

Dynamic Action Spaces for Autonomous Search Operations

by

Caleb A. Earnest

B.S. Operations Research
United States Air Force Academy, 2003

SUBMITTED TO THE DEPARTMENT OF CIVIL AND ENVIRONMENTAL ENGINEERING IN
PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN TRANSPORTATION
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2005

© 2005 Caleb A. Earnest. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part.

Signature of Author: _____
Department of Civil and Environmental Engineering
May 6th, 2005

Certified by: _____
Liyi Dai
Member of the Technical Staff, The Charles Stark Draper Laboratory, Inc.
Thesis Supervisor

Certified by: _____
Lance A. Page
Member of the Technical Staff, The Charles Stark Draper Laboratory, Inc.
Thesis Supervisor

Certified by: _____
Nicholas Roy
Assistant Professor of Aeronautics and Astronautics
Thesis Advisor

Certified by: _____
Cynthia Barnhart
Professor of Civil and Environmental Engineering and Engineering Systems
Thesis Reader

Accepted by: _____
Andrew J. Whittle
Chairman, Departmental Committee on Graduate Students

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 00 JUN 2005		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Dynamic Action Spaces for Autonomous Search Operations				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute Of Technology				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 150	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

[This Page Intentionally Left Blank]

Dynamic Action Spaces for Autonomous Search Operations

by

Caleb Aaron Earnest

Submitted to the Department of Civil and Environmental Engineering on May 6th, 2005 in
Partial Fulfillment of the Requirements for the Degree of
Master of Science in Transportation

Abstract

This thesis presents a new approach for a Navy unmanned undersea vehicle (UUV) to search for and detect an evading contact. This approach uses a contact position distribution from a generic particle filter to estimate the state of a single moving contact and to plan the path that minimizes the uncertainty in the location of the contact. The search algorithms introduced in this thesis will implement a motion planner that searches for a contact with the following information available to the decision system: (1) null measurement (i.e., contact not detected at current time), (2) time-dated measurement (i.e., clue found at current time that indicates contact was at this location in the past), and (3) bearings measurement (i.e., angular measurement towards contact position detected at current time). The results of this thesis will be arrived at by evaluating the best methods to utilize the three types of information. The underlying distribution of the contact state space will be modeled using a generic particle filter, due to the highly non-Gaussian distributions that result from the conditions mentioned above. Using the particle filter distribution and the measurements acquired from the three conditions, this thesis will work towards implementing a path planning algorithm that creates dynamic action spaces that evaluate the uncertainty of position distribution. Ultimately, the path planner will choose the path that contains the position distribution and leads to sustained searches.

Thesis Supervisor: Liyi Dai

Title: Member of the Technical Staff, The Charles Stark Draper Laboratory, Inc.

Thesis Supervisor: Lance Page

Title: Member of the Technical Staff, The Charles Stark Draper Laboratory, Inc.

Thesis Advisor: Nicholas Roy

Title: Assistant Professor of Aeronautics and Astronautics

Thesis Reader: Cynthia Barnhart

Title: Professor of Civil and Environmental Engineering and Engineering Systems

[This Page Intentionally Left Blank]

ACKNOWLEDGEMENTS

I would like to thank all of those individuals who have helped me throughout my life, including family, friends, and teachers. To the many individuals that I owe thanks for my two years at Draper Laboratory and MIT:

First, I thank Professor Nick Roy of MIT for his invaluable role in the completion of this thesis. I owe him many thanks for his dedication to my research. His expertise and insight were such a valuable source of guidance. Without the extra time he was willing to commit to my research, my thesis research would not have been the same.

I must also thank William Kreamer for his wisdom and enthusiasm in guiding me in my research. During his time as my Draper advisor, he always found time to answer any question that I had despite his workload. He introduced me to the research process and laid the foundation for the research presented in my thesis. I also thank him for finding advisors interested in my research to take his place upon his departure from Draper Laboratory.

I am equally thankful for Lance Page and Liyi Dai of Draper Laboratory. I am truly grateful for them taking on the role as my Draper supervisors. Their patience and understanding were greatly appreciated.

I want to also thank Professor Cindy Barnhart for her help. I am greatly appreciative of her willingness to donate her limited time to my thesis.

I would also like to thank the Education Office at Draper Laboratory, the Center for Transportation and Logistics at MIT, and the United States Air Force for the opportunity to continue my education at Draper Laboratory and MIT. In particular, I am grateful for Lt Col Andrew Armacost at the United States Air Force Academy and Dr. Steve Kolitz at Draper Laboratory for their efforts in helping me secure this assignment.

To all of my friends that I met here in Boston that made my two years here so enjoyable, I also say thank you. Without these friendships, I would not have grown and matured through many of the challenges and opportunities life has presented me over the past two years.

Last but not least, I would like to thank my parents. They have provided constant love and support throughout all my life. Through their faith, guidance, and encouragement, I have been able to accomplish so much.

The thesis was prepared at The Charles Stark Draper Laboratory, Inc. under Contract N00014-02-C-0191 sponsored by the Office of Naval Research.

Publication of this thesis does not constitute approval by Draper or the sponsoring agency of the findings or conclusions contained therein. It is published for the exchange and stimulation of ideas.

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or The U.S. Government.

Caleb A. Earnest, 2d Lt, USAF

Date

ASSIGNMENT

Draper Laboratory Report Number T-1515

In consideration for the research opportunity and permission to prepare my thesis by and at The Charles Stark Draper Laboratory, Inc., I hereby assign my copyright of the thesis to The Charles Stark Draper Laboratory, Inc., Cambridge, Massachusetts.

Caleb A. Earnest, 2d Lt, USAF

Date

TABLE OF CONTENTS

LIST OF FIGURES	10
LIST OF TABLES	12
CHAPTER 1	
INTRODUCTION.....	13
1.1 Problem Motivation	13
1.2 Problem Statement	16
1.3 Problem Approach	17
1.3.1 Search and Detection Theory.....	17
1.3.2 Guaranteed Detection Conditions.....	18
1.3.3 Current Search Methods	20
1.3.4 Application of Search Methods	22
1.4 Contributions	23
1.5 Organization	24
CHAPTER 2	
SUBMARINE TRACK AND TRAIL OPERATIONAL DESCRIPTION	25
2.1 Concept of Operations	25
2.2 Scenario Description	27
2.2.1 Description of Search Environment.....	28
2.2.2 Description of Contact Dynamics.....	30
2.2.3 Description of UUV Mission.....	31
CHAPTER 3	
STATE ESTIMATION	35
3.1 State Estimation Models	35
3.1.1 Motion Model.....	36
3.1.2 Observation Model	37
3.2 Recursive Bayesian Estimation Techniques.....	37
3.2.1 Kalman Filters	38
3.2.2 Particle Filters.....	40
3.2.2.1 State Parameters.....	41
3.2.2.2 Generic Particle Filter Algorithm.....	42

3.3	Comparison of State Estimation Techniques.....	46
------------	---	-----------

CHAPTER 4

CONTACT STATE ESTIMATION USING PASSIVE SENSOR TECHNOLOGY	49
---	-----------

4.1	UUV Sensor Technology.....	49
4.1.1	Range and Bearing Measurements	50
4.1.2	Time-Dated Position Measurements.....	50
4.2	Particle Filter Implementation.....	55
4.2.1	Particle Description	55
4.2.1.1	Descriptive State Parameters.....	56
4.2.1.2	Particle Weights	58
4.2.2	Contact Motion Model.....	58
4.2.3	Observation Model	60
4.2.3.1	Stage 1: Null Measurement Model	60
4.2.3.2	Stage 2: Time-Dated Measurement Model	63
4.2.3.3	Stage 3: Bearings-Only Measurement Model	67

CHAPTER 5

MOTION PLANNING FOR DYNAMIC SEARCH OPERATIONS	69
--	-----------

5.1	Dynamic Decision System Architecture	69
5.2	Motion Planning Formulation.....	74
5.2.1	Objective Function	75
5.2.2	Decision Space	78
5.3	Motion Planning Strategy.....	79
5.3.1	Action Space Generation	79
5.3.1.1	Discrete Action Space.....	79
5.3.1.2	Cluster-Based Action Space.....	80
5.3.2	Monte Carlo Simulation of Search Environment.....	83
5.3.2.1	Action Simulation	83
5.3.2.2	Particle Propagation	94
5.3.2.3	Measurement Update	94
5.3.3	Path Enumeration	95
5.3.3.1	Path Depth Selection.....	96
5.3.3.2	Enumeration Methods.....	97
5.4	Path Evaluation	100

CHAPTER 6

SEARCH AND DETECTION SIMULATION RESULTS AND ANALYSIS	101
---	------------

6.1	General Simulation Design Process	101
6.1.1	Computer Simulation.....	102
6.1.2	Simulation Scenarios	102
6.1.3	Simulation Design Factors.....	103

6.2	Search Algorithm Performance Comparisons.....	104
6.2.1	Base Case Search Algorithm Analysis	105
6.2.1.1	Base Case Search Algorithm.....	105
6.2.1.2	Base Case Performance Evaluation.....	106
6.2.2	Path Planning Analysis.....	112
6.2.2.1	Action Space Analysis	112
6.2.2.2	Search Depth Analysis	121
6.2.2.3	Resampling Frequency Analysis.....	126
6.2.2.4	Dynamic Path Planning Algorithm Conclusions.....	129
6.2.3	Time-Dated Measurement Analysis	131
6.2.3.1	Base Case Search Algorithm.....	133
6.2.3.2	Dynamic Action Space Search Algorithm	134
6.2.3.3	Time-Dated Measurements Conclusions	139

CHAPTER 7

	CONCLUSIONS AND FUTURE WORK	145
7.1	Thesis Contributions	145
7.2	Future Work.....	147
	REFERENCES.....	148

LIST OF FIGURES

FIGURE 1.1: GUARANTEED DETECTION CONDITIONS FOR SWEEP MANEUVER	19
FIGURE 1.2: GUARANTEED DETECTION CONDITIONS FOR SPIRAL-IN MANEUVER.....	20
FIGURE 1.3: EFFECT OF DETECTED CLUE	21
FIGURE 2.1: POTENTIAL SENSOR NETWORK PLACEMENT.....	26
FIGURE 2.2: SUBMARINE TRACK AND TRAIL OPERATIONS	27
FIGURE 2.3: SCENARIO MAP.....	29
FIGURE 2.4: KEY ENVIRONMENTAL CONSTRAINTS OF SCENARIO MAP	29
FIGURE 2.5: DETAILED DESCRIPTION OF CONTACT DYNAMICS	31
FIGURE 2.6: DETAILED DESCRIPTION OF INITIAL STAGES OF SCENARIO.....	32
FIGURE 3.1: KALMAN FILTER CYCLE	40
FIGURE 3.2: RECURSIVE PARTICLE FILTER ALGORITHM.....	42
FIGURE 3.3: PARTICLE FILTER ILLUSTRATION	45
FIGURE 4.1: CLUE DEPOSIT FREQUENCY	51
FIGURE 4.2: EXPONENTIAL DECAY OF CLUE SIGNAL STRENGTH.....	52
FIGURE 4.3: SENSOR DETECTION RATE - PROBABILITY OF CLUE DETECTION GIVEN THAT A CLUE EXISTS	53
FIGURE 4.4: UNCERTAINTY IN THE CLUE POSITION	54
FIGURE 4.5: PROBABILITY DISTRIBUTION OF CONTACT POSITION AFTER DETECTION OF CLUE	54
FIGURE 4.6: PARTICLE ARRAY ILLUSTRATION.....	57
FIGURE 4.7: CONTACT MOTION MODEL FOR SUBMARINE SEARCH AND DETECTION SCENARIO	59
FIGURE 4.8: ILLUSTRATION OF NULL MEASUREMENT MODEL	61
FIGURE 4.9: CLUE EFFECT ON CONTACT POSITION DISTRIBUTION	64
FIGURE 4.10: ILLUSTRATION OF TIME-DATED POSITION MEASUREMENT MODEL.....	65
FIGURE 4.11: CONTACT POSITION DISTRIBUTION IN THE AREAS SURROUNDING THE CLUE LOCATION DUE TO THE UNCERTAINTY IN THE CLUE LOCATION AND THE ASSUMED CONTACT MOTION MODEL.....	66
FIGURE 4.12: ILLUSTRATIVE EXAMPLE OF BEARINGS MEASUREMENTS	68
FIGURE 5.1: HIGH LEVEL UUV DYNAMIC DECISION SYSTEM	72
FIGURE 5.2: DYNAMIC PATH GENERATION FOR UUV DECISION SYSTEM	73
FIGURE 5.3: CLUSTERING ALGORITHM.....	81
FIGURE 5.4: ILLUSTRATION OF CLUSTERED PARTICLE FILTER DISTRIBUTION.....	82
FIGURE 5.5: CONSTRAINING ACTION SPACE ALGORITHM	85
FIGURE 5.6: ILLUSTRATIVE EXAMPLE OF CONSTRAINING ACTION SPACE ALGORITHM.....	86
FIGURE 5.7: DISCRETE CELL DESCRIPTION OF SEARCH ENVIRONMENT	88
FIGURE 5.8: NUMERICAL POTENTIAL FIELD ALGORITHM.....	89
FIGURE 5.9: NUMERICAL POTENTIAL FIELD FOR SEARCH ENVIRONMENT	90
FIGURE 5.10: NUMERICAL POTENTIAL FIELD FOR COMPUTATION OF SHORTEST PATH	91

FIGURE 5.11: PATH ENUMERATION ILLUSTRATION	95
FIGURE 5.12: BREADTH-FIRST PATH ENUMERATION ILLUSTRATION.....	98
FIGURE 5.13: DEPTH-FIRST PATH ENUMERATION ILLUSTRATION	99
FIGURE 6.1: SECTOR-BASED ACTION SPACE	105
FIGURE 6.2: CONTACT DETECTION USING SECTOR-BASED SEARCH	106
FIGURE 6.3: ILLUSTRATION OF POOR SEARCH PERFORMANCE – (3) CREATION OF MULTI-MODAL DISTRIBUTION, (6) OSCILLATION BETWEEN MODES OF A DISTRIBUTION	107
FIGURE 6.4: OSCILLATIONS BETWEEN MODES OF A DISTRIBUTION	108
FIGURE 6.5: BASE CASE ALGORITHM SEQUENCE FOR SCENARIO 2	110
FIGURE 6.6: SEARCH DEPTH NEEDED FOR TOTAL COVERAGE WITH A DISCRETE ACTION SPACE.....	114
FIGURE 6.7: LIMITED COVERAGE AREA USING A DISCRETE ACTION SPACE	114
FIGURE 6.8: AREAS WITH NO INFORMATION EXPLORED BY A DISCRETE ACTION SPACE	115
FIGURE 6.9: DYNAMIC ACTION SPACE ILLUSTRATION	116
FIGURE 6.10: INITIAL SEQUENCE OF ACTIONS USING THE DYNAMIC ACTION SPACE PATH PLANNER	117
FIGURE 6.11: CONTINUED SEQUENCE OF DYNAMIC ACTIONS LEADING TO DETECTION.....	118
FIGURE 6.12: PATH PLANNING SEARCH ALGORITHM (SEARCH DEPTH = 1)	119
FIGURE 6.13: PATH PLANNING SEARCH ALGORITHM (SEARCH DEPTH = 1)	120
FIGURE 6.14: PATH PLANNING SEARCH ALGORITHM (SEARCH DEPTH = 2), ...	123
FIGURE 6.15: PATH PLANNING SEARCH ALGORITHM (SEARCH DEPTH = 3), ...	125
FIGURE 6.16: ILLUSTRATION OF POOR SEARCH PERFORMANCE DUE TO LOW RESAMPLING FREQUENCY	127
FIGURE 6.17: IMPROVED SEARCH PERFORMANCE DUE TO INCREASED SAMPLING FREQUENCY	128
FIGURE 6.18: COMPARISON OF AVERAGE TIMES UNTIL DETECTION AND TOTAL NUMBER OF DETECTIONS	130
FIGURE 6.19: EFFECT OF DISCOVERING A CLUE	132
FIGURE 6.20: ILLUSTRATION OF PATH PLANNING SEARCH ALGORITHM WITH CLUES AVAILABLE (CLUE AGE = 20)135	
FIGURE 6.21: EXPLANATION OF PATH PLANNING RESULTS (SEARCH DEPTH = 3).....	137
FIGURE 6.22: ILLUSTRATION OF PATH PLANNING SEARCH ALGORITHM WITH CLUES AVAILABLE (CLUE AGE = 40), SEQUENCE OF SCENARIO 2 EVENTS	138
FIGURE 6.23: COMPARISON OF AVERAGE TIMES UNTIL DETECTION AND TOTAL NUMBER OF DETECTIONS FOR SCENARIO 1	140
FIGURE 6.24: COMPARISON OF AVERAGE TIMES UNTIL DETECTION AND TOTAL NUMBER OF DETECTION FOR SCENARIO 2	142

LIST OF TABLES

TABLE 4.1 : PARTICLE ARRAY STRUCTURE	57
TABLE 6.1: BASELINE DESIGN FACTOR VALUES	103
TABLE 6.2: INITIAL MOTION PLANNING PARAMETERS	104
TABLE 6.3: SIMULATION RESULTS FOR PATH PLANNING SEARCH ALGORITHM WITH SEARCH DEPTH = 1	121
TABLE 6.4: COMPARISON OF SCENARIO 1 RESULTS FOR VARYING SEARCH DEPTHS.....	124
TABLE 6.5: COMPARISON OF SCENARIO 2 RESULTS FOR VARYING SEARCH DEPTHS.....	126
TABLE 6.6: COMPARISON OF SCENARIO 1 RESULTS WITH INCREASED RESAMPLING FREQUENCIES	129
TABLE 6.7: COMPARISON OF SCENARIO 2 RESULTS WITH INCREASED RESAMPLING FREQUENCIES	129
TABLE 6.8: BASELINE DESIGN FACTORS WITH UPDATED MAXIMUM DETECTABLE CLUE AGES	132
TABLE 6.9: BASE CASE SEARCH ALGORITHM RESULTS WHEN CLUES DETECTABLE.....	133
TABLE 6.10: COMPARISON OF SCENARIO 1 RESULTS WITH CLUES	136
TABLE 6.11: COMPARISON OF SCENARIO 2 SEARCH RESULTS WITH CLUES	139

Chapter 1

Introduction

The objective of this thesis is to develop and implement new decision systems and algorithms that will enable a United States Navy unmanned undersea vehicle (UUV) to search for and track a moving contact using time-dated measurements. In other words, we assume that the UUV carries a sensor that can detect clues that indicate that the contact has passed by a position at a specific time in the past. The information available to the decision system includes: (1) null measurement (i.e., contact not detected at current time), (2) time-dated measurement (i.e., clue found at current time that indicates contact was at this location in the past), and (3) bearings measurement (i.e., angular measurement towards contact position detected at current time). The results of this thesis will be arrived at by evaluating the best methods to handle the three types of information. The underlying distribution of the contact position space will be modeled using a generic particle filter, due to the highly non-Gaussian distributions that result from the conditions mentioned above. Using the particle filter distribution and the measurements acquired from the three conditions, this thesis will work towards implementing a dynamic path planning algorithm that seeks to minimize the uncertainty in the position of the contact and ultimately predict where the contact will move based on current and past position measurements.

1.1 Problem Motivation

With recent advances in research and technology, autonomous vehicle capabilities have steadily improved. These autonomous vehicle technologies, which perform missions and tasks

without the direction of human operators, have changed the way scientists and engineers approach problems. Because these robotic devices can work without manned guidance, they can execute missions that were too difficult, dangerous, or tedious for human operators to attempt. These enhanced and expanded capabilities provided by autonomous control have quickly been incorporated into military operations. The ability of these intelligent machines to operate in hostile environments and complete difficult tasks allow the military to decrease manpower, lower costs, and limit the risk to human life. A recent example of the implementation of unmanned systems is the United States Air Force's unmanned aerial vehicle (UAV) referred to as the Predator. In recent combat operations, the Predator has been successful in supporting ground operations by providing reconnaissance and surveillance in environments that are possibly too risky for ground forces to negotiate. Although the Predator requires remote human control, it demonstrates the advantages that unmanned vehicles can bring to military operations [19].

Based on the recent success of unmanned vehicles, such as the Predator, in the battlefield, the United States government plans to expand its autonomous technologies to naval operations through the unmanned undersea vehicle. The UUV is a self-propelled submersible whose operation is either fully autonomous or under minimal supervisory control. In the past, humans have exclusively made mission planning decisions on military vehicles, but in the coming years, the United States Navy plans to shift the decision making capability from the human operator to an onboard decision system. As scientists and engineers develop these systems, the difficulty rests in the need to translate the thoughts and actions of a human operator to a set of rules and behaviors for an autonomous system to follow [7].

In the future, the United States Navy plans to place these unmanned undersea vehicles in critical roles within the battle space. The Navy needs stealthy and unmanned systems to gather information and identify contacts in areas inaccessible by traditional maritime forces [7]. Mounting threats require the United States to tighten security by closely monitoring the critical maritime locations of both the United States and its allies. Due to the vast surface areas of the world covered by water, new assets such as UUVs are needed to patrol important waterways and prevent vessels from executing hostile actions from the sea. The capability to perform these missions with UUVs in the future will allow submarine and surface combatants to expand their sphere of influence while reducing possible vulnerabilities [20].

Critical missions including intelligence, surveillance, reconnaissance, mine countermeasures, tactical oceanography, communications, navigation, and anti-submarine warfare are on the verge of being addressed with UUVs. These autonomous underwater vehicles are advantageous in these types of missions because they increase performance, lower cost, and reduce risk to human life in manned systems. These improvements can be realized because of the following operational advantages afforded by UUVs:

- Autonomy. The ability of a UUV to operate independently for extended periods creates a force multiplier that allows manned systems to extend their reach and focus on more difficult tasks. Reduced costs are also a result when sensors and weapons are operated from smaller platforms like UUVs.
- Risk Reduction. Due to the unmanned nature of UUVs, there is a reduced threat to personnel from the environment or enemy combatants.
- Low observability. As a result of the small size and engineering of UUVs, they can operate fully submerged with low acoustic and magnetic signatures. These covert features enable the UUV to put sensors in positions that previously could not be reached.
- Deployability. UUVs can be designed as flyaway packages or pre-positioned in forward areas. They can be launched from a wide variety of platforms including ships, submarines, aircraft, and shore facilities and can be recovered from a different craft than they were launched from. Recoveries may also be delayed or abandoned because of the expendability created by the low cost of the UUV.
- Environmental Adaptability. UUVs can operate in a diverse range of environments including deep to shallow water, adverse weather and seas, and tropical or arctic conditions [7].

Due to the complicated nature of these missions and the environments in which they are carried out, there exists a continued need to develop and implement new decision systems and algorithms that can handle untested situations and environments. As the number of advanced autonomous vehicles continues to grow, the United States Armed Forces must determine how to take advantage of these new technologies in order to perform various missions including, but not limited to, surveillance, reconnaissance, and anti-submarine warfare. Although the technology and industrial capacity are ready to proceed, UUV capabilities need to progress before confidently deploying them to carry out important missions for the U.S. Navy. Specifically, this thesis will look at expanding the future capabilities of the UUV by enhancing the submarine track and trail capability.

1.2 Problem Statement

According to the description in the “Navy UUV Master Plan,” the objective of the submarine track and trail capability is to patrol, detect, track, trail, and handoff adversary submarines to U.S. Forces all the while remaining undetected by the enemy. Because UUVs can be launched from safe distances to accomplish these missions in high-risk areas or water that is too shallow for larger, more conventional platforms, they are the logical candidate to perform the submarine track and trail mission. Although the advantages of using UUVs to track and trail contacts at sea are numerous, the UUV must be able to operate in hostile areas with dynamic threats without taking actions that could inadvertently advance the stage of conflict. As a result, there exists a need to design the decision system aboard the UUV to use all its resources to appropriately process any situation that it may face at sea [7].

In accordance with The Navy Unmanned Underwater Vehicle Master Plan, Charles Stark Draper Laboratory is currently developing a Maritime Reconnaissance Demonstration system. Currently, research groups are attempting to expand the search capability of the UUV system. Using prior information on the position distribution and the assumed dynamic capabilities of the mobile vessel in question, Draper Laboratory seeks to develop a search strategy that effectively searches for and detects the moving vessel while remaining covert. This system will integrate automated search strategies with closed-loop planning and control aboard in-water systems. As

the UUV maneuvers through the sea environment, the system will maintain a dynamic position distribution for use by the UUV planning system and the onboard automated controller.

The specific problem to be addressed in this thesis consists of a single UUV searching for an unknown contact in an open sea environment. Initially, we assume that the contact's location is uniformly distributed within a given space. Although we assume that the contact is contained within that space initially, it is not confined to that space as the contact can move freely in all directions. Due to the goal to remain covert, the UUV must also avoid use of traditional sensor technology and instead rely on passive sensors. Although passive sensors may provide little information about the location of contact at each time step, the objective of the UUV decision support system is to use all available information to plan the best possible path in order to detect the position of the contact. In the end, this thesis will focus primarily on how to appropriately use new sensor information to update the position distribution and how to then use the updated position distribution to plan an effective path that will maneuver the UUV within the passive sonar range.

1.3 Problem Approach

Search and Detection Theory describe classic problems for which applications can be found in the military, fishing, mineral exploration, and search and rescue. Numerous algorithms and methods have been developed to aid searchers to achieve fast and accurate detection of the unknown or missing object of interest.

1.3.1 Search and Detection Theory

In *Search and Detection*, Washburn summarizes many classic search strategies that have been applied to naval operations [25]. Many of the optimal search methods have been applied to searches for stationary contacts. In these cases, if the location of the contact is known with a given probability distribution, the searcher can guarantee detection by using maneuvers such as circle in, circle out, sweep, or movement in a manner that maximizes the probability of detection based on the predefined distributions [25]. In many cases though, contacts such as missing

persons or lifeboats may move considerable distances during the search. Dynamics such as these add considerable complexity to the search. Washburn introduces search methods for dynamic environments where the search area is confined as well as scenarios when the target's position at some benchmark in time is known with reasonable accuracy. For the purposes of this thesis, the search area is not confined and the initial location of the contact is not known with detailed accuracy. As a result, the contact position distribution will expand as time advances and detection will not necessarily be guaranteed.

1.3.2 Guaranteed Detection Conditions

In order for the sensor vehicle to guarantee detection of the target within the initial search space, one of the following two conditions must be met:

1. The sensor must be fast enough to travel twice the length of the search area (i.e., move up and back down the length of the search area) before the target can move a distance greater than the diameter of the detection region covered in one time step by the sensor. In other words, if a target moving at maximum speed along a straight path moves a distance greater than the detection diameter before the sensor can sweep back and forth across the length of the search area, the target can escape detection. If this condition is met, the sensor can sweep across the search area until the target is detected (see Figure 1.1).




Figure 1.1: Guaranteed Detection Conditions for Sweep Maneuver

2. Detection can also be guaranteed if the sensor can travel the circumference of the search area before the target can move a distance greater than the diameter of the detection area covered in one time step. If this condition is met, the sensor can spiral in until the target is detected (see Figure 1.2):




Figure 1.2: Guaranteed Detection Conditions for Spiral-In Maneuver

In reality, a sensor vehicle will probably not be able to travel at speeds fast enough to satisfy these conditions. Therefore attempts to implement these search algorithms will not necessarily result in high probabilities of detection. Consequently, alternative search algorithms must be employed by the UUV system.

1.3.3 Current Search Methods

Before looking at how to approach the search problem when detection is not guaranteed, it is good to examine how analogous scenarios are currently accomplished with human operators. Some of the best examples of dynamic search scenarios are seen everyday in search and rescue operations. Take for instance a missing hiker scenario. In this scenario, the search and rescue workers must plan a coordinated search by taking into account the following facts: the mobile

hiker is traveling randomly at a known speed and the initial position is distributed according to an assumed probability. Using these pieces of information, the key to a successful search, according to *Search is an Emergency: Field Coordinators' Handbook for Managing Search Operations*, involves minimizing the search area. The search area can be kept to a minimum by responding quickly and searching for clues not subjects. Searchers focus the search on clues (i.e., time-dated measurements) instead of the specific subject, because there are a larger number of clues and the discovery of a clue can significantly reduce the search area in which the missing contact is. Although a clue does not pinpoint the exact location of the target, it can dramatically narrow the probability space. The effect of the clue depends on the “age” of the clue and the estimated speed of the target. If the clue has recently been deposited, the discovery of this clue can greatly reduce the search area. Conversely, the search area will not be reduced to the same extent when an older clue is discovered [16].

Specifically, when a clue is found, the searcher gains additional information about the location of the target. The uncertainty in the target location is reduced, because the searcher now knows that the target passed through this point at a particular time in the past. With that information, the new search area is constrained to the distance surrounding the clue that the missing contact could have reached based on the age of the clue and the estimated speed of the missing contact (see Figure 1.3) [16]




Figure 1.3: Effect of Detected Clue

Based on the logic of these search methods, this thesis will attempt to enhance the search capability of an autonomous vehicle by modifying the objective of the searches to include potential clues left behind by a mobile subject. If clues can be discovered that reveal the position of the object of interest at a point in the recent past, the search area can be limited to the immediate area surrounding the clue. Similar discoveries of clues by an autonomous vehicle could considerably reduce the size of the search area and as a result lead to faster detection of the contact.

1.3.4 Application of Search Methods

Using the approaches applied by human searchers, this thesis will focus on how an autonomous vehicle should operate when detection is not guaranteed. Specifically, the thesis will look at how a UUV decision system processes the available information and determines where to move based on that information. The UUV sensor technology observes the environment surrounding the UUV and uses the information it gains to analyze situations. As the UUV moves through the environment, it can observe one of the following conditions: (1) null measurement (i.e., contact not detected at current time), (2) time-dated measurement (i.e., clue found at current time that indicates contact was at this location in the past), and (3) bearings measurement (i.e., angular measurement towards contact position detected at current time). Each of these observations provides information about the location of the contact and will help constrain the potential position of the contact. Eventually, the decision system of the UUV processes the current and past information in order to form a position distribution representing the likely location of the contact at the current time. The UUV is then able to enumerate and analyze the impact of future actions on the position distribution of the contact. Based on the anticipated effect of these future actions, the decision support system will choose the action that most effectively satisfies the objective of the search, and the UUV will execute the decision and move accordingly.

This thesis will concentrate on implementing a decision support system that will enhance the search capability of an autonomous vehicle. The research presented in this thesis will contribute to the Charles Stark Draper Laboratory's continued research on the Maritime

Reconnaissance Demonstration system. In order to improve the current capabilities to include the ability to detect clues, this thesis will first consider how to represent the search space probabilistically. With the available information, the probabilistic model of the search space must correctly illustrate the potential contact positions. With an appropriate probabilistic model of the search space, the remainder of the thesis will focus on designing a decision support system for autonomous vehicles. The resulting dynamic search algorithms provided in the thesis will compute the appropriate course and speed commands for the UUV to follow in order to effectively decrease the search area and minimize the uncertainty in the distribution of the contact's position.

1.4 Contributions

The work in this thesis will contribute to the future advancement of the submarine track and trail capability discussed earlier. Specifically, this thesis proposes a new search algorithm that exploits the unique structure of particle filters. By taking advantage of the sample-based distribution produced by the particle filter for path planning and controls, the dynamic decision system expands the search and detection capabilities within the following problem areas:

Complex and Dynamic Information

Due to the limited information gained from the measurements in this search problem, the probability distributions can evolve into multi-modal distributions. In addition, the assumption that the contact is in motion causes the distributions to become even more complex as they will change over time. To account for this complex and changing information, the distributions must be modeled correctly and intelligent motion plans must be designed. The work in this thesis will utilize the sample-based particle distributions to maintain the complex distributions. Additionally, the structure of particle distributions allow the information to be clustered when designing motion plans. The cluster-based action space will provide more tractable motion plans that can consider larger portions of the search space.

Time-Dated Measurements

Given the assumed capability to detect time-dated measurements (i.e., clue measurements), this search algorithm possesses a measurement model to handle these measurements within the context of the particle filter structure. This measurement model determines the information contained within a clue measurement and uses the information to properly update the probability distribution.

The advancements made in this research will lead to continued development of the artificial intelligence of Navy UUVs which could potentially be used in future real-world search applications.

1.5 Organization

The remainder of this thesis is divided into seven chapters. Chapter 2 presents the concept of operations for UUVs in submarine track and trail missions and introduces a scenario in which this future capability can be analyzed. Chapter 3 presents the background into general state estimation, while Chapter 4 investigates the use of particle filtering techniques to process specific sensor measurements and provide an estimate of the contact's position. Chapter 5 introduces a dynamic search algorithm that analyzes the position distribution of a contact and directs a UUV along paths that reduce the uncertainty in the position of the contact and ultimately place the UUV in a position to begin track and trail operations. Chapters 4 and 5 will be the major contributions made by this thesis. Chapter 6 will provide examples of how the algorithm presented in Chapter 5 can be used by a UUV during a search and detection operation. Chapter 7 provides conclusions and suggestions for future work.

Chapter 2

Submarine Track and Trail Operational Description

As mentioned in the United States Navy UUV Master Plan, the vision of future United States naval operations involves the use of unmanned underwater vehicles for many key mission areas. Specifically, the research in this thesis will attempt to enhance the future search and detection capabilities of a UUV within the framework of the submarine track and trail mission. Although it is conceivable to envision a UUV autonomously searching, detecting, tracking, identifying, and targeting hostile contacts, these application areas remain several years in the future [7]. In planning for these future operations, it is important to develop reasonable scenarios in which to test the perceived capabilities. Therefore, a feasible scenario must be designed within this thesis in order to examine potential improvements to the UUV decision support system. The remainder of this chapter will discuss the development of a submarine search and detection scenario that will appropriately assess the research presented in this thesis.

2.1 Concept of Operations

Based on previously obtained intelligence in these regions, it is assumed that certain information about the operating environment as well as the readiness and capabilities of opposition forces are known. By taking advantage of the operating environment, it is assumed that a fixed sensor network is in place near a strategic naval chokepoint or port entry (Figure 2.1) [7].




Figure 2.1: Potential Sensor Network Placement

The sensors in this network will monitor traffic through the corridor. The sensors are assumed to detect vessels moving through the network, but the precise position and course of possible adversaries remain unknown. At the moment a potentially hostile contact passes through the seaway, patrolling UUVs in the area will be informed and expected to respond [15]. Due to the limitations caused by bathymetry and a possible lack of air superiority in certain coastal regions, UUVs may be the best option for manned submarines to project power deeper into these littoral environments [7].

As the primary sensor platform for the submarine track and trail mission, UUVs will be tasked with the objective to patrol, detect, track, trail, and handoff adversarial vessels to nearby conventional forces [7]. Due to the nature of the search environment, this UUV capability is an important component in anti-submarine warfare [7]. When dealing with underwater contacts, such as submarines, it is best to be in the position of first to act [7]. Dominance is not possible in reactive submarine warfare. If hostile contacts are not identified before they reach open sea, the likelihood of detection significantly decreases. As a result, one of the more crucial elements of the submarine track and trail mission is the search and detection phase. Immediately upon being notified of a contact making aggressive actions, a UUV will launch from a substantial distance and transit into the search area [7]. Based on information such as the position distribution and the assumed speed of the adversary vehicle, the UUV will patrol the area of interest and attempt to effectively search for and detect the contact. Therefore, the UUV must be able to maneuver as

necessary to detect the contact and upon detection begin a trailing operation until the contact has been classified [7]. At this point in the mission, the UUV will communicate to auxiliary military forces the position of detection and that a trail had been initiated (see Figure 2.2) [7].




Figure 2.2: Submarine Track and Trail Operations

2.2 Scenario Description

In order to test the performance of the UUV decision support system, it is necessary to design a scenario that appropriately models the future UUV support operations. Based on the concept of operations, the modeled scenario provides a search environment and contact dynamics which are representative of likely events and produce situations that are challenging to the autonomous controllers aboard the UUV. Specifically, the scenario involves the use of one UUV to search and detect the position of a potentially hostile vessel. Upon the recognition of a hostile contact moving through a pre-positioned sensor network, the UUV in this scenario will be

tasked with locating the contact in order to begin a trailing operation. Due to the perceived threat of the contact, the naval forces in the area need to monitor the movements and possible intentions of the contact. In addition, the UUV must remain covert and undetected by the contact, and therefore active sensor technology will not be made available. Consequently, this scenario will test the ability of the UUV decision support system to use the limited passive sensor measurements available to maneuver within range and begin a trailing operation. The remainder of this section familiarizes the reader with the search environment and the contact dynamics within the context of this specific scenario.

2.2.1 Description of Search Environment

As the map in Figure 2.3 illustrates, the UUV in this scenario will be tasked with searching in both littoral and open sea environments. These different environments cause challenges to naval forces and create situations that will test the robustness of the UUV decision support system. For example, due to the constraints imposed by the littoral environment, it is easier to deduce the position of the contact, but maneuvering through the shallow water depths and obstacles is much more difficult than an open sea environment. On the other hand, the open sea provides greater maneuverability, but it also allows more freedom for the contact to move and as a result, makes it difficult to determine the location of contact.




Figure 2.3: Scenario Map

The map contains several environment constraints that challenge the decision support system aboard the UUV but can be exploited by properly designing the autonomous system. These key environmental constraints can be seen in detail in Figure 2.4.




Figure 2.4: Key Environmental Constraints of Scenario Map

These elements include the naval chokepoint, the seaway, and the channel passage. The naval chokepoint consists of a main waterway or port entrance that possesses a shipping lane that vessels in the region must travel through to reach the open sea. Due to the natural constraints imposed by the chokepoint, it seems like a logical area to closely observe through some form of surveillance such as satellite imagery or a pre-positioned sensor network [15]. Based on observations acquired through the surveillance of the chokepoint, an initial distribution of the contact position can be generated.

The other features of the littoral environment which provide natural constraints for maneuvering vessels in the region include: the seaway and channel passage. According to the map for this scenario, these two waterways are the only possible passages into the open sea due to the position of island beyond the naval chokepoint (see Figure 2.4). Because of the distinctive size and depth of the waterways, the ability to navigate the waterways varies significantly. In this particular scenario, it is assumed that the seaway is much more navigable than the channel passage. Therefore, the number of passable routes is much lower through the channel, and consequently the likelihood that the contact will choose that route is much lower than the seaway.

2.2.2 Description of Contact Dynamics

The next aspect of the scenario involves the assumed contact dynamics. For this scenario, the contact motion is assumed constant and linear with inputs along the route that correspond to maneuvers needed to avoid obstacles and proceed towards the targeted location. Given this particular motion model, the two important factors in predicting contact motion include heading and speed. The assumed speed is based on the known capability of the vessel discovered moving through the chokepoint.

The initial heading of the contact is unknown. Initially, the contact moves through the seaway or the channel passage. Based on the differences in the ability to navigate each passage, it is assumed that the contact will pass through the sea passage with a higher likelihood. After the initial decision, it is assumed that the contact moves towards one of three potential goal

locations (i.e., Goal A, Goal B, or Goal C). These potential targets are located beyond the open sea and are equally likely to be targeted by the contact (see Figure 2.3). The decision to move toward a specific goal location is assumed to be independent of the decision to travel through the seaway or channel. Therefore, regardless of the initial route taken by the contact, there exists an equal probability that the contact will head towards any one of the three potential targets. A map showing the assumed contact dynamics is featured below (see Figure 2.5).




Figure 2.5: Detailed Description of Contact Dynamics

2.2.3 Description of UUV Mission

With the description of the search environment and the contact dynamics previously introduced, this section will focus on the mission of the UUVs with respect to submarine track and trail operations. In general, the mission of unmanned undersea systems is force projection [15]. UUVs allow the U.S. Navy to expand its operational concept of “Forward...From the Sea” by extending the reach of sensors to the shallow waters that are denied to traditional submarines [15]. In other words, sensor-carrying UUVs can be launched from a safe distance and extend the detection range of traditional naval vessels without increasing the risk to manned assets [7]. To

extend the general case further, the deployed UUV would then use known chokepoints and contact movements to maneuver as necessary to classify targets and begin a trailing operation [7].

For this particular scenario, the UUV is launched from a position outside the littoral region upon notification that a hostile contact has moved through the pre-positioned surveillance network (see Figure 2.6).




Figure 2.6: Detailed Description of Initial Stages of Scenario

Due to the high threat levels within the coastal regions and the limited endurance of the UUV, it might be too risky for a UUV platform to patrol continuously within the coastal waters. Therefore, a UUV will be commanded to enter the littoral region only when necessary. Despite the increased risk, a deployed UUV must enter the littoral region in order to take advantage of the environmental constraints and narrow the position distribution of the contact. While the UUV could patrol beyond the island and wait for the contact to enter the open sea, the contact position distribution would expand beyond an area that can reasonably be searched. Reactive posture is not effective in search operations, because the search area expands as time passes [7].

After being launched from a safe distance, the UUV is ultimately tasked with maneuvering within passive sonar range where bearing measurements can be detected and used for tracking the contact. The ability to achieve this objective depends on the ability of the UUV to process information made available through the on-board sensors. Depending on the relative positions of the UUV and the contact, the UUV sensors will generate one of the following measurements: (1) null measurement (i.e., contact not detected at current time), (2) time-dated measurement (i.e., clue found at current time that indicates contact was at this location in the past), and (3) bearings measurement (i.e., angular measurement towards contact position detected at current time). The sensors aboard the UUV can only return these pieces of information due to the necessity to remain covert throughout the mission. These measurements allow the UUV to remain stealthy, because they originate from passive sensor technology and do not provide a signal that enemy vessels could detect.

[This Page Intentionally Left Blank]

Chapter 3

State Estimation

Estimation is the process of extracting information from imperfect data. Estimation methods determine the desired information from measurements while taking into consideration measurement errors, the effects of disturbances and control actions on the system, and any prior knowledge of the information [8]. In the context of state estimation, measurements, typically in the form of noisy sensor measurements, are used to acquire the relevant state variables such as position, velocity, or heading. Based on this concept of estimation, this chapter will discuss the general theory behind two common state estimation techniques, Kalman filters and particles filters. After a basic introduction of the two estimation techniques, the strengths and weaknesses of both techniques will be discussed along with the conditions that define when each technique should be utilized. Work in this chapter and the remainder of this thesis will be restricted to two-dimensional motion for convenience.

3.1 State Estimation Models

This section will outline general motion and measurement models used in common navigation and tracking applications. Before estimating the contact state, it is important to identify the appropriate assumptions and model the known information about the contact dynamics and behavior as well as the relationship between the state variables and the measurements. These assumptions and models must closely resemble the true behavior of the contact in order to generate an accurate estimate of the state variables.

3.1.1 Motion Model

Central to every navigation and tracking application is the motion model to which different model based filters can be applied. For the most part, state estimation problems are related in that they can be described by quite similar state space models, where the state vector contains the position and derivatives of position. Traditional methods are based on linearized models and Gaussian noise approximations so that the traditional filter methods can be applied [12]. In the linearized motion model, the underlying model attempts to estimate the state $x \in \mathfrak{R}^n$ of a discrete-time controlled process that is governed by the following linear stochastic difference equation:

$$x_{k+1} = A_k x_k + B_k u_k + w_k, \quad (3.1)$$

where the matrix A_k relates the state at time step k to the state at step $k+1$ and B relates the control input u to the state x . In addition, the process noise is contained within the variable w_k [26].

Due to the complexity of many real world problems, the object being estimated often does not follow a linear track. In these cases, the linear motion assumption does not hold, and the motion model must represent the more complex non-linear motion. The state vector is yet again $x \in \mathfrak{R}^n$, but now the time controlled process is governed by the non-linear stochastic difference equation [26]:

$$x_{k+1} = f_k(x_k, u_k, w_k). \quad (3.2)$$

In this case, the non-linear function $f(x_k, u_k, w_k)$ defines the relationship between the state at time step k to the state at time step $k+1$.

3.1.2 Observation Model

For the purposes of state estimation, measurements are used to provide feedback about the relative state of the object of interest. Typically, measurements relate the position of one's own platform to another object's position. Depending on the nature of the measurements, the equation relating the respective positions can take several forms. For a measurement $z \in \mathfrak{R}^m$ whose relationship to the state variables is linear, the following general equation describes the relationship between the estimate of the state and the measurement:

$$z_k = H_k x_k + v_k, \quad (3.3)$$

where the matrix H_k relates the state vector, x_k , to the measurement and v_k represents the measurement noise [26].

For other measures, such as bearing measurements, the relationship of the measurement to the estimated state is characterized by a non-linear relationship. In these situations, the measurement is governed by the non-linear equation:

$$z_k = h_k(x_k, v_k), \quad (3.4)$$

where the non-linear function $h_k(x_k, v_k)$ defines the relationship between the measurement and state at time step k [26].

3.2 Recursive Bayesian Estimation Techniques

As defined previously, estimation is the process of extracting information from imperfect data. The work in this thesis must utilize state estimation techniques to estimate the state of the

contact, because no functional form of the state exists. The state variables cannot be estimated directly and therefore must be represented probabilistically. In order to arrive at the probability distribution of the state at time step $k+1$ given the set of available observations at current time k , the posterior distribution $p(x_{k+1}|Z_k)$ must be approximated through Recursive Bayesian Estimation. Given that the set of available observations at time k is given by $Z_k = \{z_0, \dots, z_k\}$, the Bayesian solution to the posterior distribution is described by the following expression [9]:

$$p(x_{k+1}|Z_k) = \int p(x_{k+1}|x_k)p(x_k|Z_k)dx_k, \quad (3.5)$$

where the distribution, $p(x_k|Z_k)$, is defined by [9]:

$$p(x_k|Z_k) = \frac{p(z_k|x_k)p(x_k|Z_{k-1})}{p(z_k|Z_{k-1})}. \quad (3.6)$$

All state estimation techniques will be derived from these initial Bayesian equations.

While many implementations of target tracking algorithms have proven successful for state estimation, this section will examine two accepted techniques for state estimation: Kalman filters and particle filters. After reviewing these methods, this section will conclude with an explanation as to why particle filters are better suited for state estimation when using the passive sensor technology discussed in this thesis.

3.2.1 Kalman Filters

Over the last half century, Kalman filters have dominated the field of estimation. Since R.E. Kalman introduced the recursive solution to the linear filtering problem in 1960, the Kalman filter has been the subject of extensive research and application [26]. The filter is very powerful, because it can maintain estimates of past, present, and future states even though the

exact nature of the modeled system may be unknown [26]. Specifically, the Kalman filter is a set of mathematical equations that function recursively to provide efficient solutions of the least-squares method. The underlying model attempts to estimate the state $x \in \mathfrak{R}^n$ of a discrete-time controlled process that is governed by the linear stochastic difference equation as in (3.1) using a measurement $z \in \mathfrak{R}^m$ with a linear relationship with the state variables as in (3.3) [26]. This model holds under the following strict assumptions:

1. Linear process model
2. Linear measurement model
3. Gaussian probability density after every time step.

When these assumptions hold, no algorithm can perform better than the Kalman filter [26].

Often real world problems cannot be modeled with linear process and measurement models, and as a result, the discrete Kalman filter leads to inaccurate estimates. In these cases, the state vector is again $x \in \mathfrak{R}^n$, but now the time controlled process is governed by a non-linear stochastic difference equation as in (3.2) and the measurement is described by a non-linear relationship as in (3.4).

When either the process to be estimated or the measurement relationship are non-linear as in the previous equations, an adapted Kalman filter known as the extended Kalman filter (EKF) can be successfully implemented. The EKF uses linearized mathematical models for both the state error dynamics and the measurement relationship by taking partial derivatives of the process and measurement functions in order to compute estimates about the current mean and covariance [10]. These approximations allow the extended Kalman filter to recursively estimate the state vector despite the non-linear relationships.

The process and measurement models for both the Kalman filter and the extended Kalman filter are used to estimate the state variable. The filter estimates the process state at time k and then obtains feedback in the form of noisy measurements. As a result, Welch and Bishop categorize the equations for the Kalman filter into two groups: time update equations (“prediction”) and measurement update equations (“correction”). The prediction equations project the current state and error covariance estimates forward in time to obtain *a priori*

estimates for the next time step. The correction equations provide the feedback for the predicted state by including a new measurement into the *a priori* estimate to obtain an improved *a posteriori* estimate. The overall estimation algorithm cycles through the prediction and correction phases as seen in Figure 3.1 to continually update the estimation of the state variables over time [26].




Figure 3.1: Kalman Filter Cycle

For a detailed description of the specific equations used by Kalman and extended Kalman filters in target tracking refer to Bar-Shalom or Gelb [2] [8].

3.2.2 Particle Filters

In recent years, an alternate state estimation technique known as a Sequential Monte Carlo (SMC) method has been applied to numerous problems. SMC methods that are used within a Bayesian framework have been referred to as bootstrap filtering, the condensation algorithm, particle filtering, interacting particle approximations, and survival of the fittest. For the sake of simplicity, the remainder of the discussion will refer to this approach as particle filtering. For a more detailed description of particle filters refer to Doucet, de Freitas, and Gordon [5].

Although the particle filter process follows the recursive Bayesian approach to dynamic state estimation, the process is much different than the approach utilized by Kalman filters. Kalman filters use the predicted state and measurement data in conjunction with an exact

analytical expression to compute the parameters of the evolving sequence of Gaussian posterior distributions. Contrary to Kalman filtering which assumes Gaussian posterior distributions, particle filtering creates a sample-based representation of the entire probability density function [10]. The sample-based representation is constructed by using a set of random samples of the state with associated importance weights. These weighted particles can then be used to build point mass representations of the probability densities. Because the probability densities are generated from the set of random samples, this filter can be applied to any state space model (i.e. non-linear and non-Gaussian) [12]. As the probability densities are computed, estimates of the state can then be derived using the state of each particle and magnitude of its weight [22].

3.2.2.1 State Parameters

In contrast to other state estimation techniques, particle filters create a sample-based distribution from which the state of the contact can be estimated. In order to generate the sample-based distribution, particle filters simulate a set of M state samples called particles. The state estimates for each particle consist of descriptive variables, X_k^i , that would be found in typical state estimates (i.e., position, speed, heading) as well as importance weights, w_k^i , for each particle i at time k . Therefore, the state of each particle is defined as [22]:

$$S_k^i = [X_k^i, w_k^i] \quad \forall \quad i = 1 \dots M . \quad (3.7)$$

The weight defines the contribution of this particle to the overall estimate of the variable [22]. Thus, the higher a particle is weighted, the greater influence it has on the probability density function as particle i represents $\frac{w_k^i}{\sum_i w_k^i}$ of the probability mass. Collectively, these particles create a sample-based representation, $S_k = \{S_k^i\}_{i=1}^M$ which leads to the complete probability density of the state [22].

3.2.2.2 Generic Particle Filter Algorithm

This section introduces the basics of the generic particle filter algorithm. The generic particle filter algorithm operates by cycling through the following three stages: (1) prediction, (2) measurement update, and (3) resample in order to generate a numerical approximation to the posterior distribution $p(x_{k+1}|Z_k)$ (see Figure 3.2) [5]. Given the general motion models and measurement models introduced previously, the following section discusses each of these stages and explains how the particle filter generates the sample distribution that appropriately models the current state of the contact.




Figure 3.2: Recursive Particle Filter Algorithm

3.2.2.2.1 Initialization

Before implementing the particle filter algorithm, the initial probability distribution must be established. Based on historical information or current intelligence, an initial probability distribution of the state vector, p_{x_0} can be assumed. From this information, the initial particle distribution will be formed by simulating a set of M particles according to the parameters of the assumed distribution. As a result, there will be more particles within areas with higher

likelihoods of detection. Ultimately, this particle distribution defines the initial state space, $\{S_0^i\}_{i=1}^M$.

3.2.2.2.2 Prediction

Based on the *a priori* particle distribution and the set of measurements collected prior to time k , the prediction phase attempts to predict the current probability distribution as depicted in the following expression [9]:

$$p(x_k | Z_{k-1}) = \int p(x_k | x_{k-1}) p(x_{k-1} | Z_{k-1}) dx_{k-1} . \quad (3.8)$$

Using the *a priori* particle distribution and the assumed motion model, the predicted probability distribution can be found by simulating the effect on each particle. After propagating each particle one time step ahead, the new set of particles collectively forms a prediction of the state distribution at time $k+1$ [10].

For the simple case of constant linear motion, the linear motion model stated earlier can be used to propagate the particles according to the following system of equations:

$$X_k^i = \begin{bmatrix} x_k^i \\ y_k^i \\ \dot{x}_k^i \\ \dot{y}_k^i \end{bmatrix} = \begin{bmatrix} x_{k-1}^i + \dot{x}_{k-1}^i \Delta t \\ y_{k-1}^i + \dot{y}_{k-1}^i \Delta t \\ \sim N(\dot{x}_{k-1}^i, \sigma_w^2) \\ \sim N(\dot{y}_{k-1}^i, \sigma_v^2) \end{bmatrix} \quad \forall \quad i = 1 \dots M \quad (3.9)$$

where the assumed state variables consist of position and speed. This model propagates each particle according to a constant linear process model with Gaussian noise applied to the velocities of the particles to account for disturbances in the contact heading direction.

3.2.2.2.3 Measurement Update

After propagating the particles ahead one time step, data is processed in the form of a new measurement of the surrounding environment at time k . This phase attempts to update the current probability distribution based on the measurement received at time k as represented in the following expression [9]:

$$p(x_k | Z_k) = \frac{p(z_k | x_k) p(x_k | Z_{k-1})}{p(z_k | Z_{k-1})} \quad (3.10)$$

where

$$p(z_k | Z_{k-1}) = \int p(z_k | x_k) p(x_k | Z_{k-1}) dx_k . \quad (3.11)$$

This probability distribution can be approximated by updating the particle weights according to the likelihood of receiving the measurement z_k given the particle state. This likelihood function can be used to update the particle weights according to the following expression [12]:

$$w_k^i = w_{k-1}^i p(z_k | x_k^i) \quad \forall \quad i = 1, \dots, M . \quad (3.12)$$

Depending on the type of measurement received, the particle filter will apply the appropriate measurement model to properly compute the likelihood function and update the particle weights. Upon updating the weight of each particle, the weights accurately describe the complete probability density function.

3.2.2.2.4 Resampling

The final step in the particle filter algorithm involves resampling the particles. Without this step in the basic algorithm, a few particles tend to dominate after a few iterations. As this

phenomenon persists a large computational effort is devoted to updating particles whose contribution to the approximation to the posterior filtered density is almost zero. This problem is alleviated by resampling the particles during the particle filtering process [22].

The importance resampling process consists of probabilistically eliminating particles with small weights and duplicating the particles with larger weights. Specifically, M samples will be selected with replacement from the set $\{S_k^i\}_{i=1}^M$, where the probability of choosing a sample i is equal to w_k^i [12]. At the conclusion of the resampling step, a new set of particles with identical weights will be created with a higher concentration of particles around the areas of higher probability. Although this process leads to a more accurate description of the contact state, eliminating particles with lower weights causes the completeness of the distribution to be lost. The effect of importance resampling in the particle filter process is illustrated in Figure 3.3 [23].




Figure 3.3: Particle Filter Illustration

3.3 Comparison of State Estimation Techniques

Both Kalman filters and particle filters are useful tools in state estimation. The degree of success that each filtering algorithm achieves depends on the nature of the application. For applications where linear motion models, linear measurement models, and Gaussian distributions can appropriately model the real system, the Kalman filter provides the optimal estimation of the state. Even in applications with non-linear motion or measurement models, but sufficiently linear error behavior, the extended Kalman filter provides an adequate estimation of the state. Increasingly though for many complex application areas it is becoming important to include highly non-linear and non-Gaussian elements in order to accurately estimate and model the dynamics of a system. Unfortunately, the Gaussian assumptions required to apply the extended Kalman filter do not hold in many practical application areas. Given that the extended Kalman filter assumes that the probability density function is Gaussian, the true density must meet the Gaussian assumption in order for the extended Kalman filter to guarantee sufficient results. If the true density is non-Gaussian (i.e. bi-modal or skewed), then an assumed Gaussian distribution will not provide a complete description of the state space, and the extended Kalman filter will not perform as well as other filters.

On the other hand, the more non-linear the motion or measurement models or more non-Gaussian the posterior distribution, the greater the potential application of particle filters [12]. Because particle filters are not bound by Gaussian assumptions, they can approximate an optimal solution numerically based on Monte Carlo simulations of the physical model [12]. The sample-based distributions produced by the particle filter can take any form and therefore can be used to estimate all types of problems. Although the particle filter can be utilized in these complex application areas, the simulations require a significant amount of computational power as the number of particles and the sampling rate increases. Therefore, for highly non-linear and non-Gaussian models where computational power is cheap, the particle filter can be applied with high levels of success [12].

Given the complex nature of the operating environment and the information obtained with the assumed passive sensor aboard the UUV, many times the most representative posterior distributions are non-Gaussian. As a result, the assumed Gaussian distribution of the EKF will

provide a poor description of the state distribution. Therefore, it is advantageous for the UUV decision system to use particle filters for state estimation while searching for naval vessels using passive sensor technology. Without the range or bearings measurements obtained with traditional sonar, point estimates for target position will not be attainable. Instead, the only measurements available to the sensor will be time-dated position measurements. The information contained in these observations lead to distributions that are best represented with the sample-based distributions of the particle filter.

[This Page Intentionally Left Blank]

Chapter 4

Contact State Estimation Using Passive Sensor Technology

In order for the autonomous system to make intelligent decisions, it must first properly observe and model the surrounding environment and the relevant contacts. In the submarine track and trail role, a UUV must be able to estimate the states of the surrounding contacts and know how accurate those estimates are before any maneuver decision aid can direct the UUV towards the correct destination. The purpose of this chapter is to discuss how to appropriately estimate the state parameters based on the sensor measurements available to the UUV. The discussion will begin with an introduction to the information collected by the assumed sensor technologies aboard the UUV. Based on the information obtained through measurements of the surrounding environment, the rest of the chapter will detail how the decision system aboard the UUV processes information and provides an estimate of the contact state.

4.1 UUV Sensor Technology

A decision support system must be able to observe the surrounding environment and use the information it gains to analyze the situation. In order to accurately estimate the state of the contact, the UUV decision support system needs an understanding of the information that it is able to acquire with its on-board sensors. The following section will familiarize the reader with the information obtained through traditional naval sonar and introduce additional pieces of information that we assume the UUV can collect.

4.1.1 Range and Bearing Measurements

With sonar as the primary sensing technology used by the U.S. Navy, undersea platforms typically are able to detect and track contacts using range and bearing measurements. The two types of sonar used are active and passive. Active sonar operates by generating a directed sound pulse and measuring the time delay of the reflected return off objects in the environment [21]. The distance to the object or range can be calculated using knowledge of the speed of sound and the pulse duration. The direction of the object relative to the sensor or bearing can be determined by knowing where the sound pulse was directed. Although active sonar produces two measurements useful in tracking, the pulse of sound created by the sonar can be heard by other contacts at sea, revealing the UUV position [19]. In contrast to active sonar, passive sonar relies on being able to detect acoustic signals emitted by other objects [21]. When within the operating range of passive sonar, contacts can be tracked using bearings measurements. Because passive sonar listens and collects bearings measurements without transmitting a signal, it reduces the likelihood of counter-detection and thus is more effective in military operations [19].

4.1.2 Time-Dated Position Measurements

The need to remain covert in track and trail problems rules out the use of active sonar. When contacts can be detected with passive sonar, bearing measurements allow the UUV to perform track and trail operations; however, during the search and detection phase of the mission, contacts might maneuver out of the reach of traditional passive sonar technologies. Under these circumstances, the performance of the decision support system could be significantly enhanced with additional passive sensor information.

When faced with a similar scenario, typical search and rescue operations focus the search on clues instead of a specific missing subject, because numerous clues exist in comparison to subjects, and the discovery of a clue can significantly reduce the search area. Based on this reasoning used in standard search procedures, we understand that the existence of clues and the ability to detect these clues could greatly reduce the potential areas in which the targeted vessel could be located. Therefore, throughout the remainder of the thesis, we assume that as the suspected contact moves throughout the maritime environment it leaves behind clues that mark

its trail. Although the exact nature of these clues is not specified, we assume that the clues are present in a form detectable by a sensor aboard the UUV.

Specifically, as the contact continues on its track, a clue is set down periodically and remains stationary at the point of deposit. Although clues are left behind at regular intervals, the contact is not guaranteed to deposit a clue at every point in time. Instead the contact deposits clues according to an assumed deposit probability f , which is equal to the probability that a contact leaves a clue behind within the current time. The percentage of occurrences is independent and identically distributed for each time step.




Figure 4.1: Clue Deposit Frequency

Once deposited, the clues will stay fixed at the point of deposit and remain detectable for a given amount of time depending on the decay rate of the clue. In addition, throughout the entire thesis, we will assume that the strength of clue signal decays exponentially (see Figure 4.2). Consequently, as the time stamp on each of the clues grows older, the ability to detect the clue becomes increasingly difficult.




Figure 4.2: Exponential Decay of Clue Signal Strength

Finally, we assume that the UUV possesses the capability to collect the time-dated position measurements, z_k^α or “clues” that reveal the position of a contact at a specific time $k - \alpha$ in the past, where α represents the age of the clue. In particular, if the UUV passes a deposited clue within its indicated detection range, d , it can generate a measurement that identifies the position and age of the clue.

In reality, all sensor technologies have missed detection rates. Sensor missed detection rate is defined as the probability of no detection given that a clue does exist. Due to the assumed exponential decay rate of the clues, as the clues become older, the sensor detection rate (i.e., the probability that the sensor can detect the clues) exponentially decreases (see Figure 4.3).




Figure 4.3: Sensor Detection Rate - Probability of Clue Detection Given that a Clue Exists

Upon detecting a clue, the UUV receives a measurement specifying the position and age of clue. While it is assumed that these measurements are highly accurate, the exact position and age of the detected clue contains uncertainties. The effect of both these uncertainties is very similar. Because the uncertainties in both position and age expand the new search area defined by the identified clue, we will focus only on the uncertainty in position in order to simplify the state estimation. For this particular problem, we assume that the uncertainty in clue position is normally distributed in the x and y direction around the detected clue position with standard deviations denoted by (σ_x, σ_y) where these uncertainties are not necessarily the same (see Figure 4.4).




Figure 4.4: Uncertainty in the Clue Position

As a result, the potential distance in which the contact may have moved from the time of the clue deposit will increase. As can be seen in Figure 4.5, there exists a small probability that the contact is located beyond the farthest distance the contact could have traveled since the clue was deposited.




Figure 4.5: Probability Distribution of Contact Position after Detection of Clue

Using these assumptions and the information they provide, we must determine how to model the capability to detect clues and determine how these clues enhance the ability of a UUV to search for and detect a hostile contact.

4.2 Particle Filter Implementation

Many recent implementations of target tracking algorithms use particle filters for state estimation. Using techniques similar to those used in other recursive Bayesian estimation methods, particle filters attempt to track a target of interest as it evolves over time using noisy measurements. During the search and detection phase of this particular scenario, the measurements available to the decision system include: (1) null measurement (i.e., contact not detected at current time), (2) time-dated measurement (i.e., clue found at current time that indicates contact was at this location in the past), and (3) bearings measurement (i.e., angular measurement towards contact position detected at current time). This section will discuss how the particle filter algorithm will use the three measurement models to estimate the state of the contact for the maneuver decision aid algorithms discussed later in this thesis.

4.2.1 Particle Description

As stated previously, the particle filter state parameters consist of both state variables and weight; the variables of interest at time k are represented by a set of M samples (particles). Each particle consists of a copy of the state variables of interest and a weight that defines the contribution of this particle to the overall distribution [22]. The remainder of this section will discuss how the structure of the state variables and the particle weights were determined and how they will be employed in the submarine search and detection scenario.

4.2.1.1 Descriptive State Parameters

In the context of this chapter, the contact's descriptive state parameters include its position and velocity. A particle state estimate for a contact in this implementation can then be represented by:

$$X_k^i = \begin{bmatrix} x_k^i \\ y_k^i \\ \dot{x}_k^i \\ \dot{y}_k^i \end{bmatrix} \quad (4.1)$$

where the position is stored in the Cartesian north and east coordinates (x_k^i, y_k^i) in a locally flat earth reference frame and the velocity of the contact is defined by its speed in the north and east direction. These parameters were chosen to fulfill the objective of the search. Specifically, the state parameters for position in the north and east direction are necessary to localize the position of the contact. State parameters for speed in the north and east direction provide additional information about the behavior of the contact. Eventually these state parameters will be useful in predicting where the contact will move.

As stated earlier, we assume that as the contact travels through the maritime environment it deposits time stamped clues that provide information about the contact's position at some previous point in time. Based on this information, particles cannot consist of only the contact state at the current time. In order to appropriately describe the search space, the existence of clues must be considered when defining the state of each of the particles. Ultimately, the distribution must account for not only the current location of the contact but also the clues it may have deposited in the past.

Because the location of the clues is directly related to the previous positions of the contact, the search space will be characterized as arrays of particles. In other words, instead of representing the search space with independent, individual particles, the search space is composed of independent particle arrays. Each array consists of a particle for the current contact

state and a history of previous particle states representing each potential clue along the target trail (see Figure 4.6).




Figure 4.6: Particle Array Illustration

Depending on the capability of the passive sensor technology, a history of previous particle states is kept until the age, n , at which the sensor is no longer capable of detecting the presence of clues. Because the clues remain stationary upon deposit, only the particle representing the current state of the possible target will be propagated in the particle filter. As the current state of the contact is propagated, the particle history is updated and stored as shown in Table 4.1.

Table 4.1 : Particle Array Structure

Particle Number	1	2	...	m
Estimated Target Location	(x_k^1, y_k^1)	(x_k^2, y_k^2)	...	(x_k^m, y_k^m)
Newest Estimated Clue Location	(x_{k-1}^1, y_{k-1}^1)	(x_{k-1}^2, y_{k-1}^2)	...	(x_{k-1}^m, y_{k-1}^m)
...	...	...	...	...
Oldest Estimated Clue Location	(x_{k-n}^1, y_{k-n}^1)	(x_{k-n}^2, y_{k-n}^2)	...	(x_{k-n}^m, y_{k-n}^m)

4.2.1.2 Particle Weights

The final consideration for the current proposal for the search space distribution is the particle weights. In order to correctly model the search space and provide a path planner with accurate information, the particles must be appropriately weighted. Because the objective of the search is to localize the target, the particle weights will collectively describe the probability distribution for the location of the contact. Therefore, each particle array will be assigned only one weight which defines the contribution of the current particle state to the overall estimate of the true contact state. Previous particle positions found in the particle history will influence the weights of the current particle state based on the relevant probability model. Depending on the location of the UUV with respect to the locations in the particle history, the weight of the current contact state will be updated. Further explanation into updating of probabilistic particle weights will be discussed later in Section 4.2.3.

4.2.2 Contact Motion Model

One of the fundamental decisions for navigation and tracking applications is the motion model that describes the contact. Typical contact motion can be modeled with linear state dynamics as seen in equation (3.1), where the contact's motion follows a linear path except for measured inputs, u and unmeasured disturbances, w_k [12]. Similarly, for this particular search and detection scenario, the contact motion is assumed constant and linear. Based on intelligence about the assumed contact, the mean traveling speed is known with reasonable certainty. Although the distribution of the moving contact speed is known, the path traveled by contact is uncertain. Due to the geographical constraints of the environment and the uncertain objectives of the contact, the contact could maneuver on numerous courses.

Specifically, the geography of the area of interest creates two possible courses on which the contact could travel (see Figure 4.7). As a result of the island existing beyond the sensor network, the contact can either move through the channel passage or through the seaway. Because the channel passage is much narrower and more difficult to maneuver through, it is more likely for the contact to choose the wider sea passage. For the purposes of state estimation, it is assumed that there is a p_s percent likelihood that the contact will move through the seaway

as opposed to channel passage. The effect of this likelihood function will be made evident through the propagation function of each of the particles, where initially approximately p_s percent of particles will travel through the seaway and $(100 - p_s)$ percent of particles will move through the channel.




Figure 4.7: Contact Motion Model for Submarine Search and Detection Scenario

In addition to the uncertainty in the decision of the contact to travel through the seaway or channel, the particle filter motion model must also properly model the uncertainty in the objective of the contact. Based on information obtained through other forms of intelligence, we assume that contact has three potential goal destinations across the sea (i.e., Goal A, Goal B, Goal C). Upon passing through either the channel or the seaway, the contact will adjust its heading towards its objective (see Figure 4.7). Each of these destinations are assumed equally as likely, and as a result, an approximately equal number of particles will head towards Goal A, Goal B, and Goal C.

4.2.3 Observation Model

As an autonomous underwater vehicle maneuvers through the maritime environment, its sensors collect information that describes the contact's position relative to the UUV. Due to the clandestine nature of the operations, the decision support system aboard the UUV must explore with passive sensor technologies. Depending on the position of the UUV relative to the contact and the contact's trail, the sensor will be provided with one of the following pieces of information: (1) null measurement (i.e., neither contact or clue detected at current time), (2) time-dated position measurement (i.e., clue detected at current time denoting position of contact at time $k - \alpha$), and (3) bearings measurement (i.e., angular measurement towards contact position detected at current time). Therefore, instead of using sensors which provide continuous measures, the decision system in this problem must be able to relate each of measurements acquired by UUV sensors to estimate the state of the contact. The remainder of this section will explain how to appropriately model each of the measurements and describe how the particle filter will use the measurements to estimate the state of the contact.

4.2.3.1 Stage 1: Null Measurement Model

The null measurement model refers to the initial stages of the search when the contact is positioned outside the range of all available passive sensor technologies. During this initial stage of the search, no sensor measurements will be obtained by the sensor aboard the UUV. Even though no time-dated or bearing measurements were detected at this point, information about the position of the contact can be derived. For instance, the fact that the UUV sensors detected no measurements at the current UUV position tells us that there is a high likelihood that the contact is not located within the passive sensor detection range. Therefore, the probability that the contact or any of its associated clues exist within the detection region surrounding the UUV is very small. Using that information, the contact position distribution can be adjusted accordingly by greatly reducing the probability of detection in the area surrounding the UUV (see Figure 4.8).




Figure 4.8: Illustration of Null Measurement Model

As previously introduced, the null measurement model refers to the stages of the search when no measurements are obtained by the UUV sensor technologies. Given that the UUV receives no measurements at the current UUV position, there is a high probability that the contact is not located within the passive sensor detection range. Due to uncertainties in sensor performance and environmental effects, the particle weights within this region cannot be strictly set to zero. Instead the weights of the particles whose current or past positions exist within the detection region should be reduced according to the probabilistic models that define the measurement model.

In order to appropriately reduce the weights of the particles affected by the null measurement model, we must determine the probability of detecting an equivalent clue with age α given that a contact had passed through the detection region α time steps previously. Based on the assumptions of the time-dated measurements, we can use a Bayesian approach to determine the required probabilities. With the understanding that the contact passed through the detection region at time $k - \alpha$, the probability of detecting a trace of the contact's trail can be found using the following expression:

$$p(D_k^\alpha) = p(D_k^\alpha | C_k^\alpha) p(C_k^\alpha) \quad (4.2)$$

where the probability of detecting a clue, $p(D_k^\alpha)$ is equal to the conditional probability of detecting a clue given a clue had been deposited, $p(D_k^\alpha | C_k^\alpha)$, multiplied by the probability that a clue had been deposited at time $k - \alpha$, $p(C_k^\alpha)$. These probabilities can be estimated by applying the assumptions about sensor performance and clue deposit frequencies. The conditional probability $p(D_k^\alpha | C_k^\alpha)$ is equivalent to the sensor detection rate introduced earlier, where the probability of detecting a clue decreases exponentially with respect to the age of the clue, α . The specific probability can be computed using the following expression:

$$p(D_k^\alpha | C_k^\alpha) = (e^{-\alpha/S}), \quad (4.3)$$

where S is the factor controlling the rate of decrease. Furthermore, the probability that a clue has been deposited at time $k - \alpha$, $p(C_k^\alpha)$ is equal to the assumed clue deposit probability, f as seen in the next expression:

$$p(C_k^\alpha) = f. \quad (4.4)$$

These expressions can be used to update the weights of the particles that have passed through the detection region within the past n time steps. If the position of particle j or any of the past history of particle j falls within the detection region (i.e., $(x_{k-\alpha}^i)^2 + (y_{k-\alpha}^i)^2 \leq d^2$), the particle weight should be reduced according to the time since the particle passed through the region. Depending on the amount of time elapsed, the weight of the particle will be reduced by the probability of not detecting a contact or its trail that has traveled through the detection area $(1 - p(D_k^\alpha))$. In particular, the weights of the affected particles will be updated according to the following expression:

$$w_{k+1}^i = w_k^i (1 - p(D_k^\alpha)) . \quad (4.5)$$

For those particles whose current and past positions exist outside the detection region, this measurement model provides no relevant information. As a result, the weights of these particles will remain unchanged.

4.2.3.2 Stage 2: Time-Dated Measurement Model

The time-dated position measurement model refers to the stage of the search when the passive sensor aboard the UUV detects a clue. At time k , a measurement of the contact produces the following time-dated measurement:

$$z_k^\alpha = \begin{bmatrix} x_k^\alpha \\ y_k^\alpha \end{bmatrix} , \quad (4.6)$$

which consists of the position of a contact at a time $k - \alpha$ in the past. Although the clue does not provide a point estimate of position of contact at time k , it considerably reduces the contact position distribution.

Using the age of the detected clue and the assumed maximum contact speed, we can calculate the maximum distance that the contact could have traveled during the time since the clue deposit. This maximum travel distance, E , is obtained through the following equation:

$$E = \alpha V , \quad (4.7)$$

where α is the age of the detected clue and V is the assumed maximum speed of the contact. This value defines the radius of the area surrounding the clue position in which the contact could currently be located. Because the contact cannot travel beyond this maximum travel distance, we

know that contact position distribution can be constrained to positions within this radius surrounding the clue (see Figure 4.9).




Figure 4.9: Clue Effect on Contact Position Distribution

After constraining the position distribution with the information gained from the clue measurement, the distribution can be constrained further by decreasing the probability of locating the contact within the detection region surrounding the UUV. Based on the logic used before in updating the distribution with the null measurement, there is a low probability that the contact is located in the detection region, because no additional measurements were discovered near the UUV. After combining the information from the null measurement and the time-dated position measurement, the contact distribution is constrained further and the uncertainty in the distribution is greatly reduced (see Figure 4.10).




Figure 4.10: Illustration of Time-Dated Position Measurement Model

As introduced earlier, the time-dated position measurement model is applied when the passive sensor aboard the UUV detects a clue, which consists of the position of the contact at a time $k - \alpha$ in the past. Based on the position and age of the clue as well as intelligence regarding the maximum speed of the contact, the particle distribution can be updated according to the time-dated measurement model which proposes that the contact could not have traveled beyond the maximum travel distance, E . Consequently, the probability of detecting the contact outside the maximum travel distance stretching from the clue position is very low. Although the probability of detection is extremely small outside the maximum travel distance, we cannot assume that the particle weights in this outer region will be zero.

Due to the uncertainty in the position of the detected clue and the uncertainty in the traveling velocity of the contact, the circular region surrounding the clue is not strictly defined. Specifically, the probability of detection extends beyond the maximum escape distance as a result of the uncertainties. In addition, the probability of detection within the maximum escape distance is not uniformly distributed. Due to the assumed constant linear motion of the contact, there is a higher probability of the contact being positioned at the edges of the escape radius, E (see Figure 4.11).




Figure 4.11: Contact Position Distribution in the Areas Surrounding the Clue Location Due to the Uncertainty in the Clue Location and the Assumed Contact Motion Model

To correctly model the reduction of the particle weights upon the detection of a clue, we must consider the resulting position distribution in the areas surrounding the clue location. Based on the assumptions of the contact motion and the uncertainty in the age of the clue, we can model the probability distributions as a function of the distance from the detected clue location. To account for the probability distribution shown in Figure 4.11, the particle weights will be updated according to the distance between the particle and detected clue location, d_k^i which can be computed with the following expression:

$$d_k^i = \sqrt{(x_k^i - x_k^\alpha)^2 + (y_k^i - y_k^\alpha)^2}. \quad (4.8)$$

Because of the asymmetrical form of the resulting position distribution, the particle weights will be updated according to one of two updating strategies. For particles existing within the escape radius E , the particles will be updated according to the following expression:

$$w_{k+1}^i = w_k^i \left(e^{-\left(d_k^i - E\right)^2 / E^2} \right). \quad (4.9)$$

For the remaining particles which lie outside of the radius E , the weights will be updated according to the next expression:

$$w_{k+1}^i = w_k^i \left(e^{-\left(d_k^i - E\right)^2 / E} \right). \quad (4.10)$$

4.2.3.3 Stage 3: Bearings-Only Measurement Model

The bearings-only measurement model refers to the stage of the search when the UUV detects a bearings measurement using passive sonar. A bearings measurement will only be returned when the UUV moves within passive sonar detection range. When this event occurs, the following bearings measurement can be processed [5]:

$$z = [\theta] = \arctan\left(\frac{x_k - x_{sk}}{y_k - y_{sk}}\right), \quad (4.11)$$

where (x_k, y_k) is the position of the contact at time k and (x_{sk}, y_{sk}) is the position of the UUV platform at time k . The measurement value, θ , is defined as the angle from East to the line of sight between the sensor and the target in the counter-clockwise direction (see Figure 4.12) [5].




Figure 4.12: Illustrative Example of Bearings Measurements

For the purposes of this thesis, the bearings-only measurement model will not be implemented, but only discussed. Once the UUV maneuvers within the passive sonar range and returns a bearing measurement, the search and detection phase of the submarine track and trail problem will be completed. Future work will consist of transferring from the search and detection phase to the track and trail phase of the problem. The track and trail problem will consist of following the contact within the passive sonar detection range in order to successfully track the contact with bearings measurements.

Chapter 5

Motion Planning for Dynamic Search Operations

Due to the nature of moving target search operations, the available information on the position of the target is constantly changing. As the UUV and the contact maneuver through the fixed search environment, the contact position distribution maintained by the UUV decision support system evolves as a result of the available sensor readings and the dynamics of the contact. Therefore, in order for a UUV to search autonomously within this dynamic environment, the decision support system aboard the platform must be able to develop motion plans online. These plans must account for varying distributions of possible target locations as well as for obstacles in the environment in order to plan traversable paths that most effectively narrow the search area. This planning process can be complex, because sensor measurements acquired at each moment in time are dependent on the location of the sensor with respect to the contact position. For this reason, each potential action will have a different effect on the resulting posterior distribution. Based on the stated difficulties of searching in the midst of dynamic information, the remainder of this chapter will describe how to use the available information to design an autonomous navigation system for use in UUV search operations.

5.1 Dynamic Decision System Architecture

The search and detection mission offers many challenges to the decision system within an autonomous vehicle. The ability to dynamically plan vehicle paths to track maneuvering

contacts and create safe vehicle trajectories in complex water environments is extremely important for mission success and overall system survivability. In order for a functional UUV to be aware of its dynamic surroundings and assess the impact of that knowledge on its plans, the central autonomous functions within the mission controller aboard the UUV must be integrated and tuned for effective real-time operation. This integration involves elements of planning, mapping, and awareness. To facilitate the analysis of the search and detection mission, this thesis will focus solely on UUV motion planning and route control. By ignoring the functional integration within the decision system, we can direct our attention to the improvement of the path planning algorithm used to search for a contact. In particular, the integration of localization, mapping, and situational awareness will not be considered. Instead it is assumed that the UUV location and the map of the environment are known with complete certainty throughout the scenario.

Although these simplifying assumptions remove the dynamics and uncertainty from the environment, the UUV must handle the complexities and uncertainties that stem from the dynamics of the available information about the target. For instance, as the single sensor platform searches for long periods of time, the estimate of the contact location can grow increasingly complex without the support of additional autonomous vehicles or detailed intelligence on the location of the target. Due to the endurance requirements and complexity of the information available, an effective autonomous search and detection mission will involve significant levels of replanning. In order for the decision system aboard the autonomous vehicle to create motion plans and effectively update those plans when necessary, it must operate within a dynamic system framework. The remainder of this section will present an overview of dynamic systems before introducing the specific dynamic system architecture used by the decision system in this thesis.

Using Liu and Chen's definition, a probabilistic dynamic system is a sequence of evolving probability distributions $\pi_k(x_k)$, indexed by discrete time $t = 0, 1, 2, \dots, k$. In the search and detection scenario, the system is a probabilistic dynamic system, because the contact position distribution, $\pi_k(x_k)$, changes from time k to $k+1$ due to the uncertain maneuvers of the contact and the incorporation of new data into the decision system [17]. Maintaining and updating these dynamic state distributions is critical in making intelligent decisions. The

information contained in these distributions can be manipulated by autonomous decision systems to evaluate decisions and execute the plan that results in the most optimal posterior distribution.

Dynamic systems such as search and detection operations can be modeled by integrating the state estimation techniques discussed in Chapter 3. The state estimation techniques enable decision systems to perform real-time analysis on the dynamic systems by updating the estimate of the state variables. As introduced in Chapter 3, our model consists of two parts: (1) observations, Z_k and (2) unobserved states, x_k , where the distribution of the unobserved states at time k can be approximated according to the following conditional probability:

$$\pi_k(x_k) = p(x_k|Z_k). \quad (5.1)$$

Therefore as new information is processed, the chosen filtering process incorporates the new information and updates the state estimate. The UUV planner can use the probabilistic state estimate to formulate intelligent motion plans.

The decision system designed in this thesis generates path plans through two coordinated levels of dynamic systems: a primary control system and the path generation function. The primary dynamic system consists of a feedback control loop. This control system monitors the search environment through on-board sensors and forwards motion plans to the UUV guidance navigation and control system to initiate UUV movements. The information passing through this system consists of the contact position distribution. With the assistance of the sensor measurements, the position distribution flowing through the system will be updated and used as a means to develop path plans. This flow of contact state information will be modeled with state estimation techniques as discussed in Chapters 3 and 4. Specifically, for the remainder of the thesis, all state estimation carried out during the control process will be accomplished with the application of the particle filter as introduced in Chapter 4.

Using the state information provided by the particle filter, the real-time UUV motion planning system will follow the dynamic system architecture as seen in Figure 5.1.




Figure 5.1: High Level UUV Dynamic Decision System

This high-level dynamic decision system provides the path commands to the UUV guidance navigation and control system. Specifically, based on strategic inputs, such as geographical, intelligence, and localization information, the state estimation filter is provided with an initial contact position distribution. The state estimation filter will format the distribution according to the appropriate state estimation technique. This initial state information as well as the current position of the UUV is received by the path generator. Upon procuring the current state distribution and UUV position, the path generator evaluates the information available and returns the best path with respect to the stated search objective. After generation, the planned path is executed and sent to the UUV guidance navigation and control to initiate the path. After completing the path, the UUV monitors its current status by processing a new sensor reading and updating the contact position distribution through the state estimation filter. At this point, the feedback control process continues until the contact's position has been identified.

In addition to the primary control system architecture seen in Figure 5.1, a secondary dynamic system exists within the path generation function. The path generation function evaluates the information it receives through a second dynamic process in order to return the path that best meets the search objective. The dynamic path planner as seen in Figure 5.2 is necessary to produce feasible motion plans, because the available information changes during the

execution of an action. In order to manage the dynamic information, the state distribution must be resampled and evaluated before generating an additional step in the path.




Figure 5.2: Dynamic Path Generation for UUV Decision System

According to the system design, the path generation function takes the following parameters: current state distribution from the particle filter and the current UUV position from an assumed localization function. Based on the value and characteristics of the inputs, the path planner generates a feasible action space. The action space consists of the feasible maneuvers that the UUV considers at the current time. After generating the action space, the planner implements a separate “planning” particle filter to simulate one action from the available action space and predict the future distribution. Upon proposing the future distribution and UUV position, the simulated filter will resample the particles in the distribution under the assumptions of measurement model 1 (i.e., assume no measurement is observed at the current location in the

path). Depending on the desired path depth, the simulated state distribution and enumerated path will either be evaluated (if desired path depth has been reached) or will be returned to the beginning of the planning algorithm as inputs for the next level in the path (if desired path depth has not been reached). This recursive path enumeration process will continue until all paths have been enumerated and evaluated. At the conclusion of the enumeration process, the path that leads to the most desired posterior distribution (i.e., the best estimate of the contact location) at the stated path depth will be forwarded as the optimal motion plan. Once received by the UUV guidance navigation and control, the motion plan will be executed by the vehicle controls.

The dynamic path planning architecture as depicted above allows the decision system aboard the UUV to look farther into the search environment when determining motion plans. Without a dynamic process in place, a path planner would either design extended paths that fail to consider the changing information or propose limited routes that only consider the adjacent information space in hopes that the information would not change significantly during the execution of the route. Neither of these alternatives will provide intelligent motion plans. To improve the motion plans over the duration of the search, the path planner needs to consider as much of the available information as possible. By looking further into the future and considering more of the available information, the planner can make better decisions at the current time step. In other words, the UUV needs to consider future movements and contact position distributions when determining where to move next: the maneuver that is optimal after one time step may not be optimal when determining how the UUV should maneuver to optimize after several time steps. Therefore, the path planner simulates future UUV actions and evaluates what series of actions will maximize the search objective.

5.2 Motion Planning Formulation

With the path planning structure established, the following sections seek to describe how to employ the functions within the path generator to form effective paths. While the majority of this chapter will describe how to use the particle filter distribution to guide the planning process, the motion planning problem must first be formulated. In particular, this section will define the appropriate objective function in which the path generation function seeks to solve as well as the

feasible decision space which provides the options with which to evaluate the expressed objective function.

5.2.1 Objective Function

In anti-submarine operations, the initial objective involves detecting the location of a moving contact. The challenge thus becomes designing a mathematical model that locates the contact within the adversary vessel's projected area of operation. At first thought, it might seem appropriate to design a model with an objective function that seeks to maximize the probability of detection. Although this objective may lead to quick detection, it can also potentially lead to distributions with greater uncertainty when detection is not made early in the search. Therefore if the UUV is expected to conduct extensive searches, the objective function should instead be designed in a way that produces more accurate estimates of the contact's location. An objective function that minimizes the uncertainty in the location of the contact achieves more accurate estimates by containing the distribution. In other words, the UUV should travel along the path that "shrinks" the search area by condensing the possible locations in which the contact could be positioned.

Based on that logic, we must determine a mathematical model that results in a reduction in the uncertainty in the location of the contact. In order to obtain these results, we must be able to quantify uncertainty. A common method to accomplish this in information theory is to determine the belief state's entropy. Assuming we know the probability distribution of state s of a system, entropy can be expressed as [13]:

$$H(P) = - \sum_{s \in S} P(s) \log P(s). \quad (5.2)$$

The complement of entropy is referred to as information. While statistical entropy is a probabilistic measure of uncertainty, information is a measure of a reduction in that uncertainty. If information about the state of the system is obtained through observations, these observations

will reduce our uncertainty about the system's state by excluding, or reducing the probability of a number of states. Therefore, the information, I received from an observation is equal to the degree to which uncertainty is reduced [13]:

$$I = H_k - H_{k-1} . \quad (5.3)$$

Based on this reasoning, the objective function for this search and detection problem is formulated as the process of evaluating different sensing actions that the UUV can take and choosing the action, u_k , at time k that maximizes the information acquired by the UUV. The action that maximizes the information can be expressed as [17]:

$$u_k = \max_u I_{k+1} \quad (5.4)$$

By maximizing the information gain, the UUV executes actions that reduce the uncertainty in the location of the contact. The challenge in implementing this concept in practice is to develop a methodology for quantifying the expected information for different sensing actions and evaluating them in a computationally feasible manner given limited *a priori* information [17].

For extended Kalman filter applications, the expected information is contained in the estimate error covariance matrix, P . Specifically, under the assumptions discussed in detail in [17], the inverse of the error covariance P will be an estimate for the Fisher information of the system, which is [17]

$$I \approx P^{-1} .$$

Because the particle filters do not maintain an estimate of the estimate error covariance, this thesis must use the particle filter distribution to estimate the uncertainty for each state variable. With that said, the estimate for the covariance matrix at each time step can be expressed as:

$$P = \begin{bmatrix} s_{xx}^2 & 0 & 0 & 0 \\ 0 & s_{yy}^2 & 0 & 0 \\ 0 & 0 & s_{\ddot{x}\ddot{x}}^2 & 0 \\ 0 & 0 & 0 & s_{\ddot{y}\ddot{y}}^2 \end{bmatrix}, \quad (5.5)$$

where the variances are estimated with the N particle filter samples according to the equation for sample variance:

$$s^2 = \frac{\sum_{i=1}^N (x_i - \bar{x})^2}{N-1}. \quad (5.6)$$

At this point, the information is in the form of a matrix. Unfortunately, for the evaluation purposes, we require a metric to quantify the information. Therefore, we will implement the metric, $C(P)$ defined by Feder, Leonard, and Smith, which gives the total area of the error ellipse and thus is a measure of the confidence in the location of the contact [17]:

$$C(P) = \pi \sqrt{\det(P)}. \quad (5.7)$$

Using the variance estimates and formation of this cost metric, the specific objective function can be formulated and evaluated. In particular, the objective of the path generation function is to produce a set of actions u_k that minimizes the area of the error ellipse as seen in the following objective function [17]:

$$\min C(P) = \min \pi \sqrt{\det(P)}. \quad (5.8)$$

5.2.2 Decision Space

The second part of the dynamic motion planning formulation involves the selection of the decision space. For the path generation function, the decision space contains all the available paths from which to evaluate the objective function in (5.8). For the remainder of this chapter, the decision space will be referred to as:

$$u = \begin{bmatrix} R_1 \\ R_2 \\ \vdots \\ R_m \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} & \cdots & A_{1d} \\ A_{21} & A_{22} & \cdots & A_{2d} \\ \vdots & \vdots & & \vdots \\ A_{m1} & A_{m2} & \cdots & A_{md} \end{bmatrix}, \quad (5.9)$$

where each R represents a potential path that is made up of a set of d action inputs A . The number of action inputs comprising each path is dependent on the desired path depth d which will be consistent for all paths. The number of available paths m to be evaluated is restricted by the number of possible actions for each input, A . In general, each of the potential actions within a path can be defined as:

$$A = \begin{bmatrix} \theta \\ t \end{bmatrix}, \quad (5.10)$$

where θ is the UUV heading and t is the time duration of that heading command [19]. The generation of these actions depends on the search strategy being implemented. Due to environmental constraints, the UUV may need to avoid obstacles and as a result additional maneuvering may need to take place along each action sequence. Further discussion into UUV obstacle avoidance as well as the specifics of action generation functions will be introduced in the next section.

5.3 Motion Planning Strategy

With the objective function formulated, this section focuses on the applied motion planning strategy. When evaluating decisions, it is necessary to understand that different inputs can lead to more accurate estimates of the state parameters. Therefore, we must attempt to optimize the input in a manner that leads to the best possible estimates. This section will explain the details of the action generation function by describing the most advantageous choice of inputs, the simulation of those inputs, and the evaluation of each of the generated inputs.

5.3.1 Action Space Generation

One of the most important decisions for the UUV decision system during the motion planning process involves the action space selection. Due to the computational complexity involved in simulating future actions, the search depth is limited because the size of the action space grows exponentially with each additional unit of depth. Therefore, it would be beneficial to prune actions within the path generation tree that lead to little information gain in order for the path depth to be expanded further into information rich sections. Unfortunately, in many previous applications, the action space has been generated using discrete sets of actions. Often times this action space generation leads to searches into areas with no information to be gained. In response, we propose a cluster-based action space generation function that examines the complex search space for available information when defining feasible actions. In the remainder of this section, we present these two action space generation strategies: the discrete action space and the cluster-based action space. We will attempt to show how the proposed cluster-based action space generates deeper paths by taking advantage of the information available and ignoring actions with no information to be gained.

5.3.1.1 Discrete Action Space

In many current applications of adaptive motion planning, the available actions originate from a fixed, discrete set. The action set is often constrained in the degrees of freedom allowed.

For example, during a path planning process, an autonomous vehicle could be limited to a discrete number of heading directions, as seen in the following general expression:

$$A = \begin{bmatrix} n \times \frac{2\pi}{df} \text{ Radians} \\ t \text{ time units} \end{bmatrix}, n = 0, 1, 2, \dots, df. \quad (5.11)$$

According to this discrete action space generation, the number of actions generated is equal to the maximum allowable heading directions. Each of these actions will be executed for t time units before re-clustering the data and adding another unit of depth to the current path.

5.3.1.2 Cluster-Based Action Space

Due to the structure of the discrete action generation, a great deal of computation could potentially be spent evaluating areas containing no information. The action space generation function could make more efficient use of available computation power, if branches into areas of no information gain could be removed from the search tree. The following section will present an alternative action space generation function that proposes a method for considering only the actions that lead to information gain. The alternative action space constrains the actions by clustering the points of potential information gain and only planning paths towards those clusters of information. This section will detail the cluster-based action space by explaining how the clusters are created, selected, and used for the purposes of path planning.

Clustering is a division of data into groups of similar objects. For the purposes of the path generation function, we seek to group points of information gain together according to location (i.e., x-y coordinates). Due to the sample-based structure of the particle filter, geometric clustering algorithms can be applied to the positions of the samples comprising the particle filter distribution. In Figure 5.3, the clustering algorithm used to cluster the particles within the sample distribution can be referenced [3].

Clustering Algorithm

Input: Particle positions (x_k^i, y_k^i) where $i \in \{1...M\}$ at time k

A cluster size parameter $Tdist$

Steps:

1. Initialize all assignment values to one, $w_i = 1.0$
2. Set the initial number of clusters equal to zero, $numClusters = 0$
3. Iterate over all points, finding centroid of all other points lying within specified threshold of each point

For all $i \in \{1...M\}$

If $w_i > 0$

$$xtmp = x_k^i$$

$$ytmp = y_k^i$$

$$mtmp = 1$$

a. Count the number of points that lie within $Tdist$ of particle i

For all $j \in \{i...M\}$

$$\text{if } \left(x_k^j - x_k^i \right)^2 + \left(y_k^j - y_k^i \right)^2 < Tdist^2$$

$$xtmp = xtmp + x_k^j$$

$$ytmp = ytmp + y_k^j$$

$$mtmp = mtmp + 1$$

$$w_j = 0$$

end

end

b. If more than one particle associated with the current particle i

If $mtmp > 1$

Increment the number of clusters

$$K = K + 1$$

Find the centers of gravity of the associated data points

$$c_{xK} = \frac{xtmp}{mtmp}$$

$$c_{yK} = \frac{ytmp}{mtmp}$$

Save the number of associated points within the cluster

$$c_{wK} = mtmp$$

End

End

Figure 5.3: Clustering Algorithm

This clustering algorithm compares the relative distances between particles and partitions the particles according to a cluster size parameter, $Tdist$. Therefore, as the state distribution

becomes dispersed and multi-modal, a clustering algorithm can divide the particles into distinctive sets which identify the regions with high probabilities of detection. Each of these K subsets, referred to as clusters $C_{i \in \{1, \dots, K\}}$, will be represented by a centroid, (i.e., the mean of the particles assigned to the cluster) [19]:

$$c_i = \begin{bmatrix} c_{x_i} \\ c_{y_i} \end{bmatrix}, \quad \forall i \in \{1, \dots, K\}. \quad (5.12)$$

These cluster centroids represent the centers of mass within the particle distribution as illustrated in Figure 5.4. This visual representation of the clustered particle distribution demonstrates how the particles are partitioned according to relative distance. For a more detailed explanation of clustering algorithms reference Berkhin [3].




Figure 5.4: Illustration of Clustered Particle Filter Distribution

After the cluster algorithm is implemented, the computed centroids serve as the catalyst for the action generation function. These centroids establish goal locations for the autonomous vehicle to head towards. Thus, the generated actions are a function of the current UUV position and velocity as well as the position of the cluster centroids, as seen in the following expression:

$$A_i = \begin{bmatrix} \theta \\ t \end{bmatrix} = \begin{bmatrix} \arctan\left(\frac{c_{x_i} - x_s}{c_{y_i} - y_s}\right) \text{radians} \\ \frac{\sqrt{(c_{y_i} - y_s)^2 + (c_{x_i} - x_s)^2}}{v_s} \text{time units} \end{bmatrix} \quad \forall i \in \{1, \dots, K\}. \quad (5.13)$$

In this expression, we find that the number of actions generated is dependent on the number of clusters. For each action, the heading input is the angular measurement from the current UUV position to the selected cluster centroid. In addition, unlike the discrete action generation, the durations of the actions are not uniform. Instead the cluster-based action set contains inputs with different durations depending on the distance to the given cluster centroid and the speed of the UUV.

5.3.2 Monte Carlo Simulation of Search Environment

After the formation of actions, the next step in the motion planning algorithm is to evaluate the action space according to the objective function in (5.8). Due to the complexity of the system, no closed form solution exists to solve the mathematical program. Without a general formula to solve the system, the system needs to be simulated in order to estimate the effect of the action space on the state distribution. In the same way that the particle filter is used for the UUV controls, a particle filter simulation can be used for the purposes of planning. With the prior real-time distribution, the contact motion model, and measurement models, Monte Carlo simulations can be applied to the samples to obtain a simulated posterior distribution which can then be used to evaluate the objective function. To demonstrate how a particle filter can be used for path planning, the remainder of this section will detail the three main components of the simulation: the action execution, the sample propagation, and the measurement update.

5.3.2.1 Action Simulation

The first part of the planning simulation involves executing the generated actions. Each of the UUV actions must be simulated in order to determine the potential effectiveness of the

movement. While the control inputs within the action space consist of headings and movement durations, these actions may require modifications prior to their actual execution. Due to limitations in computation capabilities, constraints within the search environment, or dynamics of changing information, these initial action inputs change prior to execution. Specifically, the action space can change in one of the following ways: limiting the number of available actions, re-routing to avoid obstacles, or constraining the duration of the action to account for evolving information.

5.3.2.1.1 Cluster Selection

One of the constraining factors that must be addressed prior to action execution involves the available computation time. The process of simulating UUV movements and state estimation filters can be very computationally expensive. As the action space grows in size, the number of necessary simulations increases. In an effort to reduce the computational workload and number of simulations, limitations may be placed on the size of the action space. For discrete action space generation, the size of the action space can be easily reduced by constraining the degrees of freedom. Conversely, a cluster-based action space cannot be as easily reduced. Because the actions are not uniform in duration or angular spacing, greater thought must be put into which actions should be considered.

Although we cannot truly evaluate the effect of each action without actually simulating each of them, we must determine a way to constrain the action space in a way that minimizes the possibility of overlooking sections of available information. The constraining algorithm identifies a smaller set of clusters that covers the available information in a way that limits redundancy of actions. The redundancy of actions can be limited by spatially constraining the clusters used to create the action space. Basically, the algorithm seeks to remove any clusters that produce similar action headings and travel durations. This process is accomplished using the following algorithm (see Figure 5.5):

Cluster Constraint Algorithm

Input: Cluster centroids $c_i = \begin{bmatrix} c_{x_i} \\ c_{y_i} \end{bmatrix}$ where $i \in \{1 \dots K\}$

Steps:

1. Initialize indices

considerList = [1 : K]

validcluster = []

2. While the number of clusters in the consideration list is greater than zero

WHILE considerList *IS NOT* empty

a. Find the minimum distance from the sensor position to the cluster centroids

minClusterDist = 10^9

For all $i \in \{1 \dots K\}$,

$$\text{clusterDist} = \sqrt{\left(c_{x_i} - x_s\right)^2 + \left(c_{y_i} - y_s\right)^2}$$

If clusterDist < minClusterDist

minClusterDist = clusterDist

minClusterIndex = i

End

End

b. Add the closest cluster to the valid cluster list

push validCluster, c using minClusterIndex

c. Find the clusters within the minimum distance radius of the closest cluster

For all $i \in \{1 \dots K\}$,

$$\text{If } \sqrt{\left(c_{x_i} - c_{x_{\text{minClusterIndex}}}\right)^2 + \left(c_{y_i} - c_{y_{\text{minClusterIndex}}}\right)^2} \leq \text{minClusterDist}$$

Add cluster to the invalid cluster list

push invalidCluster, c using i

End

End

d. Remove all invalid clusters from the consideration list

For all $j \in \{\text{invalidCluster}\}$,

pull considerList, c using j

End

END

Figure 5.5: Constraining Action Space Algorithm

This algorithm constrains the action space by considering only one cluster in the general vicinity. This is accomplished by selecting the closest cluster and then disregarding all

remaining clusters that exist within a radius surrounding the cluster, where the radius is equal to the distance between the selected cluster and the UUV position. This process is continually repeated for the remaining clusters until all clusters have been eliminated from consideration or selected for use in generating the action space. An example of the process of constraining the action space is illustrated in Figure 5.6 seen below.




Figure 5.6: Illustrative Example of Constraining Action Space Algorithm

5.3.2.1.2 Obstacle Avoidance

In order to steer clear of any environmental obstructions or boundaries, the motion plans must provide obstacle avoidance. As the actions are currently stated, each movement consists of constant linear motion over the established duration time. In reality, search environments possess obstacles and boundaries which will cause collisions if proper modifications are not made to the initial actions. In the following paragraphs, we will explain how to segment the original actions to facilitate maneuvers around any obstacles. Specifically, the algorithm will

find the shortest unobstructed path to the intended destination (i.e., the location of the cluster centroid).

At this tactical point in the motion planning process, we are concerned with the problem of moving the UUV from its current position to the goal destination which has been established during the strategic development of the initial action space. Based on this problem statement, this sub-problem takes the form of a traditional static motion planning problem. Assuming that the map of the search environment is known with complete certainty, the problem is to find the shortest path from the current position to the goal destination while avoiding all impassable terrain. Extensive research has been carried out within this problem area and is better known as shortest path algorithms. While several approaches exist, this motion planner will implement a numerical potential field method in an attempt to limit the amount of computational effort required during the tactical action execution.

The use of potential functions for obstacle avoidance works by constructing a function called the potential that has a minimum at the goal destination and a high value at obstacle locations [1]. Everywhere throughout the search space, the function slopes down toward the goal configuration so that the UUV can reach the goal configuration from any other position on the map by following the negative gradient of the potential [1]. For the action execution function, the potential field can be achieved by dividing the map into discrete cells which represent the potential positions for a maneuvering UUV, as shown in Figure 5.7.




Figure 5.7: Discrete Cell Description of Search Environment

By discretizing the search area, the UUV motion can be restricted to the positions of the discrete cells. Depending on the degrees of freedom allowed, the potential function can be formed by computing the distance covered by the discrete steps taken from the current UUV position to the goal state and by giving discrete cells over obstacle positions an arbitrarily large number, M . The specific algorithm for the numerical potential field generation can be seen in Figure 5.8 [27]. This algorithm demonstrates how a numerical potential field can be created over a discrete grid.

Numerical Potential Field Algorithm

Input: N discrete cell locations, x_i where $i \in \{1...N\}$

Goal cell location, x_g

Potential field function, U

Distance matrix, D_{ij} where $i \in \{1...N\}$ and $j \in \{\text{neighbors of } i\}$

1. Initialize the value function of the numerical potential field

For all passable cells

$U(x_i) = -1$

For all impassable cells

$U(x_i) = M$

For the goal cell

$U(x_g) = 0$

2. Initialize priority queue, Q by adding the goal cell. Where Q is a priority queue sorted by the potential field function

push Q, x_g using 0

3. Update numerical potential field by adding the grid cell with the lowest potential function value from the priority queue

while Q is not empty

Remove cell with the minimum potential function value from the priority queue

$x_i = \min Q$

For all neighbors x_j of x_i

if $U(x_j) = -1$

Update the potential function values of the neighboring cells

$U(x_j) = U(x_i) + D_{ij}$

Add the neighboring cells to the priority queue

push Q, x_j using $U(x_j)$

end

end

end

Figure 5.8: Numerical Potential Field Algorithm

An application of this algorithm to the grid map seen in Figure 5.7 produces the potential field in Figure 5.9. This depiction of the numerical potential field represents the scenario map where each cell is connected to its eight neighbors. The distance needed to travel laterally in any

direction is equal to 1 and the distance needed to travel diagonally in any direction is equal to $\sqrt{2}$. Based on these travel distances, the numerical potential field value, U , within each of the cells represents the distance needed to travel from the given cell to the goal cell.




Figure 5.9: Numerical Potential Field for Search Environment

Using this numerical potential field, the shortest path from any grid cell on the map to the goal cell can be found by moving to the neighboring grid cell with a lower value. For example, in Figure 5.10 cell number 42 (see Figure 5.7) has a value of 9.071, and therefore must travel through 9.071 units prior to reaching the goal cell.




Figure 5.10: Numerical Potential Field for Computation of Shortest Path

In this example, the shortest path is formed by considering the values of the neighboring cells and moving in the direction of the smallest value as seen in Figure 5.10. Due to the structure of the numerical potential field, this negative descent guarantees a shortest path to the destination. In addition, due to the arbitrarily large values in the cells containing obstacles or boundaries, we are assured that no path will cross impassable regions.

Because the numerical potential field provides the shortest collision-free path from any point in the search environment to a specified goal position, it can be applied to the action execution function. In essence, the action generation function creates a list of goal destinations, and the numerical potential field ensures that the UUV travels on the shortest path to that goal while avoiding any obstacles in the area. With these modifications to the original actions, the action space must be expressed in a different way. Instead of consisting of strictly a heading and a duration, the actions will be segmented to account for the discrete movements towards the goal destination. The discrete actions will form a path that will be expressed as:

$$A = [a_1 \quad a_2 \quad \cdots \quad a_G], \quad (5.14)$$

where the number of discrete actions in the action set A is equal to the total number of grid cells in the path, G .

Even though now all motion is restricted to the available actions, the cluster-based action space is still much different than a discrete action generator. Because the cluster-based generator looks further into the search space to identify collections of information, it will require less computational effort than a discrete action generator to plan the same path. A discrete action generator has a high computational burden, because it considers all actions. Conversely, the use of clusters eliminates many unnecessary paths and considers only the shortest paths to the available information.

5.3.2.1.3 Branching Strategy

To account for the dynamic information within the search environment, the action sequences must be modified further before simulating an action. As the simulated actions are currently stated, the UUV will travel the entire length of the path prior to re-examining the available information. At the conclusion of the movement, the UUV will be located in the area in which a cluster of information existed when the action sequence was generated. Depending on the time to travel to the goal destination, the information could have changed significantly. Because the contact is moving according to an assumed speed and motion model, we must account for the evolving contact position distribution. Therefore the length of the current action set must be restricted, because the full action duration cannot be simulated without possibly missing information that evolves over time.

Even though an adaptive motion planning simulation is certain to miss pieces of information during a prolonged motion plan, it is necessary to consider as much information as possible to achieve a reasonable plan. Consequently, the duration of the current actions must be constrained to allow for a reevaluation of the search space. In other words, only a portion of the action sequence will be simulated before stopping to recluster the updated particle distribution

and establishing new actions based on the updated cluster locations. This reevaluation allows the simulated actions to keep pace with the changing contact position distribution.

Based on this reasoning, the final question remaining is when to reassess the environment. This reassessment point is a function of the shortest path distance to the selected cluster centroid and the relative speed of the UUV in relation to contact. In general, smaller durations lead to better results while longer durations without reevaluation increase the opportunity to miss information. On the other hand, frequent reassessment increases the computational burden on the motion planning system. As a result, a constraint must be designed that balances these two effects by producing an intermediate action that limits the amount of missed information. The proposed heuristic limits the length of the simulated action by the following ratio of the shortest path distance $\left(1 - \frac{1}{v}\right)$, where v equals the relative speed of the UUV in relation to the contact. For this problem, it is assumed that the relative speed must be greater than 1. Because the information changes with respect to the assumed dynamics of the contact, the ratio is a function of the relative speed. It ensures that a UUV with a higher speed than the contact will travel further into the initial action set before reevaluating the search space, while a lower relative speed causes reevaluation to take place earlier in the action set. With that said, the further restricted action set can be expression as the following:

$$A = \begin{bmatrix} a_1 & a_2 & \cdots & a_{G\left(1 - \frac{1}{v}\right)} \end{bmatrix}, \quad (5.15)$$

where the number of discrete actions is the original number of discrete actions G reduced by $\left(1 - \frac{1}{v}\right)$.

5.3.2.2 Particle Propagation

Upon executing the UUV action, the next step in the simulation is to propagate the particles within the simulated filter. The simulated particle distribution provides the foundation for all planning decisions and therefore must mimic the real-time particle filter used for the UUV controls. Although the simulated filter will be handled in the same manner as the real-time particle filter, the propagation step differs in the durations of the particle propagations. While the particle filter used for controlling the UUV moves particles forward according to set time increments, the duration of the simulated propagation step is dependent on the duration of the current action. The simulated propagation step can therefore be expressed as:

$$X_{k+t}^i = \begin{bmatrix} x_{k+t}^i \\ y_{k+t}^i \\ \dot{x}_{k+t}^i \\ \dot{y}_{k+t}^i \end{bmatrix} = \begin{bmatrix} x_k^i + \dot{x}_{k+t}^i \Delta t \\ y_k^i + \dot{y}_{k+t}^i \Delta t \\ \sim N(\dot{x}_k^i, \sigma_w^2) \\ \sim N(\dot{y}_k^i, \sigma_w^2) \end{bmatrix} \quad \forall \quad i=1...M. \quad (5.16)$$

During the simulated propagation step, if any particle completes the passage through the channel or seaway then it will adjust its heading towards one of three objective destinations across the sea (i.e., Goal A, Goal B, Goal C). As explained in Chapter 4, each of these destinations are assumed equally as likely, and therefore, an equal number of particles will head towards each of the three destinations.

5.3.2.3 Measurement Update

The final step in the simulation is to update the distribution based on the proposed UUV location after the action execution. At this new position, the particles are resampled according to a previously stated measurement model. Specifically, during the simulation, it is assumed that a null measurement is obtained at each simulated step. In comparing the predefined measurement models, the null measurement model provides the least amount of information. Therefore, after each simulated action and particle propagation, Measurement Model 1 is implemented and the particles are resampled accordingly. The simulated measurement model will decrease the

weights of the simulated particles within the detection region as shown in Chapter 3. By assuming a null measurement after each action, the plans are generated under the worst-case scenario (i.e., least amount of information gained). Any other measurement such as a clue deposit that is found during the course of the action can help to further reduce the uncertainty of the distribution.

5.3.3 Path Enumeration

Following the simulation of the filter and the UUV action, the motion planning algorithm continues by adding another step in the path. After evaluating the updated particle filter distribution, the path planner creates a new set of actions based on the simulated distribution. At these reevaluation points, the simulated distribution is reclustered and a new branch is inserted in the path enumeration tree to account for the latest action space generation. These reevaluation points can be seen at the branches of the search tree as depicted in Figure 5.11 below.




Figure 5.11: Path Enumeration Illustration

Using this search tree structure, the m feasible paths can be produced from the initial UUV position. Based on the tree structure, each additional depth in the path leads to an exponential expansion of the number of potential paths. Although longer paths increase the coverage of the search space, they also exponentially expand the number of paths to evaluate. Therefore, the two biggest questions in developing these paths are how deep should the paths search and how should the paths be enumerated? The remainder of this section will cover both of these topics and attempt to provide a solution strategy that produces effective paths that can be solved in a computationally feasible manner.

5.3.3.1 Path Depth Selection

The first question to be considered is how deep should the paths search? Depending on the stated path depth d , the path generation function will produce paths with d reevaluation points. Therefore for larger values of d , the paths will search farther into the search space by reassessing the environment numerous times during the enumeration process. After simulating the path and reassessing the simulated distribution d times, the search objective will be evaluated using the posterior simulated distribution. In other words, the path depth also determines when the objective function will be evaluated. Because information is gained in the search problem by covering areas of the contact position distribution, paths with greater depths that cover more of the search area will provide the most information to the decision system. Deeper paths provide better motion plans, because they consider a greater section of the search area when evaluating the impact of the path. When paths used for motion planning are relatively short, they only consider the local search space and as a result could overlook a more distant, better solution.

Generating paths with larger path depths does not come without a cost. Although it is advantageous to look deeper into the search area when evaluating the impact of an action, it also creates additional computation. Computation time is typically exponential in the depth of the paths under consideration. For this specific search application, a great deal of computational effort is required each time a particle filter is simulated. During the path enumeration process, the particle filter must be simulated at a minimum at each reevaluation point (see Figure 5.11).

Due to the complexity of the simulations, the path generation function cannot return paths with sizeable path depths. The exact number of simulations that must be processed is dependent on the path enumeration method used. After introducing this remaining factor in computation times, a comparison of attainable path depths can be seen in the computation table in the next section.

5.3.3.2 Enumeration Methods

Planning operations require searching through all the possible routes to find the most optimal, or efficient route among all the possibilities. As mentioned in the earlier sections, the path generation application must simulate the movements and environment of a UUV to a given search depth. In order for the planner to evaluate the effectiveness of the available paths, it must simulate each path and assess the posterior contact position distribution created by each set of actions. The efficiency of these simulation processes is directly tied to the number of steps in the simulation required to evaluate the paths. The number of simulations varies depending on the number of actions in each action set, A and the depth of the paths, d . If the number of actions in the action set were equivalent at each branch point, enumerating the paths iteratively would require dA^d simulations of particles. On the other hand, a recursive search algorithm would only require $\sum_{i=1}^d A^i$ simulations. In order to avoid strict enumeration, this process of enumerating paths is accomplished with conventional recursive search algorithms. The search algorithms will produce a tree structure consisting of nodes placed at choice points and edges connecting each node in the graph. A choice point is anywhere in the graph where a decision must be made as to where to go next and the edges represent the actions taken to arrive at the next choice point [11]. In Figure 5.11 above, the initial point represents the initial location of the UUV and the remaining decision points (i.e., reevaluation points) suggest potential future locations of the UUV. The search begins at the point placed at the initial location, and ends when a specified search depth has been reached (i.e., a given number of choice points have been reached). The remainder of this section will cover two algorithms used to search through a tree structure and discuss which algorithm was chosen to complete the path enumeration for this scenario.

5.3.3.2.1 Breadth-First Search

One possible search strategy that could be implemented is the breadth-first search [11]. This strategy begins with the source node and expands every one of its neighboring nodes before spreading out to their respective successors (see Figure 5.12). In other words, it generates all the actions at the current search depth before looking any further into a path. Because each action creates a unique posterior distribution and all the nodes at the current depth are expanded before an additional unit of depth is added, one obvious problem to this approach is the memory requirements for large search spaces [21]. Due to the dynamics of the available information, the posterior distributions differ at each of the decision points. In order to simulate from the correct distributions, each path must store its current distribution. As the path depth grows, the number of paths expands exponentially. With each additional unit of depth, the memory requirements for the path generation function will increase significantly.




Figure 5.12: Breadth-First Path Enumeration Illustration

5.3.3.2.2 Depth-First Search

Another search strategy is the depth-first search. This search algorithm is the complement to the breadth-first search, because it visits all of a node's descendants prior to any of its siblings rather than visiting all siblings before any children [11]. Therefore, this particular method will enumerate an entire path to the specified search depth before it goes back and expands nodes at shallower levels (see Figure 5.13). This particular enumeration method reduces the memory requirements needed by a breadth-first search, because only a single path from source to destination is stored—at most d posterior distributions will need to be stored during the enumeration process (i.e., one posterior distribution for each depth in the search). Based on this method's ability to reduce the required computation and memory necessary to enumerate the paths, the path generation function will enumerate path simulations using this recursive depth-first search.




Figure 5.13: Depth-First Path Enumeration Illustration

5.4 Path Evaluation

After the decision space has been finalized, each of the enumerated paths must be evaluated to determine the “best” motion plan. The “best” motion plan for this problem is the set of actions that leads to the posterior distribution that optimizes the objective function expressed in (5.8). Once a path has been simulated out to the stated depth, the objective function value must be computed from the posterior distribution. Because the simulated distribution consists of samples (i.e., particles), the evaluation of the distribution will begin by determining the sample covariance based on the state parameters of the particle samples. Using this sample covariance, the rest of the objective function can be computed by taking the determinant of the covariance and multiplying by π . This calculation provides the total area of the error ellipse and thus is a measure of the confidence in the location of the contact.

Following the evaluation of all the enumerated paths, the path that leads to the minimum cost function value will be chosen as the optimal motion plan. Based on the structure of the objective function, this path plan should lead to a reduction in the uncertainty in the location of the contact. After dynamically executing several maneuvers, the posterior distribution should begin to shrink and provide a more accurate estimate of the actual location of the contact.

Chapter 6

Search and Detection Simulation Results and Analysis

This chapter will present the results and analysis of the motion planning algorithm introduced in Chapter 5. This chapter will begin by discussing the implemented computer simulation and the selected parameters used to test the algorithm performance. The next section of the chapter will illustrate the usefulness of the motion planning algorithm by first introducing a simplified search algorithm and then showing the benefits of the improvements until the complete motion planning algorithm described in Chapter 5 is tested. The results of several simulations will be shown to illustrate the effectiveness of the dynamic action spaces in motion planning.

6.1 General Simulation Design Process

The objective of this section is to introduce the simulation used to compare the effectiveness of search algorithms. The first part of this section will discuss the computer simulation designed for the purposes of this research. It will document the software used to produce the results found throughout the remainder of the chapter. The next part of this section reviews the search and detection scenario presented in Chapter 2 and introduces the parameters of the scenario that will be of interest during the testing of the search algorithms.

6.1.1 Computer Simulation

The scenario introduced in Chapter 2 was implemented in Matlab® and used to track a target with the measurement models discussed in Chapters 3 and 4. The implementation of the particle filter in Matlab® maintains a sample-based distribution of the contact position. The software used to maintain and update the particle filter distribution was based on a modified version of ReBEL software. ReBEL is a Matlab® toolkit of functions and scripts, designed to facilitate sequential Bayesian estimation in general state space models. The code is developed and maintained by Rudolph van der Merwe at the OGI School of Science & Engineering at Oregon Health & Science University [24]. The original ReBEL software package was modified to handle the unique measurement models discussed in Chapter 4.

With the particle filter software providing the information on the contact position distribution, I implemented the motion planning algorithm designed in Chapter 5 in additional Matlab® software. This motion planning software uses the sample-based particle filter distribution to be used in conjunction with the modeled scenario and measurements. The software simulates the particle distributions and allows the on-board sensor to look at the probability distribution several steps into the future. The executed motion plans are the motion plans that lead to the optimal posterior distribution.

6.1.2 Simulation Scenarios

This section offers a brief description of each of the scenarios analyzed in this chapter. These simulations produce different events that may occur during a search and detection mission. The chosen scenarios seek to create situations that identify the strengths and weaknesses within the search algorithms examined throughout the rest of this chapter.

With the scenario described in Chapter 2 serving as the backdrop, the following events will be tested:

Scenario 1: Contact chooses the seaway route to sea. According to the assumed probabilistic motion model, the contact would choose this route with a probability of 75%.

Scenario 2: Contact chooses channel route to sea. This is the route with only a 25% probability according to the assumed probabilistic motion model.

Each version of the proposed search algorithms will be simulated twenty times for each of the scenarios described above. At the conclusion of these simulations, the performance of the search algorithms will be evaluated using the following metrics: average time until detection, total number of detections, and average final variance in the position distribution (for runs with no detection prior to the fixed run time).

6.1.3 Simulation Design Factors

Prior to carrying out the simulations, design factors must be identified and assigned values. For this particular scenario, the main design factors that could affect the performance of any search algorithm include: the relative speed of the UUV in relation to the contact, the detection radius of the sensor aboard the UUV (i.e., the maximum range at which a clue can be detected), and the maximum detectable clue age. Unless otherwise stated later in this chapter, the design factors will be assigned these baseline values:

Table 6.1: Baseline Design Factor Values

Maximum Detectable Clue Age	0 units
Relative Speed	1.75 units
Detection Radius	0.9 unit

By adjusting these factors up or down, this section seeks to determine the effect on search algorithms. While these factors can be set at a level in which many algorithms can be proven successful, we aim to show that the algorithm presented in Chapter 5 can maintain high levels of performance as these parameters are lowered.

After showing the advantages of the dynamic path generation function, we will introduce additional design factors that can be used to improve the performance of the motion plans

described in Chapter 5. Specifically, we will test the algorithm to see how the following planning parameters within the algorithm affect the performance of UUV decision system: path depth, replanning point, and sampling frequency. Initially these parameters will be set at the following policy:

Table 6.2: Initial Motion Planning Parameters

Path Depth	1
Replanning Point	At 1st Evaluation Point
Sampling Frequency	At Evaluation Points

6.2 Search Algorithm Performance Comparisons

The final part of this chapter will compare the performances of different search algorithms. Using the computer simulations described in the previous section, this section illustrates the differences in available algorithms and demonstrates how the motion planning search algorithm seen in Chapter 5 outperforms more conventional search algorithms. Throughout this section the results are designed to answer the question of how well the algorithm does in meeting the search objective (i.e., finding the contact). The answer to this question will be presented through the results of specific cases. These examples illustrate situations or circumstances in which the search algorithm developed in this thesis executes a more logical search strategy.

The results in this section will start by illustrating the difficulties experienced by a simplified search algorithm. Elements of complexity will be continually added to this baseline algorithm until the complete dynamic action space path planning search algorithm is tested. The section will conclude by introducing the existence of time-dated measurements and analyzing their effect on the search algorithm performance.

6.2.1 Base Case Search Algorithm Analysis

The first analysis performed in this chapter will focus on the search objective of the motion planning algorithm. The initial search algorithm presented will serve as the base case for comparisons. This algorithm seeks to recreate the intuition of how to search for a missing contact. After witnessing the shortcomings of this algorithm in practice, the remainder of this section will discuss an alternative objective that challenges the initial intuition and the basic search structure and leads to more favorable results.

6.2.1.1 Base Case Search Algorithm

When first faced with the problem of searching an unidentified contact, the initial intuition of where to search is likely to direct the search vehicle in the direction of the highest probability. Following this intuition, the decision system aboard the UUV would need to use a method for determining the direction with the highest probability of detection. This could be accomplished by heading toward the discrete cell with the greatest probability mass (i.e., cell with the highest weighted sum of particles) or heading towards the weighted mean of the contact position distribution. For the purposes of this thesis, this search objective will be achieved by separating the search area into eight sectors around the UUV position (see Figure 6.1).




Figure 6.1: Sector-Based Action Space

With this sector-based action space in place, the decision system chooses the action that corresponds to the sector with the maximum amount of particles. Based on this strategy the UUV will move towards the sector with the greatest probability at each step throughout the entire simulation.

6.2.1.2 Base Case Performance Evaluation

This section reports on the performance of the base case search algorithm in the scenarios mentioned previously.

6.2.1.2.1 Scenario 1 Analysis

When the contact chooses the route to sea with the highest probability, the sector-based search performs relatively well. Because the algorithm directs the UUV in the direction of the highest probability mass, the UUV has success in finding the contact when this event occurs (see Figure 6.2).




Figure 6.2: Contact Detection Using Sector-Based Search

As is the trend, the sector-based search algorithm tends to perform well when the position distribution is uni-modal and concentrated (in comparison to the sweep width of the sensor). When these features are present, the algorithm often leads to quick detection. Unfortunately, if the contact is not detected quickly, the algorithm cannot maintain an effective search. In the sequence of events illustrated in Figure 6.3, the problems associated with this algorithm become apparent.




Figure 6.3: Illustration of Poor Search Performance – (3) Creation of Multi-Modal Distribution, (6) Oscillation between Modes of a Distribution

In particular, scene 3 demonstrates two major problems with this search algorithm. In this scene, the sensor moves into the center of the probability mass and narrowly misses detection of the contact. By moving directly into the heart of the distribution, it maximizes the probability of detection but also splits the distribution moving through the seaway into two distinct modes. This further complicates the contact position distribution and makes it even harder for the UUV to recover and continue an effective search.

After creating an additional mode, another problem with the algorithm is made clear. This problem, referred to as oscillation, is observed when the UUV is located between modes of a distribution. In Figure 6.3, the oscillations are clearly observed in scenes 3 and 6 as the UUV appears indecisive about which area of the distribution it wants to search. As the sequence of events continues and the distribution becomes increasingly dispersed and complex, the oscillations become even more prevalent (see Figure 6.4).




Figure 6.4: Oscillations between Modes of a Distribution

Based on the results of the twenty experimental runs, the sector-based algorithm performs well in the first scenario, as the average time until detection was 173.6 time steps and the total number of detections equaled 17. These values suggest an effective search algorithm, but upon

looking at the average concluding variance for the 3 runs without detection, we discover the following high values: 13.63 (variance in the x position) and 147.87 (variance in the y position). These values support the earlier claim that the aggressive nature of the algorithm leads to quick detection, but in cases of missed detection, there is little to no chance of recovery.

6.2.1.2.2 Scenario 2 Analysis

When the contact chooses the route to sea with the lowest probability, the sector-based search performs poorly. Because there is little opportunity for immediate detection, the search algorithm experiences the same problems as introduced in the Scenario 1 discussion. It does not properly handle the distribution moving through the route of greatest probability, and therefore it has little chance to recover and reach the dispersing position distribution. As time moves forward, the distributions continue to disperse and create multi-modal distributions spread over the map. As the distributions become more complex, the performance drops significantly (see Figure 6.5).




Figure 6.5: Base Case Algorithm Sequence for Scenario 2

In the sequence of events for this example from Scenario 2, the problems with the base case algorithm continue as additional modes are created in scene 2 and oscillations are found in the remaining images in the sequence.

Contrary to the quantitative results returned during Scenario 1 testing, the base case algorithm in Scenario 2 produces unsatisfactory results. Specifically, the base case algorithm did not detect the contact at any point during any of the twenty simulation runs. The average position variances after 500 time steps were: 12.93 (variance in the x position) and 149.7 (variance in the y position). Therefore, the algorithm fails to detect the contact and furthermore leads to dispersed position distributions. These complex distributions will make it increasingly difficult to locate the contact at any point further in time.

6.2.1.2.3 Base Case Algorithm Conclusions

After observing the sector-based search algorithm attempt to guide the UUV, it is apparent that heading towards the sector with the greatest probability mass is not sufficient. Under certain circumstances, it may lead to quick detections, but other times it leads to ineffective posterior distributions that make sustained searches next to impossible. An algorithm needs to be designed to handle the complex contact position distributions and produce posterior distributions that allow for sustained searches.

The first modification to the algorithm will be changing the objective function. The current objective function, as seen in the previous results directs the UUV and its on-board sensor into the center of the probability masses. This maneuver results in bi-modal or multi-modal distribution which further increase the uncertainty in the location of the contact. Therefore, the next search algorithm implementation will use the objective function, introduced in Chapter 5 that seeks to minimize the uncertainty in the overall distribution. Instead of heading for the center of probability mass, this objective will attempt to contract the overall distribution and bound the potential locations of the contact. This will prevent the distribution from dispersing too quickly, increase the possibility of sustaining a prolonged search, and improve the probability of detecting the contact over time.

In order to implement this objective function, the structure of the search must be adapted. Without simulating the effect of an action on the distribution, we do not have an adequate estimate of the posterior distribution and the resulting information gain. The objective function will need to be evaluated over a set of simulated actions. The action that produces the best

posterior distribution will then be chosen and executed. Section 6.2.2 will provide the results and analysis of the path planning algorithms used to evaluate this objective function.

6.2.2 Path Planning Analysis

The next part of this chapter will focus on the analysis of the motion planning algorithm. After observing the weaknesses of the base case search algorithm, the objective and structure of the search algorithm were modified. This updated algorithm consists of simulating the UUV according to a set of actions and evaluating the resulting distributions at the conclusion of the UUV motion plans. The following section will provide an analysis of several methods used to create these motion plans. The first part of the analysis will look at the techniques used to select the action space. After choosing the best method for generating the action space, the remainder of Section 6.2.2 will examine the effectiveness of different motion plans by varying the path planning parameters introduced in 6.1.3.

6.2.2.1 Action Space Analysis

This section will investigate two different ways to form the action space used to evaluate the search objective function. The posterior distributions that are generated during the simulations will vary depending on the allowable movements within the action space. When determining the allowable actions within the action space, it is important to choose actions that take advantage of the given information. Failure to form an action space based on the available information can lead to path plans that provide little or no information gain. Searches into these areas waste valuable computation time and limit the size of the search area covered. The remainder of this section will describe two different ways to form action spaces and investigate which method most effectively evaluates the search objective.

6.2.2.1.1 Discrete Action Space

The first action space generation method examined will be discrete generation, which consists of a set of discrete actions as introduced in 5.3.1.1. For this analysis, the discrete actions will be generated according to:

$$A = \begin{bmatrix} n \times \frac{2\pi}{8} \text{ Radians} \\ 1 \text{ discrete cell} \end{bmatrix}, n = 0,1,2,\dots,8, \quad (6.1)$$

where the UUV has eight available actions and moves along each action for the length of one discrete cell. According to this action generation scheme, no thought is placed into where the actions are being directed. Instead the discrete action generation process merely enumerates actions uniformly around the current position of the on-board sensor. Using this expression to generate actions during the path planning process can potentially require a great deal of computation to evaluate the entire search space. The remainder of this section will reveal the computational inefficiencies associated with this method through visual illustrations.

The computation time associated with any motion planning algorithm within the current state space representation is governed in part by the number of particle updates required. As stated previously, the particles will be simulated and updated after each action, and therefore as the search depth increases, the computation time needed to evaluate the paths increases exponentially (see Section 5.3.3.2). For the discrete action generation process, the search depth needed to cover the meaningful search area requires computation times that are unrealistic. This is prevalent in dispersed position distributions. For example in Figure 6.6, the step size of the actions is small in relation to the size of the search area and therefore requires a search depth of at least nine to reach the far edges of the position distribution.




Figure 6.6: Search Depth Needed for Total Coverage with a Discrete Action Space

Due to the dynamics of the information, Monte Carlo simulations will need to be performed for planning purposes throughout a search mission. Therefore, the necessary computation for search depths that cover the entire search space might not be achievable. One possible solution to this problem would be to limit the search depth to a reasonable level. While this reduces the computation time, it also constrains the search space being considered. In Figure 6.7, the search depth is limited to four, and as a result, the motion planning process does not consider an entire mode of the distribution.




Figure 6.7: Limited Coverage Area using a Discrete Action Space

Faced with these limitations, the next question one must ask is how can the actions be generated more efficiently? To answer this question, the problem with the discrete generation function must be identified. The main problem with discrete actions is that they produce branches of a search tree into areas with little to no information to be gained. This problem is illustrated in Figure 6.8. By looking into all areas surrounding the UUV uniformly, numerous paths will be generated and evaluated in areas with little information to be gained.




Figure 6.8: Areas with No Information Explored by a Discrete Action Space

To avoid the wasted computation, a new method for generating actions needs to be developed and implemented. The next section will talk about benefits of the cluster-based action space introduced in Section 5.3.1.2.

6.2.2.1.2 *Dynamic Action Space*

An alternative action space generator discussed earlier in Section 5.3.1.2 takes advantage of the current state distribution when choosing the action space. For dispersed, non-parametric distributions, the cluster-based generation eliminates many unnecessary branches in the path enumeration trees. This method only branches UUV paths in the direction of information gain. By creating the actions dynamically, the motion plans can search deeper into the search space

with less computation. In Figure 6.9, a similar distribution as seen in Figures 6.6, 6.7, and 6.8 is clustered and displayed with the initial cluster-based action headings.




Figure 6.9: Dynamic Action Space Illustration

The clusters of information allow the motion planner to handle multiple modes and dispersed distributions more effectively than the discrete action generator. By concentrating the actions on the centers of mass within the distribution, the UUV is guided towards the available information. Because these actions more effectively cover the distributions, the dynamic action spaces will be used throughout the remainder of path planning analysis.

The remainder of this section focuses on the performance of the simplest form of the dynamic action space search algorithm (i.e., search depth equal to 1). Even in its simplest form, the dynamic action space search algorithm selects actions that guide the UUV more efficiently and produce more effective posterior distributions than the base case algorithm. In particular, while watching the path planning algorithm in a simulated search, this algorithm outperforms the base case algorithm in its ability to maneuver around obstacles and in its capacity to contain the position distributions. These traits are illustrated in Figure 6.10, where the images depict the

current UUV path and reevaluation point represented by the line connecting two triangles. The triangle at one end represents the initial position in the path, while the second triangle represents the reevaluation point at the end of the path. UUV will travel along this line until it reaches the second triangle and determines a new path at the reevaluation point.




Figure 6.10: Initial Sequence of Actions using the Dynamic Action Space Path Planner

Due to utilization of the numerical potential field, the UUV is able to tactically plan and maneuver around obstacles, as witnessed in scene 1 of the sequence of images in Figure 6.10. In the remaining scenes in the sequence, the UUV demonstrates how it searches through the available information. Unlike the sector-based search, this algorithm does not dive into the center of the probability mass. Although this will reduce the number of immediate detections, the ability to sustain an effective search over time increases significantly. The ability to sustain a prolonged search is evidenced in Figure 6.11.




Figure 6.11: Continued Sequence of Dynamic Actions Leading to Detection

In the remaining sequences of the search, the UUV attempts to contain or “corral” the distribution. By searching the distribution in this manner, the UUV reduces the uncertainty of the contact location and minimizes the number of modes in the distribution. Although this is a much more conservative approach to locating the target, it improves the chances of recovery in situations where the UUV may initially fail to detect the location of the contact.

This ability to recover can be witnessed in the experimental results from Scenario 2 testing. In Figure 6.12, the contact moves through the channel as the UUV searches through the contact distribution moving through the seaway.




Figure 6.12: Path Planning Search Algorithm (Search Depth = 1)
Initial Sequence of Events for Scenario 2

Throughout this initial series of images, the UUV attempts to contain and eliminate the mode moving through the seaway by spiraling in on the probability mass. Without the availability of clue measurements at the current speed and detection radius, it is difficult for the UUV to

eliminate the possibility of the contact moving through the seaway. As the scenario progresses, the UUV eventually eliminates the mode and makes an effort to recover (see Figure 6.13).




Figure 6.13: Path Planning Search Algorithm (Search Depth = 1)
Remaining Sequence of Events for Scenario 2

Unfortunately, given the current factors the UUV is not able to recover and detect the target moving towards Goal C. Even though the contact may not be found, this algorithm controls the uncertainty in the location of the contact by keeping the position distribution within a few modes. The overall performance of the search algorithm is reinforced by the quantitative results of the simulation runs. In Table 6.3, the simulation results for this search algorithm are found.

Table 6.3: Simulation Results for Path Planning Search Algorithm with Search Depth = 1

Scenario Number	Average Time Until Detection	Number of Runs Contact Located	Variance in X	Variance in Y
Scenario 1	313.6	13	11.79	140.29
Scenario 2	488.9	3	9.8	146.03

As assumed for Scenario 1 runs, the average time until detection increases due to the conservative nature of the search algorithm. Although the number of immediate detections drops in comparison to the base case algorithm, it is important to note that the concluding variance for the missed detections has been reduced slightly in both scenarios and the UUV was able to recover and detect the contact three times during the second scenario. These values indicate that the algorithm is succeeding in containing the distribution and providing an opportunity to recover and potentially detect the contact later in the simulation.

As the complexity of the search algorithm increases, the path planning algorithm will generate more accurate descriptions of the position distributions during the planning process. This improved information should enhance the selected actions and lead to more efficient searches. Throughout the remaining sections, we will analyze the performance of the search algorithm as its complexity is increased with added search depth and increased resampling frequency.

6.2.2.2 Search Depth Analysis

The next part of the path planning analysis involves looking at the importance of search depth. The objective of the analysis is to determine how much, if any, increasing the search depth improves the performance of the search. The following sections will examine the benefit of increasing the search depth to a computationally reasonable level. The different versions of the motion planner will be tested within the context of Scenario 1 and Scenario 2 described in the beginning of the chapter. The observations from these experiments will be detailed within the next sections.

6.2.2.2.1 Scenario 1 Analysis

During the analysis of Scenario 1 simulations, we find that the UUV path planner benefits from deeper searches. In particular, the planner will obtain a better perspective of where it needs to move at the current time in order to optimize the objective function at some point in the future. While the UUV may not choose the action that optimizes the objective function at the next time step, this strategy helps provide for more sustained searches by moving in regards to future conditions. In Figure 6.14, the series of events leading to detection are displayed. Contrary to the basic version of the algorithm discussed previously, Figure 6.14 shows how the UUV evaluates the posterior distribution after simulating several steps into the future. The path is again represented by the line connected to the current UUV position and the triangle marking the reevaluation point. The ensuing triangles represent subsequent reevaluation points in the optimal path selected during the planning phase. By increasing the length of the simulated paths and consequently considering more of the available information, the UUV can make more informed decisions which will lead to quicker elimination of modes within the probability distribution.




**Figure 6.14: Path Planning Search Algorithm (Search Depth = 2),
Sequence of Scenario 1 Events**

From a quantitative perspective, the advantage of deeper searches is confirmed in the simulation results. In Table 6.4 below, we find that as the search depth is increased, the average time until detection drops and the number of detections increases. Because these algorithms provide more conservative searches, it is not a surprise to see that the base case algorithm leads to faster average times until detection. Under the conditions of Scenario 1, an aggressive algorithm finds the contact much faster, but when either algorithm fails to detect the contact, the path planning searches do a much better job at reducing the variance in the posterior position distribution. Because the final variances as seen in Table 6.4 are lower than those produced by the base case algorithm, it is fair to say that these algorithms will have a better opportunity of recovering and detecting the contact at a later point.

Table 6.4: Comparison of Scenario 1 Results for Varying Search Depths

Scenario Number	Search Depth	Average Time Until Detection	Number of Runs Contact Located	Variance in X	Variance in Y
Scenario 1	1	313.6	13	11.79	140.29
Scenario 1	2	243.15	16	9.98	142.43
Scenario 1	3	200.6	17	10.73	142.53

6.2.2.2.2 Scenario 2 Analysis

For this section, the search depths were applied in the second scenario. After observing the simulated search operations, it is clear that these versions of the algorithm continue to perform better than the base case and possess an ability to carry out sustained searches. In Figure 6.15, we see how the UUV uses paths with three reevaluation points (i.e., search depth equal to three) to determine where to move next.




**Figure 6.15: Path Planning Search Algorithm (Search Depth = 3),
Sequence of Scenario 2 Events**

Using this strategy, the UUV is able to contain the mode moving through the seaway and attempts to eliminate the possibility that the contact chose the seaway route as soon as possible.

Because the UUV is able to contain the distribution within a few modes, it is able to sustain the search and ultimately detect the contact in scene 12 of Figure 6.15.

While the benefit of deeper searches is confirmed in the results of Scenario 1, the results are not as conclusive in the Scenario 2 analysis. Although all search depths prove much more effective than the base case algorithm, there is no significant quantitative evidence that the algorithms performance increases with added search depth (see Table 6.5).

Table 6.5: Comparison of Scenario 2 Results for Varying Search Depths

Scenario Number	Search Depth	Average Time Until Detection	Number of Runs Contact Located	Variance in X	Variance in Y
Scenario 2	1	488.9	3	9.8	146.03
Scenario 2	2	496.95	2	9.33	153.47
Scenario 2	3	493.7	2	8.93	147.07

Given the current design factors, the importance of increasing the search depth is therefore inconclusive within this scenario. Due to the relative speeds of the UUV and contact and the relative sizes of the detection region and the scenario map, it is difficult to search the entire probability distribution. Therefore, the true effectiveness of increasing the search depth within this scenario may be difficult to determine. In order to examine the value of increasing the search depth within the second scenario, the design factors need to be updated to allow the UUV to more realistically move through the probability distribution.

6.2.2.3 Sampling Frequency Analysis

The final algorithm modification to be tested involves increasing the sampling frequency during the particle filter simulation used in planning. Currently, the motion planner resamples the simulated particle distribution at the reevaluation points during the path enumeration. This section will examine whether increasing the frequency of sampling improves the search results. The motion planner should obtain better estimates of the possible information gain by increasing the sampling frequency, because additional information is gained during the path from the

current position to the reevaluation point. This additional information gained should lead to more accurate posterior distributions and as a result, more effective motion plans. Specifically, this strategy should prevent the UUV from splitting modes of the distribution. Unfortunately, as the sampling frequency is increased, the computational workload also increases. Therefore, the motion planner must balance these conflicting factors and determine how to balance the improvement in the search with the increased computational work due to increased sampling.

To justify the necessity for increasing the sampling frequency, reference Figure 6.16. Scene 2 of this series illustrates how the UUV in certain circumstances chooses a path that splits a mode of a distribution. Even though the structure of the algorithm allows it to recover from this maneuver, as evidenced in scene 3, an alternative action selection may have contained the mode and therefore further decreased the uncertainty in the location of the contact.




Figure 6.16: Illustration of Poor Search Performance Due to Low Resampling Frequency

The UUV decision system chose this action in this particular sequence, because the planning algorithm did not resample along the path that stretches across the distribution. Because the planning algorithm only samples and resamples the particles at the reevaluation points, the planner assumed that this action would result in the posterior distribution with the least amount of uncertainty. If instead the path planning algorithm sampled and resampled several times along the path, it would produce a more accurate posterior distribution. Given this information, the UUV would no longer select the action that splits the distribution but would instead maneuver to control the scattering of the distribution. In Figure 6.17, this change in trajectory is illustrated. In this same sequence of events, the planning algorithm is operating with

an increased sampling frequency and therefore uses a more accurate distribution when evaluating the paths. As a result, the UUV does not split through the mode but instead constricts the location of the particles moving through the seaway.




Figure 6.17: Improved Search Performance Due to Increased Sampling Frequency

6.2.2.3.1 Scenario 1 Analysis

Although the improved search performance can be seen in the specific example in the previous section, the quantitative results are not as convincing. After examining the results of the first scenario, the average detection time increases and the number of detections decreases (see Table 6.6). While these values might not indicate an improved algorithm during the first scenario, the increased sampling frequency allows the path planner to maintain more accurate posterior distributions. Therefore, the UUV will have more information to base its decisions and therefore will be less likely to split distributions. Because the UUV will not split the distribution, there is less probability of early detection but greater chance of sustaining a longer, more effective search.

Table 6.6: Comparison of Scenario 1 Results with Increased Sampling Frequencies

Scenario Number	Search Depth	Average Time Until Detection	Number of Runs Contact Located	Variance in X	Variance in Y
Scenario 1	1	350.45	11	12.23	144.82
Scenario 1	2	263.35	15	11.58	152.38
Scenario 1	3	313.15	12	15.2	123.33

6.2.2.3.2 Scenario 2 Analysis

While the quantitative results from the first scenario do not necessarily show improved search performance, the analysis of the second scenario reveals the benefits of increasing the sampling frequency. In the results in Table 6.7, the average time until detection decreases and the number of detections increases in comparison to the values from the previous simulations. This again supports the statement that increasing the sampling frequency creates an algorithm that more effectively contains the distribution. By minimizing the uncertainty in the position distribution, the UUV can continue searching and increase its chances of finding the contact later in the simulation as evidenced in the results in Table 6.7.

Table 6.7: Comparison of Scenario 2 Results with Increased Sampling Frequencies

Scenario Number	Search Depth	Average Time Until Detection	Number of Runs Contact Located	Variance in X	Variance in Y
Scenario 2	1	450.45	4	10.68	151.13
Scenario 2	2	468	2	9.53	156.6
Scenario 2	3	436.65	4	9.05	162.94

6.2.2.4 Dynamic Path Planning Algorithm Conclusions

After running the dynamic action space search algorithm and observing its performance, it is apparent that it is more effective than the base case algorithm. While the base case algorithm tends to detect the contact faster on average than the dynamic path planning algorithm,

it displays an inability to effectively recover and find the contact at a later time after initially missing the contact or searching near the unused passage to sea. On the other hand, each successive version of the path planning algorithm has shown an improved ability to sustain a prolonged search by more effectively searching through the search environment. These statements are supported by Figure 6.18 which compares the average detection time and the total number of detections for the each of the search algorithm variations.




Figure 6.18: Comparison of Average Times until Detection and Total Number of Detections

In addition to the differences between the two search algorithms, the path planning algorithm yields different results depending on the selection of search depth and sampling frequency. Based on observations and unique examples, increasing the search depth and the sampling frequencies appear to improve the quality of the search, but within the context of the designed scenarios the quantitative results are inconclusive. Perhaps these results are unconvincing because the UUV is not provided with the adequate speed or sensor range to cover the search area in the selected simulation time. Specifically, the UUV needs to collect sufficient information to eliminate at least one mode prior to the position distribution reaching the open sea. Eliminating the possibility of the contact moving through one of the sea corridors is important to sustaining an effective search. Otherwise, after enough time has passed for the contact to reach the open sea, no physical obstacles or boundaries constrain its movement, and as a result the position distribution will quickly disperse and become increasingly complex. In order for the UUV to eliminate a mode more quickly and have a chance to confirm its effectiveness, the design factors need to be modified or additional information needs to be provided. Therefore, the next section will provide the search algorithms with more information by activating the capability to detect time-dated measurements.

6.2.3 Time-Dated Measurement Analysis

The last analysis will look at the effect of time-dated measurements on the performances of the search algorithms. If modeled correctly, the ability to detect the clue measurements should increase the performance of the search algorithms. The presence of clues should increase the information available in the search space and therefore lead to quicker detections. We will examine two different levels of detectable clues and determine how much the presence of clue measurements helps the search performance. Specifically, this section will analyze the search algorithms with clues detectable at levels seen in the updated design factors in Table 6.8.

Table 6.8: Baseline Design Factors with Updated Maximum Detectable Clue Ages

Maximum Detectable Clue Age	20 or 40 units
Relative Speed	1.75 units
Detection Radius	0.8 unit

When the sensor possesses the ability to detect clues up to 20 or 40 time steps old, the search will be enhanced in two significant ways. The most obvious advantage brought forth by the ability to detect clues occurs when a clue is discovered. Whenever a clue measurement is detected, the uncertainty in the position distribution is significantly reduced. In Figure 6.19, the sequence of images demonstrates the effect of finding a clue. Although the contact is not within the UUV's circular detection region in scene 2 below, the UUV has crossed over a portion of the contact's trail. Because the clues were dropped recently enough to be detected, the UUV is able to reduce the uncertainty in the distribution considerably and therefore decrease the time until detection.




Figure 6.19: Effect of Discovering a Clue

The other means in which the search will be enhanced involves the null measurement model which accounts for the additional information available in the search space. Specifically, the particles are structured in a way that maintains an array of their past positions. As the UUV detection region crosses over these past particle positions, the weights of the particles drop probabilistically as seen in Section 4.2.3.1. This additional information in the state estimation process allows the UUV to shrink the probability distribution earlier and ultimately lead to more effective searches and faster detection times.

6.2.3.1 Base Case Search Algorithm

Prior to analyzing the path planning algorithm, the base case algorithm will again be tested for comparisons. When the capability to detect clues is enabled, the sector-based search algorithm performance increases as evidenced in Table 6.9. The results confirm the hypothesis that the clues provide additional information that enhances the search algorithms. Specifically, the capability to detect these time-dated measurements leads to faster detection times and an increased number of detections. In addition, the results from the base case search algorithm simulations reveal that as the sensor aboard the UUV is able to detect older clues, the search performance continues to improve.

Table 6.9: Base Case Search Algorithm Results when Clues Detectable

Scenario Number	Max Clue Age	Average Time Until Detection	Number of Runs Contact Located	Variance in X	Variance in Y
Scenario 1	20	147.35	18	12.8	151.03
Scenario 1	40	128.9	20	n/a	n/a
Scenario 2	20	401.55	8	12.29	147.07
Scenario 2	40	367.75	11	7.59	95.66

While the clue measurements improve the sector-based search algorithm, the algorithm does not account for the clues during its planning process. Rather, the planning phase of this algorithm continues to simply look for the sector that contains the greatest probability of locating the contact. In order to account for the presence of this information, the search algorithm must once again simulate the effect of an action on the distribution. By evaluating these simulated actions we can obtain an estimate of how the additional information used in the measurement update affects the overall distribution. Therefore, the next section will again analyze the performance of the path planning search algorithm.

6.2.3.2 Dynamic Action Space Search Algorithm

In Section 6.2.2, the path planning search algorithm was tested without the availability of clues. In this section, the algorithm planned its maneuvers in a manner that minimized the uncertainty in the location of the contact. By moving in this way, the UUV attempted to constrain the distribution and sustain a prolonged search for the contact. Unfortunately, the algorithm performance was reduced due to the design factors defined in Table 6.8. Therefore, this section seeks to determine how the path planning search algorithm will improve with the availability of clues. Specifically, the analysis will examine whether the path planning algorithm takes better advantage of the information contained in the clues when planning search routes.

6.2.3.2.1 Scenario 1 Analysis

The analysis of the Scenario 1 simulations reveals that the performance of the UUV planner improves significantly when the UUV possesses the ability to detect clues. The additional information gained from the clues allows the UUV to eliminate the probability distribution faster without increasing the detection range of the sensor or the speed of the UUV. With the extra information available during the planning process, the motion planner gains a better perspective of the probability distribution than the base case algorithm. This enables the UUV to move more effectively and carrying out extended searches.

The series of images displayed in Figure 6.20 shows how the UUV is able to constrict the position distribution moving through the seaway. Due to the assumption that the contact is

leaving behind detectable clues, the UUV is able to reduce the uncertainty in the distribution much faster and find the target in less time than without clues.




Figure 6.20: Illustration of Path Planning Search Algorithm with Clues Available (Clue Age = 20)

These assertions are confirmed in the following table. In Table 6.10, the results from the first scenario simulations reveal that the performance of the path planning algorithm is increased significantly for search depths equal to one or two. For these search depths, the UUV missed detection only once or twice and decreased its average time until detection greatly in comparison to the simulations without clues present.

Table 6.10: Comparison of Scenario 1 Results with Clues

Scenario Number	Search Depth	Max Clue Age	Average Time Until Detection	Number of Runs Contact Located	Variance in X	Variance in Y
Scenario 1	1	20	195.95	18	n/a	n/a
Scenario 1	2	20	195.75	19	n/a	n/a
Scenario 1	3	20	446.3	9	12.95	82.32
Scenario 1	1	40	197.75	18	n/a	n/a
Scenario 1	2	40	188.6	19	n/a	n/a
Scenario 1	3	40	369.75	14	4.59	13.36

Unexpectedly, the results for search depths equal to three reveal degrading search performance. After observing these simulations, the poor performance can be explained illustratively. In Figure 6.21, the sequence of images shows that the motion planner changes strategies when the search depth increases to a level that can reach the segment of the distribution moving through the channel. The motion planner did not choose this route when clues were not present, because the UUV could not move quickly enough to reach the distribution moving through the channel. With clues present, the UUV can make this trailing maneuver, because the particle filter is able to resample much of the distribution probabilistically when no clues are detected. By quickly eliminating this portion of the position distribution, as seen in scene 5 in the figure below, the UUV reduces the uncertainty of the distribution soon enough to recover and complete the search of the distribution moving through the seaway passage.




Figure 6.21: Explanation of Path Planning Results (Search Depth = 3)

6.2.3.2.2 Scenario 2 Analysis

In the second scenario, the results continue to show that the clue measurements significantly improve the performance of the search algorithm. In many cases the additional information allows the UUV to eliminate a piece of the probability distribution in time to recover and finish the search. In the images seen in Figure 6.22, the UUV removes the possibility of the contact moving through the seaway passage and recovers soon enough to search the distribution moving through the channel. Without the capability to detect clues, the probability distribution would have dispersed and created a search area too large for the UUV to realistically cover as time progressed.




Figure 6.22: Illustration of Path Planning Search Algorithm with Clues Available (Clue Age = 40), Sequence of Scenario 2 Events

The improved search performance is verified in the quantitative results found in Table 6.11.

Table 6.11: Comparison of Scenario 2 Search Results with Clues

Scenario Number	Search Depth	Max Clue Age	Average Time Until Detection	Number of Runs Contact Located	Variance in X	Variance in Y
Scenario 2	1	20	421.15	10	7.55	52.27
Scenario 2	2	20	405	11	7.08	59.86
Scenario 2	3	20	150	20	n/a	n/a
Scenario 2	1	40	401.22	11	5.64	21.5
Scenario 2	2	40	336.3	14	6.22	18.7
Scenario 2	3	40	160.2	20	n/a	n/a

As in Scenario 1, the results from Scenario 2 demonstrate a search algorithm that possesses the ability to sustain searches over extended periods of time. Even though this scenario consists of the contact choosing the least probable route to sea, this algorithm is able to rule out the possibility of the contact moving along the high probability route in time to continue an effective search through the remaining distribution. This is confirmed by the final variances of the simulation runs in which the UUV did not detect the contact; these variances have dropped noticeably compared to the results without clues. This indicates that the motion planner uses the assumed clues to eliminate potential contact locations and more effectively contain the position distribution.

In addition, the performance of the search algorithm with a depth equal to three appears exceptionally more effective in detecting the contact. This inconsistency in performances is due to the altered search strategy that was depicted in Figure 6.21. By choosing to search the channel route first, the UUV detects the contact rapidly within the context of the second scenario.

6.2.3.3 Time-Dated Measurements Conclusions

After enabling the UUV to detect time-dated measurements, the performance for all search algorithms increased. The results of these simulations demonstrate that the clue

measurements have been modeled correctly and that the search algorithm utilizes the additional information effectively. For both the first and second scenarios, the addition of clues decreases the average time to detection and increases the total number of detections. For the first scenario, the base case algorithm continues to perform slightly better than the path planning algorithm with or without the ability to detect clues (see Figure 6.23).




Figure 6.23: Comparison of Average Times until Detection and Total Number of Detections for Scenario 1

While the base case algorithm may have slightly better results, Figure 6.23 also shows that the path planning algorithm's marginal improvement is greater than the base case algorithm. This is likely the result of the algorithm's design which assumes the presence of clues in the search environment during the planning process. Even though the path planning algorithm with search depth equal to three does not support this claim in the first scenario, we see in the second scenario why this inconsistency occurs. Therefore, it appears that the path planning algorithm takes better advantage of the new information.

This claim is again supported in the results from the second scenario. In the graphs in Figure 6.24, the improved performances due to time-dated measurements can clearly be seen. When the clue measurements are introduced, the performance of all algorithms is greatly enhanced. In particular, the number of detections increases significantly, because the information provided by the clues allow the UUV to eliminate sections of the distribution more quickly and recover in time to continue searching a more likely search area.




Figure 6.24: Comparison of Average Times until Detection and Total Number of Detection for Scenario 2

Additionally, the results from the second scenario reveal considerable improvement in the path planning algorithm with a search depth equal to three. As mentioned in the earlier section, this improvement is due to a change in strategy for increasing search depths. When the search depth increased to three, the UUV was able to search deeper into the distribution and as a result found it advantageous to maneuver through the channel route first and eliminate any possibility of the contact choosing that path. After ruling out this route, the UUV recovers and maneuvers to search the other route. This explains the discrepancy in the results of the deeper searches for the first and second scenario. Overall this change in strategy seems to produce the best results as

this version of the algorithm yields only six missed detections when clues are detectable up to 40 time steps after their deposit.

[This Page Intentionally Left Blank]

Chapter 7

Conclusions and Future Work

This chapter summarizes the contributions of this thesis and presents suggestions for future research related to this thesis.

7.1 Thesis Contributions

This thesis presented decision-making algorithms that can assist the Navy in the advancement of the UUV submarine track and trail capability. It advanced these autonomous search operations by addressing the following planning issues: state variable representation, objective function declaration, path planning structure, and time-dated measurement depiction. By addressing each of these issues and combining the theory of the supporting topics, ultimately this thesis enables a UUV to search effectively through an environment with little sensor information and narrow in on the location of an evading contact.

Due to the assumed contact motion and measurement models introduced in this thesis, the posterior probability distributions can evolve into complex, multi-modal distributions. Given that these distributions cannot be modeled with parametric distributions such as the Gaussian distribution, the decision system aboard the UUV estimated the state of the contact by implementing a particle filter to represent the dynamic, non-parametric features of the distributions. Based on the relative success of the search algorithms, this thesis has shown that the sample-based particle distributions modeled these non-parametric distributions successfully.

Not only does the particle filter accurately depict the distribution of the state variables, but its unique structure also enabled us to construct more efficient motion planners. Specifically, this thesis introduced a new action generation function that dynamically chooses actions based on the location of available information. Due to the unique structure of the particle filter, the path planning search algorithm clustered the particles to locate sections of information when designing motion plans. The cluster-based action space provided more efficient motion plans that considered larger portions of the search space.

Using the clustered distribution to dynamically formulate action spaces, the implemented search algorithm designs intelligent motion plans that simulate the effect of actions on position distribution and chooses the set of actions that leads to the posterior distribution with the minimum uncertainty. The results of the simulations demonstrate that the objective prevents the UUV from creating more complicated distributions. Instead the objective to maximize information gain compels the UUV to maneuver in a way that contains the distribution and facilitates sustained search operations.

Additionally, the path planning search algorithm is enhanced by increasing the search depth and the sampling frequency during the planning phase. By increasing these values to computationally feasible levels, the UUV path planner is provided with more accurate representations of the position distribution while planning. With a more accurate depiction of the probability distribution, the UUV can make more informed decisions that better fulfill the objective of the search.

Finally, this thesis enhances the UUV search operations by enabling the motion planner to handle the capability to detect time-dated measurements. Given the assumed capability to detect time-dated measurements (i.e., clue measurements); the search algorithm possesses a measurement model to process these measurements in the particle filter. The time-dated measurement model determines the information within a clue measurement and uses the information to properly update the probability distribution. This additional information enhances the search algorithm and leads to a greater number of detections and lower times until detection. Specifically, the information contained in the clues or the absence of clues allows the UUV motion planner to confine the potential locations of the contact and as a result eliminate portions of the distribution at a much faster rate.

7.2 Future Work

There are several possible opportunities for future research related to the work of this thesis. The first area of research that could be explored involves transitioning from contact detection to track and trail operations. The work in this thesis involved the process of searching for and detecting an evading contact. In future work, this research could be continued by examining the best methods for tracking a moving target with bearings-only measurements. Specifically, the research would determine what actions provide the most information on the location of the moving contact. In addition, the UUV tracking algorithm would need to maneuver in a way that kept it in within passive sensor range but out of range of potential adversary sensing devices.

The next area of research that could be investigated involves enhancing the current search operations. The search algorithm could be improved by adding levels of complexity to the search, such as multiple UUVs, multiple contacts, or dynamic obstacles. The work in this thesis looked at only one UUV searching for one evading contact. Increasing the number of resources in the search area increases the complexity of the search operations and calls for more detailed algorithms. This research would examine the best methods for these multiple UUVs to work together to locate the position of one or more contacts. Additionally, dynamic obstacles could be introduced to the search environment. Currently, the search environment consists of strictly static obstacles. In future research, dynamic obstacles, such as other non-threatening vessels, could be added to the search environment. The UUV would need to determine how to plan a path that avoids any moving obstacle in the search area.

References

- [1] Barraquand, J., B. Langlois and J.C. Latombe “Numerical Potential Field Techniques for Robot Path Planning,” *IEEE Transactions on Systems, Man, and Cybernetics* , Vol. 22, No. 2, March 1992.
- [2] Bar-Shalom, Y. and X. Li. *Estimation and Tracking: Principles, Techniques, and Software*. Artech House, 1993.
- [3] Berkhin, Pavel. “Survey of Clustering Data Mining Techniques.” Technical Report, Accrue Software, San Jose, CA, 2002.
- [4] Brown, R.G. and P.Y.C. Hwang. *Introduction to Random Signals and Applied Kalman Filtering*. John Wiley & Sons, New York, 1997.
- [5] Copsey, Keith. “Tutorial on Particle Filters.” NCAF January Meeting. Aston University, Birmingham, UK., 2001.
- [6] Doucet, A., de Freitas, N., and Gordon, N., ed. *Sequential Monte Carlo Methods in Practice*. New York: Springer, 2001.
- [7] Dunn, P. “Navy UUV Master Plan”. *Proceedings of the International Unmanned Undersea Vehicles Symposium*, Newport, RI, April 2000.
- [8] Gelb, Arthur. *Applied Optimal Estimation*. MIT Press, 1973.
- [9] Gordon, Neil and David Salmond. “Sequential Monte Carlo Methods for Tracking.” UK Adaptive Signal Processing Club Discussion Meeting. 7 Mar 2001.
<<http://www.aspc.qinetiq.com/>>.
- [10] Gordon, Neil, David Salmond, and A.F.M. Smith, “Novel Approach to Nonlinear, Non-Gaussian, Bayesian State Estimation.” *IEE Proceedings-F*, Vol. 140, No. 2, April 1993.
- [11] Gray, Sarah. “What is Path Planning?” Autonomous Solutions, Inc.
<http://www.autonomoussolutions.com/research/press/what_is_path_planning.html>.
- [12] Gustafsson, Fredrik, Fredrik Gunnarsson, Niclas Bergman, Urban Forssell, Jonas Jansson, Rickard Karlsson, and Per-Johan Nordlund. “Particle Filters for Positioning, Navigation and Tracking.” *IEEE Transactions on Signal Processing*. Vol. 50, Nr. 2, 2002.
- [13] Heylighen, Francis and Cliff Joslyn. "Cybernetics and Second Order Cybernetics", *Encyclopedia of Physical Science & Technology*, Vol. 4, 3rd ed., Academic Press: New York, p. 155-170, 2001

- [14] Hwang, Yong K. and Narendra Ahuja. "Gross Motion Planning – A Survey." *ACM Computing Surveys*. Vol. 24, Issue 3, September 1992.
- [15] Johnson, Edward A., Commander, United States Navy. "Unmanned Undersea Vehicles and Guided Missile Submarines: Technological and Operational Synergies." Air University. Maxwell AFB, AL, February 2002
- [16] LaValla, Rick and Skip Stoffel. *Search is an Emergency: Field Coordinator's Handbook for Managing Search Operations*. The Emergency Response Institute, Inc: Olympia, Wa, 1987
- [17] Liu, Jun S. and Rong Chen. "Sequential Monte Carlo Methods for Dynamic Systems." *Journal of the American Statistical Association*, 1998.
- [18] McKeever, S. *Path Planning for an Autonomous Vehicle*. Master of Science Thesis, Massachusetts Institute of Technology, June 2000.
- [19] Mierisch, Andrew Peter. *Situational Awareness for a Navy Unmanned Undersea Vehicle*. Master of Science in Electrical Engineering and Computer Science Thesis, Massachusetts Institute of Technology, June 2003.
- [20] Office of Naval Research: Future Naval Capabilities. "Autonomous Operations Future Naval Capabilities." 8 June 2004. <http://www.onr.navy.mil/fncs/auto_ops/>.
- [21] Quazi, Azizul H. "An overview on the time delay estimate in active and passive systems for target localization." *IEEE Transactions on Signal Processing*. Vol. ASSP-29, No. 3, June 1981.
- [22] Rekleitis, Ioannis M. "A Particle Filter Tutorial for Mobile Robot Localization." Centre for Intelligent Machines. McGill University, Montreal, Canada. <<http://www.cim.mcgill.ca/~yiannis/Publications/index.html>>
- [23] van der Merwe, Rudolph. Arnaud Doucet, Nando de Freitas, and Eric Wan. "The Unscented Particle Filter." Technical report CUED/F-INFENG/TR 380, Cambridge University Department of Engineering, May 2000.
- [24] van der Merwe, Rudolph and Eric Wan. "ReBEL-0.2.6." OGI School of Science & Engineering at Oregon Health & Science University. 5 Feb 2004.
- [25] Washburn, Alan R. *Search and Detection*, 3rd ed. Linthicum, MD : Institute for Operations Research and the Management Sciences, 1996
- [26] Welch, Greg and Gary Bishop. "An Introduction to the Kalman Filter." 5 Apr 2004. <http://www.cs.unc.edu/~welch/media/pdf/kalman_intro.pdf>.

- [27] Williams, Brian and Nicholas Roy. “Notes on Path Planning for 16.410 and 16.413.” Principles of Autonomy and Decision Making, 2004.
<http://web.mit.edu/16.410/www/lectures_fall04/17_path-planning-notes.pdf>.