behavior that maneuvered the vehicle into spatially diverse positions such that ranging to the CNA would improve the state estimation quality. This quality was described by a metric η , which was used to dynamically change the priority weights of the two behaviors. Thus, the survey vehicle could balance its objectives and complete the survey with an appropriate quality of estimated position.

Shafer et al. [43] used MOOS-IvP to investigate cooperative search techniques for

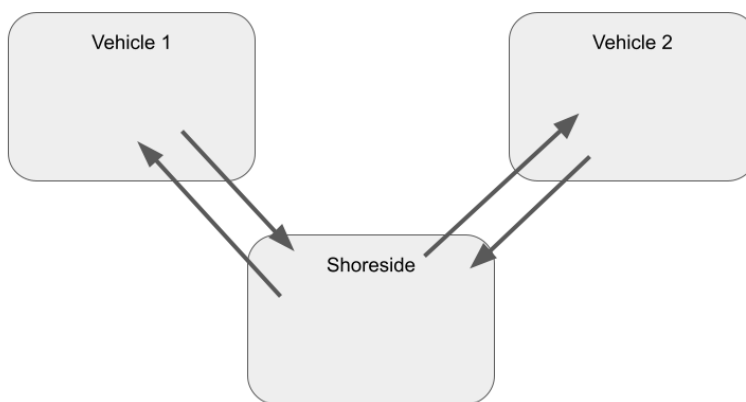


Figure 2-6: A representation of MOOS’ shoreside-vehicle topology for communication and information flow. Instead of the two vehicles communicating directly, messages flow through the shoreside community which oversees the mission and can set environmental factors for simulated conditions.

mine countermeasures (MCM). A MOOS-IvP application for simulating sensor data was created that mimicked the probabilistic nature of the sensing capabilities for a mine-hunting sensor. Several behaviors were developed to allow for a greedy search algorithm, prioritize cooperative navigation positioning, and carry out cooperative search. MOOS-IvP proved to be a powerful tool for implementing a cooperative search algorithm in both simulation and field trials by dynamically weighing the different mission priorities.

Thumma [46] utilized the MOOS-IvP simulation environment to expand the comparison of a prioritized planning algorithm defined in [35]. Implementation of a potential field planner worked well in the behavior-based system and allowed for testing with inter-agent collision avoidance as well as obstacle avoidance.

2.6 Energy Usage for Underwater Vehicles

An important factor to consider when optimizing the performance of underwater autonomous vehicles is energy usage. Bellingham et. al [5] define the energy usage of

an AUV completing a survey to be

$$E_{total} = \left[\frac{\rho C_d S v^3}{2\eta} + H \right] \tau \quad (2.24)$$

where ρ is the density of water, C_d is the vehicle's drag coefficient, S is the wetted surface area of the vehicle, v is the vehicle's speed, η is the propulsion efficiency of the vehicle, H is the hotel load from the onboard sensors, and τ is the time taken to complete the mission. The hotel load for a vehicle will depend on the vehicle's sensor suite, which is selected to meet the mission goals. Additionally the parameters C_d, S, η will depend on the specific vehicle, its hydrodynamics, and propulsion setup. For a survey mission, v and τ have an inverse relationship, creating a trade-off when attempting to design an ideal vehicle speed for optimized energy usage with Equation 2.24 [5],[11].

Willcox [49] applied energy usage considerations to a situation in which an AUV is conducting an area survey to reconstruct a spatially distributed and time evolving process through sampling. Error performance metrics to quantify the ability of the AUV to properly reconstruct these fields were balanced against energy metrics derived from Equation 2.24. Willcox was able to combine these competing interests to develop survey analysis tools for AUVs.

Singh et al. [44] explored the competing interests in designing an AUV for use in underwater data collection. More specifically they considered energy usage and design considerations for improved sensing capabilities and navigational abilities. This blend of theoretical and practical considerations allows for further refinement and a more detailed design for optimized AUV performance. By incorporating specific vehicle dynamics or parameters, the energy and performance interests can be fine-tuned to a higher degree.

This thesis will not evaluate detailed energy usage but will consider the performance of the state estimation in relation to the energy usage for each maneuvering pattern. A balance between estimation performance and energy usage can be explicitly evaluated depending on specific vehicle dynamics in future works.

2.7 Problem Formulation

Based on previous work that has shaped the field of underwater cooperative navigation, this thesis seeks to improve the state estimation capabilities of a heterogeneous marine team through the development of maneuvering strategies. The team of interest will be composed of two agents, one with low navigational capabilities and one with advanced capabilities. The former will be referred to as the follower and the latter as the leader throughout this thesis. The mission goal is for the follower to navigate to a destination while the leader maneuvers around the follower, sending range measurements at set intervals. We will compare the performance of the follower state estimation for different leader maneuvering strategies. These strategies will vary in their complexity, with some utilizing simple shapes as defined in previous works [14], [33] and others as adaptive responses to the mission progression based on concepts explored in other research efforts [2], [14].

The simulations will be designed in the two-dimensional plane as we assume both agents will be able to resolve their depth with adequate accuracy to transform range measurements from three-dimensional space to two-dimensional space. In reality, the agents may be operating underwater at different depths to maintain a separation throughout the mission.

Chapter 3

Methods

This chapter will detail the software components implemented to simulate and evaluate cooperative navigation maneuvering patterns. The basis for these methods is the work of Pelletier and O’Neill, specifically their `ros_hat` repository [37],[38]. The `ros_hat` repository is composed of Python scripts for use in the robot operating system (ROS) middleware, while the methods developed for this thesis are written in C++ for use in MOOS-IvP.

3.1 MOOS-IvP Setup

Each iteration of a MOOS-IvP simulation is called a mission or a trial and involves launching several communities that host their own set of MOOS-IvP applications. For our mission, there are two communities representing the leader and follower vehicles and the shoreside community that creates visualizations and handles communication flow. The mission goal is for the follower to arrive at a predetermined destination. The leader aids in localization by maneuvering based on the employed strategy. Section 3.5 provides an in-depth discussion of the design of the five tested strategies. The leader and follower vehicles are designed with different behaviors and separate methods of computing their estimated position in order to reflect their desired real-world counterparts.

All maneuvering for MOOS-IvP missions will take place on an XY plane for sim-




Figure 3-1: MOOS-IvP coordinate system for XY plane and angles. North is positive Y direction and zero degrees, with increasing angles proceeding clockwise to the positive X axis.

pler computation and design. In reality, the two vehicles operate at depth; range measurements are three-dimensional and the two agents can be separated by a depth offset to maintain distance. As pressure sensors can accurately resolve a vehicle’s depth, we will consider the depth to be known and available for all computation to be projected into a two-dimensional plane [36]. All forthcoming discussions will refer to operation in an XY plane but can be applied more generally to the three-dimensional ocean environment. Figure 3-1 shows the global reference frame that is used in MOOS; when referencing angles in any way throughout this thesis, we will be utilizing this coordinate frame.

We created standard mission parameters for all of our trials for ease of comparison. The follower completes a 400 meter transit in the negative Y direction with range pings to the leader every 30 seconds. The standard speeds for the follower and leader are 0.5 m/s and 1.5 m/s respectively, though increased leader speed will be an area of comparison in Chapter 4.

The main contributions of this thesis within MOOS-IvP are the creation of new applications and behaviors to implement state estimation, range measurements, external currents, and adaptive maneuvering strategies. The remaining sections of this chapter will walk through the development and purpose of these applications for use in our trials. Figures 3-2 and 3-3 illustrate the distributions of these modules across MOOS-IvP communities.

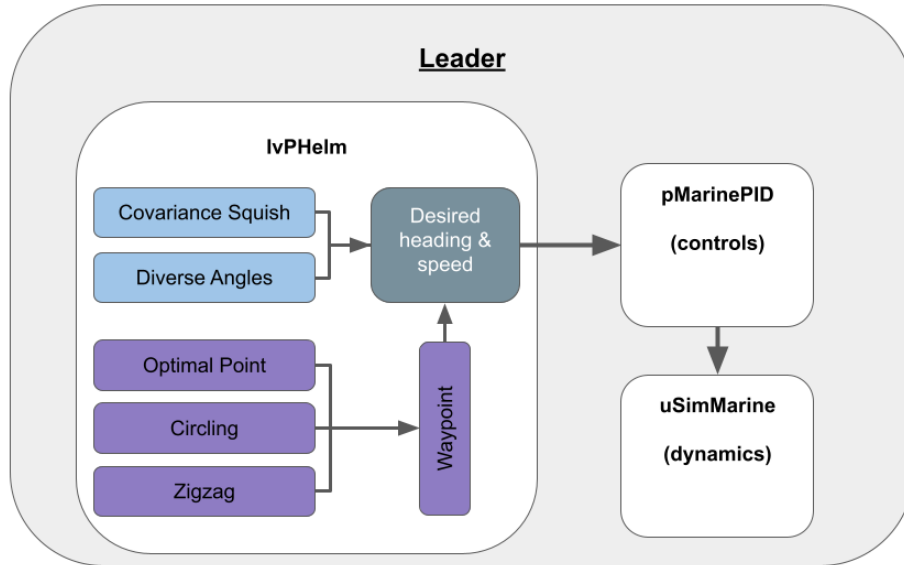


Figure 3-2: Visualization of how maneuvering strategies are integrated in the leader community. Strategies colored in blue are implemented as behaviors and directly produce speed and heading decisions. Strategies in purple generate new waypoints and are implemented as MOOS-IvP applications. A simple waypoint behavior directs the leader toward the next waypoint for strategies in purple. Speed and heading decisions are fed to other MOOS-IvP applications that handle vehicle dynamics and controls.

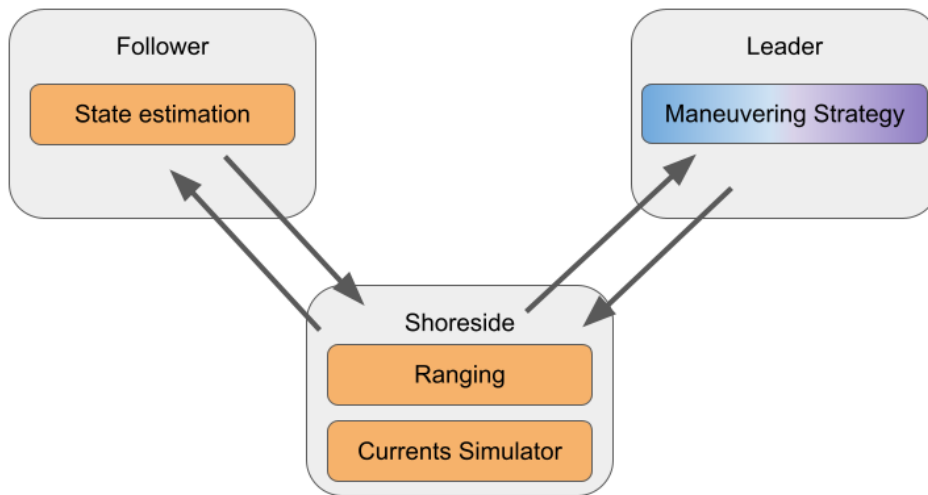


Figure 3-3: MOOS-IvP topology for our mission setup highlighting where new applications and behaviors are hosted. The follower community uses the new state estimation MOOS-IvP application pFactorGraphEst while the shoreside hosts the ranging application uCoopRangeSensor and the current simulation applications. Depending on the chosen strategy, the leader uses a new MOOS-IvP behavior or application to make maneuvering decisions, detailed in Figure 3-2.

3.1.1 Leader Vehicle

The leader’s main goal is to aid in the localization of the follower with iterative range measurements. The behaviors that are run on the leader’s IvPHelm application depend on the chosen maneuvering strategy and involve generating waypoints or a desired heading and speed. Section 3.5 will discuss how each strategy computes a new waypoint or desired heading and speed in detail. The estimated position of the leader is simulated to drift from the ground truth position as a percentage of distance traveled. The direction of this drift is chosen randomly at the start of a mission. For the simulations in this thesis, a value of 0.04% of distance traveled is used [19].

This method of generating an estimated position was designed as a simple yet realistic proxy for an AUV using an advanced INS. The covariance of the estimated position is approximated as a circular distribution in x and y proportional to the accumulated error. The estimated position of the leader vehicle is used to make heading and speed decisions in the leader’s IvPHelm, again dictated by the employed maneuvering strategy.

3.1.2 Follower Vehicle

The follower has a single behavior running for all missions that uses its estimated position to navigate to the destination coordinates. The estimated position of the follower is the output of pFactorGraphEst, a new MOOS-IvP application that will be described in Section 3.3. Heading and speed measurements are incorporated into the estimator, as well as range measurements from transmissions. In order to simulate realistic conditions, noise is added to the speed and heading measurements before incorporating them into the estimator as odometry factors. Based on the work of Pelletier [37] and O’Neill [38], the added noise was chosen to be a Gaussian distribution with standard deviations of 0.025 m/s and 2 degrees for speed and heading respectively; biases for the noise were chosen at the start of a mission from another set of Gaussian distributions with standard distributions of 0.025 m/s and 1 degree for speed and heading. A simulated current field is applied to add external disturbances

that mimic conditions in the field. The development of simulating currents will be covered in Section 3.2.

3.2 Ranging Protocol

In order to mimic the communication protocol designed in [38], a new MOOS-IvP application `uCoopRangeSensor` was created to simulate range measurements and transmit information packets between vehicles. Because this is carried out in simulation and there is no hardware involved, some simplifications are made to the ranging procedure. For example, messages are able to be passed to MOOS-IvP communities as strings and do not require encoding or decoding acoustic signals. Nevertheless, the design of the ranging application was implemented with the goal of mimicking the capabilities and error model of the WHOI Micromodem-2 [15].

In our MOOS-IvP missions, a sequence of messages between the leader, follower, and shoreside communities makes up our ranging protocol to mimic the structure of [38]. First, the follower issues a range request and is met with a range report. As soon as this report is sent, it triggers a range request from the leader which then receives its own range report. Due to the shoreside-vehicle topology, all messages are passed through the shoreside community, which hosts the `uCoopRangeSensor` application. This topology as well as the lack of an actual two-way travel time in simulation results in the sequence of information packets to be altered from the work in [37].

The ranging application `uCoopRangeSensor` is tasked with accepting range requests from vehicles and issuing range reports with the appropriate information. A range report for vehicle 1 contains the distance between the ground truth positions of vehicle 1 and 2, the latest estimated position of vehicle 2, the latest covariance matrix of vehicle 2, the MOOS-IvP time of the range measurement, and the name of vehicle 2. Both the leader and follower receive this information in their range reports. In order to properly populate the range report, the `uCoopRangeSensor` application has constant updates to the ground truth and estimated positions of both vehicles, again reinforcing the shoreside community as an overseer for simulations.

The follower uses all information in the range report to create a range factor within its estimation routine. The leader uses the latest estimated position of the follower as the most up-to-date information for planning maneuvers.

As it is currently implemented, `uCoopRangeSensor` does not incorporate noise into measurements or have a configurable blackout period to simulate dropped ranges. The noise for range measurements is added on reception from the vehicles. Dropped range measurements can be simulated through blocking range request messages for future work. Overall, this method of simulating ranges balances the realistic aspects of hardware challenges and simplifies some of the appropriate aspects for a simulation environment.

3.3 State Estimation with iSAM2

The incremental smoothing and mapping 2 (iSAM2) [24] library within Georgia Tech smoothing and mapping (GTSAM) [12] is utilized as the core of `pFactorGraphEst`, a new MOOS-IvP application that calculates the estimated position of the follower over the course of the mission. We utilize odometry factors and range factors as the main components of the factor graph, as well as priors on the follower starting position and the leader landmark positions at range transmissions.

Odometry factors are added every second from speed and heading measurements. Pelletier and O’Neill [38] used a constant speed assumption and a heading measurement from a compass to form their odometry factors in `ros_hat`. Once the noisy heading and speed measurements are acquired, the change in distance and heading is calculated from the last odometry factor. The change in side-to-side distance in the vehicle frame is considered zero for the forming of the odometry factors, as there is no measurement or guess of ocean currents. However, the noise model for the odometry factors allows for the estimated solution to incorporate this type of movement, usually after an update from a range measurement. Following the precedent set in `ros_hat` [38], our odometry noise models were increased to place more reliance on range measurements. This adjustment allowed for proper performance for our state

estimator and represents the potential influences of high current environments. An increased odometry noise model also reinforces the delineation of the follower as a low-cost, poor navigator and the leader as a reliable source of range measurements.

Every time an odometry factor is added, a new estimate is calculated and visualized. A prediction for the next pose must be added for the nonlinear least squares estimator to initialize. The last estimated position and the odometry measurements are used to compose the new pose guess through a one-step dead reckoning. Then, the new factor and initial value are added into the iSAM2 graph and optimization is carried out to calculate the most likely path considering all the factors. We use a PDL configuration for optimization, as explained in Section 2.4. An entire smoothed path is created at each time step, but only the final pose estimate is visualized in order to continually propagate the vehicle position. Smoothed paths collected along the transit can be saved and visualized. When there is a large change in the smoothed path, for example after the first ranges are incorporated, the estimated position can jump. This leads to the real-time estimated trajectory containing discontinuities while the smoothed path is free from such disturbances.

When ranges are incorporated there are two factors that must be added: the prior factor to describe the leader position and the range factor that connects the follower and leader positions at the time of the range measurement. The prior factor utilizes the information from the range report, namely the leader’s estimated position and uncertainty. Then the range factor is added, using an isotropic noise model that describes the measurement error of the WHOI Micromodem2 ranging capabilities with a standard deviation 2.9 meters as derived in [10],[38]. The first three range measurements are stored away to be added all at once when the fourth range measurement is acquired. This preempts any issues caused by spatial ambiguity. An essential issue to consider in the field is that the ranging procedure may take several seconds to measure and pass messages, meaning both agents may have moved between the time of the range request and the time the range factor is incorporated. When the follower issues a range request, it stores the current factor number, thus allowing the range factor to connect to the correct follower pose. One of the strengths of a factor graph

implementation is the ability to add information from the past and re-calculate the estimate in an efficient way.

3.4 Simulating Currents

MOOS-IvP implements external currents as a drift vector applied to the vehicle during propagation within the core uSimMarine application. A user can specify a magnitude and direction or x and y components for this drift vector and update the values by publishing messages to a vehicle's MOOSDB. For the trials in this thesis, we utilized a constant drift vector with a set magnitude and angle in the MOOS-IvP coordinate system.

Building on previous work for an advanced current simulator within MOOS-IvP [16], we also have the capabilities to simulate time- and space-varying currents. Through a MATLAB interface, a user can parameterize a noisy yet smooth vector field that updates in time. Two new MOOS-IvP applications were created to parse the vector field from MATLAB into updates for a vehicle's drift vector. This creates a more realistic current field for the vehicle to operate in. Users can input real data in the form of NetCDF files. Future work can continue to utilize the advances in current simulation within MOOS-IvP to further test for robustness before implementing these methods in the field.

3.5 Maneuvering Strategy Development

This section will detail the design of the five strategies for leader maneuvering. The strategies will produce either a desired heading and speed from a MOOS-IvP behavior or a new waypoint from a MOOS-IvP application.

3.5.1 Zigzag

Our zigzag pattern is defined as a repetitive side-to-side motion that maintains a trail distance behind the follower. The work of Fallon et al. [14] motivates the

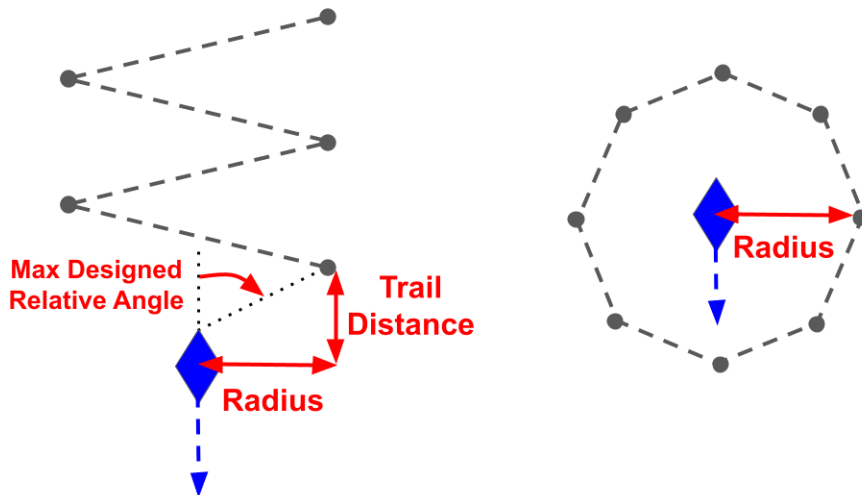


Figure 3-4: Geometric properties of the zigzag and circling strategies. The blue diamond represents the follower, which is traveling in the direction of the blue arrow. The zigzag pattern is described by a radius and a trail distance while the circling pattern is specified by a radius. The maximum designed relative angle for a zigzag pattern is a function of the radius and trail distance. Generated waypoints for the leader are shown in gray. Note that in the circling strategy, waypoints are generated in a batch of eight, but zigzag waypoints are generated one at a time.

use of a zigzag strategy. The shape of a given zigzag is parameterized by a radius and trail distance, where the trail distance is opposite the direction of the follower’s transit and the radius is orthogonal to the follower’s path. Figure 3-4 defines the geometric components of the zigzag shape. In order to create this shape, a new waypoint is generated as the leader approaches its current waypoint. Algorithm 1 is used to calculate the next waypoint based on the latest follower estimated position, the assumed speeds of the vehicles, and the desired geometry.

This process is triggered when the leader comes within 10 meters of its upcoming waypoint to allow for any necessary computation time. An issue in previous field trials involved the leader autonomy struggling when left without a new waypoint [37]; thus we are motivated to have a new point ready and update the leader’s desired waypoint queue such that the vehicle can complete its current transect and continue onto its next without a gap autonomy or a stale follower estimate.

The radius and trail distance of a given zigzag pattern effectively dictate a set of possible relative angles that can be experienced between the follower and leader.

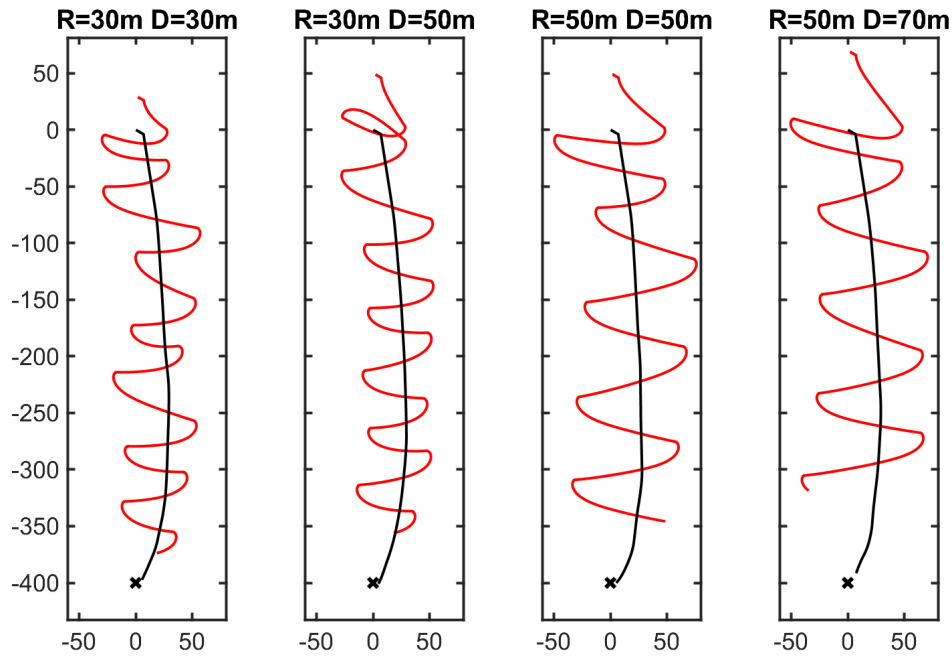


Figure 3-5: Sample trajectories using the zigzag strategy with different radii and trail distances. The leader's path is shown in red while the follower's is shown in black. An x denotes the destination coordinates. Values for the radius and trail distance ranged from 30 to 80 meters for all trials.




Figure 3-6: Relative leader trajectories for example zigzag patterns. The red line denotes the leader's path relative to the follower, which is represented as a blue diamond. The leader start position is marked with a red square. Note the effects of changing radius and trail distance, especially in regards the set of possible relative angles between the leader and follower.

Algorithm 1 Zigzag waypoint generation

Require: radius (R), trail distance (D), follower speed (v_f), leader speed (v_l)
In: Latest follower position (X_f), latest leader position (X_l)
Assume: Follower is travelling to destination (X_d)

- 1: Set look ahead time $T = 10$ seconds
- 2: **repeat**
- 3: Project new follower position (X_f^{new}) from X_f towards X_d by distance $v_f \times T$
- 4: Calculate desired leader position, (X_l^{new}) as $X_f^{new} + (\pm R, -D)$
 ▷ Sign of R is flipped between waypoints for zigzag motion
- 5: Get distance leader must travel (d) between X_l^{new} and X_l
- 6: **if** $d < (v_l \times T)$ **then**
- 7: Reachable point, **break repeat**
- 8: **else**
- 9: Increment look ahead time T by 1 second
- 10: **end if**
- 11: **until**

Out: Send next leader waypoint X_l^{new}

For example at a given trail distance, a larger radius would allow for a wider set of relative angles than a smaller radius. We can define the maximum designed relative angle as

$$\Theta_{max} = 90^\circ - \tan^{-1} \left(\frac{D}{R} \right) \quad (3.1)$$

where R is the zigzag radius, D is the trail distance, and Θ_{max} is measured in degrees.

Example zigzag trajectories are shown in Figure 3-5 for different combinations of radii and trail distances. The leader's relative path to the follower for these trials is shown in Figure 3-6, which allows for better visualization of how the maximum designed relative angle can change with radius and trail distance.

The change in relative bearings between range measurements is dependent on the geometric properties as well as the range period and vehicle velocities. Applied current can also affect the relative leader positions, especially by skewing the spread in one direction. The value of the maximum designed relative angle of a zigzag will be used as a means of comparison in Chapter 4.

3.5.2 Circling

A circling maneuver is used to allow for the full spectrum of relative angles between the follower and leader through the transit. This shape is again motivated by work from Fallon et al. [14]. A risk when circling is the possibility that the follower will drift outside the leader's circle and will produce unfavorable geometry.

The main parameter for a circling maneuver is the radius. When a set of waypoints is planned, the circle is designed such that the follower enters the circle as the leader begins its circular motion at the entry point. Our circles are composed of eight waypoints and are generated as a batch with Algorithm 2 when the leader comes within 10 meters of finishing its final waypoint from the previous batch.

Algorithm 2 Circling waypoint generation

Require: radius (R), follower speed (v_f), leader speed (v_l)
In: Latest follower position (X_f), latest leader position (X_l)
Assume: Follower is travelling to destination (X_d)

- 1: Set look ahead time $T = 10$ seconds
- 2: **repeat**
- 3: Project new follower position (X_f^{new}) from X_f towards X_d by distance $v_f \times T$
- 4: Set circle center (X_{center}) R meters ahead of X_f^{new}
- 5: Calculate leader entry point (X_l^{new}) as $X_{center} + (R, 0)$
- 6: Get distance leader must travel (d) between X_l^{new} and X_l
- 7: **if** $d < (v_l \times T)$ **then**
- 8: Reachable point, **break repeat**
- 9: **else**
- 10: Increment look ahead time T by 1 second
- 11: **end if**
- 12: **until**
- 13: Build list of eight circle waypoints W_{circ} from X_{center} , R , and X_l^{new} as entry point
 ▷ Our circles are configured to proceed in the clockwise direction

Out: Update leader waypoint queue with W_{circ}

Examples of final trajectories of the leader and follower are shown in Figure 3-7. Figure 3-8 transforms the leader trajectories into the relative paths from the follower's perspective. The trajectories of the leader are made of circles and connecting segments, while the relative motion forms a deformed circle from the view of the follower. Even though the circular motion is not perfect from this relative perspective, the original desire for changing relative angles is achieved nonetheless.




Figure 3-7: Example leader and follower trajectories for circling strategy. The leader trajectory is shown in red while the follower ground truth trajectory is shown in black. The black x denotes the destination coordinates and each figure pane illustrates a different circle radius. Radii from 20 to 70 meters were used in trials.




Figure 3-8: Relative leader trajectories for example circling patterns. The red line denotes the leader's path relative to the follower, which is represented as a blue diamond. The leader start position is marked with a red square.

3.5.3 Diverse Relative Bearings

One of the strengths of the circling strategy is the ability of the leader to transmit ranges to the follower from a diverse set of relative positions. As has been mentioned, receiving ranges at repetitive relative positions will result in poor performance from our estimation methods. Expanding on the idea of diverse relative angles, a new MOOS-IvP behavior was created to position the leader at the best relative angle to the follower for an upcoming range transmission.

The MOOS-IvP behavior was implemented to make heading decisions based on a calculated best relative bearing, which is re-computed when a range report is received by the leader. The diverse angles strategy calculates the best new relative angle based on Algorithm 4 and then maneuvers the leader with a heading decision from Algorithm 3. Minimum and maximum separation limits between the leader and follower can be dictated. This allows for a relative angle to be transformed into a segment of possible positions for the leader instead of an infinite set of positions along the relative angle. We will refer to this segment as our desired axis.

When the leader uses Algorithm 3 to maneuver, there are three modes for selecting a desired heading. In the first mode, the leader does not have enough time to reach the desired axis and will instead cut toward the vehicle; this does not happen often but is a necessary backup procedure. Alternatively, the leader may be closing in on the desired axis and thus can start to travel along the axis to maintain its relative position until the range measurement occurs. The third and final situation is when the leader has enough time to travel to the desired axis but is not yet close, so it directs its heading to the nearest point on the axis.

When a new range report is received by the leader, the MOOS-IvP behavior calculates a new best relative bearing based on diversity criteria as shown in Algorithm 4. Namely, the leader considers all possible relative bearings for the projected follower position, creates a set of reachable axes, and selects the option with the highest score. Scores are calculated as the sum of the angle differences between an option and the last five relative bearings.

Algorithm 3 Maneuvering to best relative bearing

Require: Desired relative bearing ($\Phi_{desired}$), leader speed (v_l), follower speed (v_f), ranging period (T_{range}), minimum and maximum separation (s_{min}, s_{max})

In: Leader position and heading (X_l, θ_l), latest follower position (X_f), time of last range report (t_{range})

Assume: Follower is travelling to destination (X_d)

- 1: Calculate time until next range measurement (Δt) from T_{range} and t_{range}
 - 2: Calculate distance leader can travel before next measurement ($d_{reachable}$) as $v_l \times \Delta t$
 - 3: Project new follower position (X_f^{new}) from X_f towards X_d by distance $v_f \times T_{range}$
 - 4: Form desired axis is from $X_f^{new}, \Phi_{desired}, s_{max}, s_{min}$
 - 5: Find closest point along desired axis ($X_l^{closest}$) to leader position (X_l)
 - 6: Calculate distance (d_{close}) from $X_l^{closest}$ to X_l
 - 7: **if** $d_{close} > d_{reachable}$ **then**
 - 8: Leader cannot make it to desired axis, so we calculate a catch-up point $X_{catchup}$ from $X_f^{new}, \Phi_{desired}$, and s_{min}
 - 9: **return** $\theta_{desired} \leftarrow relAng(X_l, X_{catchup})$
 - 10: **else if** $d_{close} > 3$ meters **then**
 - 11: **return** $\theta_{desired} \leftarrow relAng(X_l, X_l^{closest})$
 - 12: **else**
 - 13: Calculate separation (s) as distance between X_l and X_f^{new}
 - 14: **if** $s > s_{max}$ **then**
 - 15: $X_{cruise} \leftarrow$ project X_f^{new} along $\Phi_{desired}$ by distance s_{min}
 - 16: **else**
 - 17: $X_{cruise} \leftarrow$ project X_f^{new} along $\Phi_{desired}$ by distance s_{max}
 - 18: **end if**
 - 19: **return** $\theta_{desired} \leftarrow relAng(X_l, X_{cruise})$
 - 20: **end if**
-

Algorithm 4 Choose new desired relative bearing with diverse angles strategy

Require: leader speed (v_l), follower speed (v_f), range period (T_{range}), minimum and maximum separations (s_{min}, s_{max}), set of previous relative bearings (Φ^{prev})

In: Leader position and heading (X_l, θ_l), new Range Report

Assume: Follower is travelling to destination X_d

- 1: Get latest follower position (X_f) from Range Report
 - 2: Project new follower position (X_f^{new}) from X_f towards X_d by distance $v_f \times T_{range}$
 - 3: Calculate leader's reachable distance ($d_{reachable}$) as $v_l \times T_{range}$
 - 4: **for** all possible relative angles Φ_i **do**
 - 5: Construct possible axis from $X_f^{new}, \Phi_i, s_{min}, s_{max}$
 - 6: Find minimum distance d_{min} between leader position (X_l) and axis
 - 7: **if** $d_{min} < d_{reachable}$ **then**
 - 8: Score the relative bearing based on the last five previous bearings where
score $\leftarrow \sum_j^{j=5} |\Phi_j^{prev} - \Phi_i|$
 - 9: **end if**
 - 10: **end for**
 - 11: **return** Φ_i with maximum score as $\Phi_{desired}$
-

3.5.4 Covariance Squish

A different approach for maneuvering based on relative angles is to focus on the influence of the range measurement on the covariance of the estimated follower position [4],[50]. When the estimated position is calculated by iSAM2 throughout the transit, there is an additional output of a covariance matrix. The covariance matrix corresponds to the confidence of the estimated pose for x, y, and heading components.

Upon receiving the latest range report with the position and covariance information of the follower, the leader is able to calculate the major axis of the follower's covariance ellipse. A range measurement reduces uncertainty to the highest degree along the direction of the relative angle between the vehicles. The leader sets its best relative angle as the major axis of the follower covariance ellipse, as detailed in Algorithm 5. In this way, the leader seeks to "squish" the covariance ellipse to the maximum degree each time a range is transmitted. Mathematically, the leader seeks to maximally reduce the trace of the covariance matrix. After a new range report is received and Algorithm 5 calculates the best next relative bearing. Algorithm 3 is used to make heading decisions for the leader just as the diversity of angles strategy; for the covariance strategy the desired axis extends across the follower's path to also

allow for the mirroring of the best relative bearing.

An even simpler approach for implementing this strategy is possible due to the fact that each range measurement provides the most certainty in the direction of the relative bearing between the vehicles. Thus, the major axis of the covariance ellipse will be in the perpendicular direction to the last relative bearing. For our missions, it was sufficient to have the leader maneuver to a relative bearing that is orthogonal to its last relative bearing; commanding based on the covariance information provided the same decisions with more computation.

Algorithm 5 Choose new desired relative bearing with covariance squish strategy

In: new Range Report

- 1: Unpack follower position (X_f) and covariance matrix (P_f) from Range Report
 - 2: Get largest eigenvector (ν) of covariance matrix (P_f)
 - 3: Set desired angle ($\Phi_{desired}$) as the angle of ν
 - 4: **return** $\Phi_{desired}$
-

3.5.5 Optimal Point

The final maneuvering strategy that was developed and explored for this thesis is the optimal point strategy, which builds on the work of Fallon et al. [14] and utilizes the observability criteria as defined in Equation 2.3. Similar to geometric maneuvering strategies, it is implemented as a waypoint generator for the leader vehicle. However, for this strategy, the timing of the range measurement is taken into account explicitly, as the leader aims to arrive at its optimal transmission point for the next range measurement. Once a new range report is received by the leader, a new optimal point is selected using Algorithm 6.

When selecting the next optimal point, the following steps are taken. First, a set of reachable points is generated by considering the leader speed, the expected time until the next range, and a resolution for grid spacing. The follower position is propagated forward in time using the latest estimated position and its desired speed. All reachable leader points are evaluated and separated into an acceptable and unacceptable set; a potential new leader waypoint is considered unacceptable if

it fails the observability criteria (or the determinant is close to zero), if the resulting relative bearing to the follower is too similar to the last measurement, if the next range distance is too similar to the last measurement, or if the distance between the leader and follower violates the minimum and maximum separation bounds. All of these acceptability criteria come from the discussion in [14] around the derivation and use of the observability criteria. From the set of acceptable points, scores are assigned and the point with the best score is selected as the new waypoint. Scores are higher for points that are more closely aligned with the current vehicle heading and for points that create a more diverse set of relative bearings and distances compared to the last few measurements. Additionally, higher scores were assigned for points that were an appropriate distance away from the leader to minimize loitering while still choosing reachable points.

Algorithm 6 Optimal point waypoint generation

Require: follower speed (v_f), leader speed (v_l), range period (T_{range}), minimum and maximum separations (s_{min}, s_{max}), grid spacing (Δg)

In: Leader position and heading (X_l, Θ_l), Range Report

- 1: Get last range (D^{last}) and follower position (X_f^{last}) from Range Report
 - 2: Calculate last relative bearing ($\Phi_{relative}^{last}$) between follower (X_f^{last}) and leader (X_l)
 - 3: Project new follower position (X_f^{new}) from X_f towards X_d by distance $v_f \times T_{range}$
 - 4: Calculate leader's reachable distance ($d_{reachable}$) as $v_l \times T_{range}$
 - 5: Construct set of reachable points ($X_l^{reachable}$) with grid spacing Δg
 - 6: **for** X_i in $X_l^{reachable}$ **do**
 - 7: Calculate relative bearing (Φ_i) and range (D_i) between X_i and X_f^{new}
 - 8: Calculate $\det(O)$ from Equation 2.3 using X_i, X_f^{new}
 - 9: **if** all thresholds obeyed for $\det(O), |D_i - D^{last}|, |\Phi_i - \Phi^{last}|$, and separation **then**
 - 10: Add X_i to acceptable set ($X^{acceptable}$)
 - 11: Score X_i based on $\Phi_i, D_i, \|X_l - X_i\|_2, \theta_l$
 - 12: **end if**
 - 13: **end for**
 - 14: **return** X_i from acceptable set ($X^{acceptable}$) with maximum score
-

Some configuration variables for this strategy that required tuning included the thresholds for acceptable change in relative bearing and range distance, the threshold for observability criteria, and the failure strategy for when there are no acceptable points. The current failure strategy chooses a default point for the leader to navigate

to, with the new goal of getting into a position that will allow for acceptable points to be generated for the next range measurement. To accomplish this, the leader chooses a default point that is 90 degrees from its last measurement and twice the minimum separation away from the projected follower position. This brings the leader closer to the follower in hopes that more future points will be acceptable.

3.6 Summary

To facilitate simulations in MOOS-IvP, we created several modules that mirror the implementation of `ros_hat` [38], including functionality for ranging, state estimation, and currents. Furthermore, we developed five different maneuvering strategies that generate new waypoints or desired headings for the leader. The design of each strategy was based on previous work in cooperative navigation [14],[38]. The circling and zigzag strategies are considered our geometric strategies, and the covariance squish, diverse angles, and optimal points strategies are grouped as our advanced strategies. With all of these tools in place, we can collect data to compare the localization performance across different maneuvering strategies.

Chapter 4

Data Analysis

This chapter will walk through the comparative simulations and subsequent results that were carried out based on the methods described in Chapter 3

4.1 Summary of Trials

We deployed five different maneuvering strategies as outlined in Chapter 3 with several mission configurations. We will first provide an in-depth analysis of each strategy for our standard mission configuration as described in Table 4.1, then compare energy usage and performance between strategies. In later sections we will explore the impact of increasing current magnitude, changing the angle of current, and allowing the leader to travel at a higher speed. These other mission configurations are summarized in Table 4.2.

Our quantitative metrics will be based on the performance of the follower's state estimation and the leader's energy usage. Performance metrics include the average and maximum values of the real-time and smoothed error for the follower. Real-time error is the distance between the ground truth position and the estimated position from pFactorGraphEst at every time step during the mission. The smoothed error is the distance from ground truth to the final smoothed path, which is outputted by pFactorGraphEst at the end of the mission. In the forthcoming discussion performance metrics and error metrics will be used interchangeably, but it is important

Standard Mission Configuration	
Parameter	Value
Leader speed	1.5 m/s
Follower speed	0.5 m/s
Current magnitude	0.1 m/s
Current angle	120°

Table 4.1: Parameters for standard mission configuration.

Configuration Name	Parameter Changed	Tested values
Increased current magnitude	Current magnitude	0.15 m/s, 0.2 m/s
Increased leader speed	Leader speed	2.0 m/s, 2.5 m/s
Varied current direction	Current angle	45°, 240°, 330°

Table 4.2: Summary of the different mission configurations that were run and will be compared throughout this chapter.

to note that better performance is marked by lower values of the error. The cubed average speed of the leader is used as our energy metric from Equation 2.24 for quantitative comparison.

4.2 Circling Standard

This section will use the standard mission configuration as a basis for discussion on the performance of the circling strategy, especially in regard to how the circle radius is impactful.

Figure 4-1 combines our four error metrics for all circle radii. For circle radii of 60 meters or larger, there is an increase in maximum real-time error. Average error was less affected by the circle radius, and the best performance was achieved in the middle of our tested range of radii.

Even with a current magnitude of 0.1 m/s, the 25 and 30 meter circles can miss the follower on their initial circling maneuvers because the initial error has not been resolved from the first set of ranges yet. While the estimation can recover throughout the rest of the mission, it is not desirable for the geometry of the team to be lost. Small radii have this qualitative disadvantage that is not necessarily shown in our error metric results.




Figure 4-1: Follower estimation performance as a function of circle radius for the standard mission configuration. Red circles show the average and maximum error for the real-time estimated position while blue circles show the average and maximum error for the final smoothed follower trajectory.

When considering energy usage for a circling maneuver, we find that there is increasing energy usage for larger circle radii. Figure 4-14 shows the relationship between our energy metric and circle radius along with the other strategies.

In order to understand the potential trade-off between localization performance and energy usage, Figure 4-2 shows our performance metrics against cubed leader speed. Within our circling strategy, there is not a strong relationship between localization performance and energy usage. The one exception would be for our highest levels of energy usage; there is an increase in maximum error corresponding to a similar trend for circles with larger radii.

4.3 Zigzag Standard

For our other geometric pattern, we conducted a similar analysis of performance based on radius and energy usage. Figure 4-3 breaks down the performance of the zigzag

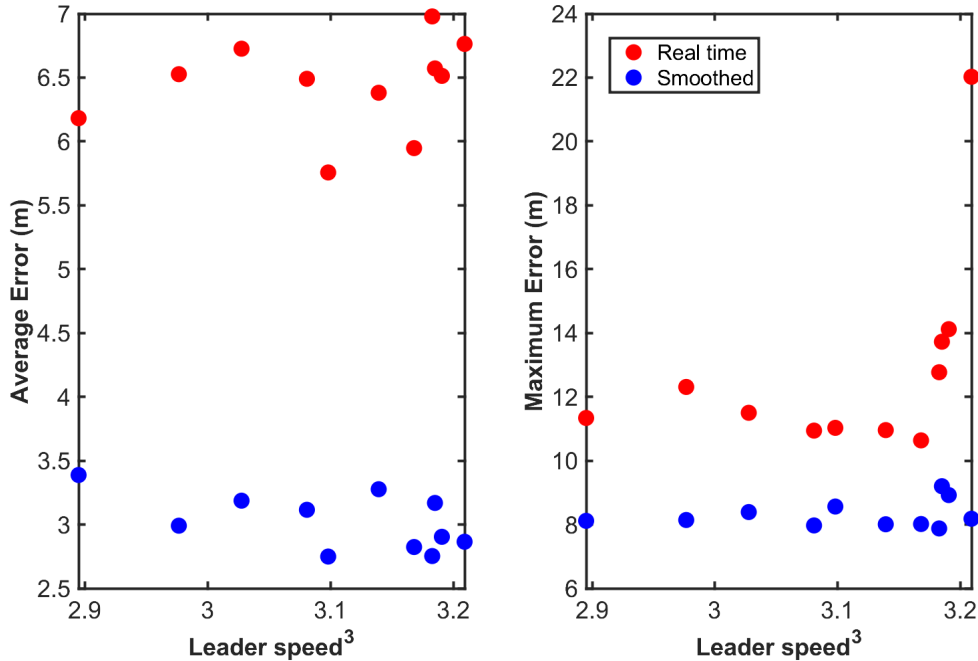


Figure 4-2: Follower estimation performance as a function of cubed leader speed for the standard mission configuration. Red circles show the average and maximum error for the real-time estimated position while blue circles show the average and maximum error for the final smoothed follower trajectory.

maneuvering strategy as a function of radius while Figure 4-4 does the same by cubed leader speed. Zigzag patterns have a small spread of performance for all radii, but the 30 meter radius performed worst in all error metrics. When performance is mapped as a function of energy usage, there is more of a downward trend, but still minimal spread to the data.

Zigzag energy usage was found to be a function of radius as shown in Figure 4-5. By visualizing our energy metric across all radii and trail distances, we can observe minimal variation in the trail distance dimension, but a significant positive relationship between radius and cubed leader speed.

Figure 4-6 maps our performance metrics across the two dimensions of the zigzag shape. The spread of performance varies for different combinations of zigzag radius and trail distance. The worst performance occurs for patterns with small radii and large trail distances. Better performance is achieved for larger values of zigzag radius paired with shorter trail distances.

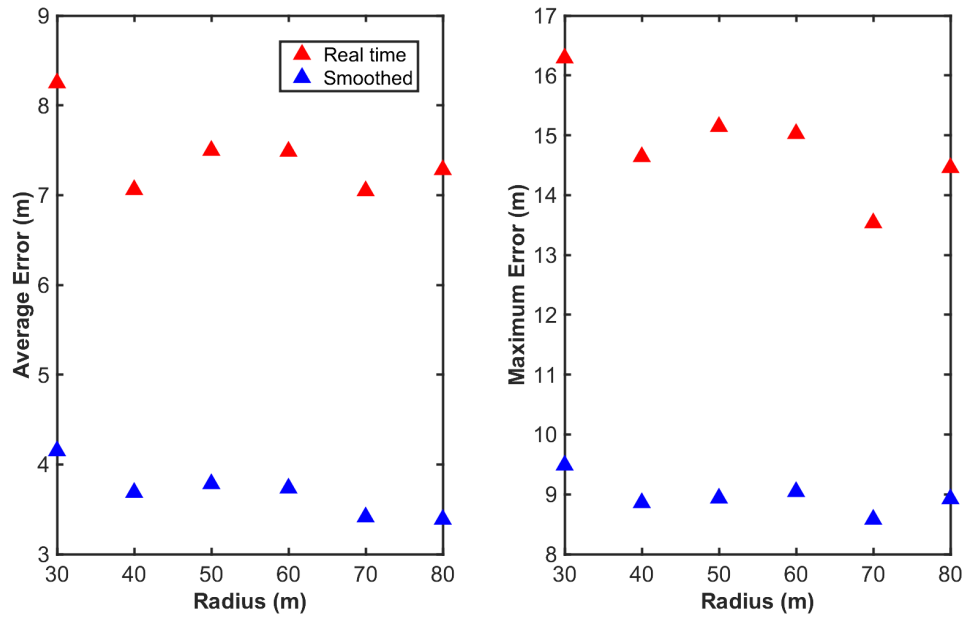


Figure 4-3: Follower estimation performance as a function of zigzag radius for the standard mission configuration. Red triangles denote the average and maximum error for the real-time estimated position while blue triangles denote the average and maximum error for the final smoothed follower trajectory.

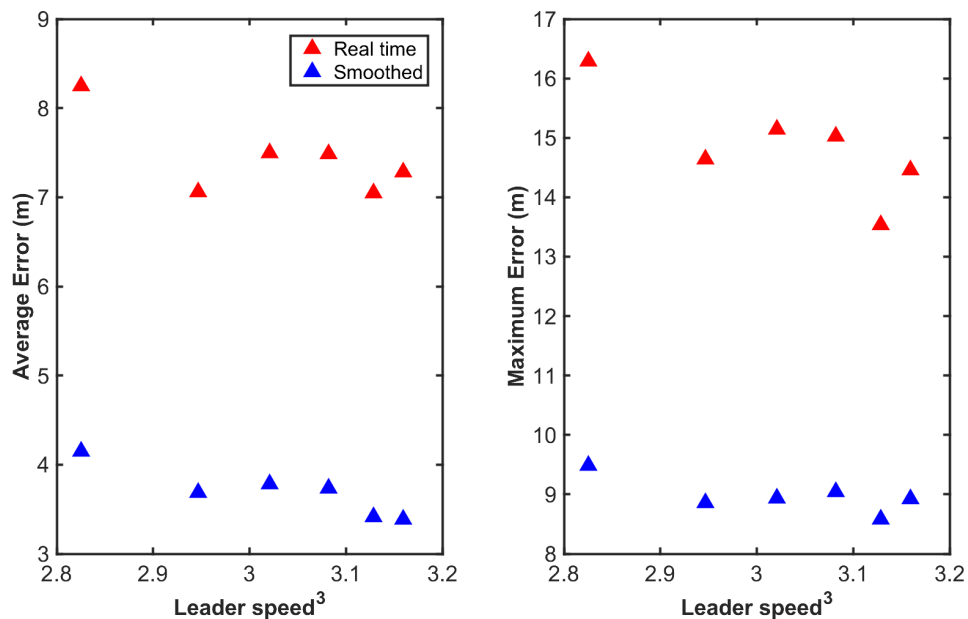


Figure 4-4: Performance metrics as a function of cubed leader speed for the standard mission configuration. Red triangles denote the average and maximum error for the real-time estimated position while blue triangles denote the average and maximum error for the final smoothed follower trajectory.

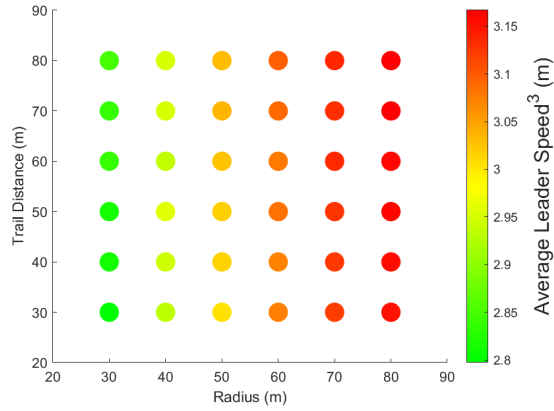


Figure 4-5: Cubed leader speed is shown for all combinations of radius and trail distance. We can state that our energy metric is a function of radius and is not dependent on trail distance.

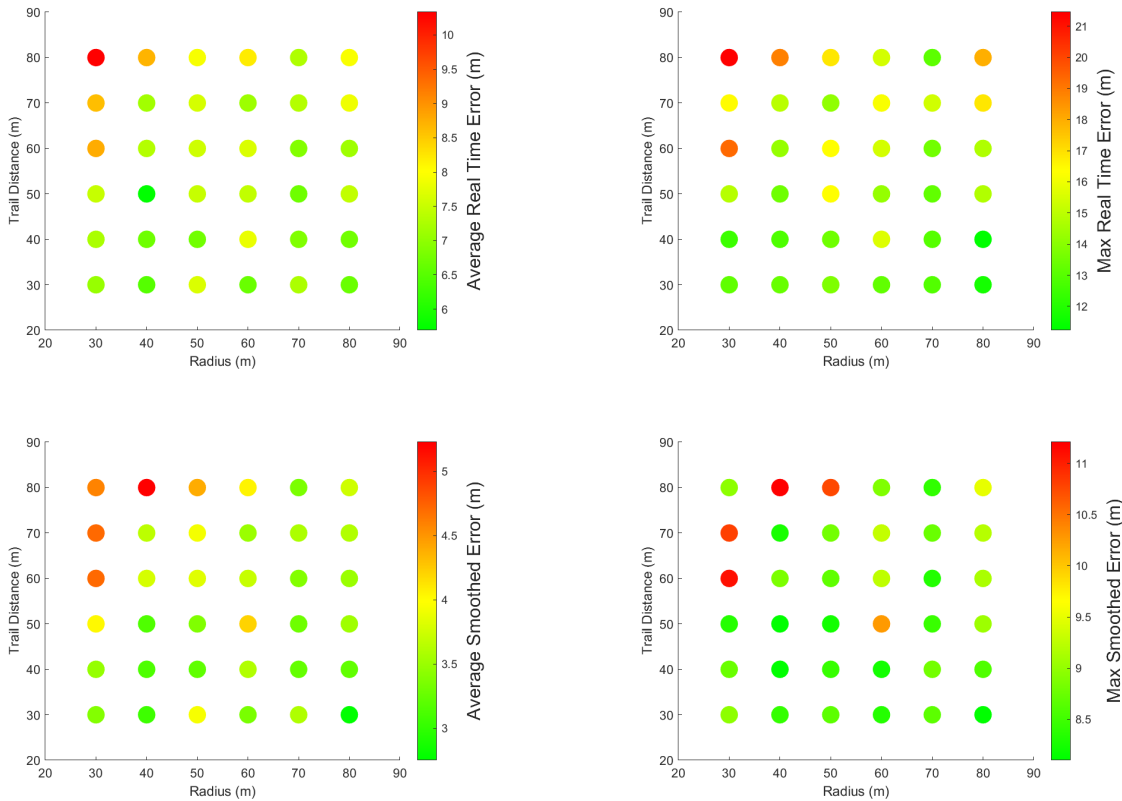


Figure 4-6: Performance metrics for the zigzag strategy with standard mission configuration, broken out by radius and trail distance. Note the different scales on the color bars.




Figure 4-7: A visualization of the maximum designed relative angle of zigzags for all radii and trail distances combinations. The maximum angle increases with radius and decreases with trail distance.

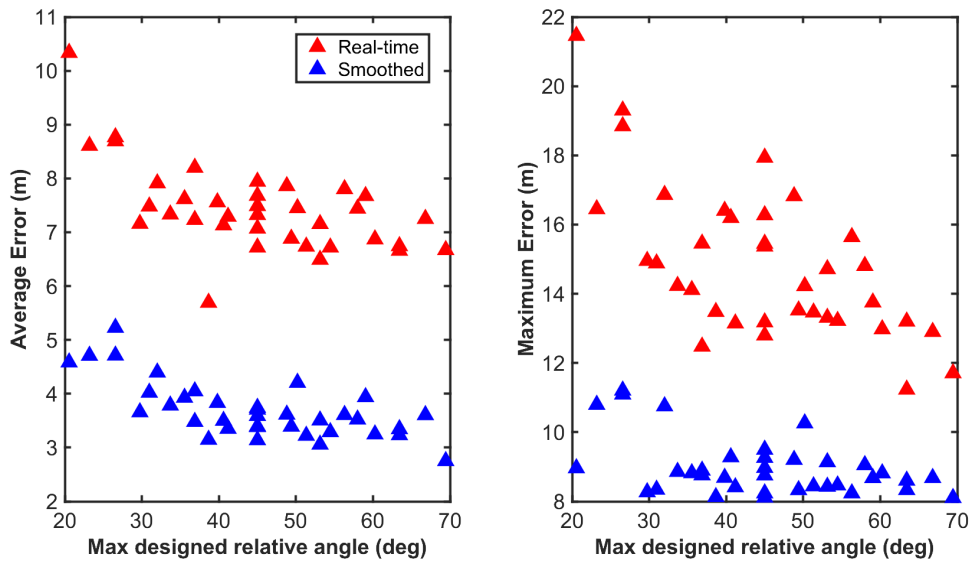


Figure 4-8: All follower error metrics as a function of maximum designed relative angle for the zigzag strategy at the standard mission configuration. Note the downward trend in the error metrics from 20 degrees to 50 degrees. For angles above 50 degrees, there is a flattening of the curve for all metrics except maximum real-time error.

To better represent the trends in our performance metrics, we can map the data across the maximum designed relative angle for all zigzag shapes. Recall this angle increases with radius and decreases with trail distance as defined in Equation 3.1 and visualized in Figure 4-7. Figure 4-8 displays our performance metrics as a function of the maximum designed relative angle. We can see a larger spread of performance when considering all possible zigzag shapes, as opposed to only considering radius. We observe that lower angles give worse performance; increasing this angle allows for our error metrics to flatten gently as we proceed to the largest maximum designed relative angle.

For our standard mission configuration, the performance of the zigzag maneuvering strategy is most reliant on the maximum designed relative angle and the energy usage is most reliant on radius.

4.4 Advanced Strategies Standard

We will consider the covariance squish, diversity of relative bearings, and optimal point strategies as our advanced strategies. As they are not parameterized by any geometry like circling or zigzag patterns, we will compare them together in this section. For brevity, their names may be shortened to the covariance, diverse, or optimal strategies in the following discussion and figures.

For the covariance strategy, we found that the resulting leader trajectories had some components similar to a zigzag path, with looping motions created as the leader attempted to travel along the major axis of the follower's covariance ellipse. Figure 4-9 illustrates some example trajectories for different leader starting positions.

One of the failure points of the covariance strategy occurs when the leader attempts to cross over the follower just before a range measurement. When in close proximity to the follower, the leader's relative angle can vary greatly depending on the exact timing of the measurement. If this happens in successive range measurements, the follower estimated position suffers due to poor geometry. If the leader makes a range measurement between the follower's ground truth and estimated positions,

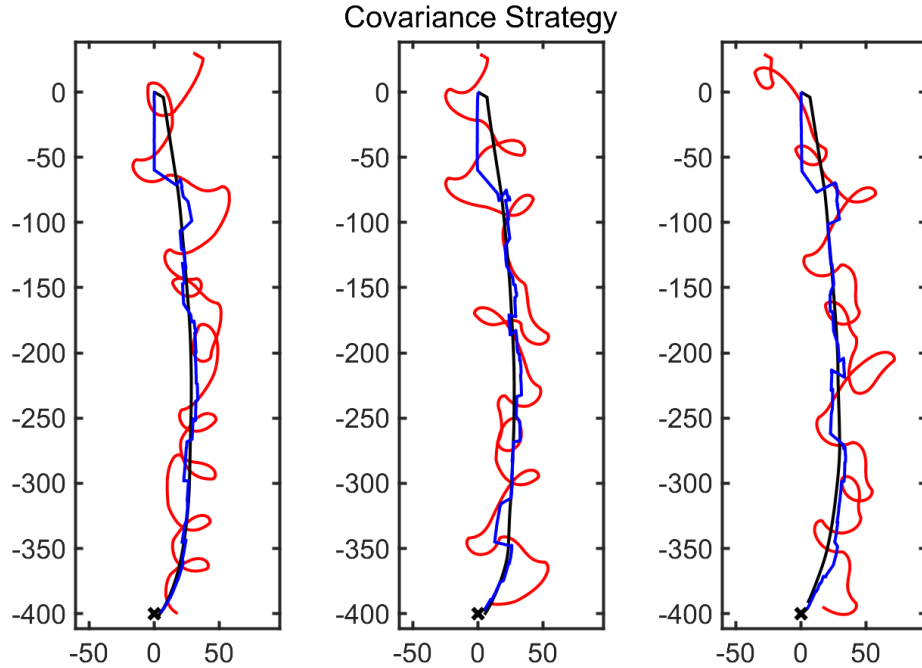


Figure 4-9: Example leader and follower trajectories for the covariance squish strategy at standard conditions. The leader’s ground truth path is shown in red while the follower’s ground truth and real-time estimated trajectories are shown in black and blue respectively. The black x marks the destination. Three different starting points were used for this strategy. Leader trajectories feature many looping motions.

spatial ambiguity can arise and deteriorate the estimate.

The diverse relative bearing strategy suffered from similar issues and resulted in leader trajectories as shown in Figure 4-10. Leader path shapes tended to switch between oscillations as seen in the covariance strategy and sweeping arcs. The performance of this strategy lagged behind the other advanced maneuvering strategies and the geometric patterns; however, its original goal to create a diverse set of relative bearings for range measurements was successful, as seen in Figure 4-13. We can further see that the covariance strategy created relative angles that form a wedge similar to the zigzag strategy, while the optimal point strategy’s relative ranging positions have less order than the other strategies.

The optimal point strategy resulted in leader trajectories as shown in Figure 4-11. The paths are characterized by small loops and periods of zigzag-like motion. The main issue with this strategy arose when no reachable points fulfilled the acceptable

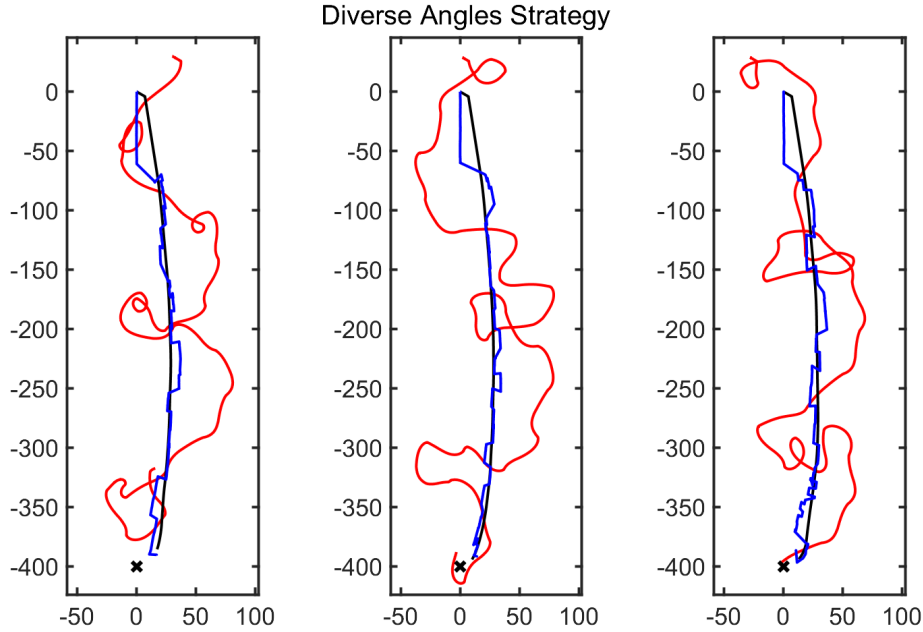


Figure 4-10: Example leader and follower trajectories for the diverse angles strategy with the standard mission configuration. The leader’s ground truth path is shown in red while the follower’s ground truth and real-time estimated trajectories are shown in black and blue respectively. The black x marks the destination. Three different starting points were used for this strategy. Leader trajectories feature smaller loops as well as sweeping arcs around the follower’s path.

criteria; the mandate for a large change in relative angle between range measurements greatly reduced the number of acceptable points. When the set of acceptable points was empty, a default point was generated for the leader to maneuver to. For our standard configuration, this occurred for roughly 30% of the waypoints during a mission.

To summarize the effectiveness of our advanced strategies, Figure 4-12 compares the mean values of our performance metrics for the three strategies. The ranking of the three strategies is consistent across all metrics, with the covariance strategy performing the best followed by optimal point and then diversity of angles.

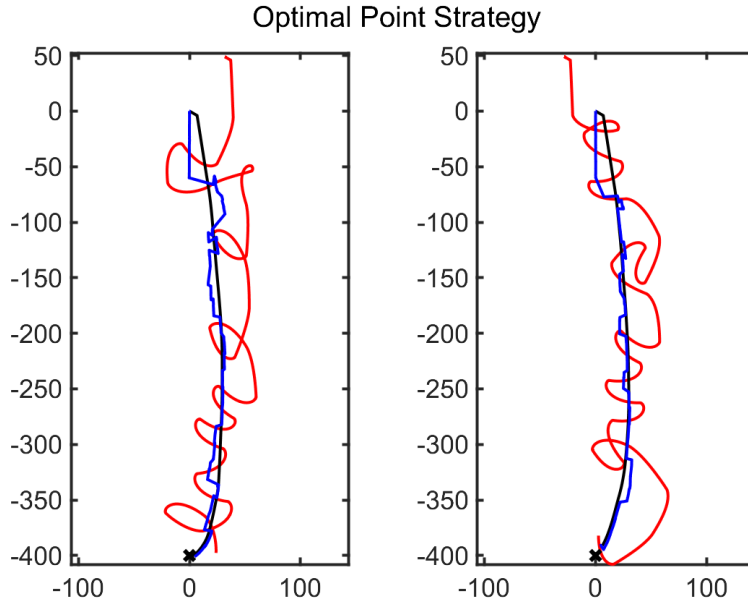


Figure 4-11: Example leader and follower trajectories for the optimal point strategy with the standard mission configuration. The leader’s ground truth path is shown in red while the follower’s ground truth and real-time estimated trajectories are shown in black and blue respectively. The black x marks the destination. Two different starting points were used for this strategy. Leader trajectories feature many looping motions and switchbacks similar to a zigzag.

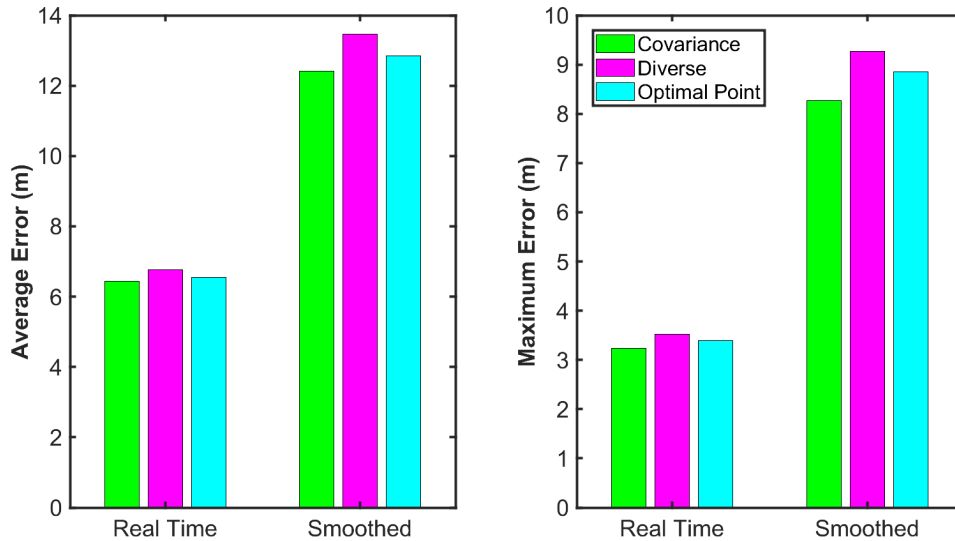


Figure 4-12: Comparison of all performance metrics for the three advanced strategies at the standard mission configuration. Values are averaged from 28 trials for each strategy. The covariance strategy has the best performance for all metrics, while the diverse angles strategy has the worst performance for all metrics.

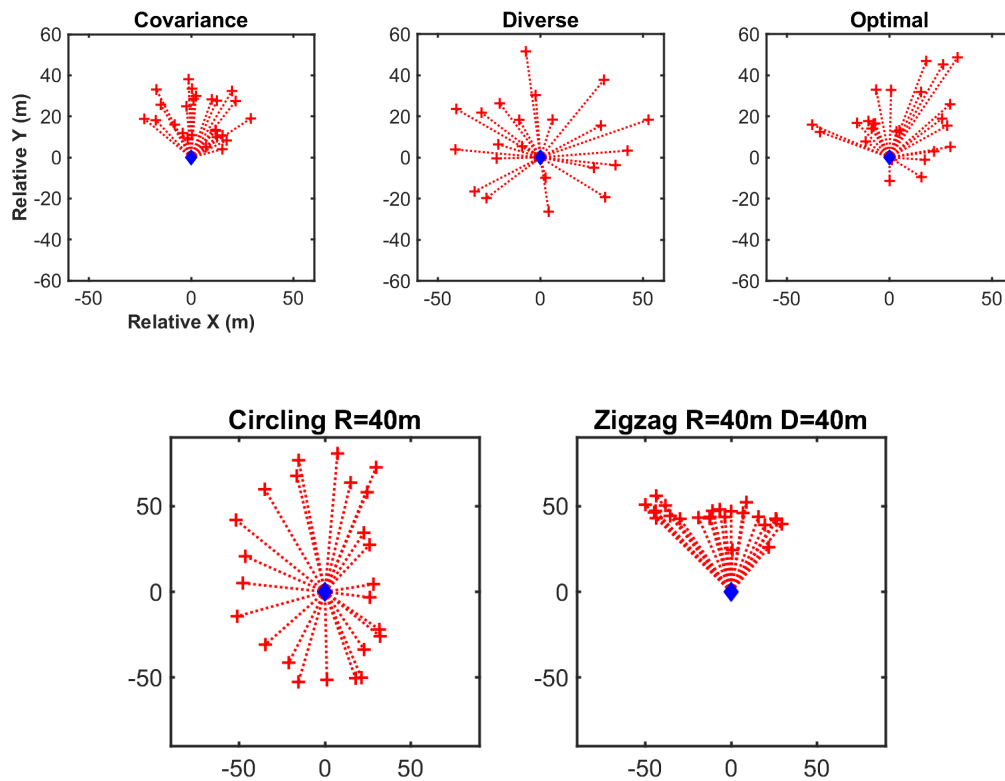


Figure 4-13: Relative leader positions at range measurements for all strategies with the standard mission configuration. The follower is represented as a blue diamond at the origin, and relative leader positions are marked with red plus and connected to the leader to emphasize the relative angle. Note the differences in the distributions of relative angles to the follower, especially between the advanced strategies in the upper row and the geometric strategies in the lower row.

4.5 Comparing Standard Results

Now that we have discussed the trends, benefits, and disadvantages within each strategy, we will compare energy usage and performance across all five strategies for the standard mission configuration.

4.5.1 Energy Usage

The geometric patterns both have a positive relationship between radius and energy usage, with circling resulting in greater energy usage than zigzagging for a given radius. Figure 4-14 also shows that all three advanced strategies operated with a lower energy metric for any circling or zigzag shape. The most energy-efficient strategy according to our energy metric of cubed leader speed is the covariance strategy.




Figure 4-14: Energy usage across strategies for the standard mission configuration, broken down by radius when appropriate. Cubed leader speed is used as the energy metric from Equation 2.24.

4.5.2 Performance

To compare performance between strategies, we can first keep the geometric patterns broken out by radius. Figure 4-15 displays the four error metrics for the five strategies together. The circling strategy outperforms all other strategies at certain radii. Zigzag maneuvering has weaker performance when compared to other strategies. The

advanced strategies perform between or similar to the geometric strategies, typically beating zigzag but lagging or matching circling performance.




Figure 4-15: Performance metrics across strategies based on radius when applicable. The circling strategy performs better than the other strategies, especially for radii around 50 meters. The zigzag strategy performs worse than circling for a given radius but is better than the advanced strategies for some radii.

When mapping performance into the energy usage domain, we find similar comparisons between the strategies. While the advanced strategies provide lower energy usage, the circling strategy offers better performance. The zigzag strategy offers poorer performance than circling at the same energy usage, making circling the more favorable strategy among the geometric patterns.

To compare the five strategies in simpler terms, we use the mean values of the performance metrics for all strategies in Figure 4-33 and Table 4.3. The covariance strategy has the lowest mean value for three of the four error metrics. Circling ties for first in real-time average error and has the lowest value for average smooth error. This reasserts our finding that the covariance strategy is the best of the advanced strategies and that circling is the best of the geometric patterns, with their performances remaining comparable. The covariance strategy is better than circling in terms of energy usage. It should be noted that there are several circling radii that provide better performance than the covariance strategy, even though the mean values of circling error metrics are higher than those of the covariance strategy.

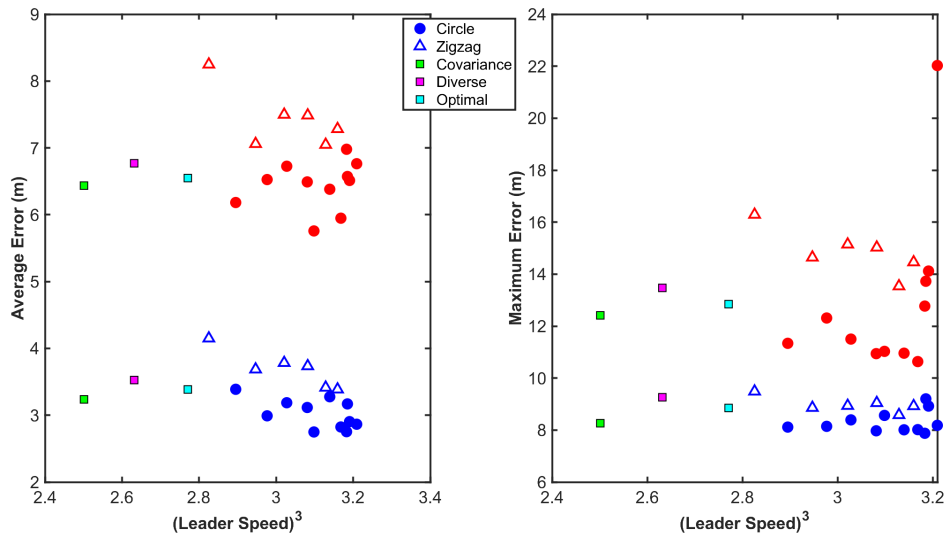


Figure 4-16: Performance across strategies based on energy usage for standard mission configuration. Zigzag has better energy usage when compared to circling but does not match its performance. All adaptive strategies resulted in reduced energy usage but produced performances between that of zigzag and circling.

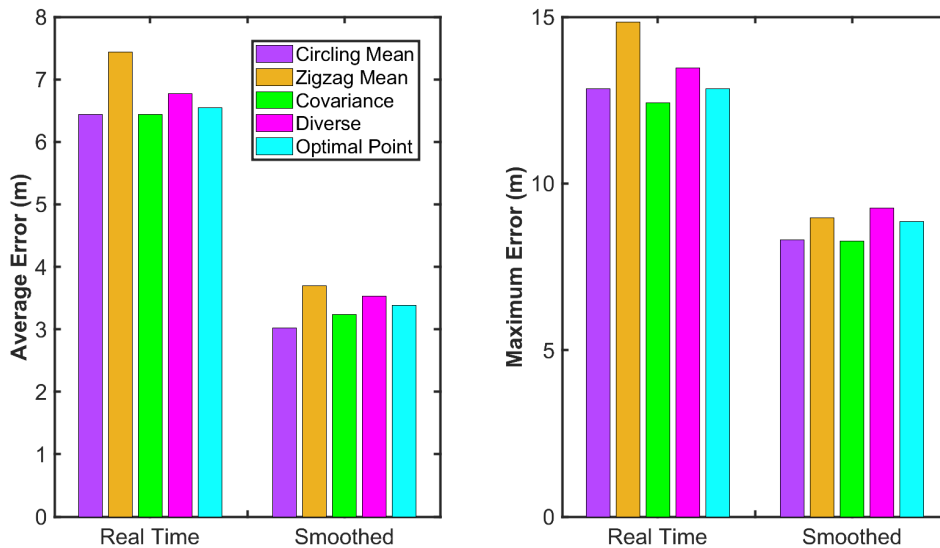


Figure 4-17: Comparison of mean values of performance metrics for all strategies with the standard mission configuration. The circling and covariance strategies perform the best while the zigzag and diverse strategies result in larger error statistics.

Strategy	Real-time Avg	Real-time Max	Smooth Avg	Smooth Max
Circling	6.44	12.85	3.02	8.31
Zigzag	7.44	14.85	3.69	8.97
Covariance	6.44	12.41	3.24	8.26
Diverse	6.77	13.47	3.53	9.26
Optimal	6.55	12.85	3.39	8.85

Table 4.3: Comparing mean values for all performance metrics across all five strategies for the standard mission configuration. The covariance strategy outperforms the other strategies in all but smooth average, where it is trumped by circling. Circling also ties for minimum real-time average error. Note from Figure 4-15 the circling and zigzag strategies have performance that varies with radius; for some radii, circling outperforms covariance. This table is the mean of all radii and effectively mirrors Figure 4-17

4.6 Increased Current Magnitude

In order to test the resilience of these strategies to environmental disturbances, we applied a greater magnitude of currents to both vehicles. An ideal strategy may still suffer reduced performance in these tougher conditions, but we seek to choose a strategy that will limit the effects of increasing current magnitude

One of the most obvious observations of increasing current magnitude is that there is more drift in the follower real-time error between range measurements. The odometry measurements that update the estimated position during this period do not incorporate the impact of currents on the follower. With higher error on the follower’s estimated position, the leader may set up a geometric pattern that is no longer appropriately positioned around the follower’s ground truth position. Figure 4-18 illustrates this issue for a circling strategy with a 30 meter radius.

For radii smaller than 30 meters this issue is seen to an even worse degree and can be categorized as a qualitative disadvantage for circling with smaller radii. Predictability in the geometry of the team is desired for safety during the mission.

To quantitatively show how the radius of a circling maneuver is related to the performance drop from increasing currents, Figure 4-19 shows the change in error metrics from the standard mission trials to increased currents trials.

For zigzag patterns, the effect of increasing current magnitudes is seen most clearly




Figure 4-18: Leader and follower trajectories for a 30 meter circling strategy with increasing current magnitudes. Red lines denote leader trajectory while black and blue denote follower ground truth and real-time paths respectively. For increased current magnitude there is more bend in the follower ground truth trajectory and more instances of erratic real-time error after the follower path is no longer encircled by the leader.

through the lens of the maximum designed relative angle. Figure 4-20 shows all performance metrics for the zigzag strategy at different current magnitudes. Note the large values of maximum real-time error in high current trials for bigger angles. For real-time error metrics, there is no longer a flattening at higher angles but an uptick when stronger currents are applied.

To compactly compare all our maneuvering strategies, we can use the mean values for our performance metrics. Figure 4-21 groups the data by applied current magnitude. For the advanced strategies, there is a greater rise in maximum real-time error. In the standard mission configuration, the covariance strategy performed well; however, with stronger currents, its performance suffers to a greater extent than the other strategies.

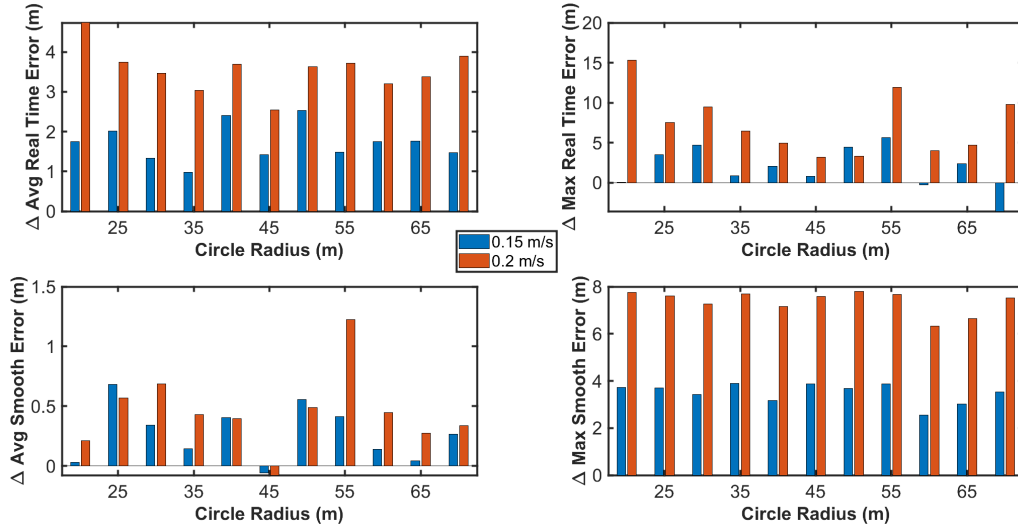


Figure 4-19: Differences in performance metrics for all circling radii as a result of increased current magnitude. Note the largest increases at low radii, especially in real-time error.

Based on this analysis of performance for our maneuvering strategies at higher current magnitudes, the circling strategy is the more robust and better-performing strategy. However, it should be noted that circling at small radii with strong currents present may cause unfavorable geometry; planning for adverse conditions by using a large enough radius is necessary for mission success.

4.7 Varied Current Direction

We carried out several more rounds of trials with a 0.1 m/s current at angles other than the standard 120 degrees. We are interested in how the direction of travel versus the current direction may be related to performance. Namely when the current is almost orthogonal to the follower’s forward progress versus when the current is along or against the follower’s direction of travel. In the MOOS-IvP coordinate system, the follower is attempting to travel at 180 degrees. After collecting runs for all strategies at 45, 240, and 330 degrees, we compared performance and highlighted the main deviations from our standard mission configuration.

For energy usage, varying the current angle did not affect the relationship between

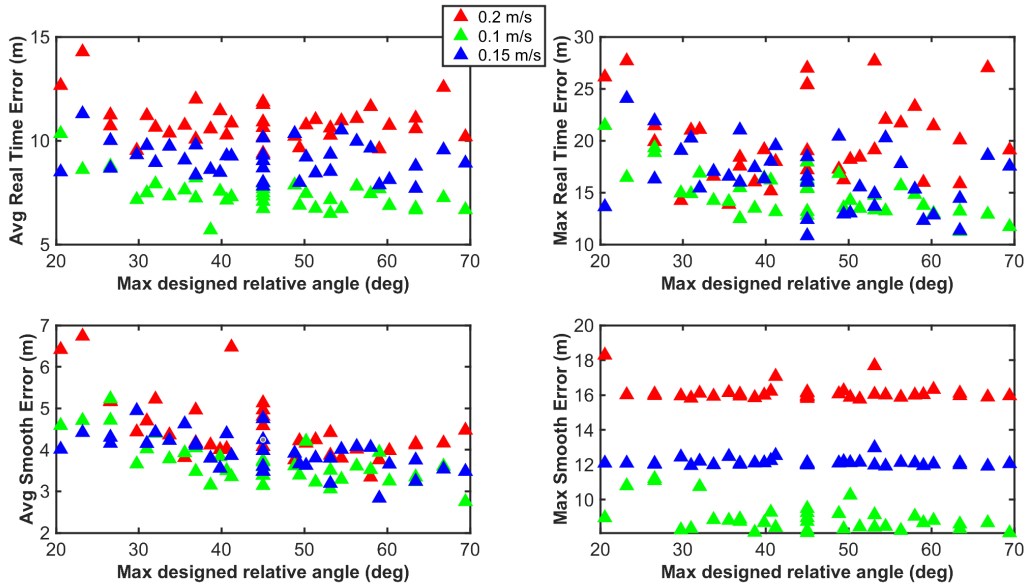


Figure 4-20: Performance of zigzag strategy as a function of maximum designed relative angle for varying magnitudes of applied current. Note the higher values of maximum real-time error for angles above 50 degrees when a 0.2 m/s current is applied.

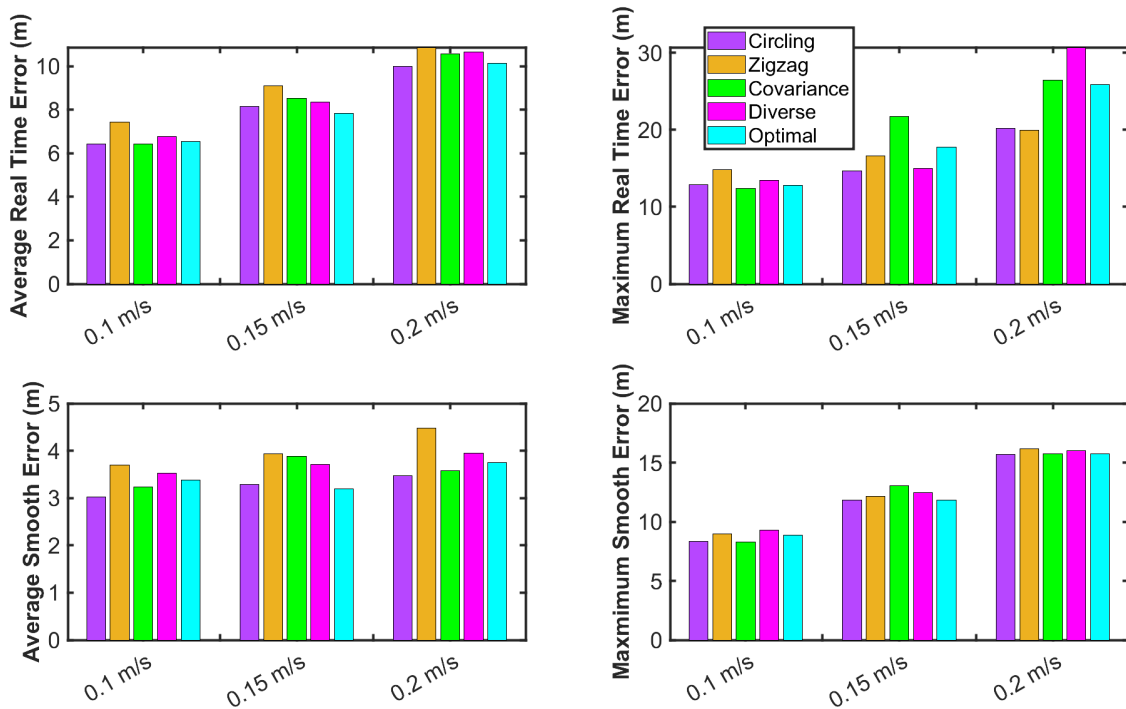


Figure 4-21: Performance of all strategies with increasing current magnitude. Note the larger increase in real-time error for advanced strategies for higher current magnitude.

strategies. Figure 4-22 shows that circling remains the top energy-intensive strategy, followed by zigzag, optimal point, diverse angles, and finally covariance squish. Energy data was not collected for 240 degrees and is not included in this comparison. For our geometric strategies, we have a deviation in energy usage for 120 degrees for lower radii; however, the relationship between energy usage and radii for each geometric pattern is maintained.




Figure 4-22: Energy usage for all strategies with different angles of current applied.

When viewing performance across all strategies for our four current angle options in Figure 4-23, we can see that trials with 330 degree currents had higher maximum errors, especially for the covariance and diverse strategies. The 330 degree configuration has the greatest component of current magnitude against the direction of follower motion. We see the most similar performance across the board between the 120 degrees and 240 degrees configurations. This is expected because the current vectors in these configurations are a reflection of each other across the direction of travel; thus they have the same magnitude of current along and orthogonal to the follower's path.

To consider solely the circling strategy performance for the various current angle configurations, Figure 4-24 provides a more detailed breakdown. For large circle radii, there is worse real-time performance with a 330 degree current vector. Most clearly,

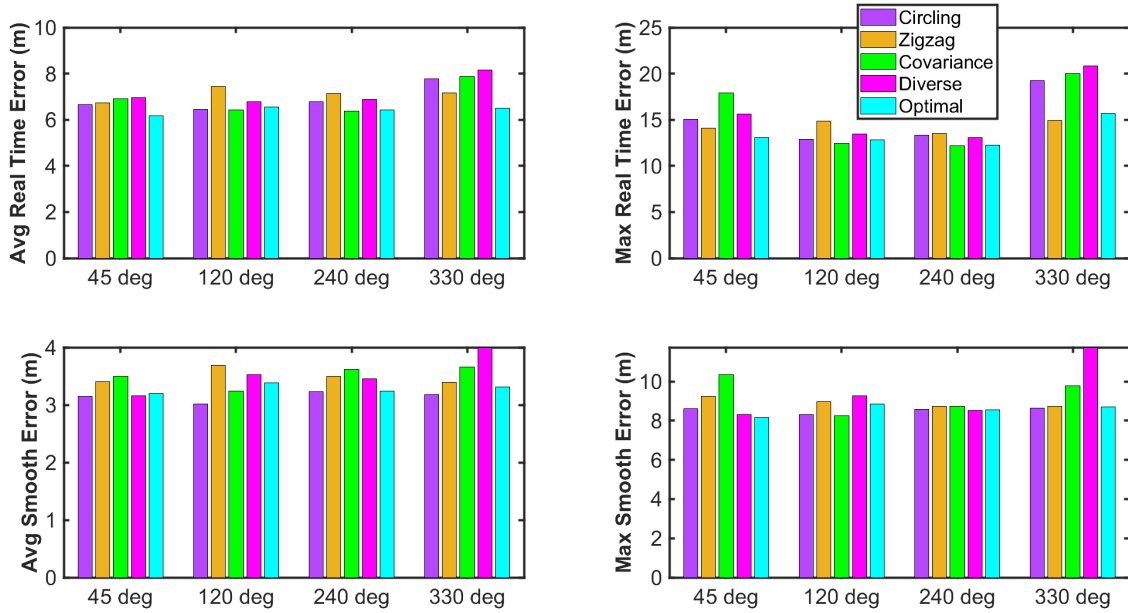


Figure 4-23: Mean values of performance metrics for all strategies with different angles of currents applied. Note marginally worse performance for 330 degree applied current and the very similar spread of performance between the 120 and 240 degree configurations.

the maximum real-time error for radii above 40 meters is similar for 120 and 240 degrees, worse for 45 degrees, and worst for 330 degrees. This pattern can again be attributed to the amount of current magnitude applied along the follower path versus orthogonal to the path.

For a similar analysis of zigzag performance by radius, we find less of a stratification by applied current angle. While there are slight variations in the relationships between radius and performance for the various configurations, there is not a clear pattern.

Two specific deviations from the standard mission configuration were found for zigzag patterns at certain current angles. First, for a 240 degree current angle, a more dramatic relationship between maximum designed relative angle and performance was observed as shown in Figure 4-26. The standard mission configuration with 120 degree current resulted in a spread of error metrics that decreased with larger angles, but flattened after 45 degrees. For the 240 degree currents, there continues to be a

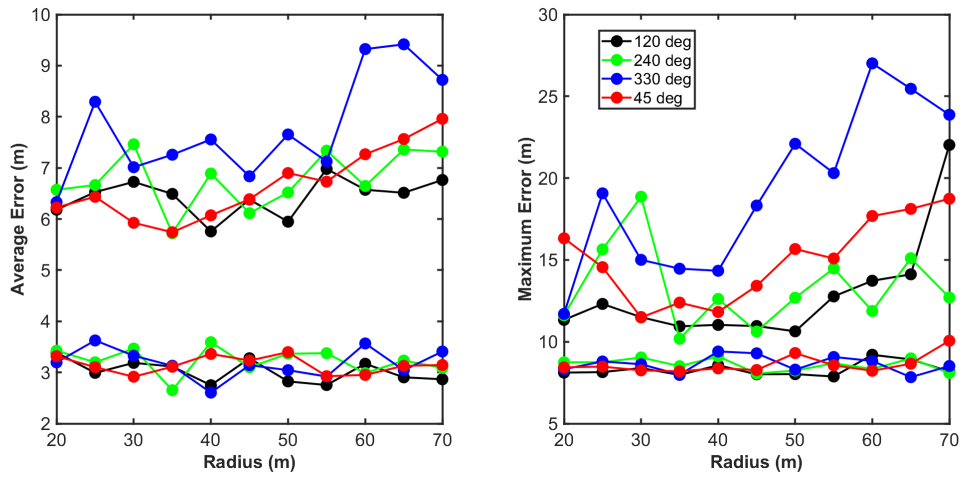


Figure 4-24: Performance metrics for circling as a function of radius with different angles of current applied. Note the uptick in real-time error for larger radii with a 330 degree applied current.

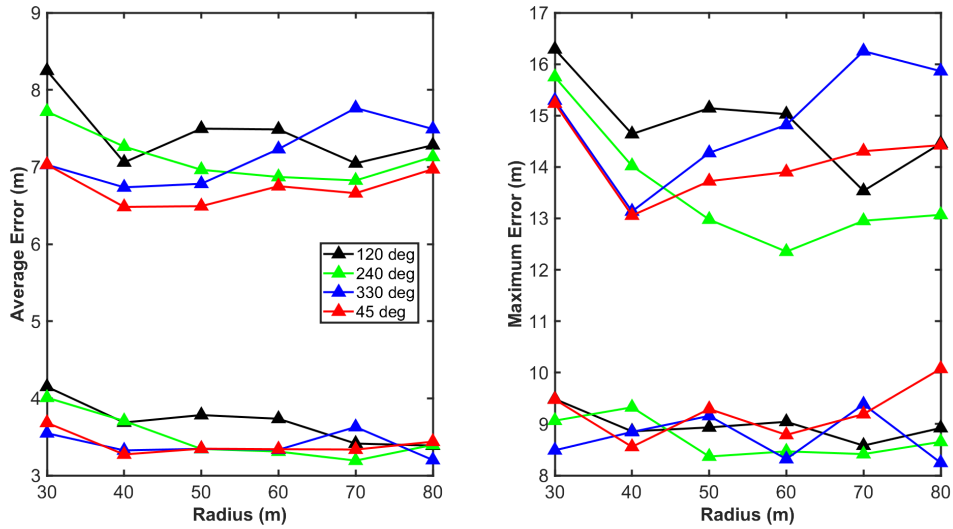


Figure 4-25: Performance metrics for zigzags as a function of radius with different angles of current applied.

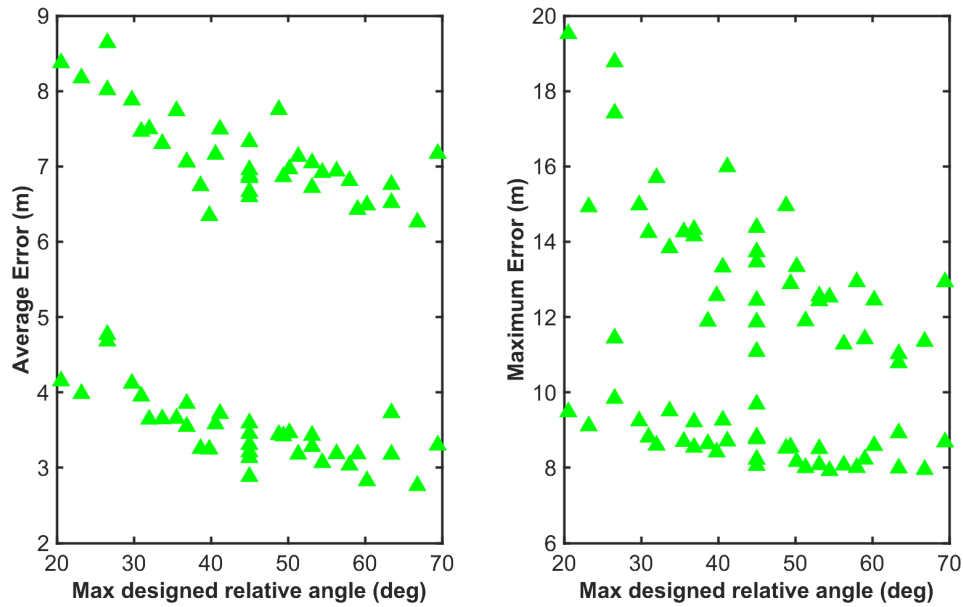


Figure 4-26: Zigzag performance by maximum designed relative angle for 240 degree current applied. Note a stronger and more consistent downward trend when compared to Figure 4-8, especially for angles larger than 45 degrees.

decrease in error metrics for larger angles.

The other deviation from the standard configuration was found for the 330 degrees applied current. Namely, when broken out by trail distance and radius, there was worse performance for larger trail distances, as shown in Figure 4-27.

Overall, changing the angle of applied current created a small shift in the performance of our maneuvering strategies, especially circling, covariance squish, and diverse angles. When planning a mission, it is important to note how the direction of current is oriented in relation to the direction of travel. For the same magnitude, a current vector pointing along the direction of travel may have more of an impact; however, for our geometric patterns the orthogonal component can create unfavorable relative positions if left unchecked.

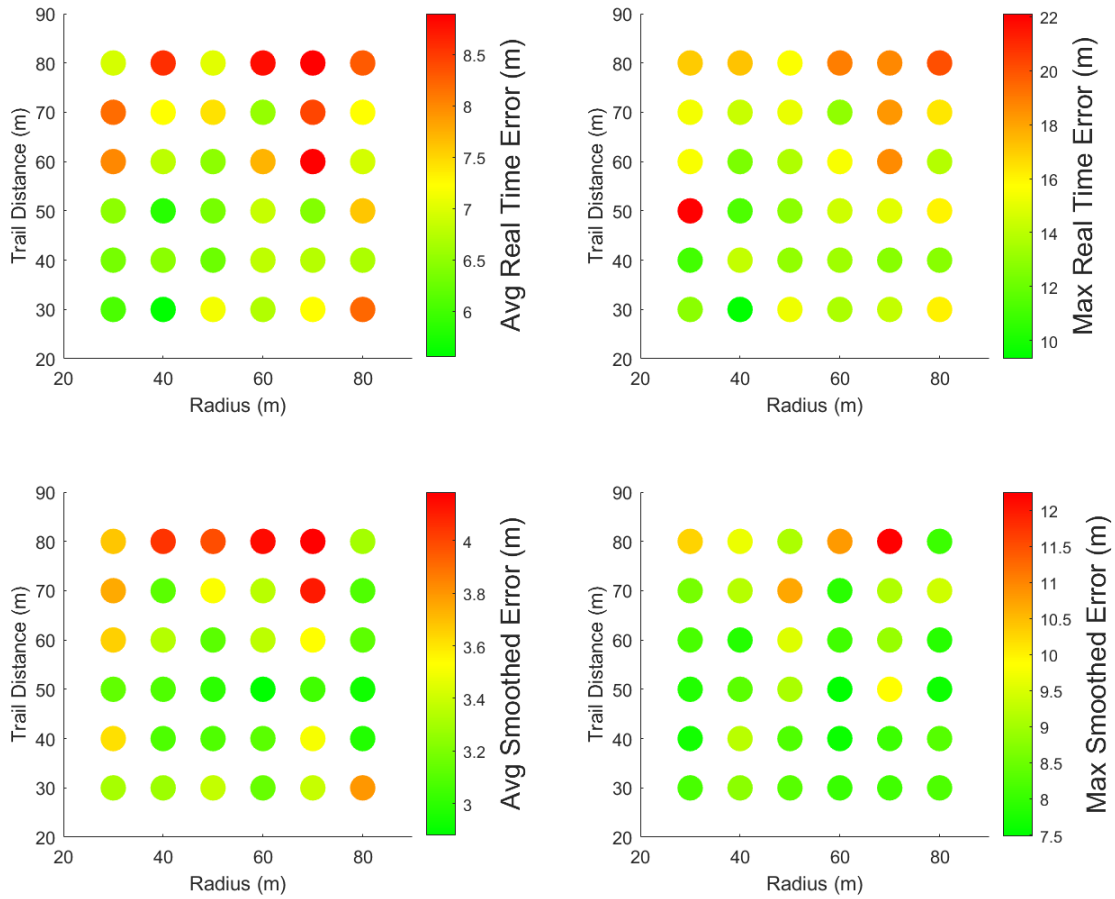



Figure 4-27: Zigzag performance by radius and zigzag trail distance for applied current at 330 degrees. Patterns with a trail distance greater than 60 meters had worse performance for this mission configuration.

4.8 Increased Leader Speed


To expand upon the discussion of energy use versus performance, we allowed our leader vehicle to travel at higher speeds to determine if there was an improvement in the estimation performance. The leader executed the same five strategies, but with a desired speed of 2.0 m/s and 2.5 m/s instead of the standard 1.5 m/s.

For geometric strategies, this allowed the repeating patterns to be more tightly packed in space, as the leader was able to outpace the follower to a higher degree. Figure 4-28 shows leader trajectories at different speeds for circling at a radius of 20 and 70 meters, with markings for leader positions at range measurements. At higher speeds, the leader can complete more circles, but the density of ranging positions around each circle decreases. For a radius of 20 meters, there can be as few as two range measurements per circle. When timed poorly this can lead to a repetitive set of relative ranging positions, effectively cancelling out the benefit that circling brings as a maneuvering strategy. On the other hand, a circle with a radius as large as 70 meters still sends several range measurements along its circumference. A higher speed allows for the change in relative bearing between successive measurements to become larger for circling patterns. Circles with large radii struggled in the standard configuration because there was not enough change in relative bearing between range measurements. Thus, a greater speed impacts the circling strategy positively for large radii and negatively for small radii. Figure 4-29 maps the change in error metrics for the circling strategy from the standard speed of 1.5 m/s to higher speeds. Real-time error metrics best display the relationship between circle radius and speed-based performance changes.

For our advanced strategies, increasing leader speed resulted in similar trajectories with more elongated features. Figure 4-30 compares example trajectories for the covariance, diverse angles, and optimal point strategies for all speeds. The covariance strategy retains its repetitive loops but now travels for a longer distance away from the follower. A similar change is seen for the diverse angles strategy, which has more components of its trajectory that resemble a zigzag path. At a higher speed, the



(a) 20 meter radius



(b) 70 meter radius

Figure 4-28: Distribution of leader ranging locations with increasing speed. Example trajectories for 20 and 70 meter radius circles. With higher leader speed, there is more distance travelled between range locations, leading to less change in relative angle between measurements. For a small radius, this can be a detriment but for a larger radius it can be helpful.

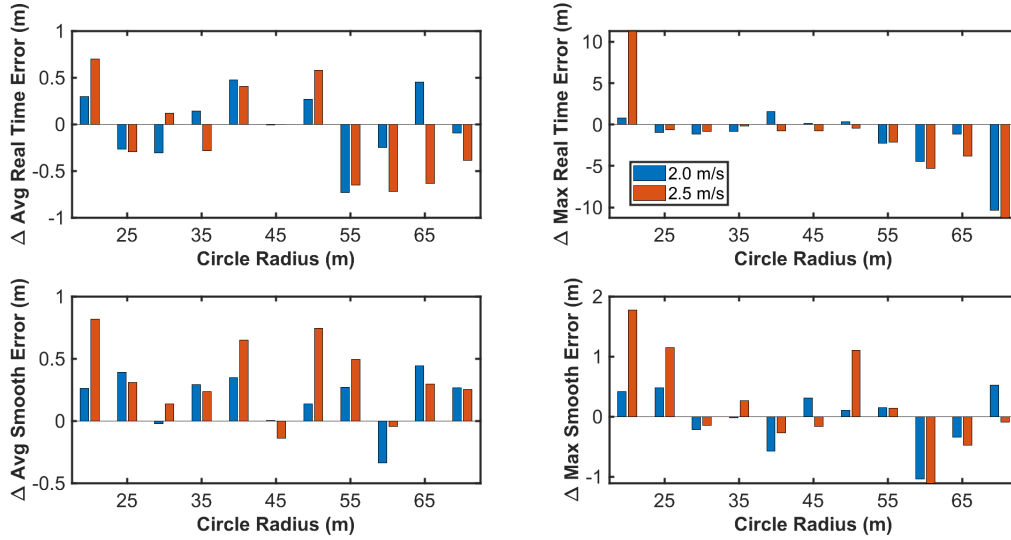



Figure 4-29: Differences in performance metrics for circling strategy with increased speed for all radii. Improved real-time performance was achieved at radii of 55 meters or greater, while smaller radii suffered performance loss to varying degrees.


leader is able to achieve its best relative bearing from the covariance strategy, while the diverse angle strategy can be more ambitious with how many possible relative angles it can consider reachable. The optimal point strategy creates trajectories that cut away in various directions. Since the leader can travel a longer distance between range measurements, there are more reachable points to consider for the optimal point selection.

Figure 4-31 shows the energy metric for all strategies, broken out by radius when possible. The spread of energy usage among the strategies is wider at higher speeds but there is more overlap between the geometric and advanced strategies. Circling remains more energy-intensive than zigzag for any given radius and the ranking of the advanced strategies is not altered. Combining all strategies, we can compare the relationship between energy usage and performance in Figure 4-32.


To compare all strategies against one another simply, we use the mean values of our performance metrics as shown in Figure 4-33. Most notable is the worse performance for the diverse angles strategy with increasing speed. While a higher speed allows the leader to travel farther between range measurements, it also allows for the diverse angle strategy to choose axes that are much farther away.



(a) Covariance strategy



(b) Diverse angle strategy



(c) Optimal point strategy

Figure 4-30: Evolution of leader and follower trajectories for increased leader speed with advanced maneuvering strategies. The leader's ground truth path is shown in red while the follower's ground truth and real-time estimated paths are shown in black and blue respectively. A black x marks the destination.




Figure 4-31: Energy usage metric for all strategies for different desired leader speeds of 1.5, 2.0, and 2.5 m/s. Circling and zigzag are broken out by radius and denoted by circles and triangles respectively. Red icons and dotted lines represent a leader speed of 2.5 m/s, blue icons and dashed lines represent a leader speed of 2.0 m/s, and black icons and solid lines represent a leader speed of 1.5 m/s. The spread of energy usage is greater for higher speeds, but the ordering of the strategies' energy usage is consistent.




Figure 4-32: All strategies' performance metrics as a function of cubed leader speed. Note the larger spread in performance for circling and zigzagging at higher speeds; this signals that the geometric design of these strategies becomes more important for increased speed configurations.

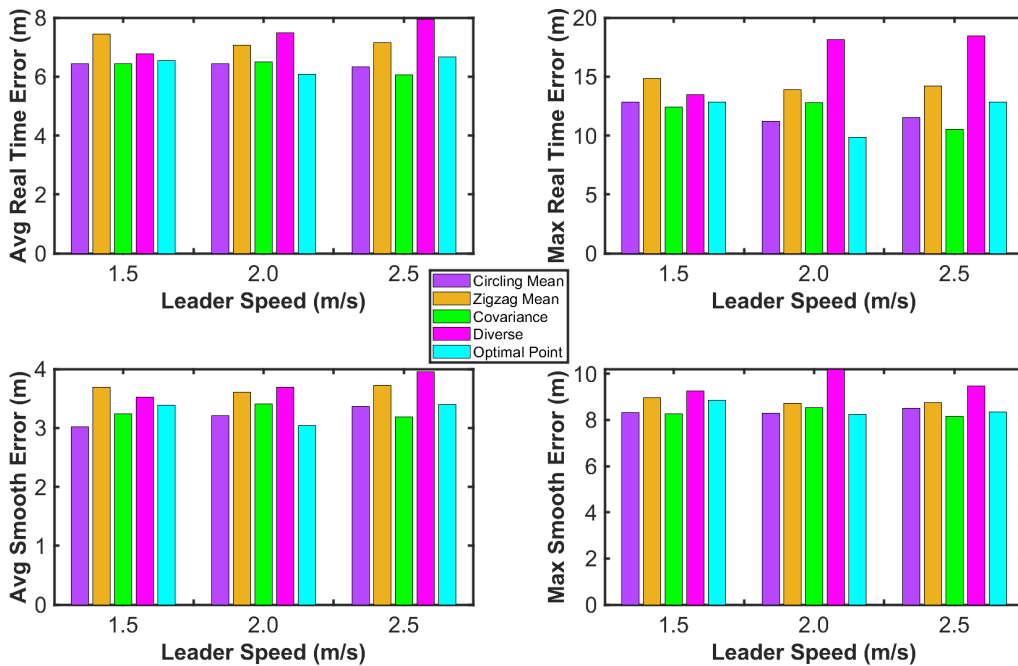


Figure 4-33: All performance metrics for the five strategies with increased leader speed. The diverse angles strategy suffered performance loss with higher speeds, especially in maximum real-time error.

4.9 Other Strategy Considerations

Depending on the mission priorities and specifications, the balance of energy usage and state estimation performance will vary. Beyond the performance and energy usage comparisons between strategies, there are other important considerations for selecting the best strategy for a given mission.

Between the geometric and advanced strategies, the former provides a level of predictability for leader maneuvering; this is especially important if the team involves human divers or if a significant depth offset is not feasible between teammates. The environment and mission goals can also impact the maneuvering strategy selection. For a use case in which the leader is equipped with sensors to monitor the environment, it is advantageous for the leader to be ahead of the follower. The opposite is true if a heterogeneous team is operating in a volatile environment and the lower-cost follower should experience any negative effects instead of the more advanced leader.

Other limitations that may narrow the field of options could include the turning radius of a vehicle. The advanced strategies produced leader trajectories that involved tight loops and turns, which may not be possible for a given vehicle.

Chapter 5

Conclusions

5.1 Contributions

This thesis evaluated several maneuvering strategies for cooperative navigation through data-driven simulation results. After building up a set of simulation tools in MOOS-IvP for state estimation and ranging, we designed five different maneuvering strategies based on previous research efforts. The performance of the follower's state estimation was the main basis for comparison, along with the energy usage of the leader. The circling strategy had the best overall performance but used more energy than other strategies; the radius of circling was also an important consideration for both quantitative and qualitative performance. This evaluation is maintained through comparative simulations with stronger currents and rotated currents. When the leader was allowed to travel at a higher speed, geometric configurations were affected depending on how the relative leader position changed between range measurements.

Even though circling performed the best among our strategies, there is not a great performance loss if the use of a zigzag shape or covariance squish strategy is desired. These other strategies may be better-suited to the mission goals or limitations of the team. Overall, there are several options that provide sufficient cooperative navigation performance and can be tailored to different use cases.

5.2 Future Work

This work can be expanded in the future for use in field trials and further testing in simulation. Field trials would help to highlight the practical limitations of the various maneuvering strategies. Additionally, the conditions in the field will incorporate more of the underwater environmental challenges and further test the robustness of all strategies.

Future efforts could benefit from exploring more advanced current fields in simulation. By incorporating real-world data into our simulation, we can test to see if there is a degradation in performance before deploying in the field. In addition, further ablation testing featuring dropped range messages or different ping frequencies would be meaningful in evaluating the resilience of these methods in adverse conditions.

Ongoing work from O'Neill [33] involves utilizing range updates in the follower's estimated position to derive a current vector in real time. For applications with HAT or followers without a method of measuring current, this is a helpful alternative to reduce the error growth between range measurements.

Future research efforts are necessary to properly tune the noise model for the follower's odometry factors. Depending on the agent acting as the follower and the target environment, care must be taken to appropriately represent the follower's sensors, motion, and external influences.

While the advanced strategies were successfully implemented according to the original design, there are still improvements to be made. Incorporating the effect of currents into the follower's estimated speed can cut down on errors when propagating the follower's position. Additionally, the leader could plan several waypoints ahead to optimize its path over a longer time horizon.

Bibliography

- [1] A. Bahr. *Cooperative Localization for Autonomous Underwater Vehicles*. PhD Thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, February 2009.
- [2] A. Bahr, J. J. Leonard, and M. F. Fallon. Cooperative Localization for Autonomous Underwater Vehicles. *Intl. J. of Robotics Research*, 28(6):714–728, 2009.
- [3] Alexander Bahr and John Leonard. Minimizing Trilateration Errors in the Presence of Uncertain Landmark Positions. January 2007.
- [4] Alexander Bahr, John J. Leonard, and Alcherio Martinoli. Dynamic positioning of beacon vehicles for cooperative underwater navigation. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3760–3767, October 2012. ISSN: 2153-0866.
- [5] J.G. Bellingham and J.S. Willcox. Optimizing AUV oceanographic surveys. In *Proceedings of Symposium on Autonomous Underwater Vehicle Technology*, pages 391–398, June 1996.
- [6] M. Benjamin, M. Grund, and P. Newman. Multi-objective optimization of sensor quality with efficient marine vehicle task execution. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pages 3226–3232, May 2006. ISSN: 1050-4729.
- [7] M. R. Benjamin, J. J. Leonard, H. Schmidt, and P. Newman. An Overview of MOOS-IvP and a Brief Users Guide to the IvP Helm Autonomy Software. Technical Report CSAIL-2009-028, Computer Science and Artificial Intelligence Laboratory, MIT, June 2009.
- [8] Michael Benjamin, Paul Newman, Henrik Schmidt, and John Leonard. Extending a MOOS-IvP Autonomy System and Users Guide to the IvPBuild Toolbox. August 2009.
- [9] Michael R. Benjamin. The Interval Programming Model for Multi-objective Decision Making. Technical report, September 2004. Accepted: 2005-12-19T23:28:02Z.

- [10] J. Curcio, J. Leonard, J. Vaganay, A. Patrikalakis, A. Bahr, D. Battle, H. Schmidt, and M. Grund. Experiments in moving baseline navigation using autonomous surface craft. In *Proceedings of OCEANS 2005 MTS/IEEE*, pages 730–735 Vol. 1, September 2005. ISSN: 0197-7385.
- [11] Thomas Curtin, James Bellingham, Josko Catipovic, and Doug Webb. Autonomous Oceanographic Sampling Networks. *Oceanography*, 6(3):86–94, 1993.
- [12] Frank Dellaert and GTSAM Contributors. borglab/gtsam, May 2022.
- [13] Frank Dellaert, Michael Kaess, and others. Factor graphs for robot perception. *Foundations and Trends[®] in Robotics*, 6(1-2):1–139, 2017. Publisher: Now Publishers, Inc.
- [14] M. F. Fallon, G. Papadopoulos, J. J. Leonard, and N. M. Patrikalakis. Cooperative AUV Navigation using a Single Maneuvering Surface Craft. *Intl. J. of Robotics Research*, 29(12):1461–1474, October 2010.
- [15] E. Gallimore, J. Partan, I. Vaughn, S. Singh, J. Shusta, and L. Freitag. The WHOI micromodem-2: A scalable system for acoustic communications and networking. In *Proc. of the IEEE/MTS OCEANS Conf. and Exhibition*, pages 1–7, September 2010.
- [16] Shir Goldfinger. sgoldfin/moos-ivp-shirg, 2022.
- [17] J. Carl Hartsfield. *Single Transponder Range Only Navigation Geometry (STRONG) applied to REMUS autonomous under water vehicles*. Thesis, Massachusetts Institute of Technology, 2005. Accepted: 2007-10-19T20:26:40Z.
- [18] O. Hegrenaes, E. Berglund, and O. Hallingstad. Model-aided inertial navigation for underwater vehicles. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1069–1076, May 2008.
- [19] iXblue. Phins Compact Series.
- [20] Finn B. Jensen, William A. Kuperman, Michael B. Porter, and Henrik Schmidt. *Computational Ocean Acoustics*. Modern Acoustics and Signal Processing. Springer, New York, NY, 2011.
- [21] Jeff Kaeli, Gwyneth Packard, Carl Fiester, Dara Tebo, and Mike Purcell. A Range-Only Acoustic Localization Technique for Marine Robotics Applications. In *OCEANS 2019 MTS/IEEE SEATTLE*, pages 1–5, October 2019. ISSN: 0197-7385.
- [22] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert. iSAM2: Incremental Smoothing and Mapping with Fluid Relinearization and Incremental Variable Reordering. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, May 2011.

- [23] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert. iSAM2: Incremental Smoothing and Mapping Using the Bayes Tree. *The International Journal of Robotics Research*, 31(2):217–236, February 2012.
- [24] M. Kaess, A. Ranganathan, and F. Dellaert. iSAM: Incremental Smoothing and Mapping. *IEEE Trans. Robotics*, 24(6):1365–1378, December 2008.
- [25] D. B. Kilfoyle and A. B. Baggeroer. The Current State-of-the-Art in Underwater Acoustic Telemetry. *IEEE J. Ocean Engineering*, 25(1):4–27, 2000.
- [26] J. J. Leonard and H. F. Durrant-Whyte. Mobile Robot Localization by Tracking Geometric Beacons. *IEEE Trans. Robotics and Automation*, 7(3):376–382, June 1991.
- [27] Tapovan Lolla, Pierre F. J. Lermusiaux, Mattheus P. Ueckermann, and Patrick J. Haley. Time-optimal path planning in dynamic flows using level set equations: theory and schemes. *Prof. Lermusiaux via Angie Locknar*, September 2014. Accepted: 2015-09-01T13:34:45Z Publisher: Springer-Verlag.
- [28] I. Masmitja, S. Gomariz, J. Del Rio, B. Kieft, and T. O’Reilly. Range-only underwater target localization: Path characterization. In *OCEANS 2016 MTS/IEEE Monterey*, pages 1–7, September 2016.
- [29] Nikola Mivskovic, Antonio Pascoal, Marco Bibuli, Massimo Caccia, Jeffrey A Neasham, Andreas Birk, Murat Egi, Karl Grammer, Alessandro Marroni, Antonio Vasiljević, and others. CADDY project, year 3: The final validation trials. In *Oceans 2017-aberdeen*, pages 1–5. IEEE, 2017.
- [30] M. Morgado, Paulo Oliveira, and Carlos Silvestre. Tightly coupled ultrashort baseline and inertial navigation system for underwater vehicles: An experimental validation. *Journal of Field Robotics*, 30:142–170, February 2013.
- [31] Đula Nađ, Christopher Walker, Igor Kvasic, Derek Orbaugh, Nikola Miskovic, Iain Anderson, and Ivan Lončar. Towards Advancing Diver-Robot Interaction Capabilities. *IFAC-PapersOnLine*, 52:199–204, January 2019.
- [32] P. M. Newman. MOOS - A Mission Oriented Operating Suite. Technical Report OE2003-07, Department of Ocean Engineering, MIT, 2003.
- [33] Brendan W. O’Neill, Jesse R. Pelletier, S. Calvert, Alan Papalia, John J. Leonard, Lee Freitag, and Eric Gallimore. Loosely-Coupled Human-Robot Teams for Enhanced Undersea Operations. In *OCEANS 2022, Hampton Roads*, pages 1–8, October 2022. ISSN: 0197-7385.
- [34] Geōrgios Papadopoulos. *Underwater vehicle localization using range measurements*. Thesis, Massachusetts Institute of Technology, 2010. Accepted: 2011-04-25T16:14:36Z.

- [35] Alan Papalia, Nicole Thumma, and John Leonard. Prioritized Planning for Cooperative Range-Only Localization in Multi-Robot Networks. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 10753–10759, May 2022.
- [36] Liam Paull, Sajad Saeedi, Mae Seto, and Howard Li. AUV navigation and localization: A review. *Oceanic Engineering, IEEE Journal of*, 39(1):131–149, 2014. Publisher: IEEE.
- [37] Jesse Pelletier. *Human-Autonomy Teaming for Improved Diver Navigation*. Thesis, Massachusetts Institute of Technology, February 2022. Accepted: 2022-06-15T13:02:56Z.
- [38] Jesse R. Pelletier, Brendan W. O’Neill, John J. Leonard, Lee Freitag, and Eric Gallimore. AUV-Assisted Diver Navigation. *IEEE Robotics and Automation Letters*, 7(4):10208–10215, October 2022. Conference Name: IEEE Robotics and Automation Letters.
- [39] M. J. D. Powell. A New Algorithm for Unconstrained Optimization. In J. Rosen, O. Mangasarian, and K. Ritter, editors, *Nonlinear Programming*, pages 31–65. ap, 1970.
- [40] Jake D. Quenzer and Kristi A. Morgansen. Observability based control in range-only underwater vehicle localization. In *2014 American Control Conference*, pages 4702–4707, June 2014. ISSN: 2378-5861.
- [41] Anwar Quraishi, Alexander Bahr, Felix Schill, and Alcherio Martinoli. A Flexible Navigation Support System for a Team of Underwater Robots. In *2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, pages 70–75, August 2019.
- [42] S. A. T. Randeni P, A.L. Forrest, R. Cossu, Z. Q. Leong, and D. Ranmuthugala. Estimating flow velocities of the water column using the motion response of an Autonomous Underwater Vehicle (AUV). In *OCEANS 2015 - MTS/IEEE Washington*, pages 1–6, October 2015.
- [43] A. J. Shafer, M. R. Benjamin, J. J. Leonard, and J. Curcio. Autonomous cooperation of heterogeneous platforms for sea-based search tasks. In *OCEANS 2008*, pages 1–10. IEEE, 2008.
- [44] H. Singh, D. Yoerger, and A. Bradley. Issues in AUV design and deployment for oceanographic research. In *Proceedings of International Conference on Robotics and Automation*, volume 3, pages 1857–1862 vol.3, April 1997.
- [45] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. The MIT Press, Cambridge, MA, 2005.

- [46] Nicole Thumma. *Distributed Potential Field Method for Cooperative Range-Only Localization in Multi-Robot Networks*. Thesis, Massachusetts Institute of Technology, February 2022. Accepted: 2022-06-15T13:02:53Z.
- [47] J. Vaganay, P. Baccou, and B. Jouvencel. Homing by acoustic ranging to a single beacon. In *IEEE Oceans*, pages 1457–1462, 2000.
- [48] J. Vaganay, J. J. Leonard, J. A. Curcio, and J. S. Willcox. Experimental Validation of the Moving Long Base Line Navigation Concept. In *Autonomous Underwater Vehicles, 2004 IEEE/OES*, pages 59–65, June 2004.
- [49] Jeffrey Scott Willcox. *Oceanographic surveys with autonomous underwater vehicles : performance metrics and survey design*. Thesis, Massachusetts Institute of Technology, 1998. Accepted: 2009-12-10T18:57:21Z.
- [50] K.X. Zhou and S.I. Roumeliotis. Optimal motion strategies for range-only distributed target tracking. In *2006 American Control Conference*, pages 6 pp.–, June 2006. ISSN: 2378-5861.