THE PENNSYLVANIA STATE UNIVERSITY

The Graduate School

College of Engineering

MULTI-UUV COOPERATIVE SEARCH FOR TIME-CONSTRAINED MCM RECONNAISSANCE

A Thesis in

Computer Science and Engineering

by

Ryan D. Prins

© 2009 Ryan D. Prins

Submitted in Partial Fulfillment

of the Requirements

for the Degree of

Master of Science

August 2009

The thesis of Ryan D. Prins was reviewed and approved* by the following:

Mahmut Kandemir

Associate Professor of Computer Science and Engineering

Graduate Program Chair

Thesis Adviser

Shashi Phoha

Professor of Electrical and Computer Engineering

Bhuvan Urgaonkar

Assistant Professor of Computer Science and Engineering

*Signatures are on file in the Graduate School

Abstract

Naval mine counter-measures (MCM) are a focal priority of defense technology, due to the increasingly asymmetric form of underwater hazards to the operation and passage of naval ships. A major element of the MCM mission is reconnaissance, in which hazards are distinguished from clear space, identified and classified. Underwater unmanned vehicles (UUV) are the most important component of MCM operation, due to their low cost and the robustness compared to manually executed missions, as well as their ability to collect and process data in parallel to compile data and quickly neutralize threats. Communication among UUVs allows them to work cooperatively to assist each other in choosing the paths and actions most conducive to mission success. This work addresses several problems related to cooperative UUV search, including time-constrained search optimization, ideal communication strategies, cooperation among heterogeneous vehicles, and use of unmanned aerial vehicles (UAV) to enhance connectivity in a UUV network and distribute UUVs to increase MCM mission efficiency.

In the first study, an algorithm is presented that optimizes a time-constrained multi-UUV search by directing each UUV to choose its immediate search task to maximize the information obtained about its environment. Several communication strategies are presented that demonstrate the trade-off between the amount of bandwidth and power consumed by communication among UUVs and the level of cooperation achieved in the UUV mission, as well as the subsequent effect on MCM mission performance. Another study describes the motivation for using UUVs of different classes, optimized for different search tasks and topologies, and defines an algorithm for cooperation between two types of vehicles, demonstrating that adding vehicles of one class increase the efficacy of the other's actions. The last study introduces the use of a hybrid team of UAVs and UUVs, presenting an algorithm that utilizes the UAVs to expand the communication range of the UUVs and allocate UUVs to interesting search regions to optimize the results of the MCM mission.

Finally, several potential expansions of this work are presented for future study. Specifically, some short-term goals are presented, including implementing a turning penalty and modeling acoustic transmitters and information priority in vehicle update, as well as long-term goals, including the addition of a localized control-based algorithm, mine placement pattern recognition, and translation to anti-submarine warfare (ASW).

TABLE OF CONTENTS

List of Figures ……. vi

List of Tables ……. viii

Chapter 1. INTRODUCTION……. 1

1. MCM ……. 1
2. UUV ……. 1
3. Cooperative Search…….. 2
4. Time-Constrained Search……. 2
5. Communication Strategy…….. 3
6. Vehicle Heterogeneity…….. 3
7. Hybrid UUV/UAV Search and Redistribution of UUVs…….. 3

Chapter 2. RELATED WORK …….. 5

1. Mine Counter-measures (MCM) …….. 5
2. Unmanned Underwater Vehicles (UUV)……. 5
3. Communication ……. 13

Chapter 3. TIME-CONSTRAINED OPTIMIZATION OF MULTI-UUV COOPERATIVE MINE DETECTION …….. 16

1. Introduction …….. 16
2. Algorithm …….. 18
3. Simulation …….. 24
4. Results ……..…….. 27
5. Discussion ……..…….. 35
6. Conclusion ……..…….. 35

Chapter 4. COMMUNICATION STRATEGIES IN MULTI-UUV COOPERATIVE MCM RECONNAISSANCE 36

1. Introduction 36
2. Experiment 36
3. Test Results 42
4. Discussion 51
5. Conclusion 52

Chapter 5. COOPERATIVE MCM FOR HETEROGENEOUS UNMANNED VEHICLES UNDER CONSTRAINED TIME 53

1. Introduction 53
2. Classification, Identification, and Neutralization 53
3. Experiment 55
4. Test Results 57
5. Discussion 65
6. Conclusion 67

Chapter 6. UAV-INSTIGATED REDISTRIBUTION OF UUVS IN NAVAL MCM RECONNAISSANCE 68

1. Introduction 68
2. Formalization 70
3. Experiment 76
4. Results 78
5. Conclusion 86

Chapter 7. FUTURE WORK 87

1. Immediate Goals 87
2. Long-Term Goals 89

References 93

LIST OF FIGURES

Figure 3.1. Three-dimensional view of underwater field and sample cell …….. 17

Figure 3.2. Example task assignment priority for vehicle V. …….. 19

Figure 3.3. Task assignment algorithm. …… 21

Figure 3.4. Example task conflict resolution for vehicles V and V'. …….. 22

Figure 3.5. Conflict resolution algorithm. …… 23

Figure 3.6. Block diagram of the inputs and outputs of the algorithm. …….. 24

Figure 3.7. Confidence after 100 iterations in control condition. …… 28

Figure 3.8. Confidence as algorithm runs to completion. …….. 28

Figure 3.9. Confidence with side length varied. …… 29

Figure 3.10. Total probability factor increase with side length varied, 10 vehicles. …….. 30

Figure 3.2. Confidence with Number of Vehicles Varied. …….. 31

Figure 3.12. Confidence with D_MAX Varied. …… 31

Figure 3.13. Mine score after n iterations.…….. 33

Figure 3.14. Total probability factor increase with mine frequency varied. …… 34

Figure 3.15. Mine score with mine frequency varied. …….. 34

Figure 4.1. Example Vehicle Update Range…….. 39

Figure 4.2. Context for Full and Minimal Update Strategies …… 39

Figure 4.3. Context for Partial Update Strategy. …… 42

Figure 4.4. Immediate and Periodic Updates. …… 42

Figure 4.5. Percent of Field Searched with D_SEE Varied. …….. 46

Figure 4.6. Percent Confidence after n Iterations, Full Update. …… 46

Figure 4.7. Percent of Field Searched, Locations per Update Varied. …… 47

Figure 4.8. Percent Confidence at Completion, Locations per Update Varied. …….. 47

Figure 4.9. Percent of Field Searched with D_MAX varied, Immediate vs Periodic. …….. 49

Figure 4.10. Number of Location Updates with D_MAX varied, Immediate vs Periodic. …… 49

Figure 4.11. Percent of Field Searched with D_MAX Varied, Immediate vs Periodic. …….. 50

Figure 4.12. Number of Location Updates Sent with D_MAX Varied, Immediate vs Periodic. …… 51

Figure 5.1. Percent confidence as number of LWVs is varied. …….. 59

Figure 5.2. Percent mines identified and neutralized as number of HWVs is varied ……. 60

Figure 5.3. Percent of mines identified and neutralized as confidence after detection stage is varied. …… 63

Figure 5.4. Percentage of mines neutralized for one-pass and two-pass strategies. ……. 65

Figure 5.5. Adjusted percentage of mines neutralized for one-pass and two-pass strategies. …….. 65

Figure 6.1. Region types as defined by the Draper study …….. 69

Figure 6.2. Connectivity of UUVs within a region. …….. 71

Figure 6.3. Example UAV plan to support vehicles performing search tasks. …….. 71

Figure 6.4. Example support decision for a UAV. ……. 72

Figure 6.5. Example allocation of UUVs. ……. 73

Figure 6.6. Example path planning for allocating UUVs. ……. 73

Figure 6.7. Algorithm to assign UAVs to support region tasks. ……. 75

Figure 6.8. Average region compared to "hot" region. …… 76

Figure 6.9. Confidence in "Very Hot" region over 1000 iterations. …….. 82

LIST OF TABLES

Table 3.1. Total Probability Factor (TPF) Increase in Control Condition …….. 28

Table 3.2. Total Probability Factor (TPF) Increase with Side Length Varied …… 29

Table 3.3. Total Probability Factor (TPF) Increase with Chaos Q Varied …….. 32

Table 3.4. Mine Score with Chaos Q Varied ……. 32

Table 4.1. Control Condition …..…… 44

Table 4.2. Varying Communication Range ……. 45

Table 4.3. Total Messages and Data Sent …… 48

Table 5.1. Control Condition Test……. 58

Table 5.2. Percent Mines removed and Neutralization Rate as Number of LWVs is Varied …… 59

Table 5.3. Percentage of mines removed and Neutralization Rate as Number of HWVs is Varied ……. 61

Table 5.4. Percentage of mines removed and confidence with previously visited cells weighted. ……. 64

Table 6.1. Initial and final confidence, control condition. ……. 79

Table 6.2. Initial and final confidence by region with a "hot" region but no redistribution. …… 80

Table 6.3. Sample initial and final vehicle configuration by region with "hot" region. …… 81

Table 6.4. Initial and final confidence by region with redistribution to the "hot" region. …… 81

Table 6.5. Initial and final confidence by region with "very hot" region. …… 82

Table 6.6. Sample initial and final vehicle configuration by region with "very hot" region. …… 83

Table 6.7. Average vehicle moves for adjusted initial probability. …… 83

Table 6.8. Sample final vehicle configuration by region for adjusted initial probability. …… 84

Table 6.9. Initial and final confidence by region with adjusted initial probability. …… 84

Table 6.10. Initial and final confidence by region with number of UAVs varied …… 85

CHAPTER 1

INTRODUCTION

1. MCM

Naval mine counter-measures (MCM) require the ability to effectively detect underwater mines, and take appropriate actions to ensure a safe operating area or transit route for naval vehicles. The primary steps in an MCM mission include reconnaissance, clearance, and sweeping of mines [17]. An effective MCM operation must be flexible to combat asymmetric hazards including surface, floor, floating, and moored mines. Additionally, clandestine operations are important, especially in littoral regions hostile to MCM tactics. The United States Navy [52] has named mine counter-measures as a high-priority mission in unmanned technology.

2. UUV

There are several advantages to using unmanned underwater vehicles (UUV) as the primary medium through which an MCM operation is performed [15].

- Low-cost, easily replaceable UUVs mitigate the inherent danger in MCM operations. In the somewhat likely event of sensor failure in mine neutralization, a new UUV may quickly be deployed to complete a mission task.
- UUVs reduce the risk to human workers. Underwater mines create an extremely hazardous environment to military personnel, and commanding UUVs remotely allows the human operator to be physically far removed from the hostile area.
- A team of UUVs allows missions to be carried out in parallel, completing a mission faster than a team of human experts would be able to.

For these reasons, a long-term goal for UUV technology is “to have the capability to: (1) deploy or retrieve devices, (2) gather, transmit, or act on all types of information, and (3) engage bottom, volume, surface, air or land targets” [52].

3. Cooperative Search

Cooperative search intends that a team of UUVs tasked with searching a hazardous area is able to communicate information in such a way that the team gathers more information about the area than would be obtained by the same number of UUVs, isolated from each other, searching in parallel [55]. This can involve (1) the communication of sensory data from one vehicle to another to prevent redundant search of an area, (2) the communication of patterns of data to improve the efficiency of search by all vehicles, (3) redistribution of UUVs from clear areas to areas in which more interesting results are being discovered, (4) reassignment of UUVs to ensure thorough coverage in the event of failure of one of the UUVs, and other constructive behaviors. Cooperation may be somewhat achieved by an intelligent mission controller, which assigns tasks in such a way as to optimize the search patterns of the UUVs, given some initial information. However, search results will more often dictate a reconfiguration of the UUVs to achieve optimal mission performance. In this case, the communication among vehicles becomes central to the MCM mission, since a single mission controller rarely has contact with all of the UUVs in the operation.

4. Time-Constrained Search

One problem that requires constant mission re-planning is that of a time-constrained search environment. Path algorithms may be used to optimize a complete search when the time allocated to perform a mission is sufficient to search an entire area. However, operation under a time constraint is more realistic, considering that (1) a hostile environment may require termination of a mission at any time, and (2) asymmetric threats [52] may sometimes involve a changing environment that renders a comprehensive search impossible, and an original mission plan ineffective. In this study, an algorithm is presented which results in each UUV choosing, at any point, the most interesting next task available to it. That is, the UUV chooses a task that maximally increases the confidence of the presence or absence of a mine in the search area. The algorithm presented was more effective in increasing the information known about the search environment than two other greedy algorithms presented in previous studies.

5. Communication Strategy

As previously mentioned, communication is fundamental to a cooperative UUV MCM mission [17]. If UUVs cannot sufficiently share data to optimize the team's search pattern, then no UUV may be performing its search tasks optimally. Underwater communication is complicated by power limitations in extended missions [32], the vastness of the search area compared with the limited range of acoustic sensors, and the minimal bandwidth available to acoustic communication channels [29]. Under such constraints, it is critical to minimize the amount of data shared between UUVs while maximizing the utility of that data to all parties involved in the communication [55]. This study presents several underwater communication strategies to take advantage of a UUV network, while demonstrating the tradeoff between message bandwidth usage and cooperative performance.

6. Vehicle Heterogeneity

The asymmetric nature of MCM missions motivates an equally versatile fleet of UUVs. Thus, an important contemporary research area is in cooperation among UUVs of different vehicle classes, to allow any search function to be performed by a UUV ideally suited to it. In [52], the U.S. Navy has defined four separate vehicle classes to be used in various underwater topographies, for different search tasks; heavier, faster vehicles may be used in open waters or for quick mine detection, while smaller, more precise instruments may be more effective in surf-based search and mine neutralization. In this study, two different vehicle classes are defined, and are used together in a single mission to determine the influence that the presence of one vehicle class has on the other's performance. It was discovered that the algorithm presented allowed a lighter, more precise UUV to attain a performance benefit from the presence of a heavier, faster UUV class by communicating information among the vehicles in the search area. Extension of this algorithm to incorporate more UUV classes may result in still better cooperation and specialization in MCM tasks.

7. Hybrid UUV/UAV Search and Redistribution of UUVs

Some recent studies [19, 56] have addressed the need for increased underwater network connectivity and asymmetric search by proposing a hybrid model involving both UUVs and aerial unmanned vehicles (UAVs). In this model, a UAV acts both as a search instrument, performing overhead reconnaissance

scanning of the target area, and as a communication relay, transmitting and receiving sensory data to and from UUVs in different search regions. The UAVs are able to order reconfiguration of the UUV search team as the environment requires it, and the UUVs are able to more quickly relay sensory data back to the host through the UAV. This study allowed UAVs to command redistribution of UUVs from an area of less interest to an area of more interest according to the distribution of confidence in the information known about the search regions. It was found that redistribution under appropriate circumstances significantly improved the confidence in the interesting search region, as well as improving the result of the overall MCM operation.

CHAPTER 2

RELATED WORK

1. Mine Counter-measures (MCM)

Gooding [17] has laid out a thorough introduction to the primary factors of the naval mine counter-measures (MCM) mission. Contemporary naval forces, particularly littoral, face the hazard of water-borne mines in hostile territories. In an MCM mission, a team of specialists, human or robotic, enter an area ahead of naval vessels to detect and clear the mines to eradicate the threat to the vessels and "rapidly establish large, safe operating areas and transit routes" [52]. The United States Navy has defined four types of MCM missions [52]: reconnaissance (detection, classification, identification, and localization), clearance (neutralization and breaching), sweeping (mechanical and influence) and protection (spoofing and jamming), which Gooding has simplified to mine reconnaissance and clearance.

2. Unmanned Underwater Vehicles (UUV)

2.1. Multi-UUV Search

The primary mechanism for naval MCM in the future has been defined by the U.S. Navy to be a team of underwater unmanned vehicles (UUV), performing tasks remotely from human operation [52]. Gage [15] has offered several advantages to a many-robot approach to MCM:

- Using a UUV team limits danger to human personnel by allowing the human operators to be physically removed from the hazardous area.
- Using multiple low-cost UUV also improves the robustness of the search effort, allowing a MCM mission to be completed even if one or more UUV fails during the mission
- Multiple targets may be processed simultaneously, allowing a parallel approach to MCM

In addition to Gage's qualitative analysis, Hayes [21] has offered quantitative insight into the optimality of a multi-UUV search team, evaluating the cost and benefit of adding vehicles to a search team and

demonstrating a method by which an optimal number of UUV may be discovered. Both studies have shown the benefits of a multi-UUV search effort.

2.2. Cooperative Search

A cooperative search model is a special case of multi-UUV search, in which the vehicles accomplish a mission more successfully by communicating sensory data and objectives than they would accomplish independently. This cooperation allows each vehicle to be more knowledgeably assigned to its search task; it prevents multiple vehicles from performing the same task. It may also aid navigation, as vehicles may inform each other of known "clear" areas, through which they may travel quickly, or hazardous areas, where speed must be decreased.

Vincent and Rubin [55] addressed the question: "How should a team of cooperating sensing platforms be employed to search for...targets in a hazardous environment?" They then propose five goals: 1) to maximize the probability of the detection of targets, 2) to minimize the time of detection of targets, 3) to minimize the number of vehicles employed, 4) to ensure robustness of the mission in case of failure of one of the vehicles, and 5) to minimize the amount of communication that must be exchanged for the vehicles to cooperate and adapt as the search progresses. They develop a search methodology that prevents multiple vehicles from being assigned to the same search region, limits redundant visits to search areas, and facilitates dynamic modification of search assignments in the event of vehicle failure. Vincent and Rubin explore both predetermined and non-predetermined flight patterns, recognizing that even in the case of predetermined flight patterns, cooperation among vehicles is essential to ensure efficient search in the case of vehicle failure. Their proposed cooperative search algorithm employs predetermined swarm patterns for the vehicles, but restricts the number of patterns that the vehicles are allowed to travel, as well as the way the vehicle configurations are shifted. Their limited communication entails periodic messages to ensure connectivity between vehicles and reconfiguration messages if the swarm pattern needs to be updated.

Ryan et al. [43] defined cooperation as a number of unmanned vehicles performing a cooperative task by traveling in a group. They noted that a changing environment required a change in group formation, since connectivity in a group is often affected by hazards or obstacles encountered during reconnaissance. In particular, the construction of a group of vehicles should be determined by the

communication range limitations of the search environment to ensure effective cooperation and improve mission performance.

Rajala and Edwards [39] explore the problem of sharing mine maps among the UUV team. UUVs are limited in the amount of information they can communicate due to low bandwidth, and therefore only critical information can be used in cooperation. This necessitates careful planning of location description, sometimes conceding some resolution in mine location in order to ensure full communication between UUVs of important environment elements.

El-Abd and Kamel [14] use cooperation among UUV to achieve optimality in particle swarm optimization (PSO), using vehicles as the particles. They divide the search space into sub-grids, and restricts a separate team of vehicles to each sub-grid, optimizing their movement within the smaller grid. The overall optimal swarm is obtained by choosing the most successful particle from each sub-grid swarm. They found that increasing the number of cooperating swarms improved performance, so long as suitable communication was maintained. They also found that the type of communication, sharing information globally vs. sharing only with two neighbors to provide a communication "circle," had an effect on search performance, with the "circle" method proving more effective as the number of cooperating swarms was increased.

2.3. Initial Information

Most previous path planning algorithms for a UUV search team have assumed no initial knowledge of the search environment. Accordingly, the optimization of the search has either been to optimize the time in which a comprehensive search is completed or to greedily choose the nearest unsearched location. However, recent progress in three-dimensional imaging [31, 38] has introduced the potential for a UUV team to have some knowledge of the search field at the beginning of the search.

Based on geometric models of mines, Rao et al. [38] have been able to use sidescan sonar images to generate sets of statistical information that are used to classify objects as mine-like or non-mine-like. This information can be used to establish a three-dimensional representation of the underwater field, with an established probability that certain objects are or are not mines. Such knowledge can be used to effect a more efficient UUV search.

2.4. Mission Structure

The structure of a UUV MCM mission is comprised of task allocation and execution, making necessary adjustments determined from data gathered by the UUV in the course of the mission.

2.4.1. Task Allocation

One of the most important facets of the search optimization problem is allocating specific tasks to individual vehicles or to groups of UUV. Most of the previous work in this regard has been based on a greedy search, wherein each vehicle is assigned an optimal search pattern. Often, vehicles may be assigned to the same task; some works have proposed a negotiation between vehicles, such as an "auction" approach, whereby the task conflict is resolved.

Hsieh et al. [23] greedily assign tasks to vehicles in order of proximity. In their approach, each vehicle is capable of completing the entire mission. All tasks are initially "available" to all vehicles, with targets later marked as "selected" or "done." At any point, each vehicle also holds its next target location. During operation, each vehicle selects the closest available target as its next target and broadcasts this to the other vehicles in range; if no other vehicle shares this next target, it is upgraded to "selected." The vehicles also broadcast target lists for targets they have marked as "done."

Sujit et al. [49] propose a negotiation scheme to efficiently allocate tasks for multiple vehicles. Four scenarios are presented for a decision-making vehicle; 1) having no target and no neighbors, 2) having no target but having neighbors, 3) having a target but no neighbors, and 4) having a target and neighbors. In scenarios (1) and (3), the vehicle acts independently, but in scenarios (2) and (4), it acts as a negotiator with its neighbors. The negotiation cycle involves a vehicle i) sending/receiving a proposal, ii) processing any received proposals and sending a decision to accept or reject the received proposal, iii) computing its own route decision and iv) implementing a decision. An individual vehicle computes its own route based on which proposals were accepted or rejected by its neighbors; a route is not chosen unless it was accepted by all neighbors. The risk of deadlock is mitigated by using potential loss of information by the vehicle that would need to switch to its next-best task to determine which vehicle should accept its neighbor's proposal, or by using a token system to determine which vehicle should defer based on past resolutions.

In the model proposed by Sariel et al. [44], each vehicle maintains finite state machines (FSM) for each task in the mission. The states are free, auctioned, in execution, uncertain, and invalid. Each vehicle receives a list of definitions of mission tasks, and selects the most suitable candidate task. Conflicts between two vehicles that choose the same task are resolved by "auction." Each UUV presents a "bid" which is defined as an overall solution to the conflict, based on an evaluation of the value of each vehicle's completed task in the resolved task allocation set. This evaluation is some measure of the cost of each vehicle carrying out its allocated task, in combination with the benefit achieved from task completion. The best "bid" presented during the auction process is accepted as the resolution to the task assignment conflict.

2.4.2. Uncertainty

Hall and Benokraitis [20] emphasized the need for dynamic mission planning at a high level, even for a UUV team controlled through artificial intelligence (AI) methods. Early AI systems only generated a plan as a chain of tasks; when a single task needed to be changed, the entire chain needed to be recomputed. Thus, they demonstrated that a hierarchical mission structure allowed for flexibility while operating within the constraints of the mission. In this hierarchical model, details of specific mission steps are generated only for a temporal interval within the overall mission. When a factor within the search environment changes, the mission controller determines whether the temporal goal is still valid, and only recomputes immediate tasks if the goal can still be reached. Factors of uncertainty include the inaccuracy of sensory data, simple ignorance of the environment (or inadequate resolution of knowledge), and ignorance of the consequences of a vehicle's actions. The study therefore encourages designers to account for such uncertainties to combat the risk of a single point of failure within the tasks of a mission.

2.5. Centralized vs. Distributed Mission Control

One element of the UUV MCM model that has been debated in previous works is a centralized vs. distributed task allocation to the search vehicles. Using a centralized controller is generally a simpler model to implement, and allows global optimization of search time in a relatively static environment. However, the communication limitations of an underwater search environment make it difficult for this model to adjust to new situations or threats, or to the failure of instruments or vehicles. A decentralized

mission control model requires more intelligence on the part of the vehicles, which presents some difficulty in preventing redundancy in search tasks; but this model is more flexible in the presence of the usually prevalent uncertainty in an underwater search area. Therefore, it has generally been accepted that the advantages of the distributed control model outweigh its few shortcomings.

Roeckel et al. [42] proposed an intelligent mission controller, which would interface with individual vehicles. This mission controller performs all of the autonomous control for the UUV team, sending a command record to each UUV, which in turn responds to the controller with status reports.

Baras et al. [3] note that the long-range communication demands of a centralized control system, as well as its susceptibility to single-point failure, rendered it inflexible in a real-world MCM mission. They propose a decentralized path generation model, wherein each vehicle made its moving decisions based on an optimization function of its position, potential targets, neighbor vehicles, and other threats. They present several candidate functions to govern UUV movement, and use vector field analysis to optimize the function parameters. Their study concluded that the primary advantage of the decentralized approach was the relative simplicity of the information required to generate vehicle paths, while decentralization also provides flexibility and robustness in a MCM mission.

Karim and Heinze [27] determined that search tasks would be best accomplished by unmanned vehicles that were based on human cognitive models. This type of decentralized coordination emphasized communication and coordination between vehicles to make mission control decisions in real time. Karim and Heinze implemented both a centralized approach and a cognitive modeling approach. The centralized design was a basic autopilot system with mission control layers, using a traditional feedback loop method to send commands to and receive data from the various agents. The human cognitive agents, on the other hand, followed a model more similar to the cyclic computational model; the agent would observe and assess its situation, then make a decision regarding a course of action and carry out the action. They found that while the minimalism of the centralized, purpose-built design has often been preferred by unmanned vehicle researchers, the human cognitive model allows better flexibility and environmental adaptability, rendering it more translatable to the increasingly asymmetric MCM operation.

Sariel et al. [44] recognized the vast amount of uncertainty in an underwater search environment, and noted that a decentralized approach to mission control was better able to make adjusted decisions in the presence of noisy communication, position uncertainty, and the likelihood of vehicle failure. Their Distributed and Efficient MultI Robot – Cooperation Framework (DEMIR-CF) emphasized robustness in performing simultaneous search tasks with diverse requirements. The DEMIR-CF used a distributed task allocation combined with dynamic task selection, as well as precaution schemes to allow task switching in the event of robot failure. Each vehicle in the framework determines the most suitable task for itself, and negotiates with surrounding vehicles to find the most suitable set of tasks to assign to all of the vehicles. The precaution schemes are designed to deal with such situations as vehicle failure, change in estimated task cost or task definition, the introduction of new tasks or new vehicles to the environment, and manual intervention. It was discovered that the DEMIR-CF achieved near-optimal solutions even in the presence of an ever-changing and uncertain environment.

2.6. Vehicle Heterogeneity

A more recent area of UUV research has been in the area of vehicle heterogeneity. This may involve either UUVs of different classes – in [52], the U.S. Navy has defined four: man-portable, light-weight vehicle, heavy-weight vehicle, and large vehicle – or a combination of UUVs and unmanned aerial vehicles (UAV). The four different UUV classes are specialized to search different types of regions; they also have different payloads and operating speeds. The UUV/UAV model is particularly interesting because it can use the UAVs not only for additional reconnaissance, but also for enhancing connectivity among the UUVs by relaying information between them.

Draper Labs [56] has developed the Risk-Aware Mixed-Initiative Dynamic Replanning (RMDR) system, a system designed to control task planning for a heterogeneous team of UUVs and UAVs. Task planning consists of three steps: 1) Preprocessing (generating goal points for each search location), 2) Search task optimization, and 3) Communication coverage planning. The RMDR system divides a search area into different regions based on their topology and risk level, including regions comprised entirely of land, entirely of water, or overlapping a coastline, as well as regions where vehicles may not enter, may not surface, or may not communicate, or incur a penalty for entering due to elevated risk. The UUVs are each assigned to a specific region based on time windows when particular regions must

be searched, due to lighting, currents, or other factors. This search assignment is optimized by solving the Vehicle Routing Problem with Time Windows (Lau, Tan, etc.). The UAVs may be given one of two tasks: 1) Search or 2) Communication Relay. UAVs may be assigned to any region for Search tasks (whereas UUVs are limited to water-only regions). Additionally, two types of Communication Relay are defined: 1) at location and 2) support vehicles. "At location" relay requires a UAV to travel in a specified pattern to support connectivity within one or more specified regions, while "support vehicles" relay requires a UAV to travel in a "racetrack" pattern between a set of specified vehicles. Search tasks always take precedence over Communication Relay tasks for UAVs, which are not available for Communication Relay while performing a Search task. Finally, Search tasks may be reorganized in two cases: 1) sensor failure or 2) a vehicle finishes its task early.

2.7. Simulation Environment

Considerable attention has been given to an appropriate environment to model a real-world MCM mission. The search environment has a great impact on the performance of various heuristics; for example, a littoral mission may be executed very differently from a deep-water mission. While the resolution of the test environment has varied through simulations developed in previous work, mission simulations used for test purposes should be designed to accurately reflect vehicle sensor capabilities, enemy mine placement, communication range, and various other factors.

De Sousa and Gollu [11] devote a study to defining the organization of a simulation environment for coordinated multi-UUV operations. They design a structured vehicle model with a systems perspective, aimed at accurately reflecting maximum interaction between multiple vehicles and devices. Their Generalized Vehicle (GV) control architecture is organized in five layers: the physical (consisting of sensors and devices), abstraction (representation of physical behavior), functional (guidance and navigation), coordination (both intra-vehicle and inter-vehicle command directives), and organization (supervision of the execution of a mission plan). The test environment itself was modeled in software using object-oriented definitions of both continuous and discrete behavior in an MCM mission, specified in the SHIFT language, developed to describe dynamic networks of hybrid automata.

Roeckel et al. [42] developed a simulation environment to model everything external to a single intelligent mission controller. This model represented vehicle hydrodynamics, sound propagation, and

other factors by generating Monte Carlo environment results from random seeds. They abstracted the communications and vehicle controller interface modules that interacted with the mission controller so that no changes would be required in the mission controller when the simulated vehicle controller was swapped for an actual in-water controller. They also modeled vehicle movement through an autopilot simulation and a probabilistic model of a GPS antenna, which interfaced with the vehicle controller module.

3. Communication

Gooding [17] rightly identifies communication as "the bedrock capability of a multiple-vehicle system." He adds that "without intra-system communication, the benefits of employing multiple assets are reduced to the trivial case of cloning vehicles to reduce mission time." Effective communication among vehicles allows for the sharing of sensory data to prevent redundant search, as well as provide added information by which heterogeneous vehicles may make better-informed subsequent search decisions. The impact of the level of communication of UUVs on their level of cooperation is highlighted in [29]. Localization, relay, and optimal location of vehicles can all enhance connectivity in a network, and all of these options have been explored to some extent in previous works.

Partan et al. [32] underscore several practical challenges to underwater sensory communication. First, underwater communication range is more limited than terrestrial communication, and the search area generally more vast. This means that underwater sensor networks, such as a UUV search, are generally sparser than land-based networks. The majority of underwater communication is over an acoustic channel [29], which has a very limited bandwidth, due to strong attenuation with increasing frequency. Multi-hop networking is therefore more important to improve communication in an underwater network. Relative to radio communication, acoustic communication also carries a much higher probability of bit error [5], as well as a longer propagation delay. Second, underwater communication is generally acoustic, as opposed to terrestrial radio-based communication. This shifts the power demands, which are dominated by transmission in acoustic networks. Optimization of energy use needs to be evaluated with this factor in consideration. In summary, Partan notes that "issues at the physical layer can drive topology, affecting routing, medium access, and even applications."

Chandrasekhar et al. [7] explore various localization algorithms for placing static acoustic underwater sensors, but many of their findings translate equally well to the mobile network of a UUV search team. They note that the heavy attenuation of radio waves underwater makes GPS and RF technologies infeasible for communication in an underwater network; the most significant factors in physical location of devices are network topology, signal propagation models, and energy requirements. Like Partan, their study notes that the primary difficulty in underwater communication stems is related to the inherent troubles with acoustic communication.

Akyildiz et al. [1] lay out some design concepts for underwater communication protocols in light of the difficulties of acoustic communication. Besides highlighting the cost and power requirements, as well as the sparsity of underwater networks, they propose that data caching may be more important in underwater network nodes than in terrestrial nodes because of the potential transience of the communication channel. An underwater protocol must account for transmission loss, due to attenuation and geometric spreading, noise, degradation of an acoustic signal over multi-path propagation, and high transmission delays. They describe a sensor network with UUVs as requiring such network coordination algorithms as adaptive sampling (varying the communication rate by commanding the UUVs to locations where their data will be most useful) and self-configuration (in the event of changes in connectivity). Finally, they outline the state of the art in protocols with regard to underwater sensor networks, including UUV networks.

Benton [4] and Haag [19] have developed the Autonomous Undersea Systems NETwork (AUSNET) protocol, which implements dynamic source routing (DSR) optimized for ad hoc networks with low bandwidth. The AUSNET protocol is designed for relatively small networks, which allows its discovery and routing mechanisms to be simplified compared to the DSR methods. It also takes advantage of several attributes unique to UUV networks, such as low transit rates and predictable changes in navigational paths, to address the specific constraints of underwater communication. There is a high emphasis on real-time network reconfiguration as vehicles pass in or out of range of each other, or their communication devices fail.

Gerla and Yi [16] note that a major problem in network connectivity in underwater networks is the limited ability of existing protocols to scale to changing network size and mobility. They present a

Multicast-Enabled Landmark Ad hoc Routing (M-LANMAR) protocol to exploit team mobility. The UUV team is divided into subgroups; within each subnet of vehicles that share connectivity and similar tasks, one vehicle is chosen as a "landmark." In Gerla and Yi's scheme, each subnet subscribes to a multicast group, with the landmark node acting on behalf of the subnet. This helps to meet the low bandwidth requirements of underwater communication.

Seah [47] and Draper [56] both build on the concept of vehicle heterogeneity by using UUVs or UAVs to increase connectivity in an underwater network. Seah's work uses UUVs to act as dynamic nodes in a network of stationary underwater sensors; the UUVs are placed as communication links as environmental changes cause disturbances in the sensor network. Draper's RMDR system shares the emphasis on location of bridge nodes, expanding the model to include UAVs. In both works, location is primary to the goal of network connectivity, requiring highly mobile relay nodes within the network to ensure robust communication among all of the UUVs and UAVs in a search team.

CHAPTER 3

TIME-CONSTRAINED OPTIMIZATION OF MULTI-UUV COOPERATIVE MINE DETECTION

1. INTRODUCTION

1.1 Contributions

A primary contribution of this work is to expand cooperative movement and search optimization, in the context of MCM, to three dimensions. Some past efforts to optimize cooperation of vehicles for search purposes have primarily been applied to land-based or ocean floor-based searches. However, since mines may be placed in a three-dimensional underwater space to prevent vehicles from passing through, it is necessary for detection to occur in three dimensions as well. Appropriate UUVs, then, are capable of movement and detection in three dimensions. Much work has already been accomplished in the development of three-dimensional mine detection, specifically through sonar imaging [31, 38]. This study builds on this work by presenting distributed task assignment for UUVs traveling in three dimensions.

Another contribution of this work is to use probability estimates to search those areas about which the least amount of information is known. Many previous works have generally assumed a uniform probability distribution, and have therefore applied greedy algorithms to always assign a vehicle to search the nearest locations. However, a preliminary reconnaissance scan using sonar imaging [31, 38] may produce valuable initial information about the presence of mines in an area. In this case, assigning a vehicle to search the closest location is not necessarily the most profitable action for the vehicle, as it may be assigned to a location for which the presence or absence of a mine is already known with great confidence. The algorithm presented in this study takes advantage of initial information [31, 38] regarding an underwater environment, using this to execute an efficient MCM search.

1.2 Algorithm Framework

One of the most effective strategies for optimizing the mine detection is to optimize the cooperative aspect of the execution patterns of the UUVs. The UUVs should traverse the field in a manner that best

takes advantage of their search capabilities by limiting their interference with each other while directing their searches toward those areas of the field that have the highest doubt regarding the presence of a mine. To establish a framework for an algorithm, several basic characteristics of the problem environment are presented, followed by a definition of the time constraint to be optimized, and a brief synopsis of the optimization strategy.

An effective search strategy requires careful examination of the problem environment. The search occurs in a three-dimensional underwater field, represented as a collection of smaller cubic unit cells, as shown in Figure 3.1. Each of these cells has a unique probability of the presence of a mine, represented by a number $0 <= P_{i,j,k} <= 1$. The probability of each cell may be initialized by either a preliminary scan of the field, usually performed by aerial surveillance, or by a specified default expected probability. Each vehicle is assigned a task, to further investigate a single cell. It is assumed that these tasks will be assigned by an omniscient controller, which will have information about each vehicle's location and the probability $P_{i,j,k}$ of each cell, as well as the search history of each cell, such as whether it has yet been searched or whether a vehicle is currently assigned to it. The task assignment problem is further defined by the time constraints of the mission.

Different search constraints may require different optimizations of the UUV search patterns. For instance, a mission under a known time constraint may search the entire field comprehensively if the time constraint permits it. However, if the allotted mission time is unknown, the search algorithm must emphasize the proximity of a UUV to its target search location as well as the probability of the presence of a mine in that location, such that the UUV will always make the best immediate decision under the assumption that its mission will be terminated after its next action.

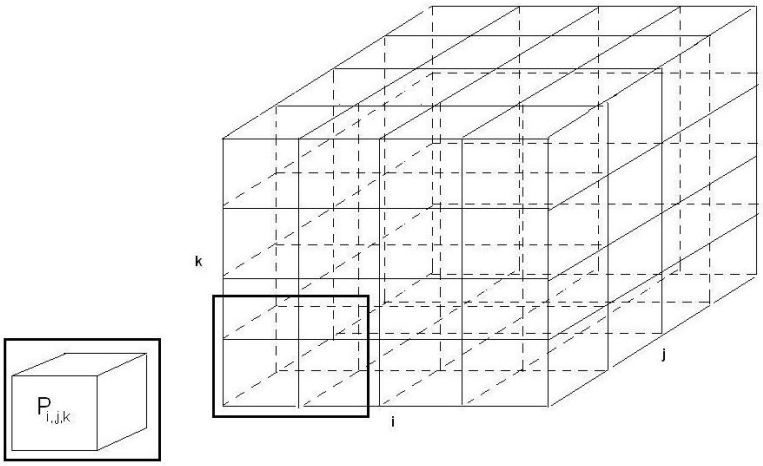

Figure 3.3. Three-dimensional view of underwater field and sample cell

The optimization strategy seeks to assign vehicles at any moment in time in such a way as to maximize the increase in information of their collective search effort. To accomplish this, it is desired that each vehicle be assigned to a unique cell. It is also important that no vehicle should search a cell that has been previously searched by another vehicle. Additionally, a vehicle should search a cell that is located relatively near to the vehicle, such that the increase in information regarding the presence of a mine is worth the time expended to travel to that cell. Ultimately, the optimal task for any vehicle is the task that minimizes distance from the task while maximizing the increase in information from the completion of that task.

1.3 Roadmap

The next section describes the algorithm used to optimize the cooperative search efforts of the UUV team. It first defines the variables used by the algorithm to describe the search environment. This is followed by a description of the rules used to assign a task to a vehicle. Lastly, a method of conflict resolution is presented, by which two vehicles assigned to the same task will resolve this situation in a manner that optimizes the collective search assignment of the UUV team.

2. ALGORITHM

The algorithm to optimize the cooperation of the UUV team is more specifically defined in terms of the relationship of many objects and smaller functions to the goal of an optimal assignment of tasks to the vehicles. First, the relevant objects and variables themselves are defined. A discussion of the rules of default task assignment is followed by a definition of the action taken when a vehicle's optimal task is

already assigned to another vehicle. After task assignment, all of the vehicles are commissioned to their final assigned tasks, and appropriate action is taken to update the state of the problem area as vehicles report completion of their tasks.

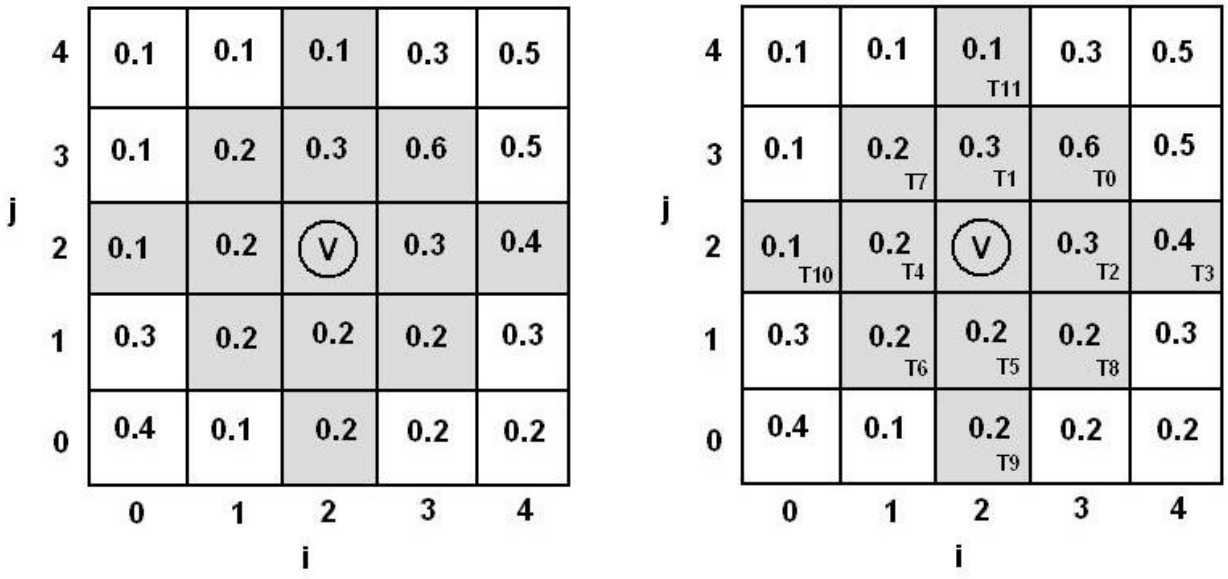

Figure 3.2. Example task assignment priority for vehicle V. Sample initial probabilities are shown for each cell in a small 2-dimensional space surrounding V. The shaded cells are those within a predefined distance D_MAX = 2, and are the only cells considered when assigning a task to V. The probability map is shown on the left, with the task assignment priority added on the right in order of assignment preference, with T0 as the highest preference, followed by T1, etc.

2.1 Variables

The algorithm uses several variables with frequency. Let *V* refer to the set of vehicles used in the search. Each vehicle holds its current task t (a location in three dimensions), a set of "reachable" tasks T, its state, and its current location. F refers to a three-dimensional underwater field to be searched, divided into unit cells which serve as search targets for each UUV. The three-dimensional array S describes these cells in terms of their search status; a 1 indicates that the cell has already been searched, while a 0 indicates that a search of the cell has not yet been completed.

An assignment array A describes the cells in terms of their assignment status; a 1 indicates that the cell has been assigned to a vehicle to be searched, while a 0 indicates that no vehicle has been assigned to search the cell. A probability array P describes each cell in terms of the probability, a number between 0

and 1, of finding a mine there. Finally, each vehicle holds a constant D_MAX, which indicates the maximum allowable distance for the vehicle to be assigned to travel to its next target (based on the vehicle's travel speed), to ensure that reasonable target proximity is maintained.

2.2 Task Assignment

Task assignments are adjusted as each vehicle is addressed. As the algorithm iterates through the set *V*, each vehicle V that is idle is assigned a task. For each of these vehicles, every reachable task that has not yet been searched is added to its set of potential tasks T. A reachable task is defined as a task that is less than or equal to D_MAX distance units from the current location of vehicle V. After every reachable task has been added to the set T, the set T is sorted by the weight factor of each task, from least to greatest, as shown in Figure 2, with only two dimensions shown for clarity. The weight factor is defined as the distance d of the task from the current location of the vehicle, multiplied by the ideal probability deviation X, where X is the absolute difference between the probability of the presence of a mine in task T_i and 0.5, the ideal target probability. The weight factor, then, is w = Xd, where $d = \sqrt[3]{((T_i.x - V.currentLocation.x)^2 + (T_i.y - V.currentLocation.y)^2 + (T_i.z - V.currentLocation.z)^2)}$ and $X = |P_{Ti} - 0.5|$. The default proposed task is set to T[0], the reachable task with the lowest weight factor. Figure 2 shows an example of the priority order of task assignment given in two dimensions. Vehicle V begins in location (2, 2). With D_MAX = 2, there are 12 reachable tasks. Of these, the closest is located in (3, 3), with a weight factor of $w = |0.6 - 0.5|(\sqrt{2}) = 0.1414$. The next lowest weight factors are cells (2, 3), (3, 2), and (4, 2), with w = 0.2. Note that though cells (4, 4) and (4, 3) have a weight factor of 0, they are unreachable because the distance exceeds D_MAX. Note also that the cell with the lowest weight factor is not necessarily the nearest cell. Our task assignment algorithm is shown in Figure 3. If the task assigned to the current vehicle V causes no conflict with a task previously assigned to another vehicle V', V holds this assignment until the next vehicle is assigned.

2.3 Conflict Resolution

If a conflict is created, the conflict must be resolved between the current vehicle V and the conflicting vehicle V'. One potential resolution is to change the assigned task of V to the task with the lowest weight factor within its set T that is unassigned, its next best-task choice. The other potential resolution is to change the assigned task of V' to its next-best task choice. In each of these cases, a ripple-down effect may occur, causing more potential conflicts between the next-best task choices and more vehicles.

All of these lower-level conflicts are resolved before the highest-level conflict between V and V'. Only two potential resolution sets of assigned vehicle tasks are available for any conflict; the first set (R1) is formed by using the already optimized existing set of vehicles and their respective assignments, and

```
V = set of vehicles (Vehicle[])
F = field of cells (int[][][])
S = field of searched cells (boolean[][][])
A = field of assigned tasks (boolean[][][])
D_MAX = maximum distance threshold to assign a task (double)

ASSIGN_TASK

For each vehicle V in V
        If V.state = idle Then
                For each potential task t in field F     //task is a location (x, y, z) to be searched
                        If DISTANCE(V.currentLocation, t) < D_MAX and S[t.x][t.y][t.z] = 0 Then
                                Add t to potential tasks V.T[]
                Next t
                Sort V.T[] by weight factor WEIGHT(V, T[index]).
                Proposed task t = V.T[0]
                For each vehicle V' in V
                        If V'.t = t Then                //Task assignment conflict – resolve
                        RESOLVE(V, V')
                        break                           //Impossible to conflict with two separate vehicles
                Next V'
        Else do nothing
Next vehicle

For each vehicle V in V
        Commission(V)
        V.state = active
Next vehicle

End Assign_Task

WEIGHT(V, t)

Ideal probability deviation X = |P_T – 0.5|
Distance d = DISTANCE(V.currentLocation, t)   //where current location of V is (x0,y0,z0) and location of t is (x1,y1,z1)
Weight factor w = Xd

Return w

DISTANCE(L1, L2)

Distance d = ∛((L2.x – L1.x)² + (L2.y – L1.y)² + (L2.z – L1.z)²)
Return d
```

Figure 3.3. Task assignment algorithm. The algorithm loops through the vehicles, assigning tasks to any that are not currently executing one. A task is defined as a location to be searched, with coordinates (x, y, z). Potential tasks are those locations that are within a distance D_MAX of the vehicle being assigned and that have not already been searched. The tasks are then sorted by weight factor to produce a task list ordered by preference for the vehicle. The proposed task to assign to the vehicle is the task with the lowest weight factor. If this causes a conflict with another vehicle assigned to the same task, the sum of weight factors for all vehicles' assigned tasks is compared for each potential resolution. The set of tasks with the lowest total weight factor is the set of tasks assigned to the vehicles.

adding the new vehicle assigned to its next-best choice of task. The second set (R2) is formed by adding the new vehicle assigned to its optimal task, and reassigning the other vehicles to their next-best task arrangement to accommodate it. The conflict is resolved by choosing the set that produces the lower sum of the weight factors of all assigned tasks, and assigning these tasks to their respective vehicles. An example conflict and resolution is shown in Figure 3.4.

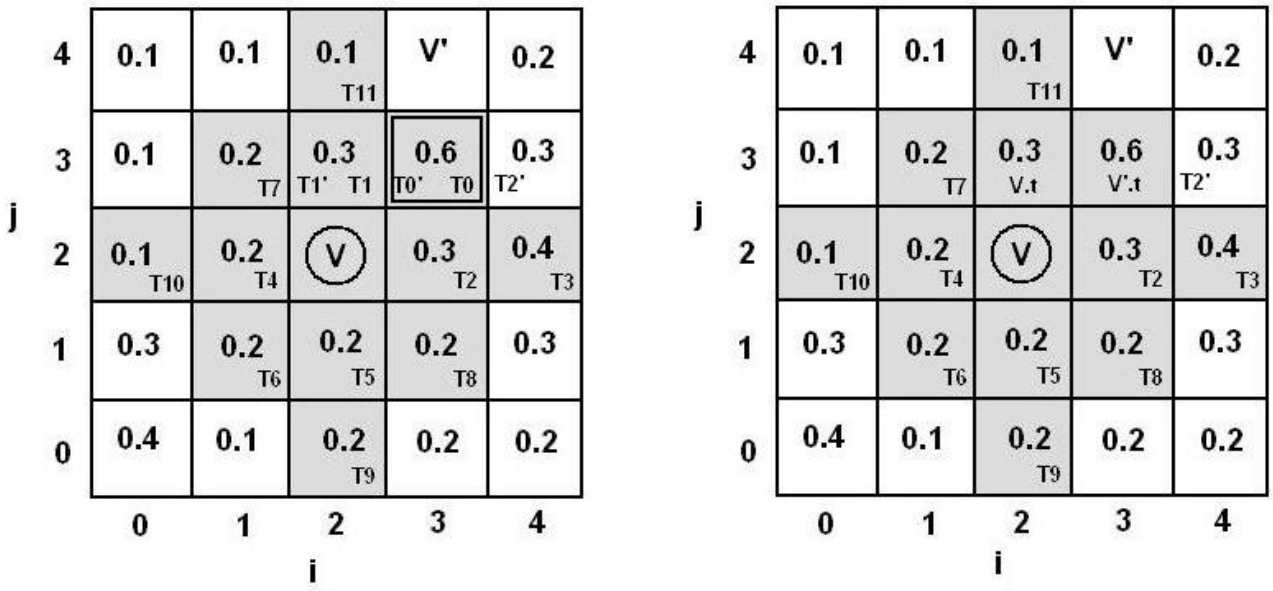

Figure 3.4. Example task conflict resolution for vehicles V and V'. The initial task assignment is shown on the left: both vehicles V and V' have been assigned to the same task, where T0 = T0'. Two possible resolutions exist: reassigning V to T1 or reassigning V' to T1'. The better resolution assigns V to T1 and V' is assigned T0'. The final task assignment is shown on the right, where V's assigned task is V.t and V' is assigned V'.t.

In an arrangement similar to that in Figure 3.2, a second vehicle V' is introduced, located at (3, 4). Assume that V' and V are the only two vehicles in the set. V' is a vehicle that is already assigned to task (3, 3). As vehicle V attempts to be assigned (3, 3), a conflict occurs with V'. In the resolution set R1, V' maintains its assignment to (3, 3), and V is assigned its next-best choice, (2, 3). In the resolution set R2, V is assigned (3, 3), and V' is reassigned to its next-best choice, (2, 3). R1 has a total weight factor of $|0.3 - 0.5|(1) + |0.6 - 0.5|(1) = 0.3$. R2 has a total weight factor of $|0.6 - 0.5|(\sqrt{2}) + |0.3 - 0.5|(\sqrt{2}) = 0.424$. Since $0.3 < 0.424$, R1 has the lower weight factor and is taken as the solution set. V' is assigned the task (3, 3) and V is assigned the task (2, 3). The algorithm for conflict resolution is shown in Figure 3.5.

```
RESOLVE(V, V')

Resolution R1 = V            //Each resolution will default to keeping all the vehicle assignments the same
Resolution R2 = V
For i = 0 to Length(R1) – 1 Then        //Find V in R1
        If R1[i] = V Then break
Next i
V_R1 = R1[i]
For j = 0 to Length(R2) – 1 Then        //Find V' in R2
        If R2[j] = V' Then break
Next j
V.t = ALTERNATE(R1[i])                  //Set alternate choices and compare weights
V'.t = ALTERNATE(R2[j])
w1 = w2 = 0
For i = 0 to Length(R1) – 1
        w1 += WEIGHT(R1[i], R1[i].t)
Next i
For j = 0 to Length(R2) – 1
        w2 += WEIGHT(R2[j], R2[j].t)
Next j
If w1 < w2 Then                                                 //R1 is better solution
        For k = 0 to Length(R1) – 1
                A[V[k].t.loc.x][V[k].t.loc.y][V[k].t.loc.z] = 0      //Reset assignment vector
                V[k] = R1[k]                                            //Use alternate vehicle assignments
                A[V[k].t.loc.x][V[k].t.loc.y][V[k].t.loc.z] = 1      //Set assignment vector
        Next k
Else                                                            //R2 is better solution
        For k = 0 to Length(R2) – 1
                A[V[k].t.loc.x][V[k].t.loc.y][V[k].t.loc.z] = 0      //Reset assignment vector
                V[k] = R2[k]                                            //Use alternate vehicle assignments
                A[V[k].t.loc.x][V[k].t.loc.y][V[k].t.loc.z] = 1      //Set assignment vector
        Next k
Return
```

Figure 3.54. Conflict resolution algorithm. The algorithm uses two sets of tasks: R1 and R2. R1 represents the set of tasks resulting from assigning the new vehicle V to the highest preference task available without reassigning the other vehicles. R2 represents the set of tasks resulting from assigning the new vehicle V to its highest preference task regardless of availability, and resolving the conflict by reassigning the other vehicles as necessary. The total weight factors of the two sets are compared, and the set with the lowest total weight factor is accepted as the set of tasks to assign to the vehicles. The assignment vector is updated appropriately.

2.4 Commission

When all idle vehicles have been assigned, the vehicles are commissioned with their tasks, and their status is set to active. The assignment array A is updated and the value of every cell currently assigned to a vehicle is set to 1. When a vehicle completes its assignment, it reports its progress and the searched array S sets its value for that cell to 1, while the assignment array A resets its value for that cell to 0. The vehicle's state is set to idle and it awaits its next task assignment.

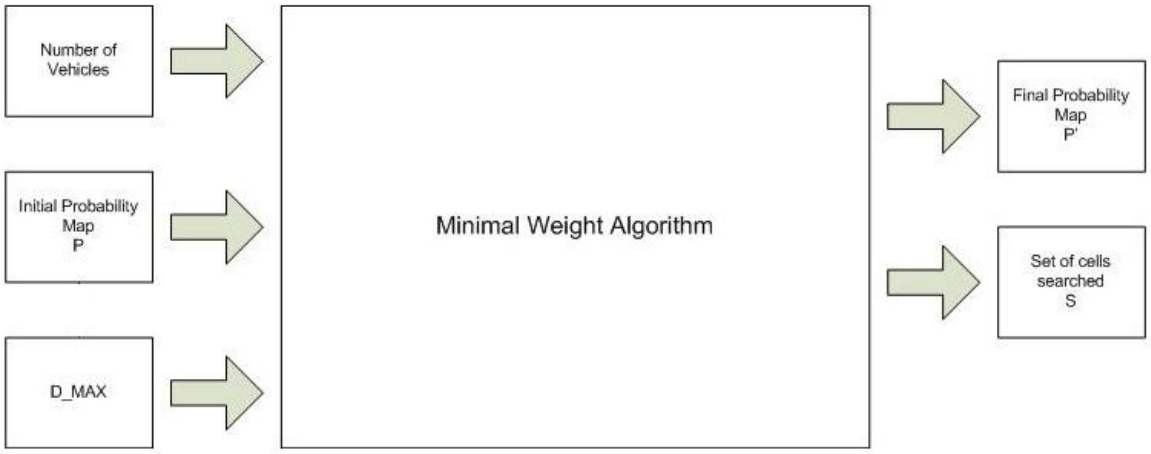

Figure 3.6. Block diagram of the inputs and outputs of the algorithm.

3. SIMULATION

Simulation of the algorithm used a test program run under controlled initial conditions such as the number of vehicles, initial probability, D_MAX, and field dimensions.

3.1 Other Algorithms

Two other algorithms were also implemented for comparison, using the same initial conditions. The first comparison algorithm always selects the closest incomplete task to assign to a vehicle, regardless of the initial probability of the task. The second comparison algorithm always assigns a vehicle to the task of the highest doubt, that is, the task with probability closest to 0.5; this second algorithm does not consider the distance between the vehicle and its assigned task, nor is the vehicle allowed to search intermediate cells. For the remainder of this chapter, our proposed algorithm will be referred to as the Minimal Weight Algorithm (MWA), the first comparison algorithm as the Minimal Distance Algorithm (MDA), and the second comparison algorithm as the Best Probability Algorithm (BPA).

3.2 Experiment Procedure

3.2.1 Control Condition

In the control condition, the field was initialized as a perfect cube, with each cell given a randomized initial probability between 0 and 1. In some tests, initial probabilities were more tightly controlled, to

test performance of the algorithm under specific conditions. Vehicles were initially positioned randomly within a plane along one edge of the cube, with a default D_MAX = 2. The interrupt nature of progress reports from the vehicles were not perfectly modeled for two reasons. First, this was not a real-time simulation. Second, vehicles should not be reassigned immediately upon their progress report because this undermines the conflict resolution of the MWA, as it is unlikely that two vehicles in the same region will be simultaneously idle after the initial assignment of tasks to all of the vehicles. Therefore, for simulation purposes, vehicles were assumed to have a maximum distance that they could travel in one iteration of the algorithm; this distance was called DV_MAX and was also set to 2. If a vehicle was within DV_MAX of its task, it was assumed that the task would be completed in this iteration of the algorithm and the vehicle's status returned to idle. If the vehicle was not within DV_MAX of its task, its distance to task was decreased by DV_MAX and its status remained active. DV_MAX was not varied in the simulation, since the relevant factor is actually the ratio of D_MAX to DV_MAX – the maximum period of time (or, under the simulation conditions, number of iterations) for the vehicle to reach the task.

3.2.2 Performance Metric

The results were analyzed by comparing the total confidence or total probability factor of the field for each of the three algorithms. An ideal result would be a field in which the probability of each cell is either 0 or 1. This condition is defined as 100% confidence. Let the probability factor of a cell i be the difference between that cell's probability and 0.5, or $f_i = |p_i – 0.5|$. Total confidence is measured by the ratio of the sum of the probability factors of all to the sum of all cells multiplied by 0.5, or $C = (\sum f_i) / (0.5*n)$.

3.2.3 Modeling Probability Update

The test program was designed to model imperfect sensory equipment; that is, completion of a task would not update the location's probability to either 0 or 1. Rather, completion of a task updated the resulting probability of the location as a function of its initial probability. Tasks with an initial probability $p_i <= 0.5$ were given a resulting probability of $p_f = p_i^2$, while tasks with an initial probability $p_i > 0.5$ were given a resulting probability of $p_f = \sqrt{p_i}$. This produces the desirable result that information always increases as the result of UUV action, and information increase is greater when the initial state of the task has greater doubt. For example, searching a cell with probability $p_{i1} = 0.5$ should

be more profitable toward the end goal of the mission than searching a cell with $p_{i2} = 0.1$, and this model produces $p_{f1} = 0.25$ and $p_{f2} = 0.01$, resulting in increases of $|\Delta f_{i1}| = 0.25$ and $|\Delta f_{i2}| = 0.09$, respectively.

3.2.4 Initial Probability

The tests can be grouped into two categories: random initial probability and designed initial probability. For simple simulations designed to show the effect of such factors as number of vehicles, field size, and D_MAX on the performance of the algorithm, random initial probability was used, with each cell initialized to a random probability between 0 and 1. Designed initial probability assumes that there is some definable order to the placement of the mines. This is a more realistic scenario than a random assortment of mines scattered throughout the field. Therefore, in this group of tests, a simple pattern of mines, such as a straight line or a square, was generated. The mines' proximity to each other was forced to be less than D_MAX, as it is assumed that D_MAX is one variable that could easily be assigned after some basic analysis of initial surveillance. To test performance under the condition of designed probability, two factors needed to be introduced to the initial probability map: chaos and mine frequency.

3.2.5 Chaos

Chaos represents the distinction between a cell containing a mine and a cell that does not contain a mine. This factor could be used to model such situations as initial surveillance of murky waters, or shallow water in which it is difficult to distinguish between a mine and a rock or plant on the floor. For simulation purposes, chaos Q was defined on a scale of 0 to 1 as a constraint for probability of cells that either did or did not contain a mine, with one endpoint being 0 for clear cells and 1 for cells containing mines. As a cell containing a mine was likely to cause greater doubt regarding a mine's presence, based on an initial scan, than a cell not containing a mine, chaos for cells containing mines was defined as the square root of chaos in clear cells, or $Q_M = \sqrt{Q_C}$. For example, introducing a chaos factor of 0.7 into the test system has the effect of assigning all probabilities of clear cells values between 0 and 0.7 (a range of 0.7), while assigning cells containing mines probabilities between $(1 - \sqrt{0.7})$ and 1, or 0.163 and 1 (a range of 0.837). Q was varied in one test to demonstrate the effect of the precision of initial surveillance equipment on the performance of the algorithms.

3.2.6 Mine Frequency

Mine frequency defines the probability that any particular cell will be assigned by the system to contain a mine. Note that this factor applies only to the scenario of designed initial probability, and is furthermore meaningless in a situation with Q = 1, as there is essentially no distinction in this case between probability assigned to a clear cell and that assigned to a cell designed to contain a mine. Therefore, mine frequency becomes increasingly important to the simulation as chaos decreases. For test purposes, frequency is varied with Q held constant at a reasonable value to demonstrate the MWA's proficiency at detecting mines based on an initial pattern, compared to the MDA and BPA. Similarly, tests of the effect of Q will hold mine frequency constant at a reasonable value.

3.2.7 Mine Score

The performance of the algorithms under designed initial probability is measured both by total probability factor increase and by a new measure called the mine score. This score is the percentage of total mines that were found by a search algorithm within the given number of iterations. It must be remembered, however, that it is a combination of these two measures that truly gives an indication of the performance of the algorithm. For instance, it is more desirable to visit a cell with high doubt that turns out to be clear than to visit a cell that is known with great confidence to contain a mine.

4. RESULTS

4.1 Control Condition

The control condition for random initial probability involved a cubic field with a side of length 20 units, 20 UUVs, and 100 iterations of the algorithms. This was necessary to achieve a reasonable execution time, as the BPA is designed to search the entire three-dimensional field and thus has $O(n^3)$ algorithmic complexity. While the other two algorithms are also $O(n^3)$, they only operate at their worst-case complexity when a vehicle is unable to find an available task within distance D_MAX and must be assigned a “distant” task, searching the entire field. As shown in Figure 3.7, the MWA outperformed both the MDA and BPA under the control condition. Table 3.1 demonstrates that MWA provided a 11.07% increase in confidence from the initial state, while MDA provided an increase of 7.02% and BPA an increase of 4.83%. Figure 3.8 shows the performance of the three algorithms as time grows sufficient to search the entire grid. In this test, the MDA completed searching the field before the

MWA, thus demonstrating that the MWA best applies to the condition of less than sufficient time to search the field, given the field's size and the number of available vehicles.

Table 3.1. Total Probability Factor (TPF) Increase in Control Condition

	MWA	MDA	BPA
Initial TPF	1985.41	1985.41	1985.41
Final TPF	2205.13	2124.83	2081.36
% Increase	11.07	7.02	4.83

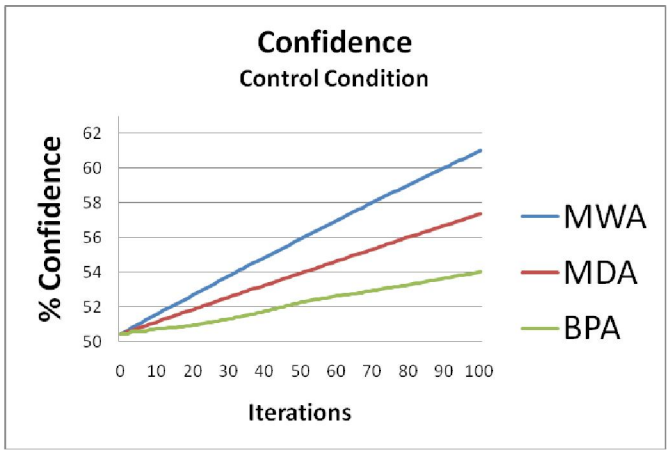

Figure 3.7. Confidence after 100 iterations in control condition. 100% confidence indicates that all cells contain a mine probability of 0 or 1. Confidence is calculated by $C = (\sum f_i) / (0.5*n)$, where $f_i = |p_i - 0.5|$. The advantage is initially minimal, because very few search tasks have been completed. However, the MWA quickly outperforms the other two algorithms, thanks to a more effective default search strategy, as well as communication between vehicles to avoid redundant search.

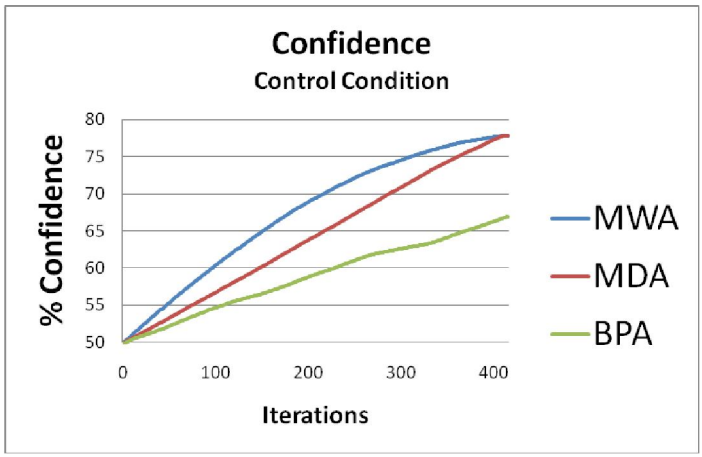

Figure 3.8. Confidence as algorithm runs to completion. The algorithm took 415 iterations to run to completion. Note that even a complete search of the field does not produce 100% confidence, based on the assumption of imperfect sensory equipment. The advantage of the MWA becomes small as the search nears completion.

4.2 Varying Field Size

Varying the field size demonstrated that the MWA provided better performance (by overall increase in information) compared to the other algorithms for larger field sizes. To allow the field size to shrink, the number of vehicles was decreased for this test from 20 to 10. Since each initial probability map was necessarily different, end performance was determined by measuring the difference between the initial

and final total probability factors produced by the MWA and the other two algorithms. The results showed that greater increase in total probability factor for the MWA occurred at larger field sizes. There was also a decrease in confidence as field size increased, as shown in Figure 3.9. However, in this test the number of searched cells remained roughly constant (since the number of vehicles and iterations remained the same), while the total number of cells in the system increased exponentially. Therefore, the ratio of searched to total cells grew less, resulting in smaller returns of increased confidence. Thus, Figure 3.10 illustrates the increase in total probability factor for the three algorithms as the field size is varied. Table 3.2 demonstrates that the MWA offers a total probability increase greater than that of either the MDA or the BPA for the entire range of side length, and that the difference increases as the side length increases. Therefore, the MWA is generally better applied to a large field (in terms of the number of cells) than to a small one.

Figure 3.9. Confidence with side length varied. Confidence was calculated after 100 iterations of the algorithm had been performed, with a fixed vehicle set of 10 vehicles.

Table 3.2. Total Probability Factor (TPF) Increase with Side Length Varied

Side Length	MWA	MDA	BPA
10	135.03	134.34	81.34
12	182.31	143.37	87.84
14	200.78	137.14	84.04
16	210.24	138.36	94.25
18	213.37	136.7	96.93
20	217.43	138.06	95.95
22	219.72	139.21	91.98
24	217.79	135.76	96.94
26	222.05	138.67	102.05
28	223.55	142.68	106.01
30	219.96	137.47	104.42

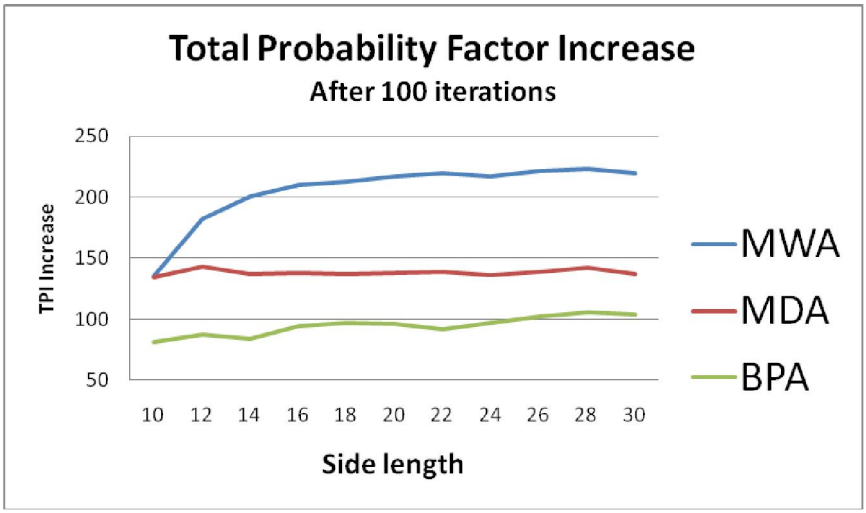

Figure 3.10. Total probability factor increase with side length varied, 10 vehicles. The probability factor is defined as the distance between the probability of the existence of a mine in a cell and 0.5; a higher probability factor is desirable, as a mine probability of 0 or 1, the desired condition, produces a maximal probability factor of 0.5. Total probability factor is the sum of the probability factors of all cells in the field. This figure demonstrates that, while 100 iterations of the algorithm do not greatly increase the overall confidence in a large field (as shown in Figure 8), the MWA is still producing similar gains in terms of the total increase in information about the presence of mines in the field.

4.3 Varying the Number of Vehicles

The number of vehicles was tested to measure the performance of the algorithms under conditions ranging from very crowded vehicles with many conflicts to sparsely distributed vehicles with few to no conflicts. With the field the same size as the control condition, the minimal number of vehicles to completely search the field in 100 iterations is 80. However, a case with 90 vehicles was also tested, as 80 vehicles are not likely to completely search the field in 100 iterations, given that some tasks will likely require more than one iteration to complete. For 90 vehicles, the entire field was searched by the MWA in 93 iterations rather than 100. The results of this test demonstrated that the MWA outperformed both the MDA and BPA over the entire test range, however the difference between the MWA and MDA grew small at both extremes. This is because the effect of optimizing cooperation is greater when more vehicles are used. However, as the field is searched more completely, prioritizing the cells with greater doubt becomes less important. Thus, the MWA best applies to a situation with many vehicles that do not have sufficient time to search the entire field. Figure 3.11 shows a graph of the final confidence as a function of the number of vehicles.

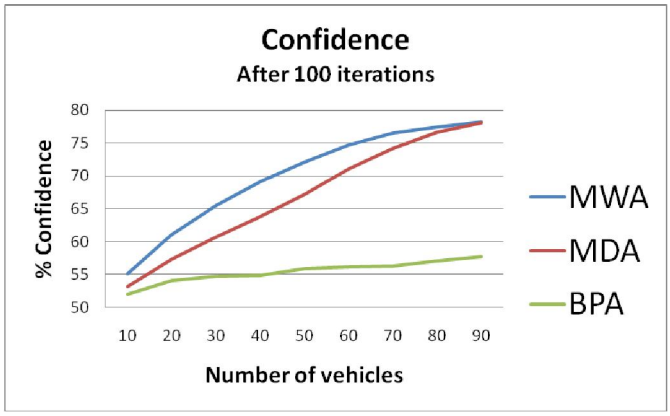

Figure 3.11. Confidence with Number of Vehicles Varied. Confidence was calculated after 100 iterations of the algorithm had been performed, with side length 20.

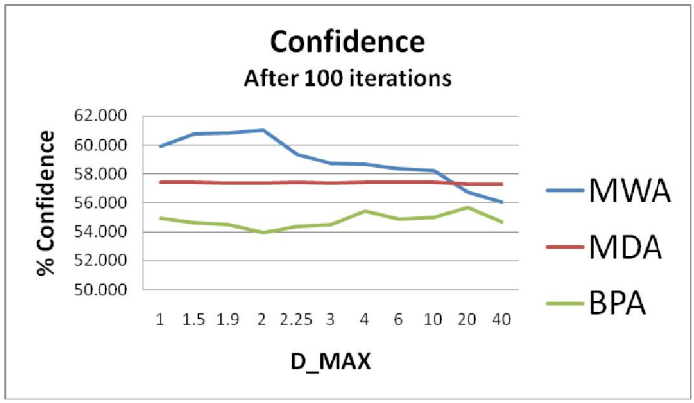

Figure 3.12. Confidence with D_MAX Varied. Confidence was calculated after 100 iterations of the algorithm had been performed, with a fixed vehicle set of 20 vehicles and side length 20. This demonstrates that choosing a D_MAX that will take the vehicle a long time to reach is detrimental to the performance of the MWA.

4.4 Varying D_MAX

D_MAX was varied to demonstrate that the maximum allowable distance from an assigned task to a vehicle affects the performance of the MWA. In a real-time scenario, this factor is more likely to be a function of the speed of the vehicles. In the test, it was demonstrated that setting D_MAX too low will lower performance because this prevents the vehicle from being assigned to tasks with a high probability factor, while setting D_MAX too high will result in large amounts of time wasted in travel between tasks while not gathering information as often as possible. Figure 3.12 shows the end confidence of all three algorithms as a function of D_MAX, and proves that D_MAX only affects the MWA. Thanks to its consideration of both distance and probability, the MWA still outperforms the MDA and BPA in most cases even for poor choices of D_MAX; however, the improvement becomes insignificant in the worst case. The optimal choice for D_MAX in the simulation was to set it equal to DV_MAX, the maximum distance the vehicle was able to travel in one iteration.

4.5 Varying Chaos

Chaos was introduced to simulate the performance of the algorithms in conditions ranging from a confident initial probability map (representing clear waters and/or precise surveillance equipment) to a meaningless initial map. Figure 3.13 shows a comparison of the mine score of the three algorithms under Q = 0.1 and gives a good indication of their respective strategies. The BPA generally attains its

final mine score very quickly when chaos is low because, according to the simulation strategy, there is generally more uncertainty in a cell containing a mine than in a clear cell. However, it is noteworthy that the MWA performs nearly as well after 100 iterations. Tables 3.3 and 3.4 show the increase in total probability factor and mine scores, respectively, of the three algorithms after 100 iterations. It is obvious that the MWA has a much higher increase in total probability factor than the BPA, while the MDA performs poorly by both standards. This test demonstrates that while the BPA is the first algorithm to reach the areas of highest doubt, it does not generally search quickly enough to provide much information beyond this in a short amount of time. Therefore, the BPA may be the best algorithm of the three if an initial probability map has a high level of confidence with only a few cells needing further analysis, but the MWA performs better if much more information is needed about the field beyond the initial map.

Table 3.3. Total Probability Factor (TPF) Increase with Chaos Q Varied

Q	MWA	MDA	BPA
0	0	0	0
1	64.42	41.9	26.77
2	124.02	82.61	50.44
3	176	114.6	74.35
4	219.79	142.21	86.52
5	242.03	162.07	101.38
6	223.02	166.08	87.11
7	220.11	164.46	94.54
8	219.15	160.47	96.21
9	218.12	149.68	90.19
10	217.3	137.88	94.56

Table 3.4. Mine Score with Chaos Q Varied

Q	MWA	MDA	BPA
0	0.12	0.15	0.02
1	0.61	0.09	0.69
2	0.48	0.13	0.57
3	0.49	0.14	0.52
4	0.33	0.12	0.34
5	0.17	0.15	0.07
6	0.12	0.13	0.03
7	0.1	0.1	0.04
8	0.1	0.1	0.04
9	0.1	0.16	0.04
10	0.15	0.09	0.06

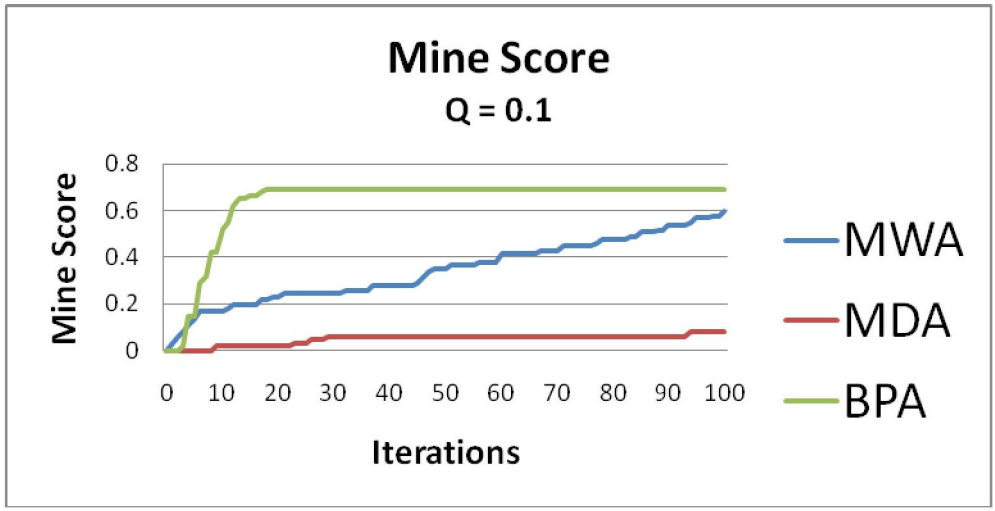

Figure 3.13. Mine score after n iterations. Mine score is the ratio of mines discovered by the algorithm to the total mines present in the field. This figure shows the mine score after n iterations with chaos Q = 0.1, indicating that all clear cells were given an initial probability between 0 and 0.1, and all cells containing mines were given an initial probability between (1 – √0.1) and 1, or 0.684 and 1. The observed upper limit on the mine score is due to the fact that some cells containing mines have very "confident" initial data, for instance a mine probability of > 0.9, and will not be given search priority over clear cells that have a similar degree of doubt. Note that while the BPA tends to discover the most mines earliest (since it disregards the distance of the vehicle from the task), the MWA makes quick and steady progress, nearly equaling the BPA in this case after only 100 iterations.

4.6 Varying Mine Frequency

Mine frequency was varied to simulate conditions from clear water to moderately populated with mines. Variations were by 0.1% from 0 to 2.5% of the map. As this was another case of designed initial probability, each 0.1% variation in frequency resulted in a difference of one "cluster" of 8 mines (either a straight line or a small circle). As shown in Figure 3.14, varying the frequency of mines had little or no effect on the comparative performance of the algorithms. After 100 iterations, the MWA produced greater confidence than either of the other two algorithms at all ranges, with a mine score comparable to that of the BPA and much greater than that of the MDA (see Figure 3.15).

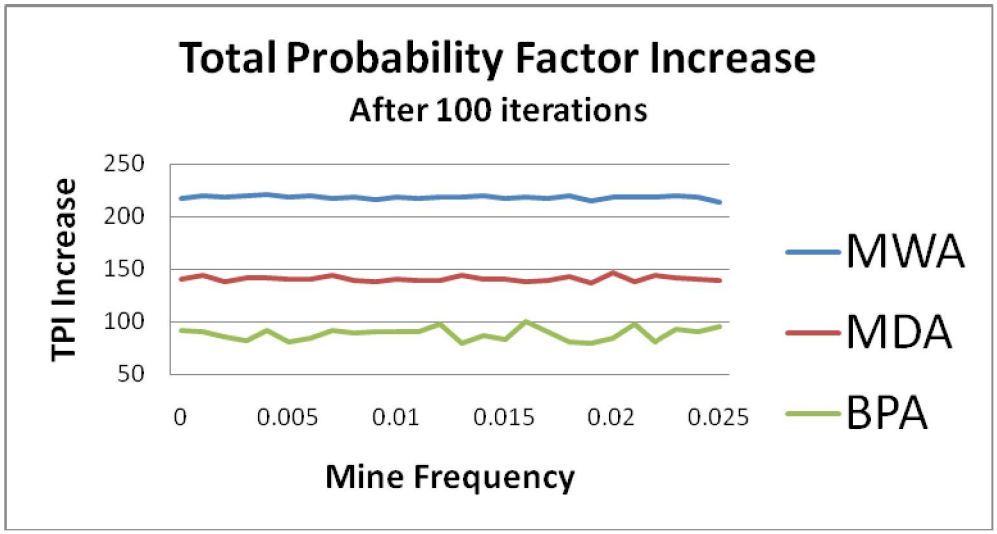

Figure 3.14. Total probability factor increase with mine frequency varied. This demonstrates that, while comparing the performance of the algorithms based on mine score may give mixed results as mine frequency is varied, the MWA still outperforms the MDA and the BPA in terms of the increase in information about the confidence regarding presence or absence of mines in the field.

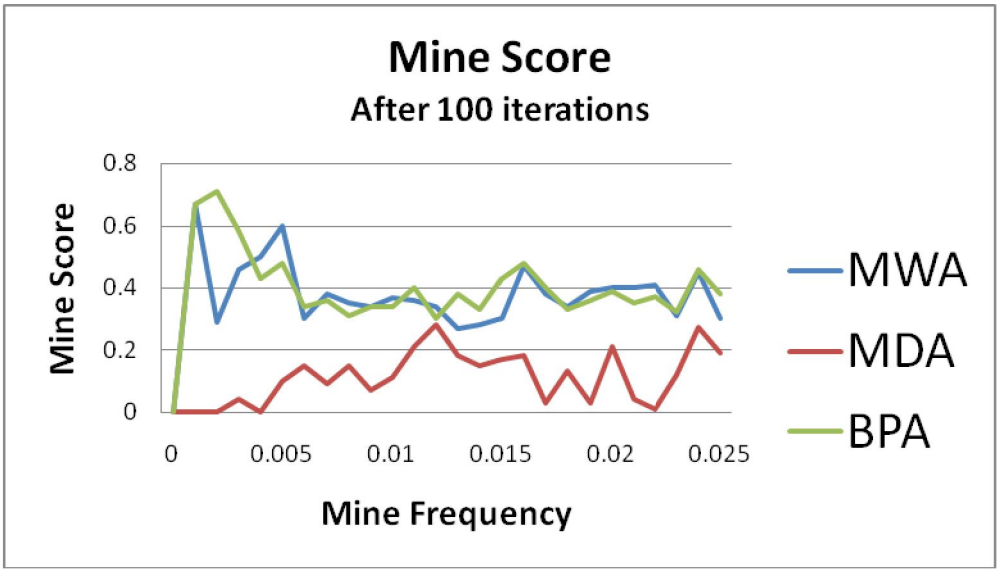

Figure 3.15. Mine score with mine frequency varied. Mine score is calculated after 100 iterations of the algorithms.

5. DISCUSSION

Under the test conditions, the MWA consistently outperformed both the MDA and the BPA. The only case in which either of the other algorithms performed measurably better than the MWA according to any metric was when mine score was the defining measurement of performance, the number of mines was known, and very few iterations were performed; the BPA visited the mine cells sooner than the MWA in this case. However, in this case and others, the MWA produced a considerably more confident resulting map than either of the other two algorithms when a search of the field did not run very near to completion.

6. CONCLUSION

The major contribution in this chapter is a time-constrained algorithm presented to optimize the information gained by a team of UUVs working in cooperation to map a three-dimensional mine field under an unknown time constraint, given an initial probability grid. The algorithm was implemented in a test program along with two other algorithms proposed by past research. In the simulated search, this algorithm provided better results than either of the others under almost any of the controlled conditions.

CHAPTER 4

COMMUNICATION STRATEGIES IN MULTI-UUV COOPERATIVE MCM RECONNAISSANCE

1. INTRODUCTION

This chapter demonstrates the effect that limiting communication has on the performance of a cooperative UUV team. An UUV team is considered to be cooperative if an UUV's decision regarding its task assignments is influenced by information obtained from other UUVs. One such cooperative algorithm, the Minimal Weight Algorithm (MWA), is used here to demonstrate that the degree of communication between vehicles can significantly affect the performance of any algorithm that attempts to take advantage of cooperation among vehicles to optimize the information gained from a search mission. We quantify the effect of the degree of communication on search results to emphasize the importance of efficient communication.

The primary contribution of this chapter is to present several information sharing strategies, each seeking to optimize a facet of communication efficiency, and to define the contexts in which each strategy is most appropriately used. Some existing work has explored the limitations of communication but has offered little advice on how to compensate for them to aid UUV cooperation. The problem addressed by this chapter is, given a set of environmental conditions (particularly the size of a field), an ad-hoc network containing limited UUV and communication resources, and some initial information regarding the presence of mines in that field, to choose an optimal communication strategy for that environment that will result in most efficiently obtaining information from the UUV mission.

2. EXPERIMENT

2.1. Appropriateness of the MWA

The Minimal Weight Algorithm presented in Chapter 3 is a representative of the set of cooperative multi-UUV control algorithms. It has several facets that make it a particularly useful subject for testing the effect of ad-hoc network limitations on the performance of a cooperative algorithm. The initial probability map, the tendency to minimize travel distance, and the increased likelihood of proximity of vehicles to each other all contribute to make this algorithm ideal for this study.

2.1.1 Initial Probability

By using an initial probability for each cell, the input to the MWA renders some tasks more profitable than others, and thereby provides incentive for UUVs to communicate with each other to prioritize completion of the most important tasks. This environment is preferable to a field with uniform probability distribution, which would be most efficiently searched by minimizing interaction between vehicles with each other in order to reduce the potential for assignment conflict.

2.1.2 Precedence for Closer Vehicles

The MWA tends to give precedence to the vehicle closer to a task, increasing the benefit of communication to the performance of the search. When two vehicles are assigned the same task, the closer vehicle will generally be assigned to the task because, all other tasks held equal, this arrangement will result in a lower total weight for the vehicle assignments. Such conflict resolution requires communication between vehicles, and therefore provides a good test case for evaluating the effect of limiting range or varying amount of information shared between vehicles.

2.1.3 Increasing Proximity of Vehicles

Using weight as a primary factor of assignment precedence has the tendency to cause vehicles to drift nearer to each other. The initial probability map contains certain cells that are inherently more likely to be searched by vehicles; these cells will tend to be clustered in areas of particular interest within the field. Each cell in these areas will likely have a low weight factor, which will tend to draw the attention of vehicles toward these areas and away from areas which have little doubt as to the presence of mines. As vehicles congregate in the important areas of the map, the degree of communication and shared information increases. Therefore, the Minimal Weight Algorithm provides an ideal environment for testing the effect of communication on the level of cooperation, because it tends to increase cooperation where possible.

2.2 Modifications for an Ad-hoc Network

A few modifications were made to adapt the Minimal Weight Algorithm to an ad-hoc environment. The Minimal Weight Algorithm for three-dimensional search was designed with the assumption of an omniscient central controller. For test purposes, this algorithm needed to be adjusted to reflect the more realistic scenario of an ad-hoc UUV network with imperfect communication. In such a network, each

UUV may or may not see any other UUV, and it can only communicate with other vehicles within range. This had several implications for implementation of the MWA.

2.2.1 Distributed Probability Maps

The most important modification was that each vehicle now must carry its own version of the probability map P and the searched location map S. Also note that the assignment map A is now eliminated, as there is no way to store this map globally and it holds little meaning for the vehicles, which must communicate with each other directly anyway to determine any assignment conflicts. These maps were previously held only by the omniscient controller; however, each vehicle must now be responsible for choosing a new task, as well as keeping track as best it can of which tasks have been assigned or completed. The accuracy of these maps may be compromised by their distributed nature; thus, extra communication is necessary to improve the degree of consensus among the vehicles regarding the environment. These measures are discussed later in the explanation of the information strategies tested.

2.2.2 D_SEE

Another modification to the simulation model involved the addition of a new variable, D_SEE. This variable represents the communication range of each vehicle. For simulation purposes, it is assumed that each vehicle will have the same communication range, though this could be easily modified in the algorithm to support vehicle heterogeneity. This modification affects the commission step of the algorithm, as only other vehicles within D_SEE units of the updating vehicle will receive the update, as shown in Figure 4.1.

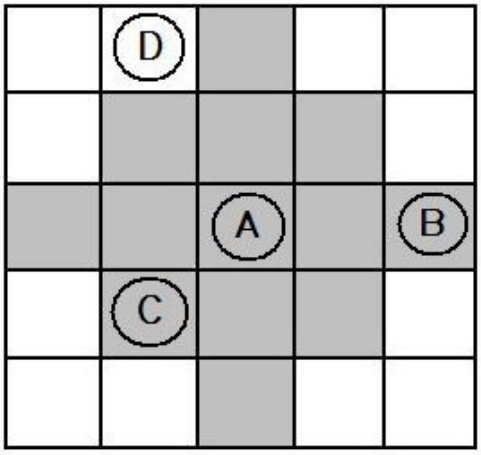

Figure 4.1. Example Vehicle Update Range. In this example, with D_SEE = 2, vehicle A has just completed a task and is broadcasting the results to its neighbors. Vehicles B and C will receive the update because they are within 2 units of vehicle A; however, vehicle D will not receive the update.

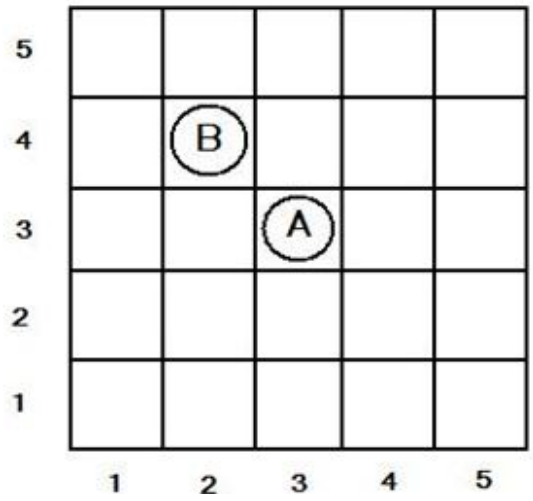

Figure 4.2. Context for Full and Minimal Update Strategies. In a full update strategy, after completing task (2, 2), A sends B a probability and timestamp for each cell in its own map, i.e. {$P_{1,1A}$, $T_{1,1A}$; $P_{1,2A}$, $T_{1,2A}$; … ; $P_{5,4A}$, $T_{5,4A}$; $P_{5,5A}$, $T_{5,5A}$}. In a minimal update strategy, A sends B only {$P_{0,0A}$, $T_{0,0A}$}.

2.3 Information Sharing Strategies

The strategy for sharing information among vehicles can have a great effect on the end performance of any search algorithm that seeks to take advantage of cooperation. If no information is shared between vehicles, there is nothing to prevent every vehicle from searching the same cell from being searched by every vehicle in the UUV team, completely destroying all benefit of using multiple UUVs. However, depending on environmental constraints such as battery life, wireless network capabilities, or the threat of enemy detection of excessive communication, sharing all information may prove prohibitive. Various strategies for sharing information are presented here, offering different performance priorities.

2.3.1 Full Update Strategy

A full update strategy results in every vehicle sharing all of its knowledge about the map. In each update, an UUV will broadcast a message to all other vehicles containing a probability value and a timestamp for each cell in the map, as shown in Figure 4.2. The probability value may change either as a result of the vehicle searching a cell and updating its perceived probability of containing a mine accordingly, or as a result of information communicated by another vehicle that has searched that cell.

The timestamp determines whether or not a vehicle receiving the message should update its own probability map with the information. If the vehicle receiving the message has a more recent probability value for a particular cell, it would be counterproductive to update its own probability map with obsolete information.

An alternative strategy that produces an equivalent performance is to modify the full update strategy to broadcast only those locations that have been updated since the original commission of the vehicles. This would initially result in very little data being passed between vehicles, as it is useless for a vehicle to communicate knowledge that is held in common with every other vehicle. However, as the search nears completion, the messages sent will grow just as lengthy as in the unmodified full update strategy. The expected result of a full update strategy is a high degree of cooperation and good performance, but also high message overhead.

2.3.2 Minimal Update Strategy

In contrast to the full update strategy, a minimal update strategy causes a vehicle to communicate only information regarding a single task. The benefit is that no vehicle will ever receive information that is not useful; each vehicle will only send messages regarding cells that it has personally searched, and will only send such a message once. Of course, this requires that any vehicle must be within communication range of the vehicle sending the update in order to receive the information. Unlike the full update strategy, the minimal update strategy offers no way to pass "secondhand" information between vehicles. In a small environment or in the presence of excellent communication range, this may affect performance only negligibly; however, the full update strategy will generally result in better cooperation between vehicles. A minimal update strategy is expected to provide messages of a much more manageable quantity and length than the full update strategy, but with compromised algorithm performance.

2.3.3 Partial Update Strategy

The minimal update strategy is a special case of a partial update strategy. Partial update seeks to communicate only a limited number of tasks at a time; for instance, with a number of tasks to update n = 10, a vehicle broadcasts its 10 most recent completed tasks, as shown in Figure 4.3. This allows many more vehicles to benefit from the updates, as a receiving vehicle must be within range of the sending

vehicle for only 1 of 10 tasks to receive a full update. Practically, the partial update strategy does not fully solve either the full update strategy's problem of redundant information sharing or the minimal update strategy's lack of secondhand communication; however, it offers an interesting potential for compromise that is worth study. The partial update strategy with $n > 1$ should result in improved algorithm performance over the minimal update strategy and messages of a restricted length, which is preferable to the full update strategy.

2.3.4 Periodic vs. Immediate Update

With periodic updates, a vehicle would broadcast messages at constant intervals, as opposed to immediate update, in which a vehicle broadcasts messages upon completion of a task. Either strategy could be used with full, partial, or minimal updates, but a periodic update strategy would seem to have the greatest benefit under a full update strategy. If the environment offers a high degree of communication among vehicles, a periodic update could be used with a long update interval in place of immediate update, to reduce the number of messages sent and alleviate overuse of resources. On the other hand, if the environment offers only a very limited communication range, periodic updates could be used more frequently than immediate updates in an effort to share information at every possible opportunity. These benefits do not come without consequence, however; if the interval is set too long, a vehicle may redundantly visit a cell because it did not receive an information update soon enough, while setting the interval shorter than immediate updates will result in redundant messages. Immediate updates offer less flexibility than periodic updates, but with an inherently more efficient and event-driven message-sending strategy.

2.3.5 Other Strategies

Many other update strategies can be conceived that are not implemented or tested here. For instance, a vehicle might transmit only the updates that have not already been received by neighboring vehicles. Or, a vehicle might send updates whenever a new vehicle comes in range. Either of these could be combined with a periodic or immediate update strategy in an effort to offer maximum communication. However, these require much communication independent of information updates, such as information acknowledgments or periodic self-identification broadcasts, thereby complicating the analysis of message overhead and placing such strategies beyond the scope of this study.

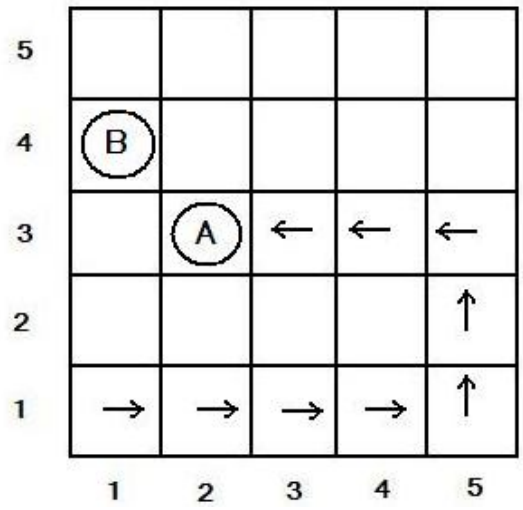

Figure 4.3. Context for Partial Update Strategy. In this example, with n = 10, the arrows represent the search path of vehicle A. In partial update, after completing task (2, 3), A sends B a probability and timestamp for each of its last 10 cells searched, {$P_{1,1A}$, $T_{1,1A}$; $P_{2,1A}$, $T_{2,1A}$; $P_{3,1A}$, $T_{3,1A}$; $P_{4,1A}$, $T_{4,1A}$; $P_{5,1A}$, $T_{5,1A}$; $P_{5,2A}$, $T_{5,2A}$; $P_{5,3A}$, $T_{5,3A}$; $P_{4,3A}$, $T_{4,3A}$; $P_{3,3A}$, $T_{3,43}$; $P_{2,3A}$, $T_{2,3A}$}. In a minimal update strategy, A would send B only {$P_{2,3A}$, $T_{2,3A}$}.

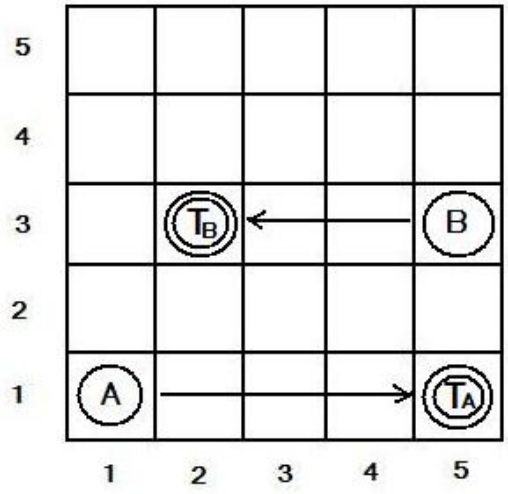

Figure 4.4. Immediate vs. Periodic Updates. In immediate update, A will only send B an update on reaching T_A, when B will be out of range. However, in periodic update with period 1, A will send B its recent updates at each unit on its way to T_B, so B will receive more information from A.

3. TEST RESULTS

3.1 Evaluation Metrics

The factors evaluated in the experiment seek to measure the effect of the cooperative efforts in terms of the volume searched and messages sent between vehicles. The evaluation of a search algorithm performance under various communication ranges or update strategies is different from an evaluation that seeks to demonstrate one such algorithm's improvement over another. Any algorithm that seeks to use cooperation among multiple vehicles to improve a search will be affected similarly by varying the amount of cooperation among those vehicles. An ideal cooperative search is one that maximizes the volume searched, while minimizing the number of iterations required, the number of redundant visits to individual cells, and messages and data required to communicate results between vehicles.

3.1.1 Maximize Volume Searched

All other factors held equal, it is desirable to maximize the volume searched by the UUV team. When comparing one search algorithm against another, the algorithm that searches fewer cells may indeed produce a smaller information gain than the other algorithm; however, when only considering one search algorithm, it is always beneficial to search more cells. A higher amount of cooperation between vehicles will result in a larger volume searched, because any vehicle will be less likely to search a cell that has already been searched by another vehicle, leaving it free to search an unsearched cell. Therefore, one measure of the amount of cooperation under a particular communication range or information sharing strategy will be the total percent of the field searched by the UUV team.

3.1.2 Maximize Confidence

Confidence refers to the level of assurance in the map as to the presence or absence of mines. 100% confidence is defined as the environment in which the probability of each cell containing a mine is either 0 or 1. 0% confidence represents an environment in which the probability of each cell containing a mine is 0.5. The total confidence at the end of a search is the sum of the deviation of each cell's probability from 0.5 ($|P_i - 0.5|$), divided by the sum of all cells with an ideal deviation of 0.5, i.e. confidence $C = \sum|P_i - 0.5| / \sum 0.5$ over all cells.

3.1.3 Minimize Number of Iterations in Full Search

Optimal cooperation among vehicles should result in a minimal number of iterations of an algorithm required to fully search the field. Under any algorithm, vehicles should accomplish a full search more quickly if they are more aware of which cells have already been searched by other vehicles. This awareness is heightened by a higher degree of information sharing or communication with other vehicles. Another measure of the amount of cooperation under one search algorithm, then, will be to run the algorithm to completion and find the number of iterations required for a full search.

3.1.4 Minimize Redundant Visits to Individual Cells

The ideal cooperative effort should result in no redundant visits to any cell. Assuming that all vehicles have identical search capabilities, no second vehicle visiting a particular cell will obtain more information than did the first, rendering a second search useless. Ideally, each vehicle should then communicate its task completion to every other vehicle, and no vehicle should ever search a cell that has

previously been searched by another vehicle. The number of redundant visits will therefore be another factor to measure in evaluating the cooperative efforts of the UUV team.

3.1.5 Minimize Number of Messages and Data Sent

Ideal communication should seek to minimize the number of messages sent to update other vehicles. Optimally, no unnecessary messages or data should be sent; vehicles should communicate only to prevent redundant visits to cells. Thus, one factor used to compare the performance of various information sharing strategies will be the number of messages and the length of the messages used to communicate information between vehicles, with message length measured in terms of the number of locations being updated within the message.

3.2 Control Case

Test results were measured against an initial simulation using an immediate full-update strategy. The dimensions of the field were 20 x 20 x 20 units (a total of 8000 locations to search), and 20 vehicles were used. The Minimal Weight Algorithm was run for 1000 iterations, with D_MAX set to 6 and a communication range D_SEE of 4 units. Each vehicle was given a uniform speed of 1 unit per iteration. The test results for the control case are shown in Table 4.1. The full update strategy in the control case sent a total of 11759 messages with 18088747 location updates, an average of 1538 location updates sent per message, searching 67% of the field.

Table 4.1. Control Condition

Volume searched	**67%**
Final confidence	69.3%
Messages sent	11759
Updates sent	18088747
Longest message sent	3891
Average updates per message	1538

3.3 Varying Communication Range

Communication range of the vehicles was varied to measure the effect of the communication range on performance. Each test measured the percentage of the field visited by the vehicles, as well as the percentage of the field visited multiple times. The expected result was that increasing the communication range of the vehicles would increase the cooperation among the vehicles and the sharing

of information. This should increase the area searched while ultimately decreasing the percentage of the area that is searched more than once, because the vehicles should target unsearched locations more frequently as they learn more information about the results obtained by other vehicles.

As expected, increasing communication range increased the percentage of the field searched, as shown in Figure 4.5. D_MAX was set to 10 for this particular test, in order to result in a more gradual performance curve, better demonstrating the behavior of the algorithm as D_SEE increases. Under this set of dimensions and vehicles, increasing D_SEE beyond 20 resulted in a complete search of the field; however, improvement in performance was still measurable by the number of iterations required for a complete search, as shown in Table 4.2. Quantitatively, performance was roughly doubled with each increase of 10 units in range, until the entire field was within range. This suggests that the increase in performance may be exponential with respect to an increase in wireless range, up to a point. The percentage of the field searched multiple times did not decrease until the range was reached at which a search was completed. However, it is noteworthy that the volume searched increased faster than the volume searched multiple times; this suggests that as range was increased, search efficiency was monotonically increasing.

Table 4.2. Varying Communication Range

D_SEE	Iterations	% searched	% multiple visits
0	1000	28	23
2	1000	28	23
4	1000	32	27
6	1000	38	31
8	1000	45	36
10	1000	54	43
12	1000	69	52
14	1000	80	58
16	1000	94	66
18	1000	99	72
20	815	100	75
22	584	100	41
24	466	100	22
26	428	100	8
28	417	100	6
30	414	100	2
32	412	100	0
34	410	100	0
36	410	100	0

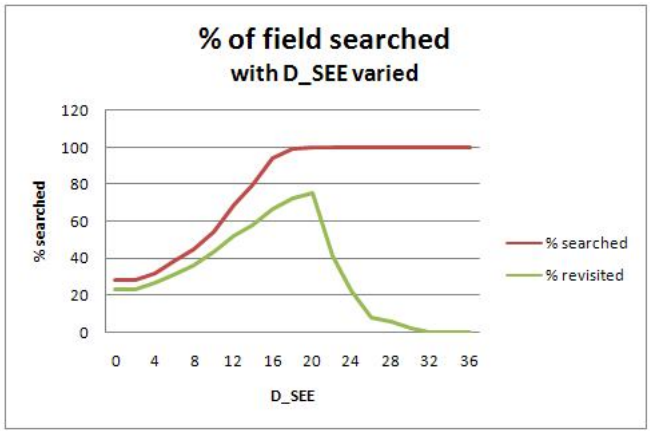

Figure 4.5. Percent of Field Searched with D_SEE Varied. It is notable that the % searched and % revisited both increase steadily as the communication range D_SEE increases, until the point at which the entire field is searched, after which increasing D_SEE serves to decrease search redundancy.

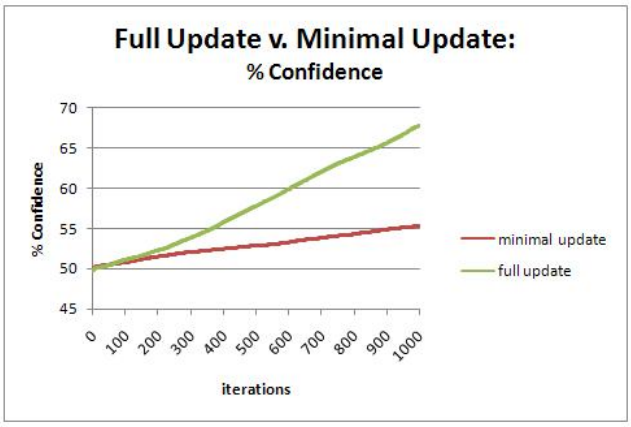

Figure 4.6. Percent Confidence after n Iterations, Full Update. Confidence is measured under full update and minimal update strategies. As can be seen, a full update strategy provides significantly improved algorithm performance over a minimal update strategy.

3.4 Minimal Update

Minimal update simulation conditions were D_MAX = 6 and D_SEE = 4 (communication range). A minimal update strategy, as expected, offered inferior performance to that of the full update strategy. While the full update strategy produced a 45.5% search of the field in 100 iterations in the control condition, the minimal update strategy produced only an 11.3% search under the same conditions. The performance curve is shown in Figure 4.6. Similarly, while the full update strategy produced a 30.0% redundant search rate (65.9% of total searched units), the minimal update strategy produced a 9.50% redundant search rate (84.1% of total searched units). However, minimal updates resulted in only 8568 messages sent, with only one location updated in each. This is 0.000474 of the data sent in the full update strategy, and this ratio would grow more extreme as the dimensions are expanded.

3.5 Partial Update

As in the minimal update, the partial strategy conditions were D_MAX = 6 and D_SEE = 4. The partial update strategy offered a performance between that of the minimal update and full update strategies, with interesting insight into the tradeoff between performance and resource consumption. Figure 4.7 shows the performance in terms of the percentage of the field searched over a number of partial update searches with varying update count, along with the minimal and full update cases, as well as the special

case of partial update with n = 0, representing a case with no communication among the vehicles, and thus no cooperation. Figure 4.8 shows results from the same test as Figure 4.7, in terms of the % confidence obtained by the search.

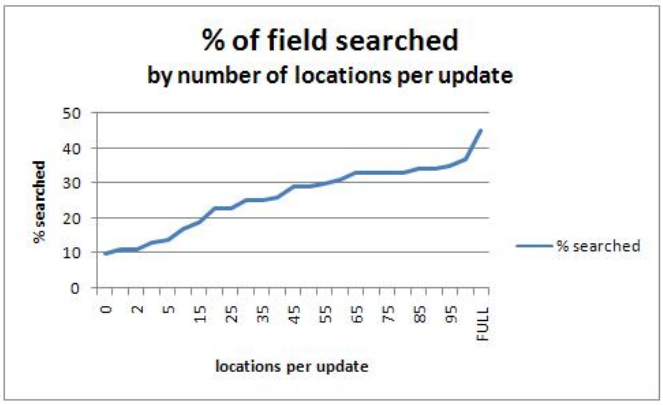

Figure 4.7. Percent of Field Searched, Locations per Update Varied. Percent is measured after 1000 iterations of the algorithm, with number of locations per update varied. The 0 locations per update condition is the condition with no communication between the vehicles. The minimal update condition has 1 location per update. The full update condition has a variable number of locations per update, which eventually exceeds 100 locations per update, the maximum tested for the partial update strategy.

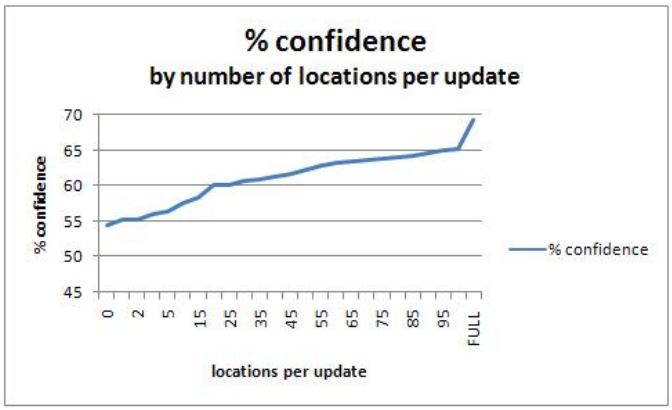

Figure 4.8. Percent Confidence at Completion, Locations Per Update Varied. Confidence is measured after 1000 iterations of the algorithm, with number of locations per update varied. The 0 locations per update condition is the condition with no communication between the vehicles. The minimal update condition has 1 location per update. The full update condition has a variable number of locations per update, which eventually exceeds 100 locations per update, the maximum tested for the partial update strategy.

The advantage of a partial update over a full update is that the number of locations sent in an update is limited, and so does not grow unmanageably as dimensions are increased. Table 4.3 shows the total amount of data (in number of locations sent) transferred by the range of partial update tests, as well as minimal and full update. Note that even sending the last 100 locations updated results in a >1600% increase in data transmitted but only a 26.6% increase in the performance. In general, the relationship between the amount of communication and the performance of the cooperative algorithm was demonstrated to be logarithmic.

Table 4.3. Total Messages and Data Sent

Locations per update	Messages sent	Data sent
0	**0**	**0**
1	**8535**	**8535**
2	**8325**	**16650**
3	**8379**	**25137**
5	**9175**	**44040**
10	**8721**	**87210**
15	**9048**	**135720**
20	**9480**	**189600**
25	**9395**	**234875**
30	**9771**	**293130**
35	**9843**	**344505**
40	**9893**	**395720**
45	**9576**	**430920**
50	**9847**	**492350**
55	**10017**	**550935**
60	**10151**	**609060**
65	**10134**	**658710**
70	**9982**	**698740**
75	**10370**	**777750**
80	**10276**	**822080**
85	**10581**	**899385**
90	**10278**	**925020**
95	**10500**	**997500**
100	**10367**	**1036700**
FULL	**11759**	**18088747**

3.6 Periodic Update

3.6.1 Minimal Periodic Update

The minimal periodic update was tested with D_SEE = 4 and varied D_MAX. D_MAX was varied so that the impact of the periodic updates could be seen as the vehicles had less opportunity for immediate update. Other variables were held equal to the control condition. Minimal periodic update tests were compared to minimal immediate update tests at the same D_MAX, and performance results are shown in Figure 4.9.

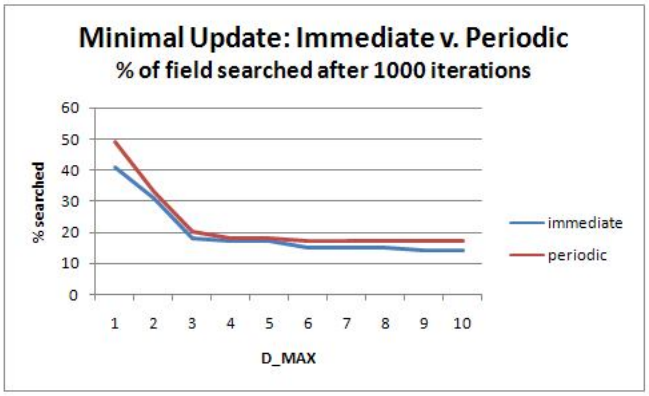

Figure 4.9. Percent of Field Searched with D_MAX varied, Immediate vs Periodic. Percentage is measured after 1000 iterations under immediate and periodic minimal update strategies. Periodic update (period = 1) consistently offered a minor improvement over immediate update in the minimal update strategy.

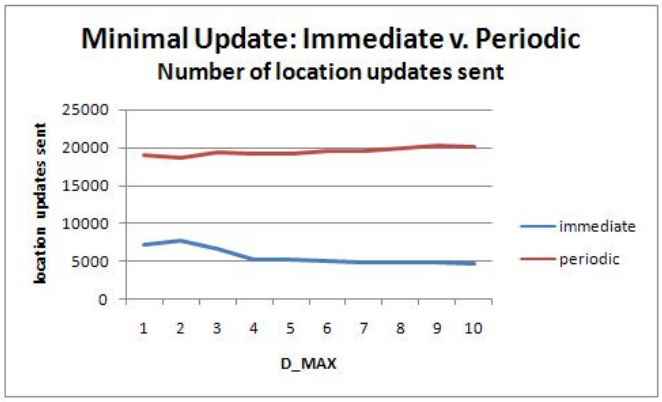

Figure 4.10. Number of Location Updates with D_MAX varied, Immediate vs. Periodic. Location updates are measured after 1000 iterations under immediate and periodic minimal update strategies. Periodic update (period = 1) resulted in a significant increase in the number of messages sent. It is noteworthy that the number of periodic updates actually increased slightly as D_MAX was increased. This is because periodic updates were only performed when another vehicle was in communicating range and increasing D_MAX tends to emphasize the probability of mines over proximity in task selection, drawing vehicles toward more concentrated areas and therefore closer to each other. As a result, vehicles tended to be in range of each other slightly more often when D_MAX was higher, providing more opportunity for periodic update.

Performance is very similar under both of these strategies. A slight performance improvement can be seen in the minimal periodic update strategy. This is because there is slightly more opportunity for a vehicle to broadcast the completion of a task to other vehicles that may not have been in range when the task was finished. However, the improvement in performance is likely not worth the significant increase in communication, as shown in Figure 4.10.

3.6.2 Full Periodic Update

The full periodic update conditions were the same as the minimal periodic update. Full periodic update tests were compared to full immediate update tests at the same D_MAX, and performance results are shown in Figure 4.11.

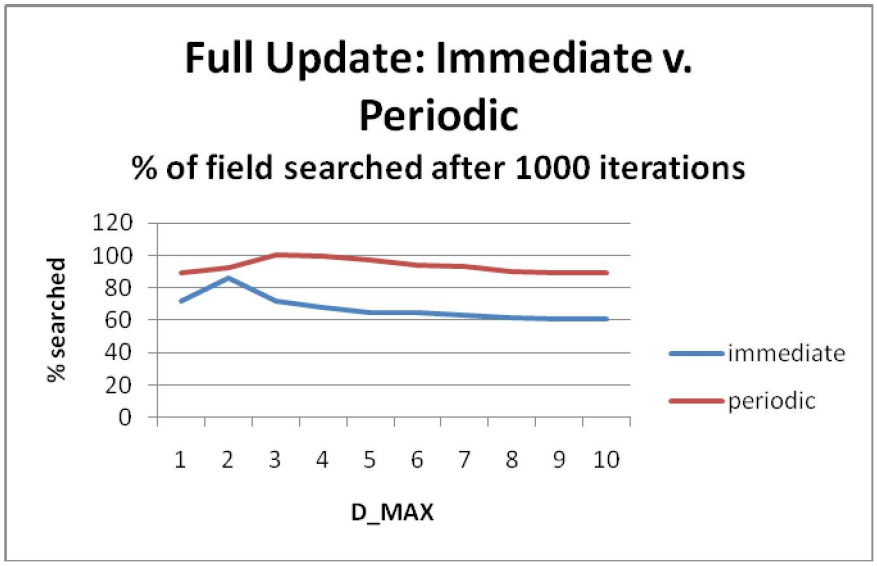

Figure 4.11. Percent of Field Searched with D_MAX Varied, Immediate vs Periodic. Percent is measured after 1000 iterations under immediate and periodic full update strategies. Periodic update (period = 1) offered a marked improvement over immediate update in the full update strategy.

As Figure 4.11 demonstrates, performance decreases more sharply as D_MAX is increased with full immediate update than with full periodic update. This would be expected, as update frequency in a periodic update strategy is not dependent on the frequency of completion of tasks. The result is that a periodic update improves performance as the vehicles travel greater distances to complete tasks. But this improvement is more dramatic in a full update strategy than in a minimal update strategy.

The number of updates decreases under a full immediate update as D_MAX is increased. This is because vehicles complete tasks less often, and therefore have less opportunity to submit immediate updates of these tasks. However, under a full periodic update, as in the minimal periodic update, the number of updates remains roughly the same, as shown in Figure 4.12, since vehicles always broadcast updates at the same rate under a periodic update strategy regardless of their frequency of task completion, and the completion of tasks also has little effect on how frequently the vehicles pass within communication range of each other. Thus, the improvement of a full periodic strategy over a full immediate strategy grows as the vehicles complete tasks less frequently, without significantly increasing the amount of communication between the vehicles.

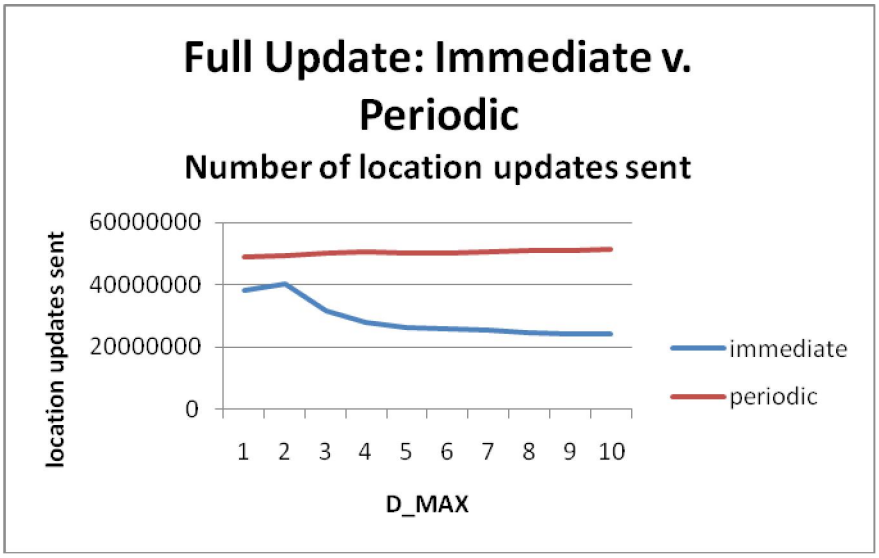

Figure 4.12. Number of Location Updates Sent with D_MAX Varied, Immediate vs Periodic. Location updates are after 1000 iterations under immediate and periodic full update strategies. Periodic update (period = 1) resulted in a significant increase in the number of messages sent. Number of updates sent varies very little in the periodic full update strategy, compared to the immediate update strategy, as D_MAX is increased, because a periodic update strategy is not dependent upon task completion.

4. DISCUSSION

Under the test conditions, varying the communication range had the most dramatic effect on algorithm performance. This suggests that the most important element of a cooperative search is the extent to which vehicles are able to communicate.

A full update strategy produced the best performance results, but with an enormous penalty in consumption of resources. A partial update strategy offered similar performance results without a prohibitive increase in communication over the course of the search, suggesting that this strategy would be preferable if power or network resources were a concern, or if the environment demanded the mission emphasize covertness. A minimal update strategy, as a special case of the partial update strategy, offered a significant performance increase over a strategy with no cooperation, but offered very poor performance compared to the other cases of partial update or full update.

Periodic updates generally offered an improvement over immediate updates, but this improvement is qualified by the increase in communication, offering similar concerns to those regarding a full update

strategy. The performance increase grew more dramatic as the updates grew more frequent compared to the rate at which the vehicles completed tasks, suggesting that in a map with many distinct locations of interest, a periodic update strategy could be very beneficial. However, in a field of relatively uniform probability distribution, or in which the vehicles are not generally separated by great distances, the significant communication penalty of a periodic update strategy would likely not be worth the improved performance.

5. CONCLUSION

Search teams of multiple UUVs are becoming increasingly important to MCM reconnaissance missions. However, algorithms that take advantage of vehicle cooperation to efficiently assign tasks to the vehicles are less effective when communication between vehicles is limited. We identified and measured the effect of such communication limits on one cooperative algorithm, the Minimal Weight Algorithm. We presented several strategies for sharing information between members of an UUV search team and measured their impact on the performance of the Minimal Weight Algorithm in a simulation of a realistic environment. Results showed that the level of communication between vehicles has a dramatic effect on the performance of a cooperative search, and that an ideal information sharing strategy is dependent on the goals and constraints of the mission. Finally, several suggestions were made as to what types of missions particular strategies would be best suited to.

CHAPTER 5

COOPERATIVE MCM FOR HETEROGENEOUS UNMANNED VEHICLES UNDER CONSTRAINED TIME

1. INTRODUCTION

1.1 Contributions

This study extends the algorithm developed in Chapter 3 to accommodate heterogeneous vehicles with different mission directives based on their assigned tasks to achieve an efficient MCM operation under constrained time. This requires separate heuristics for mine reconnaissance and neutralization. A primary goal of the study is to produce a strategy whereby each class of vehicle may accomplish its own task efficiently, while also assisting another class to better complete its task. We attempt to demonstrate the effectiveness of this strategy by showing improved results from the vehicles of one class as vehicles of another class are added.

We also address the claim that a single pass is not an ideal mission plan by testing the effectiveness of a single pass and multiple passes under constrained time. Specifically, we investigate whether the improvement of sensory data truly warrants a multiple-pass strategy when faced with an unknown time constraint. Finally, we discuss implications of this study for further heterogeneity of unmanned vehicles in the future development of MCM.

2. CLASSIFICATION, IDENTIFICATION, AND NEUTRALIZATION

Under the conditions of the Minimal Weight Algorithm, an initial map divides the target field into individual cells in three dimensions. This map contains a probability for each cell that the cell contains a mine. The map is supposed to have been obtained by an aerial reconnaissance sweep. Therefore, in the model of this study, the aerial sweep is presented as the detection stage of the MCM operation. As such, we do not address the detection stage, but assume some initial knowledge of the problem environment. The remaining steps are classification, identification, and neutralization of mines. In practice, these tasks could be allocated in three separate passes over the target area. Classification and identification both aspire to only obtain information about the environment, while neutralization is the removal of the threat. For simplicity, this study allocates the classification task to the HWVs, and identification and neutralization both to the LWVs.

2.1 Classification

In this study, the classification step of the MCM operation is assigned to the HWVs. The desired goal of this task is to obtain a maximum amount of information about detected obstacles in a target field. Thus, it is desirable to first visit the cells which have the greatest doubt as to the presence of a mine, while giving lower priority to searching cells that either very probably contain or very probably do not contain a mine. Under constrained time, it is also important to visit cells in close proximity to the search vehicle, to most efficiently use the time available to the UUV team. The Minimal Weight Algorithm (MWA) attempts to meet both these goals, and is therefore an appropriate strategy for the HWV operation. It relies on cooperation between vehicles to optimize the collective data obtained by the UUV team. The algorithm is described in its entirety in Chapter 3.

2.2 Identification

2.2.1 LWV Identification/Neutralization Tasks

The identification and neutralization steps of the MCM operation are assigned to the LWVs. The priority of these tasks is different from the classification step; while vehicles assigned to classification are most interested in those cells of greatest doubt, identification and neutralization prioritize the cells which most probably contain a mine. The goal of this pass is to achieve as high a rate of neutralization per search time as possible. Of course, much of the strategy behind the Minimal Weight Algorithm still applies; an LWV should still be assigned to a nearby task to make efficient use of time, and if two LWVs are assigned to the same task the conflict must be resolved, similar to the HWVs. Therefore, a modified version of the MWA is proposed to command the LWVs.

2.2.2 Modifications to the MWA

2.2.2.1 Task Assignment

The key modification lies in the task assignment step of the algorithm. After reachable tasks have been added, the tasks are sorted by a different weighted product than that used by the HWVs. A low weight is still desired; however, the deviation X is defined by $X = |P_{Ti} - 1.0|$, instead of $X = |P_{Ti} - 0.5|$. Hence, a weight of 0 is only achieved when there is 100% probability that a target cell contains a mine. The conflict resolution strategy does require alteration; it relies only on the priority of the tasks to obtain an

optimal solution of task assignment, and the priority that each LWV gives to its reachable tasks is now based on the optimal criteria for identification and neutralization of mines. Additionally, the LWVs are given priority over the HWVs in task assignment. This is because the LWVs are assumed to have more precise equipment for the identification of mines, and thus any information obtained by an HWV would be redundant if the cell were investigated by an LWV. One further note is that while an HWV will avoid visiting a cell previously searched by an LWV, an LWV does not disregard cells that have been previously searched by an HWV. In fact, in many cases, upon classification of a cell, the HWV may raise the probability that the cell contains a mine; thus, an LWV will actually be more likely to search the cell. This is the primary method by which adding more HWVs to a mission can actually improve the task performance of the LWVs.

2.2.2.2 Commission

The commission of the LWVs is identical to the HWVs. However, after a task is completed, the probability that a cell contains a mine is set to 0. This occurs whether or not the identification step determined that the cell contained a mine; the assumption in this study is that the LWVs identify mines with 100% accuracy and that the neutralization equipment has a 100% success rate.

3. EXPERIMENT

3.1 Experiment Format

This experiment has two goals: (1) to recognize and measure the success of both the classification and identification/neutralization tasks and (2) to compare the one-pass and two-pass strategies under constrained time. Toward meeting the first goal, we establish the following test scenario.

An underwater field is initialized to a constant size of 20 x 20 x 20 units, with a mine frequency of 0.05 but initialized probability ranging uniformly from 0 to 1. The probability map does not accurately reflect the mine frequency, but it is convenient to use a binomial probability distribution centered at 0.5 because the trends reflected in the test results are more clearly defined. In the control case, 20 HWVs are used and 20 LWVs are used, with a speed of 4 units per iteration for an HWV and 1 unit per iteration for an LWV. Both classes have D_MAX = 4, and the algorithm is run for 100 iterations, with results are averaged over 5 tests. The factors evaluated are: (1) amount of information gained (measured in terms

of overall confidence of presence or absence of mines in the field), (2) percentage of mines removed in the operation, and (3) rate of identification/neutralization per task by the LWVs. The amount of information gained is measured as a % confidence, which is determined by the sum of the differences of each probability from 0.5 over the entire field, or $C = (\sum f_i) / (0.5*n)$, where $f_i = |p_i - 0.5|$ and n is the total number of cells in the field. The rate of identification/neutralization is found by the number of mines neutralized divided by the total tasks completed by the LWVs.

3.2 Assumptions

Various assumptions are made in the test environment that simplify the experiment and help to isolate the factors being evaluated. The first assumption is perfect collision avoidance; it is assumed that the vehicles have the capability to avoid each other when traveling through the same cell on the way to a search task. The second assumption is perfect communication. This is not a realistic assumption; the effect of imperfect communication on a cooperative search algorithm is significant, and those concerns have been addressed in detail in Chapter 4. However, these effects are not pertinent to the goals of this study; they affect the end performance of the MCM operation, but are not affected by the constructs being varied in this experiment. A third assumption is imperfect classification. An HWV does not assign a probability of 0 or 1 to a cell upon completion of a task; rather, it raises or lowers the probability of a mine in that cell, depending on its determination that the cell does or does not likely contain a mine. On the contrary, for simplicity it is assumed that the identification and neutralization task has a 100% success rate by the LWVs. While this is also an unrealistic assumption, it demonstrates the effects of variations in the experiment clearly, and the overall trends would not change as long as the identification and neutralization stages result in improved information, as is assumed by the Navy UUV Master Plan.

3.3 Executed Tests

The tests run were (1) varying the number of LWVs used, (2) varying the number of HWVs used, (3) varying the effectiveness of the initial detection step to establish initial information, (4) varying the weight given to cells previously visited by an HWV when assigning tasks to LWVs, and (5) comparing the performance of one-pass vs. two-pass strategies. In the first test, the number of HWVs was held constant while the number of LWVs was varied, and the results were evaluated, with particular attention to final confidence in the environment. The second test held the number of LWVs constant while

varying the number of HWVs and especially evaluated the final number of mines remaining in the field as well as the rate of identification/neutralization of the LWVs. The third test drew a greater distinction in the initial probability map between mines and clear space. The fourth test attempted to determine whether it is profitable when assigning the LWVs to give priority to cells already visited by an HWV and determined likely to contain a mine, as opposed to a cell determined likely to contain a mine solely by the detection stage. The final test evaluates the benefit of running the same number of iterations of the algorithm in two passes over one pass, and also compares the results under time constraints, when the number of iterations available to each vehicle class is cut in half.

4. TEST RESULTS

As previously stated, the goals of the experiment were (1) to recognize and measure the success of both the classification and identification/neutralization tasks and (2) to compare the one-pass and two-pass strategies under constrained time. The first goal was met by four tests, as described below. The first two tests, varying the number of vehicles of each class, did demonstrate a trend that increasing the number of vehicles of one class led to improved task performance by the vehicles of the other class. Varying the initial confidence in the detection stage was shown to have a great affect on the identification/neutralization result obtained by the LWVs, but not on the relative ability of the HWVs to perform their tasks. Giving a higher priority to cells already visited by an HWV, however, did not seriously affect the rate of success of the LWVs. The second goal of the experiment was achieved by comparing the performance of the one-pass and two-pass strategies at various iterations of the search. This test showed that the two-pass strategy does offer improved performance over the one-pass strategy; however, because a two-pass search requires twice the time to execute, this improvement is not enough to give a two-pass search a performance advantage under a time constraint.

4.1 Control Condition

The control condition was a test field of 20 x 20 x 20 units, with a team of 20 HWVs and 20 LWVs. The probability that a particular cell contained a mine was 0.05, with uniform probability distribution from the detection stage. A one-pass strategy was used, with the LWVs assigned before the HWVs. Test results were taken after 100 iterations, averaged over 5 unique scenarios. The control condition began with a confidence of 50% and produced a final confidence of 60.76%. The LWVs averaged 1066

tasks, removing 55.4 of 410 mines, or 13.51% of the mines in the field, with a neutralization rate of 0.052 mines removed per task completed.

Table 5.1. Control Condition Test

HWVs	20
LWVs	20
Iterations	100
Mine frequency	0.05
Initial Confidence	50%
Final Confidence	60.76%
Tasks Completed	1066
Percentage of mines removed	13.51%
Neutralization rate	0.052

4.2 Varied Number of LWVs

The number of LWVs was varied from 0 to 40 vehicles. The percentage of mines removed in 100 iterations increased as the number of LWVs increased, from 0% at 0 LWVs to 40.89% at 40 LWVs. As shown in Table 5.2, while the percentage of mines removed increased, the rate of neutralization decreased, from 0.0822 for 2 LWVs to 0.0665 for 40 LWVs. This shows that there are diminishing returns for adding more LWVs to the mission, since identification/neutralization tasks of decreasing mine likelihood must be assigned as more tasks are completed.

Table 5.2. Percent Mines removed and Neutralization Rate as Number of LWVs is Varied

Number of LWVs	% mines removed	Neutralization rate
2	1.56	0.0822
4	1.90	0.0784
8	4.08	0.0801
12	6.34	0.0771
16	9.73	0.0743
20	13.51	0.0756
24	18.89	0.0711
28	22.15	0.0641
32	29.74	0.0628
36	36.15	0.0668
40	40.89	0.0665

The final confidence also generally increased with a greater number of LWVs, as shown in Figure 5.1, ranging from 58.53% at 0 LWVs to 69.01% at 40 LWVs. This demonstrates that increasing the number of LWVs, which are assigned the identification/neutralization task, improves the performance of the classification task performed by the HWVs. This is not surprising, since a visit from an LWV does improve the confidence of an individual cell, since it either identifies and neutralizes a mine or confirms that there is no mine present, in either case giving that cell a probability of 0 and confidence of 100%. However, it is more interesting to note the improvement in performance by the LWVs as the number of HWVs is increased.

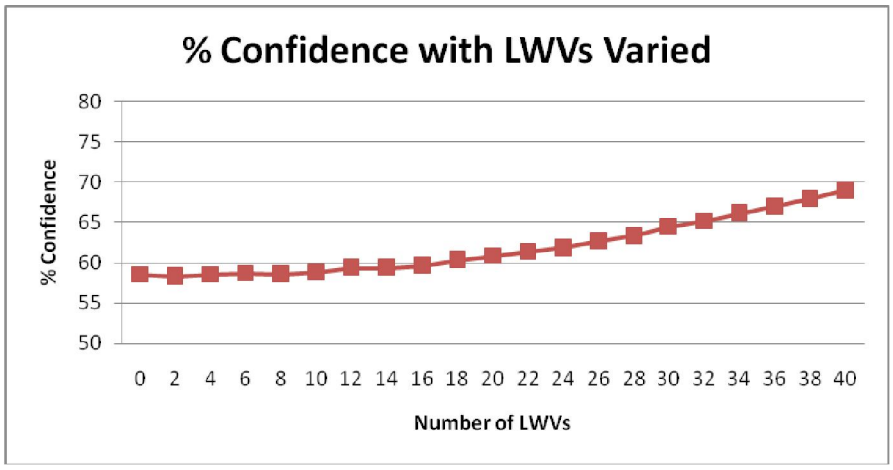

Figure 5.1. Percent confidence as number of LWVs is varied. This shows that besides completing their own task of neutralization, the LWVs are also aiding the HWVs in increasing the information in the field.

4.3 Varied Number of HWVs

As in the previous test, the number of HWVs was varied from 0 to 40. As expected, the final confidence improved as more HWVs (which primarily seek to increasing confidence of information) were added to the experiment, with a final confidence of 52.29% at 0 HWVs and 62.94% at 40 HWVs.

The percentage of mines removed in 100 iterations improved with increasing HWVs, with 14.18% of mines removed at 0 HWVs, ranging to 20.31% of mines removed at 40 HWVs. Interestingly, this improvement was very subtle for fewer than about 30 HWVs. To better establish the trend, the test was extended to 70 HWVs. These results are shown in Figure 5.2. The percentage of mines removed grew faster as the HWVs became more abundant, increasing to 36.88% at 70 HWVs. Specifically, increasing

the number of HWVs from 0 to 30 improved overall mine neutralization by only 21.2%; however, increasing the number of HWVs from 30 to 40 improved neutralization by 18.2%, an increase from 40 to 50 gave a 18.2% improvement, from 50 to 60 improved 29.4%, and increasing from 60 to 70 HWVs gave a 18.7% improvement in mine neutralization.

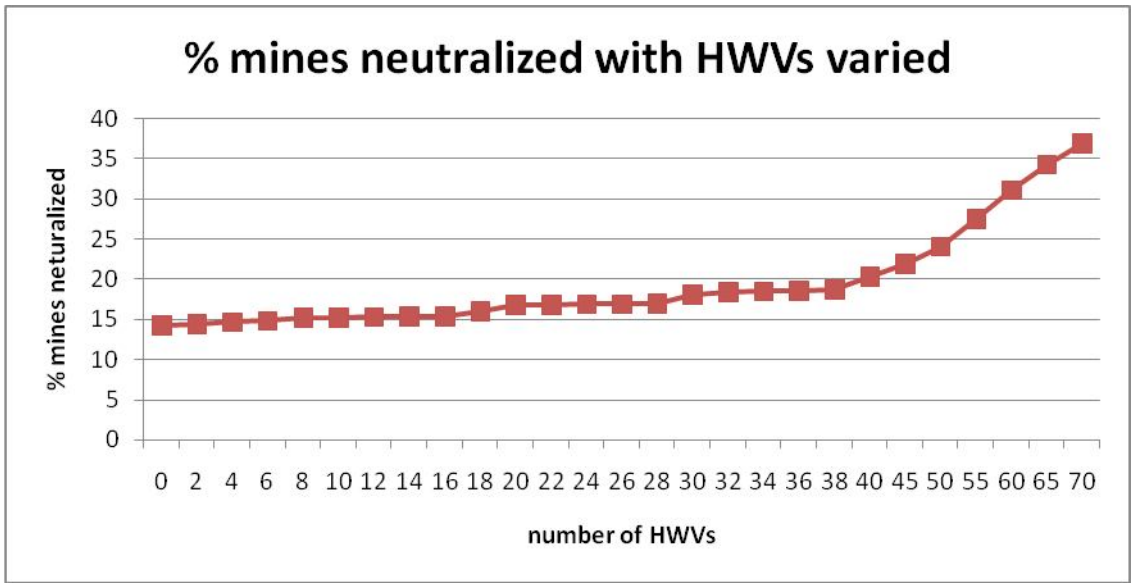

Figure 5.2. Percent mines identified and neutralized as number of HWVs is varied. This shows that there is a trend that increasing the number of HWVs aids the LWVs in their identification/neutralization task; however, the benefit is not significant for few HWVs, during this relatively short test of 100 iterations.

This trend is a result of improved communication of mine presence from the HWVs to the LWVs, as is also evidenced by the vastly improved rate of neutralization by the LWVs with increased HWVs, shown in Table 5.3. With 0 HWVs, the LWVs averaged 1059 tasks completed and neutralized 53 mines, a rate of 0.0500, identical to the mine frequency. At 40 HWVs, this rate improved to 64 mines in 1046 tasks, for 0.0612, and at 70 HWVs the LWVs neutralized 112 mines in 1051 visits for 0.1066. Thus, it is demonstrated that this heuristic does encourage cooperation between vehicle classes by improving the performance of both the classification and the identification/neutralization tasks when vehicles of either class are added to the mission.

Table 5.3. Percentage of mines removed and Neutralization Rate as Number of HWVs is Varied

Number of HWVs	% mines removed	Neutralization rate
0	14.18	0.0500
4	14.60	0.0520
8	15.17	0.0486
12	15.27	0.0519
16	15.32	0.0530
20	16.78	0.0515
24	16.83	0.0552
28	16.94	0.0525
32	18.34	0.0554
36	18.49	0.0630
40	20.31	0.0628
45	21.82	0.0660
50	24.00	0.0736
55	27.48	0.0781
60	31.06	0.0849
65	34.23	0.0888
70	36.88	0.1066

4.4 Varied Initial Confidence

Under the control conditions, the initial confidence of the field was, on average, 50%. This is because the probability produced by the detection stage was uniformly distributed between 0 and 1, producing a binomial distribution of initial confidence over all the tests, centered at 50%. For this test, we increased the initial confidence by assigning cells that contain a mine X times closer to 1 than cells that do not contain a mine. That is, $Pr(b) = 1 – (1 – Pr(a)) / X$, where Pr(a) is the average probability assigned by the detection stage to a cell that does not contain a mine, and Pr(b) is the average probability assigned to a cell that does contain a mine. This test was to determine whether an improvement in the detection stage offers an improvement in the relative performance of the classification and identification/neutralization stages. X was varied by 0.5 from 1 to 10.

The results showed little difference in the classification stage. An initial confidence factor X = 1 (control condition) produced a final confidence of 60.68%, while increasing confidence to X = 10 produced a final confidence of 61.91%. The reason the difference in classification was very slight is that the probability distribution in clear cells after the detection stage was unchanged; the confidence in

cells that contained a mine was greater, but with a mine probability of 0.05, this did not greatly affect the overall confidence of the field after the detection stage.

The identification/neutralization stage, however, showed a vast improvement, as shown in Figure 8. At X = 1, the percentage of mines cleared was 13.23%; however, at X = 10, this percentage skyrocketed to 81.09%. Note that this percentage could approach 100% as we let X increase without bound; at X = 10, there is still a 10% chance that a cell without a mine is given a greater probability from the detection stage than a cell with a mine, and will therefore be visited first by an LWV. This prevents the rate of neutralization from approaching 100%. However, the rate of neutralization is also held lower by assigning a finite D_MAX. Recall that D_MAX is intended to force a vehicle to search more locally, to make efficient use of time. This means that an LWV will sometimes attempt to identify/neutralize a clear cell that is near rather than a cell that actually contains a mine but is far from the LWV. Removing the D_MAX constraint might further improve the rate of neutralization, but may decrease overall performance; this phenomenon is explained more fully in Chapter 4. In this test, we observed that the rate of neutralization increased from 5.02% at X = 1 to 30.86% at X = 10.

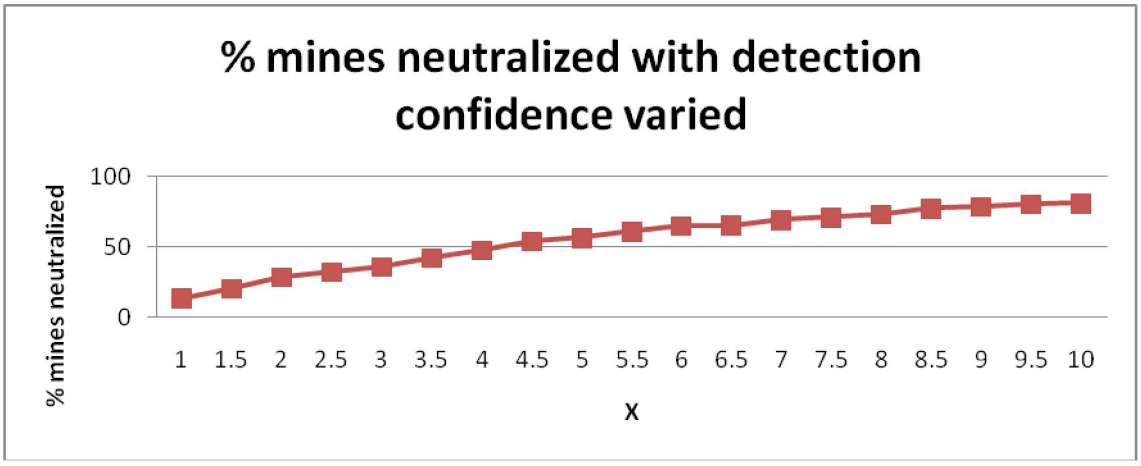

Figure 5.3. Percent of mines identified and neutralized as confidence after detection stage is varied. X denotes a multiplier, where cells that contain a mine result in a probability assigned from the detection stage that is X times closer to 1 than cells that do not contain a mine. Pr(b) = 1 – (1 – Pr(a)) / X, where Pr(a) is the average probability assigned by the detection stage to a cell that does not contain a mine, and Pr(b) is the average probability assigned to a cell that does contain a mine. X is varied from 1 (control condition, least confidence) to 10 (most confidence).

4.5 Varied Weight of Previously Visited Cells

A test was executed to determine whether giving priority to cells already visited by an HWV when assigning an LWV would improve neutralization rate. As mentioned in previous test results, the initial mine probability map obtained from the detection stage is generally less reliable than the subsequent probabilities produced by the classification stage. In particular, when we allow the detection stage to assign higher probabilities to those cells that do, in fact, contain a mine, we expect that the classification stage would help to reject some of the "false positives" that occur in the detection stage. In turn, an LWV that gives higher precedence to a cell already visited by an HWV should perform better because the information that indicates the cell may contain a mine is more reliable. Thus, we would expect a better neutralization rate from an LWV that gives precedence to previously searched cells than from an LWV that gives no such precedence.

The test applied a weight multiplier, varied from 0.1 to 1, to cells that had been visited by an HWV. The expectation was that the performance of the LWV task (mines neutralized and neutralization rate) would increase as the weight multiplier was decreased. The results in fact did not show the trend expected. Over the domain of multipliers, the percentage of mines neutralized ranged from 12.31% to 15.59%, with the control condition (multiplier = 1) neutralizing 13.02% and the most extreme condition (multiplier = 0.1) neutralizing 13.84% of mines in the field. Similarly, final confidence ranged only from 60.36% to 60.94% and showed no distinguishing trend. These results are shown in Table 4 below. The rate of neutralization per task completed also did not significantly change over the course of the test.

Table 5.4. Percentage of mines removed and confidence with previously visited cells weighted.

Weight applied to previously visited cell	% mines removed	% Confidence
1	13.02	60.66
0.9	14.89	60.69
0.8	13.59	60.56
0.7	14.54	60.87
0.6	13.90	60.47
0.5	15.59	60.71
0.4	12.31	60.75
0.3	13.50	60.63
0.2	14.13	60.64
0.1	13.84	60.47

4.6 Comparison of One-Pass and Two-Pass Strategies

To compare performance under the one-pass strategy and the two-pass strategy, we ran two tests, each varying the length of the test from 50 iterations to 400 iterations. The first test used the one-pass strategy, with the LWVs being assigned before the HWVs and performing tasks concurrently. The second test used the two-pass strategy, with HWVs performing their classification task and LWVs beginning the identification/neutralization task after the HWVs have completed a predetermined number of iterations.

As shown in Figure 5.4, the two-pass strategy did outperform the one-pass strategy at each test length. However, the test length of the two-pass strategy is deceptive; the two-pass strategy really requires twice the length of the one-pass strategy, because the LWVs begin their iterations only once the HWVs have finished theirs. Figure 5.5 shows a more accurate representation of performance over time expended in the search. From this figure it is seen that while the two-pass strategy offers a slight performance improvement over the one-pass strategy, under any given time constraint the one-pass strategy actually gains more information and neutralizes more mines. The two-pass strategy is more task-efficient; it has a higher percentage of mines neutralized per iteration and per task completed by the LWVs. However, it is less time-efficient than the one-pass strategy. Therefore, under an unknown time constraint (for which the cooperative algorithm used was optimized), the one-pass strategy is more useful than the two-pass strategy.

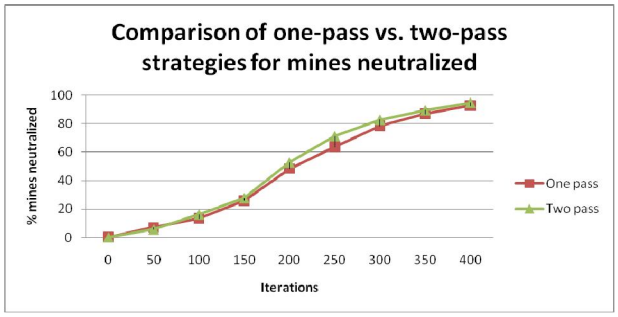

Figure 5.4. Percentage of mines neutralized for one-pass and two-pass strategies. The two-pass strategy does slightly outperform the one-pass strategy, as concluded by the Navy UUV Master Plan.

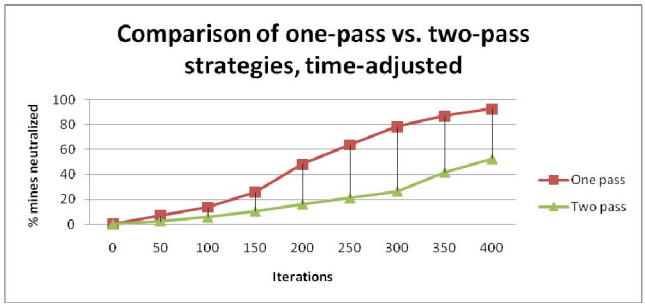

Figure 5.5. Adjusted percentage of mines neutralized for one-pass and two-pass strategies. Adjusted for the amount of search time allotted for each the HWVs and the LWVs in a two-pass strategy. Since the vehicle classes search separately in a two-pass strategy, under constrained time it is obvious that the one-pass strategy obtains better performance.

5. DISCUSSION

5.1 Expected Trends

In general, most of the expected trends in this study were realized in the test results. Increasing the number of LWVs increased the percentage of mines neutralized, while also increasing the final confidence of the information. Increasing the number of HWVs increased the final confidence of the information, while also slightly increasing the mine neutralization rate and percentage of mines neutralized. Raising initial confidence in the detection stage did increase the mine neutralization rate and the percentage of mines neutralized, while affecting the total confidence very little. The two-pass strategy was shown to be more task-efficient, but less time-efficient than the one-pass strategy. The only expected trend that was not realized was that we expected that, in LWV assignment, raising the priority of visiting cells already searched by an HWV would improve LWV performance. On the contrary, LWV mine neutralization remained largely unchanged.

5.2 Effect of Increasing HWVs

One interesting result was that increasing the number of HWVs at first had very little effect on LWV performance. The difference between LWV task performance at 0 HWVs and at 20 HWVs was negligible, with only a very slight improvement occurring at 30 HWVs. However, as the number of HWVs was increased to 40, 50, 60, and 70, the increase in mine neutralization rate grew more

significant. Intuitively, it would seem that adding HWVs would be accompanied by a diminishing return of performance enhancement per HWV; after all, this diminishing return was experienced when adding LWVs while holding the number of HWVs steady. The reason for this result is the increased information gained from the classification task.

The LWVs initially have relatively few good prospects for search; since the results of the detection stage were set to be a uniform probability distribution from 0 to 1, there were few cells that contained a probability very close to 1. An HWV, on discovering a cell that likely contains a mine, adjusts that cell's probability closer to 1, increasing the likelihood that an LWV will search the cell. In only 100 iterations of the search, the information gained by a small team of HWVs does not accumulate quickly enough to have a great effect on LWV performance. As the number of HWVs is increased, however, the LWVs become more likely to have "good" information in the early stages of their search. It is this effect that causes the greater margins in performance as the number of HWVs grows.

5.3 Weighted Previously Searched Tasks

We attempted to improve LWV performance by raising the priority of visiting cells already searched by an HWV in task assignment. This was done under the assumption that information from the HWVs was inherently better than information obtained solely from the detection stage. Indeed, from the test model, this should have been the case. An HWV that discovers a cell likely to contain a mine shifts that cell's probability closer to 1. Therefore, while a cell with a high probability result from the detection stage still only truly has a 0.05 chance of containing a mine, that chance is higher if the probability has been obtained by an HWV in the classification stage. As a result, it would seem that an LWV that searches any cell with a high probability of containing a mine is more likely to identify and neutralize a mine if the cell was previously visited by an HWV. To verify the results found by the experiment, we modified the model to reflect a more realistic map of probabilities obtained by the detection stage; however, there was still no significant difference in LWV performance under task assignment weighted by previous visit from performance under unweighted task assignment. This experiment should be developed further in a future study to determine whether the theory or the experiment was inaccurate.

6. CONCLUSION

The increasing sophistication and complexity of MCM missions require a growing set of diverse and complex tasks. These tasks are best accomplished by a heterogeneous set of unmanned vehicles. An important objective, then, is to take advantage of cooperation among these vehicles, both within a single class and across multiple vehicle classes, to optimize the performance of these tasks. In this study, we have designed a heuristic to obtain the best success for a set of UUVs of two different classes in constrained time. Our experiment analyzed task performance under varied environments according to several metrics that evaluated the level of cooperation among vehicles and success of the overall mission. We demonstrated that increasing vehicles of one class improved the task performance of vehicles of the other class, and that improving the information from one stage of the mission benefited all subsequent stages of the mission. We also tested the conclusion of the Naval UUV Master Plan that making multiple passes over the field with different MCM tasks gave improved performance, and we determined that this assertion does not hold under constrained time.

CHAPTER 6

UAV-INSTIGATED REDISTRIBUTION OF UUVS IN NAVAL MCM RECONNAISSANCE

1. INTRODUCTION

1.1 Aerial/Underwater Unmanned Vehicle Cooperation

This chapter discusses the heterogeneous UAV/UUV cooperative MCM search, in light of such elements as communication relay and redistribution of UUVs. As mentioned in [27] and [56], a recent development in MCM mission planning is to utilize one or more unmanned aerial vehicles (UAV) to enhance connectivity in a UUV search team. Under such a strategy, a UAV visits each UUV or group of UUVs periodically, receiving the most recent sensory data gleaned from the UUV search as well as sharing any new information or commands with the UUVs. This sensory information may then be shared with a mission controller, which maintains a global map of the search field and information specific to each UUV.

1.2 Draper Study

Draper Labs [56] has implemented this model with their Risk-Aware Dynamic Replanning (RMDR) system. In their system, a UAV supports a team of UUVs assigned to separate search regions, while also sometimes performing its own search tasks. The UAV may be assigned generally to offer connectivity within a specific set of regions, or more directly to support particular vehicles in the field.

1.3 Contributions

While Draper has offered a good overview of the heterogeneous UAV/UUV search problem, they have left several issues to be addressed.

1) The regions as presented by the Draper study (shown in Figure 6.1) are defined with very little resolution. Several different types of regions are specified; however, within a single region, there is no precision given to search tasks other than that the UUVs should search the entire region.

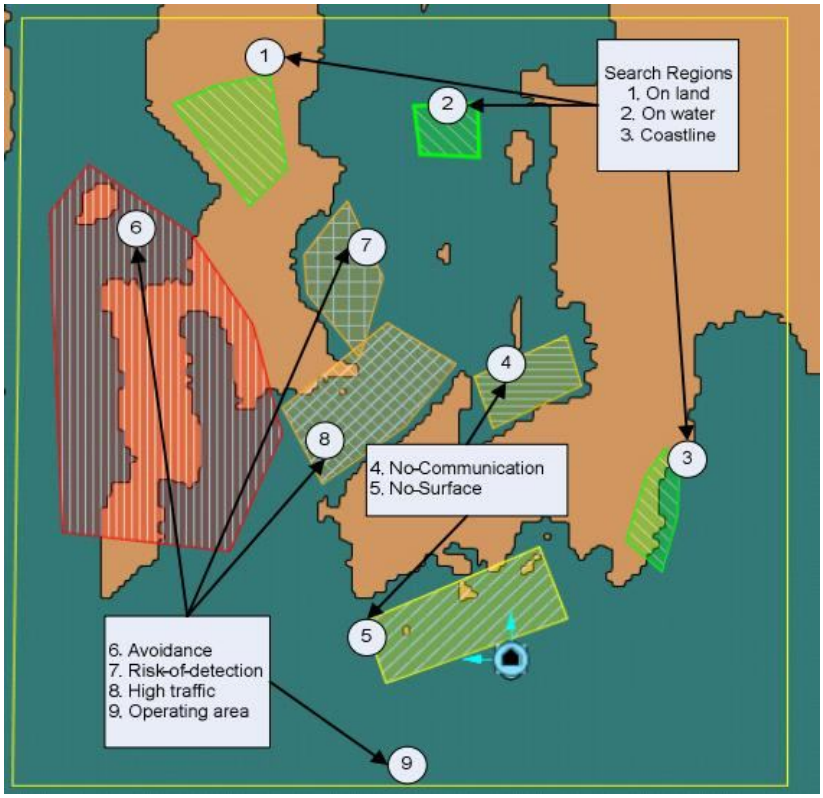

Figure 6.1. Region types as defined by the Draper study (Source: [56])

2) The Draper study offers little flexibility or scalability in UAV support for UUVs. An ideal communication relay path for the UAV is found by brute-force search over the tasks of the UUVs, optimizing the path with mixed-integer formulation. This is a feasible solution to the problem defined by Draper, because it is assumed that the mission is limited to no more than 8 regions. However, a wider search, or a search with more precise definitions of regions, cannot use this method to assign a UAV to communication relay tasks. Additionally, this solution holds only when the tasks of all UUVs are planned deterministically, and the positions of the UUVs throughout the mission are clearly known. Thus, in normal operation the UAV has no influence on the mission of the UUVs, beyond shuttling sensory data.

3) Similarly, there is no provision made for redistribution of the UUVs except in the event of failure or completion of a region search. If a particular region is found to be “interesting” by UUV search, it makes sense from a mission planning perspective to allocate more UUVs to that region.

The contribution of this chapter is to address these shortcomings, attempting to formalize the problem initially presented by the Draper study, as well as adding to the resolution of the search environment. This chapter also adds the element of redistribution of UUVs to different regions in an effort to optimize the information gained in an MCM search mission.

2. FORMALIZATION

2.1 Problem Formalization

As in previous chapters, we define a three-dimensional underwater search field, divided into unit cells. We additionally define the search area in two dimensions, for the purposes of allocating UAVs to rendezvous points with UUV groups. Furthermore, the search field is divided into a number of regions, each comprising its own stand-alone three-dimensional field. These regions are placed at two-dimensional coordinates, indicating their location in the global search area. To simplify our study, we have defined all of the regions to be identical in their spatial dimensions, however this need not be the case and does not influence the effectiveness of the search algorithms.

Each region is defined by a three-dimensional array of mine probability, with each cell being initialized with a value that indicates the probability that the cell contains a mine, as previously established by an initial UAV search. Within each region, a team of UUVs is assigned tasks using the algorithm presented in chapter 3, which optimizes the immediate search task of any UUV, based on the task's distance and "interest" (the amount of doubt as to whether that task will result in the discovery of a mine). The UUVs within a region share information under the periodic full update method described in chapter 4. These UUVs are represented by one UUV which reports sensory information and mission status to the UAV assigned to communication relay (Figure 6.2). The UAV's communication relay mode is similar to the "Support Vehicles" mode defined by Draper, with the qualification that the vehicle set to be supported is comprised of the representative UUV from each search region (Figure 6.3).

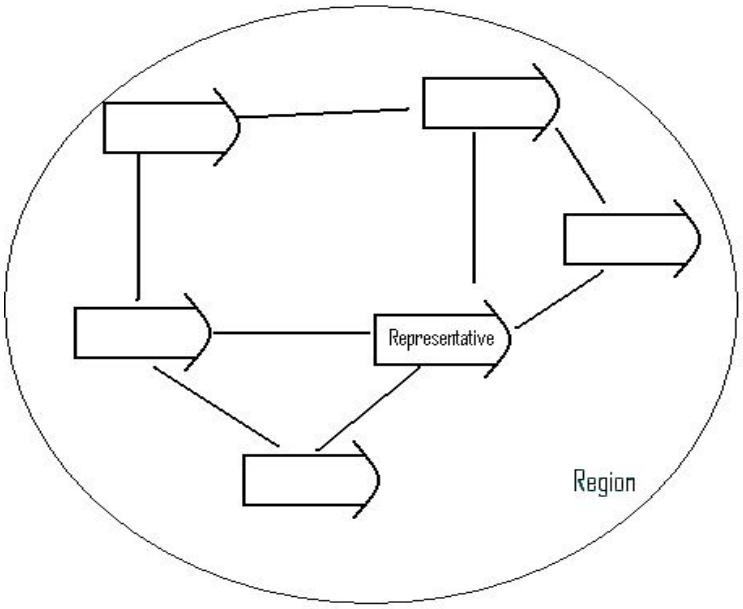

Figure 6.2. Connectivity of UUVs within a region. All of the vehicles are able to communicate with a single representative of the region, which collects sensory data from the other UUVs and transmits it to the UAV upon rendezvous. The representative also transfers any commands from the UAV to the other UUVs in its region.

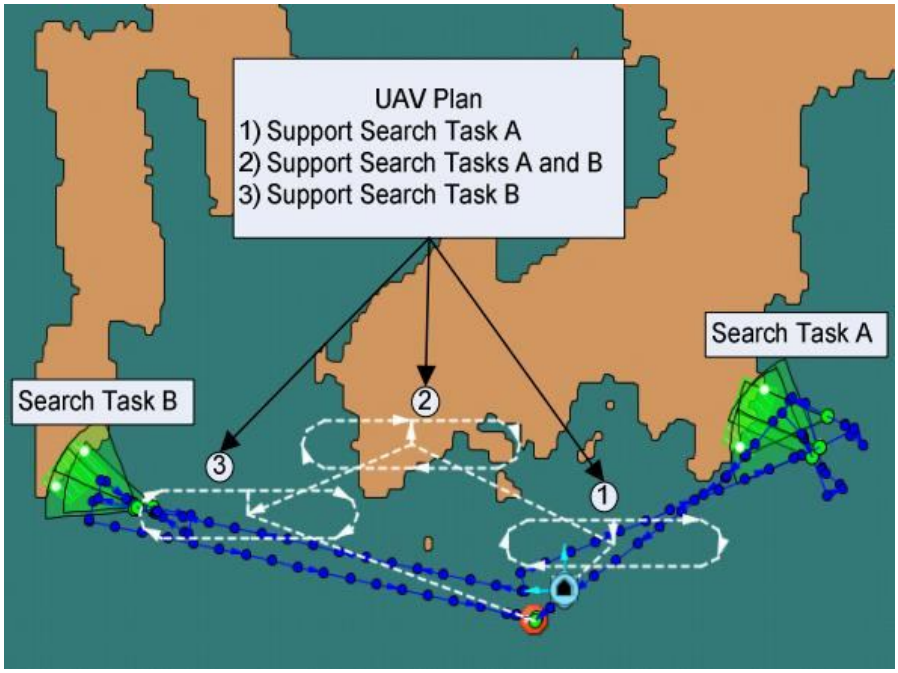

Figure 6.3. Example UAV plan to support vehicles performing search tasks (Source: [56])

2.2 UAV Communication and Redistribution

The UAV's communication relay pattern is not deterministic. Rather, the UAV chooses its next UUV to visit in a manner similar to the UUV's task selection. The UAV maintains both the distance to each UUV and the time since its last visit, and attempts to visit a UUV that is near, and has not been recently visited (Figure 6.4). As the UUVs travel through their search area, the distances between the UUVs that the UAV must support may change, affecting the optimality of a preplanned communication relay pattern.

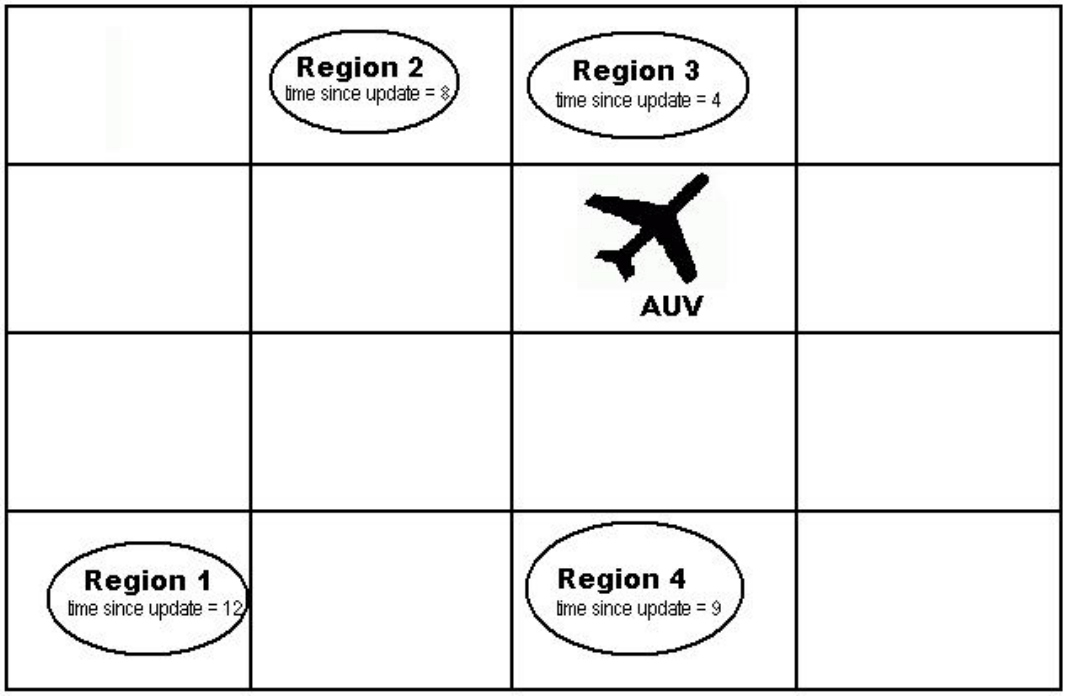

Figure 6.4. Example support decision for a UAV. Region 1 has not been updated for 12 time units, Region 2 for 8 units, Region 3 for 4 units, and Region 4 for 9 units. The weight for region 1 is the distance / time = $\sqrt{(2^2 + 2^2)} / 12 = \sqrt{8} / 12 = 0.2357$. The weight for region 2 is $\sqrt{(1^2 + 1^2)} / 8 = \sqrt{2} / 8 = 0.1768$. The weight for region 3 is $\sqrt{(1^2 + 0^2)} / 4 = 1 / 4 = 0.25$. The weight for region 4 is $\sqrt{(2^2 + 0^2)} / 9 = 2 / 9 = 0.2222$. Thus, the next region chosen for support by the UAV will be region 2. It is neither the closest, nor the least recently updated, but the best combination of the two.

Additionally, the UAV may issue commands to redistribute UUVs based on the results obtained from communication relay. If one region is of particular interest, the UAV will order a UUV from a nearby region to be reassigned to the interesting region. After each information update, the UAV totals the

confidence (as defined in chapter 3) of each region, and allocates UUVs to each region proportional to the inverse of the confidence in that region; that is, the regions of least confidence will receive the most UUVs (Figure 6.5).

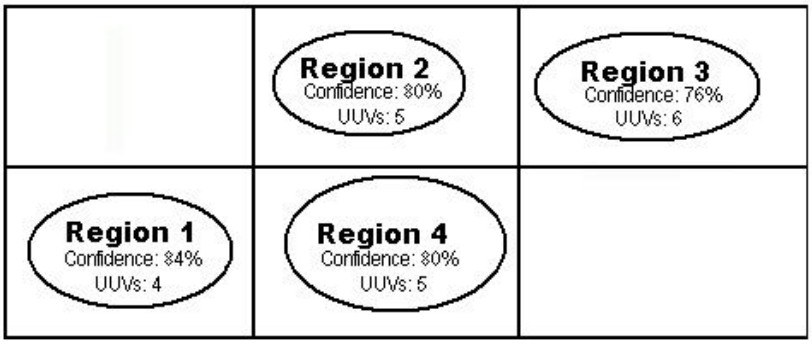

Figure 6.5. Example allocation of UUVs. The total doubt is (1 – 0.84) + (1 – 0.80) + (1 – 0.76) + (1 – 0.80) = 0.8. There are 20 total UUVs to allocate. Region 1 gets (0.16 / 0.8) * 20 = 4 UUVs, Regions 2 and 4 each get (0.20 / 0.8) * 20 = 5 UUVs, and Region 3 gets (0.24 / 0.8) * 20 = 6 UUVs.

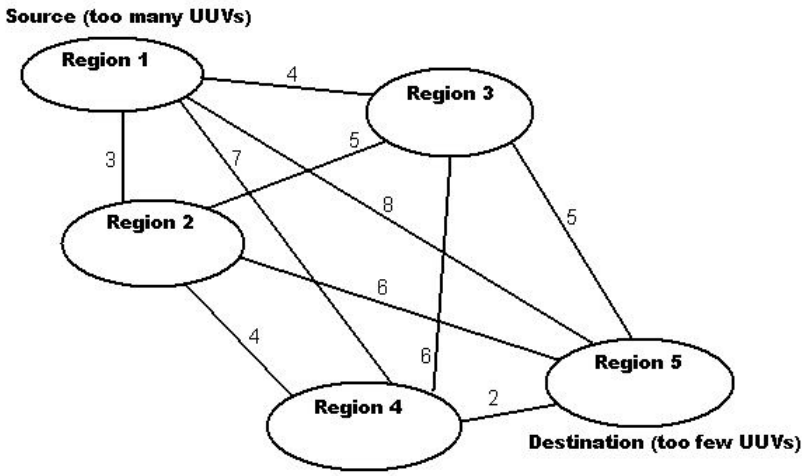

Figure 6.6. Example path planning for allocating UUVs. Region 1 has too many UUVs and Region 5 has too few UUVs. Thus, a UUV must be redistributed from Region 1 to Region 5. It is desired that the time that any UUV is not available for search is minimized; thus, a path is chosen that minimizes the distance of any UUV being allocated from one region to another. The path chosen in this case is 1 → 2 → 4 → 5, where the maximum distance between regions is 4 (between region 2 and 4). So, a UUV from region 1 is sent to region 2, a UUV from region 2 is sent to region 4, and a UUV from region 4 is sent to region 5.

To reassign a UUV to a region of interest, the UAV chooses a "source" region that has more UUVs than the UAV has allocated to it. The UAV then determines a "path" to the destination region, using the shortest possible distances between regions. The UUV being allocated will not necessarily, itself,

traverse the entire distance to the new region; rather, it will travel to the next region on the shortest path, which will at the same time allocate one of its vehicles to the next region, and so on until the destination region has received an extra vehicle (Figure 6.6).

As the UAV visits the representative UUVs of the regions along the path, it issues commands to reallocate vehicles to different regions. It should be noted that reallocation of vehicles can result in several UUVs being unable to participate in search tasks for a short time; therefore, the benefit of reassigning vehicles must be above a certain "threshold" value for redistribution to occur.

The problem of the heterogeneous UAV/UUV search mission is defined as follows. Given a set A of UAVs and a set R of regions, each containing a set V of UUVs, optimize the communication relay pattern of the UAVs and the distribution of UUVs to regions to result in the maximum amount of information gained about the search environment.

If a UAV *a* is idle, it chooses a regional representative to support with its next task. It sorts its set of supported vehicles (defined as supported regions), and chooses a region *r* such that distance(*a*, *r*) / *r*.timeSinceUpdate is minimized. Meanwhile, all UUVs in *r* perform their current task or assign a new task according to the MWA as designed in chapter 3. If the UAV is active, it travels toward the rendezvous point of *r*. When *a* reaches *r*, it updates *r* with any new information not yet shared with *r*, and *r* updates *a* with any new information that a does not yet have. Based on the information provided by *r*, *a* updates the allocation of UUVs, and if *r* has UUVs that need to be reassigned, *a* issues a command to redistribute a vehicle to another region. *r* then chooses a UUV to move to the new region, and issues the command to the UUV. This algorithm is shown in Figure 6.7.

```
A = set of AUVs (AUV[])
F = field of cells (int[][][])
R = set of regions
N = set of UUVs (UUV[])

ASSIGN_TASK_AUV

for each r in R
   for each n in r.N
      ASSIGN_TASK(n);          //task assignment occurs as in the MWA in chapter 3
   r.timeSinceUpdate++;
for each a in A {
   compute_allocations(a.R);  //determines how many UUVs each region should have
   if (distance(a, a.next_region) = 0 and a.state == "active")
      UPDATE_REGION(r);
      r.timeSinceUpdate = 0;
      a.state = "idle";
   else if (a.state == "idle")
      SORT_REGIONS(a);
      a.next_region = a.R[0];
      a.state = "active";

UPDATE_REGION(r)

if (a.needsToGiveVehicle(r))
   commandToGiveVehicle(r, destination)          //commands r to send a vehicle to destination
if (a.needsToAcceptVehicle(r))
   commandToAcceptVehicle(r, source)  //informs r it will receive a new vehicle from source
a.transmit_data(r)                    //give r any new data
a.receive_data(r)                     //accept any new data from r

SORT_REGIONS(a)

for each r in R
   r.weight = distance(a, r) / r.timeSinceUpdate;
sort_by_lowest_weight(R);
```

Figure 6.7. Algorithm to assign UAVs to support region tasks. The algorithm includes calls to the ASSIGN_TASK algorithm developed in chapter 3, and also includes commands from the UAV to various region representative UUVs to redistribute their resources to other regions or to accept newly joining UUVs into their search network.

3. EXPERIMENT

3.1 Parameters

The simulation environment used a fixed set of 10 regions. Each region was assigned a (x, y) coordinate pair, indicating its position in the global search area. Each region had dimensions of 10 x 10 x 10 units, and was initially allocated 5 UUVs. Simulations were run for 1000 iterations. UUVs had a maximum communication distance of 4 units (D_SEE = 4) and a maximum task assignment distance of 2 units (D_MAX = 2). A UUV's maximum speed was 1 unit per iteration within its region, or 1 unit per iteration between regions (a higher speed was assumed while a UUV was not performing a search task), while the UAVs were able to travel 10 units per iteration within a region, or 5 units per iteration between regions, with perfect communication among UAVs. As in previous chapters, after a visit from a UUV, the probability of a cell containing a mine was moved closer to 0 or 1; for cells at probability 0.5 or lower, the probability was squared, while for cells above probability 0.5, the square root of the probability was used, thus modeling a greater benefit from visiting cells of greater doubt. This simulation differed from previous simulations in that visits from multiple vehicles were allowed; multiple visits to a particular cell resulted in increased confidence. Different simulations included the control case (no redistribution of UUVs, no region of particular interest), a "hot" region with more mines than usual (Figure 6.8) but no redistribution of UUVs, a "hot" region with redistribution, an especially "hot" region which increased doubt for the first few UUV visits to each task, a field with lower initial probability of containing mines (clearer water), and varying the number of UAVs.

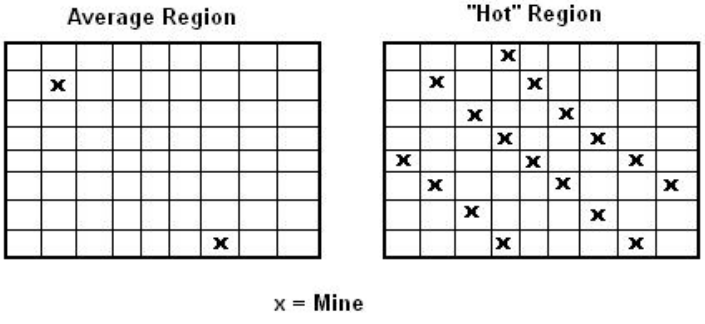

Figure 6.8. Average region compared to "hot" region. The "hot" region contains many more mines, or at least a much greater doubt as to the presence of mines, than the average region. The "hot" region will have a much lower degree of confidence in the information known about it.

3.2 Test Cases

In the control condition, the field parameters of the experiment were as mentioned above. 2 UAVs were used, however they only collected sensory data from the UUVs and did not redistribute them to different regions.

The first test case modeled one region (region 4) as more interesting than the others. Within this region, the probability of a cell was only updated after being visited by more than 5 UUVs. Thus, this region's confidence was not allowed to increase unless more vehicles were reassigned to search it; in the first test case, UUVs were not redistributed, so the confidence of the region was not updated. This case was used to establish a base against which the effectiveness of redistribution was later tested.

The second test case used a "hot" region, just as in the first test case. However, in this case, redistribution of UUVs was allowed. As discussed above, redistribution was only allowed to occur if it resulted in a benefit above a threshold value. Therefore, redistribution occurred only when a UAV determined that the "hot" region required at least 2 more UUVs than other regions, based on the confidence. It was expected that this case would produce a much higher confidence in the "hot" region, with slightly lower confidence in other regions (due to fewer vehicles allocated to search them), resulting in a slight increase in overall confidence.

The third test case used a "very hot" region. In this case, confidence was actually decreased for the first 5 UUV visits to any cell, setting the probability of a mine in that location to 0.5. This case was intended to model a region that was actually more interesting than first established by the initial probability grid, and to compare the effect of redistribution in this "very hot" case to the effect of redistribution in a less extreme case. It was expected that this case would offer a more dramatic improvement in confidence than previous tests when more vehicles were allocated to the region of interest.

The fourth test simulated generally clearer water by lower the overall initial probability of each cell containing a mine. Instead of choosing a random probability between 0 and 1, this random probability was raised to either the power of 2 or 4. The expectation is that this results in a greater difference between searched areas containing a threat and searched areas containing no threat, thereby increasing

the difference in confidence between interesting and uninteresting regions and increasing the benefit of allocating vehicles.

The final test varied the number of UAVs, from a low of 2 UAVs to a maximum of 10 UAVs. It is expected that, since the speed of the UAVs is high relative to that of the UUVs, this variation should only marginally affect the overall performance of the search. However, its primary benefit is in redistribution of vehicles; since the UAVs have perfect communication, increasing the number of UAVs means that redistribution can be done in parallel, instead of a single UAV issuing redistribution commands separately to several regions. As a result, the time required for reallocation should be slightly decreased.

3.3 Evaluation Metric

The primary metric used to evaluate simulation results is described as "confidence." Confidence is defined as the total percentage of the difference of the probability of each cell from 0.5 (worst-case doubt), that is $100 * [\Sigma(p_i - 0.5) / \Sigma(0.5)]$, where p_i is the probability of each cell. For example, a set of cells with probability 0.45, 0.71, and 0.80 has a total confidence of $(|0.45 - 0.5| + |0.71 - 0.5| + |0.80 - 0.5|) / (0.5 + 0.5 + 0.5) = 37.3\%$. The goal of a search mission is not necessarily to discover a maximum number of mines, but to optimally increase the overall confidence in the search area.

The other metric used to evaluate results is the number of vehicle moves. A vehicle move is defined as any number of UUVs being reallocated in such a way as to change the overall distribution of UUVs to all regions. For example, a redistribution that results in 1 UUV being moved from region 7 to region 5 and 1 UUV being moved from region 5 to region 4 is recorded as 1 vehicle move, rather than 2. A higher number of vehicle moves indicates that more redistribution of vehicles is taking place; it is determined from the resulting search confidence whether the redistribution was effective.

4. RESULTS

4.1 Control Condition

The control condition initialized the probability array with a random probability between 0 and 1 assigned to each cell, an average 0.5 probability. Each region contained 10 x 10 x 10 = 1000 cells, and

had 5 UUVs searching the cells, with revisits allowed, for 1000 iterations. As shown in Table 6.1, the initial confidence of all regions was 2494.15 / 5000 = 49.9%, and the final confidence of all regions was 3840.42/5000 = 76.8%. The minimum final confidence for any region was 76.4% and the maximum was 77.2%. Because no redistribution was allowed to occur, the final configuration was still 5 UUVs in each region.

Table 6.1. Initial and final confidence, control condition.

Region	Initial Confidence (%)	Final Confidence (%)
1	49.8	76.7
2	50.0	76.8
3	50.0	76.8
4	50.0	77.1
5	49.7	76.4
6	50.3	77.2
7	49.5	76.5
8	49.6	76.6
9	49.8	76.8
10	50.2	77.1
Total	**49.9**	**76.8**

4.2 "Hot" Region With No Redistribution

In the first test, region 4 was designated a "hot" region, wherein probability of a cell was not updated until more than 5 vehicles visited the cell. Since this test did not involve redistribution of UUVs, region 4 did not have a change in confidence, and the initial and final configuration of the UUVs was once again 5 in each region. The overall initial confidence was 49.9%, with a regional confidence ranging from 48.8% to 51.0%, and the overall final confidence was 74.1%. However, final regional confidence ranged from 76.5% to 77.3%, with the exception of region 4, which remained at 48.8%; the overall confidence of the other regions was 76.9%, a negligible difference from the control condition. Individual region confidences are shown in Table 6.2 below.

Table 6.2. Initial and final confidence by region with a "hot" region but no redistribution.

Region	Initial Confidence (%)	Final Confidence (%)
1	49.9	77.0
2	51.0	77.3
3	49.8	76.6
4	48.8	48.8
5	49.9	77.0
6	49.7	76.7
7	50.4	77.1
8	49.5	76.5
9	50.3	77.2
10	49.4	76.5
Total	**49.9**	**74.1**

4.3 "Hot" Region With Redistribution

The second test allowed redistribution of UUVs to the "hot" region. The redistribution was commanded by 2 UAVs which supported one representative UUV in each region. This resulted in an average of 5.2 moves per simulation, with a sample initial and final configuration shown in Table 6.3. As shown in Table 6.4, initial confidence was was once again 49.9%, with a final confidence of 75.3%, or 1.2% higher than in the second test. The "hot" region (region 4) experienced an increase in confidence from an initial 49.8% to a final 60.1%. This is not as high an increase as experienced by other regions, due to the restriction that the probability of an individual cell was not updated until the cell was visited by more than 5 vehicles. However, it is a much better performance than the case in which no UUVs were reallocated to the "hot" region. Interestingly, excluding region 4, the field had a final confidence of 76.9%, the same as in the control condition and the first test. This means that removing the vehicle resources from uninteresting regions did not significantly harm the performance of the search mission in those areas.

Table 6.3. Sample initial and final vehicle configuration by region with "hot" region.

Region	Initial vehicles	Final vehicles
1	5	5
2	5	5
3	5	5
4	5	8
5	5	3
6	5	5
7	5	5
8	5	4
9	5	5
10	5	5

Table 6.4. Initial and final confidence by region with redistribution to the "hot" region.

Region	Initial Confidence (%)	Final Confidence (%)
1	49.4	76.6
2	50.1	76.7
3	49.7	76.8
4	49.8	60.7
5	50.3	77.0
6	50.3	77.0
7	49.8	76.7
8	50.5	77.6
9	49.7	76.8
10	49.7	77.0
Total	**49.9**	**75.3**

4.4 "Very hot" Region

The third test configured region 4 to be even more dramatically of interest than the first and second tests. This was simulated by setting a searched task's probability to 0.5 until it was searched by more than 5 vehicles. This test simulated a scenario in which the UUVs discovered that the region was more interesting than initially supposed by aerial reconnaissance. In the third test, when no UUVs were reallocated, this in fact resulted in a great drop in confidence, from an initial region 4 confidence of 49.5% to a final confidence of only 10.6%. The overall confidence in the field still rose, from 49.9% to 70.2%, as shown in Table 6.5. Table 6.5 also shows the initial and final confidence of each region when UUVs were redistributed to region 4. The redistribution still resulted in a drop in confidence in region

4, due to the dramatic initial increase in doubt imposed by the simulation, but it was a much more mild fall, from 50.4% to 47.4%; the overall confidence rose from 50.0% to 74.3%, a 4.1% increase over the situation without allocation. This implies that the greater the difference between regions of greater and less interest, the greater the benefit in redistribution of UUVs, as commanded by UAVs. It is also noteworthy that the confidence in region 4 initially fell, but was steadily rising by the conclusion of the test, as shown in Figure 6.9. This test also resulted in more vehicle moves, an average of 17.6 per simulation, and Table 6.6 demonstrates that the final configuration of the vehicles allocated more vehicles to the interesting region than in the second test.

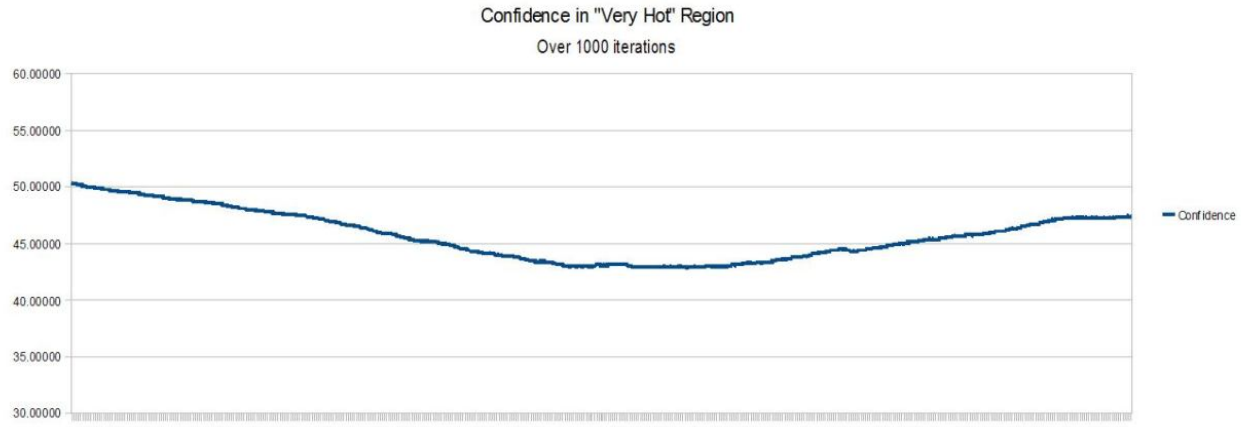

Figure 6.9. Confidence in "Very Hot" region over 1000 iterations. Note that while the overall confidence dropped over the duration of the search, this was caused by the immediate drop in confidence when any cell in the "very hot" region was initially searched. In fact, the confidence was steadily climbing at the completion of the search, and this trend offers significant benefit over the case where no extra UUVs were allocated to the "very hot" region.

Table 6.5. Initial and final confidence by region with "very hot" region.

Region	Initial Confidence % (no redistribution)	Final Confidence % (no redistribution)	Initial Confidence % (with redistribution)	Final Confidence % (with redistribution)
1	49.7	76.7	49.6	77.2
2	50.4	77.0	49.9	77.8
3	50.6	77.2	49.6	76.8
4	49.5	10.6	50.4	47.4
5	49.4	76.5	50.0	77.3
6	50.7	77.3	50.7	77.6
7	49.5	76.6	50.3	77.9
8	49.8	76.7	49.7	77.1
9	50.1	76.9	50.1	76.8
10	49.4	76.4	49.4	77.1
Total	**49.9**	**70.2**	**50.0**	**74.3**

Table 6.6. Sample initial and final vehicle configuration by region with "very hot" region.

Region	Initial Vehicles	Final Vehicles
1	5	5
2	5	4
3	5	5
4	5	10
5	5	5
6	5	4
7	5	4
8	5	5
9	5	4
10	5	4

4.5 Lower Initial Probability

To simulate clearer waters, the fourth test raised the initial probability of each cell to an exponent: 1 (the control case), 2, or 4. As in the second test, region 4 was designated a "hot" region, in which the probability of a cell containing a mine was not updated until the cell was searched by more than 5 vehicles. As the difference between simulations allocating vehicles as opposed to not allocating vehicles had been thoroughly visited in the previous three tests, the fourth test compared the performance of the search with only the initial probability varied.

Table 6.7 shows the difference in vehicle moves between the three simulations in this test. The number of moves increased as the distinction between clear water and interesting areas was increased by lowering the initial probability of each cell. However, it did not seem that more vehicles were necessarily allocated to region 4 as the initial probability was lowered. Table 6.8 shows a final configuration from each simulation, with the initial configuration allocating 5 vehicles to each region.

Table 6.7. Average vehicle moves for adjusted initial probability.

Initial Probability Exponent	Vehicle moves
1	5.2
2	9.0
4	29.2

Table 6.8. Sample final vehicle configuration by region for adjusted initial probability.

Region	Final vehicles (exp = 1)	Final vehicles (exp = 2)	Final vehicles (exp = 4)
1	5	5	5
2	5	5	5
3	5	5	4
4	8	8	8
5	3	4	5
6	5	5	5
7	5	4	5
8	4	5	5
9	5	4	5
10	5	5	4

The initial and final confidence for each simulation are shown in Table 6.9. As expected, the increase in confidence in the "hot" region was greater (13.3%) in the case when initial probability was squared than in the case (10.9%) when initial probability was not modified. In the case when initial probability was raised to the power 4, the increase in confidence in the "hot" region was slightly less (9.8%), but this may be attributed to the fact that confidence was nearing 100% for the field, and completed tasks were not resulting in as dramatic a confidence increase. Indeed, when viewed in terms of "reduced doubt," the case with power = 1 reduced doubt from 50.2% to 39.3% (a 21.7% reduction), while the case with power = 2 reduced doubt from 39.4% to 26.1% (reduced by 33.8%) and the case with power = 4 reduced doubt from 24.6% to 14.8% (a 39.8% reduction). Thus, the contention of the fourth test, that increasing the distinction between clear and ambiguous areas increases the benefit of redistribution, holds from the simulation results.

Table 6.9. Initial and final confidence by region with adjusted initial probability.

Region	Initial Confidence % (exp = 1)	Final Confidence % (exp = 1)	Initial Confidence % (exp = 2)	Final Confidence % (exp = 2)	Initial Confidence % (exp = 4)	Final Confidence % (exp = 4)
1	49.4	76.6	61.5	84.3	75.1	90.5
2	50.1	76.7	60.5	83.8	75.1	90.9
3	49.7	76.8	62.1	85.0	75.3	91.8
4	49.8	60.7	60.6	73.9	75.4	85.2
5	50.3	77.0	61.3	83.8	74.3	90.2
6	50.3	77.0	61.4	84.1	75.2	91.4
7	49.8	76.7	61.0	84.7	75.2	90.5
8	50.5	77.6	60.6	83.9	75.5	90.7
9	49.7	76.8	62.1	84.8	74.6	90.6
10	49.7	77.0	61.8	84.8	75.2	90.8
Total	**49.9**	**75.3**	**61.3**	**83.3**	**75.1**	**90.3**

4.6 Varied Number of UAVs

The final test varied the number of UAVs involved in communication relay and redistribution command. This was done to determine whether increasing the number of UAVs would increase communication sufficiently to improve the benefit of redistribution of UUVs. As in some previous tests, region 4 was designated a "hot" region, in which probability was not updated until after more than 5 unique UUV visits. The number of UAVs was varied by 2 between 2 and 10. Table 6.10 shows the initial and final confidence of each simulation. It can be seen from these results that the number of UAVs does not greatly affect the connectivity. This is probably due to the relative high speed of the UAVs as compared to the speed of the UUVs; any delay in a redistribution assignment due to the flight patterns of a supporting UAV is not significant compared to the time required for the UUV to cross to another region. It could be argued that more UAVs may be beneficial if the regions are further from each other; however, this case would also require a higher relocation penalty for UUVs to remain efficient, likely resulting in infrequent redistribution anyway.

Table 6.10. Initial and final confidence by region with number of UAVs varied

Region	Initial Conf % 2 UAV	Final Conf % 2 UAV	Initial Conf % 4 UAV	Final Conf % 4 UAV	Initial Conf % 6 UAV	Final Conf % 6 UAV	Initial Conf % 8 UAV	Final Conf % 8 UAV	Initial Conf % 10 UAV	Final Conf % 10 UAV
1	49.4	76.6	50.2	76.7	50.1	76.9	50.2	77.2	49.6	76.5
2	50.1	76.7	49.3	76.3	50.4	77.2	50.7	77.7	49.7	76.4
3	49.7	76.8	50.2	77.3	50.1	77.0	50.0	76.8	50.4	77.0
4	49.8	60.7	49.7	59.3	49.7	60.6	50.0	60.5	49.8	59.8
5	50.3	77.0	49.1	76.2	50.4	77.0	50.3	77.0	49.9	76.8
6	50.3	77.0	49.8	76.6	50.3	76.7	50.8	77.3	50.1	76.9
7	49.8	76.7	49.9	76.6	50.7	77.7	50.2	77.0	50.1	76.9
8	50.5	77.6	50.0	76.8	49.5	76.8	50.3	77.7	50.6	77.6
9	49.7	76.8	50.0	76.8	50.1	77.0	49.2	76.5	49.8	76.6
10	49.7	77.0	49.5	76.9	50.8	77.8	50.3	77.0	50.3	77.0
Total	**49.9**	**75.3**	**49.8**	**75.0**	**50.2**	**75.5**	**50.2**	**75.5**	**50.0**	**75.2**

5. CONCLUSION

This chapter has presented a strategy for redistribution of UUVs by UAV command, based on search data. This allows UUVs to be reallocated from a less interesting region to a more interesting region of the search area, effecting a more efficient use of search resources, and improving the effectiveness of an MCM mission. Tests were conducted, which demonstrated the benefit of redistribution of UUVs in the presence of an unevenly distributed threat. It was determined that the greater the distinction between a more and less interesting search region, the greater the benefit of redistribution of UUVs. The number of UAVs used to enhance network connectivity did not perceptibly affect the test results; this was attributed to the high speed of the UAVs relative to the UUVs. In general, the redistribution algorithm proved to be very effective in increasing the end confidence in the presence or absence of mines in the search field at termination of the search.

CHAPTER 7

FUTURE WORK

The future of UUVs in mine counter-measures operations offers great research potential. The U.S. Navy Master Plan for UUVs [31] has an overall goal to "Deliver UUV Capability…and Begin Using It!" To meet these goals, there are several areas that warrant further progress from the research conducted in this work. These prospects are divided into two fundamental categories: immediate and long-term goals. Immediate goals are those that merely introduce an additional element into one or more of the studies previously presented, and would constitute similar but more realistic or comprehensive works. Long-term goals would likely require a re-definition of the MCM mission as it has been presented, and while many of the concepts presented in earlier work can serve as building blocks to construct a context for these future works, they each warrant their own separate study.

1. Immediate Goals

1.1 Turning Penalty

The algorithm presented in Chapter 3 (and used in later chapters) oversimplifies the movement of UUVs. Specifically, turning is a significant problem in UUV navigation; UUVs typically have a very large turn radius. Vertical motion is also non-trivial for a UUV; UUVs cannot generally dive directly downward. Furthermore, underwater currents can make some turns impractical or even impossible for a UUV. Therefore, the algorithm as initially proposed is flawed in that it weights the distances of each cell in the UUV's acceptable task range equally. The cell directly behind a UUV is certainly not as quickly or easily reached as the cell directly in front of it, or even a cell distantly in front of it. Therefore, a more accurate version of this algorithm would account for the turning difficulty by applying a turn penalty to cells that do not lie in the direction of the vehicle's motion. Additionally, the potential tasks likely would not include those behind the vehicle, or those requiring an infeasible turn motion to reach. More likely, the candidate tasks available to a UUV would be those easily reachable, given its position, orientation and speed, as well as any environmental factors such as current. Such an effect would require only a slight modification of the algorithm, and study could be conducted as in Chapter 3.

1.2 Information Priority

As mentioned in related work, as well as in Chapter 4, bandwidth is very limited on an acoustic channel. Chapter 4 discussed several possible strategies for limiting the amount of information communicated between UUVs. An additional factor in information transfer could be information priority. Information might be shared between vehicles only if it meets a particular standard, according to the mission objective. For instance, if the emphasis of a mission is to identify all mines, it may be desirable for multiple vehicles to visit a suspicious location, but not a location that has been deemed safe. Therefore, a UUV may share information regarding the completion of tasks that did not result in the discovery of a mine, but might withhold information when its search has ambiguous results, causing other UUVs to visit that location. This could be studied as an supplement to the work presented in Chapter 4. There are several reasons why this might be a desirable property of a mission. A mission supervisor could use information from several vehicles regarding a particular location to determine whether the location contains a mine. This information could also be used as a "sanity check" to establish whether one of the UUVs has faulty sensory equipment; under a uniform information sharing strategy, one vehicle may incorrectly communicate to other UUVs that it has discovered a mine at a particular location, preventing them from searching that location and making an accurate assessment. Finally, vehicles could take advantage of knowledge of their environment to increase speed when traveling through ostensibly clear water, slowing down through unknown waters; it would not be necessary to communicate a mine's presence to other UUVs to force them to slow through an otherwise unknown area, and sharing this information would be unnecessary. Thus, prioritizing certain types of information to share between vehicles could reduce unnecessary communication while improving performance in an appropriate mission and context.

1.3 Acoustic Transmitters

One technique used to increase the efficiency of UUV travel is to deploy a small acoustic transmitter near a point of interest on the ocean floor. The transmitter emits a weak signal, intended to direct a second vehicle toward that location for further investigation. To improve network communication, such a device might be exchanged for one that is able to transmit more information on request by a visiting UUV. This equipment could prove particularly useful in a heterogeneous UUV search, such as that

studied in Chapter 5. Suppose a team of HWVs, initially scanning the search area for mines, drops small acoustic transmitters when tasks are completed that result in a cell likely containing a mine. These transmitters could be loaded with the locations of any other known transmitters that have been deployed by any of the HWVs. In this case, an LWV could travel directly from one transmitter to another, quickly neutralizing mines without requiring the completion of the HWV's detection portion of the MCM operation. The transmitter could also be loaded with information regarding other cells in its area. A UUV that enters the region without encountering other vehicles that have previously searched it would then be able to obtain this information from the transmitter, thus decreasing search redundancy in the mission.

2. Long-Term Goals

2.1 Localized Control

A search field with high task resolution may be prohibitive for individual vehicles to store a master probability map. A possible solution to this would be to use a representative UUV (presented in Chapter 6) as a local controller within its search region. This UUV would have a larger memory store and higher processing capabilities than the other vehicles, and would maintain the probability map for the region. It would be responsible for task assignment, as well as reporting progress to a UAV or to a mission host with which it has connectivity. This approach could render a large field more manageable to search and update; however, it certainly comes with its own complications. Local networks, in particular, need to be carefully configured, such that no vehicle can simultaneously be a part of two networks. Otherwise, the vehicle may receive multiple conflicting assignments from different controllers. A vehicle will additionally need to periodically return to within range of the controller to receive its next task, or a short series of tasks. As noted in Chapter 6, a controller may, on occasion, need to reassign a vehicle to a different region, and must therefore also maintain knowledge of the locations of other controllers. Modeling this mission configuration requires further definition of the tasks and behavior of the local controller, as well as adjustments to the vehicles, which would no longer have the responsibility to select their own tasks. Further study would need to be conducted to determine the amount of bandwidth and power that could be saved by this approach, as well as the effect on mission performance.

2.2 Pattern Recognition

While these studies have generally assumed that mines may be placed anywhere in an underwater search field, this is not practically true. The natural methods of mine placement tend to result in patterns of mine locations, particularly in straight lines, separated at regular intervals. Many of these patterns have been documented, and future UUV missions should seek to take advantage of this information to more efficiently identify locations that are likely to contain a mine.

In light of this, one possible improvement would be to use preprocessing to perform pattern analysis on the initial probability map. When a high mine probability in several locations in the initial map closely resembles a known mine placement pattern, the pattern should be projected onto the probability map; that is, given the pattern, those locations expected to contain a mine should receive an initial increase in the probability known to the UUVs. This projection should also allow for some tolerance in the exact locations of the mines; a "rubber" mat approach could be used, wherein those locations most closely matching the projected pattern receive the greatest increase in probability, while locations lying near the projected pattern receive a smaller probability increase. This helps to account for the drifting of certain hazards due to natural ocean currents or the momentum of the vessel deploying the mines.

The UUVs would also be responsible to modify the probability of locations in the vicinity of completed tasks. Under this tactic, the discovery of a mine at a particular location, for example (3, 3, 3), might result in the projection of a mine at such adjacent locations as (4, 3, 3) or (3, 2, 3). In addition, if multiple mines are discovered that match a known pattern, for instance at (1, 3, 3), (3, 3, 3), and (5, 3, 3), other locations that match further locations of mines according to that pattern (such as (7, 3, 3) and (9, 3, 3)) would receive an increase in probability. This would require modification to the algorithm, which currently does not change the probability of mines in other cells based on the completion of a single task.

The studies presented here did not generally specify initial locations for mines, with the exception of the simulations related to chaos and mine score presented in Chapter 3, as well as the neutralization of mines by the LWVs in Chapter 5. The test environment needs to be modified to initialize probability maps with patterns of mines that match, or nearly match, known common patterns of underwater mine placement. The initial patterns should also account for the effect of underwater currents on mine

location, given the propensity of unmoored mines to drift in a strong current. The changes to both the algorithm and the test environment would require a separate future study.

2.3 Mobile Threats and Anti-Submarine Warfare

One motivation for the development of UUV technology is the ability of unmanned vehicles to be programmed to counter asymmetric threats. The studies presented here have hitherto been concerned with mine counter-measures. However, another major use of UUVs in naval military operations is in anti-submarine warfare (ASW). UUVs are coveted for ASW operations due to their low profile and their ability to enter shallow waters where larger naval vessels cannot. ASW operations differ from mine counter-measures in that the hazards encountered by UUVs are mobile. This has many implications for UUV search patterns.

First, a UUV search is never completed. A previously searched location may later be occupied by an incoming mobile threat. Thus, control algorithms that optimize UUV movement to quickly perform a comprehensive MCM search of an environment are unlikely to translate sufficiently to ASW reconnaissance. In this regard, the time-constrained MCM search applies more readily to ASW than the full search.

The data reported by a UUV would now also need to determine whether a threat is mobile or immobile. If mobile, a UUV would need to update a probability map regularly to reflect the projected movement of the potential hazard, sharing the object's expected path with other UUVs in the object's vicinity.

This also implies that the probability map that is the basis for UUV movement is likely to change over time, independent of completed UUV tasks. In the context of ASW, there may be other reasons besides known mobile threats to dynamically update a probability map. For example, enemy units may be known to travel “in formation,” which would relate the ASW reconnaissance problem to the aforementioned pattern matching problem in MCM. Another factor may be the strategic value of particular locations at particular times; for instance, due to an enemy's desire for clandestine operations, the likelihood of an incoming threat may increase in poor lighting conditions.

The weight of a task in allocation may depend on factors besides distance and probability. All other factors held equal, it is certainly desirable when searching for mobile threats to visit locations that have

not recently been visited. Thus, an additional factor in the weight of a task may be the time since it was last known to be visited by a UUV.

Additionally, flexibility is paramount in the configuration of an ASW UUV network. Besides sensor failure or loss of network connectivity, UUVs may also need to be rearranged if a vehicle is given a new mission directive. In ASW, it may become necessary for a UUV, upon detection of a mobile threat, to monitor, or "trail," the threat, to determine its potential hazard level. Under such a condition, another UUV may need to be allocated to that vehicle's region to fulfill its search duties; alternately, additional resources might be allocated to tracking the threat, or engaging non-lethal or lethal weaponry.

Developing search algorithms for ASW would require major modification of the MCM search algorithms presented in this study. In particular, the complexity of the weight function would be increased, with the likely addition of the time passed since the last visit to the cell. The probability map would also be updated much more frequently, to reflect the movement of threats, as well as the strategic value of the location from an enemy's perspective.

Inter-vehicle communication may be decreased in the presence of hostile vessels with communication capabilities, replaced by communication either with "waypoints" along the border being searched by UUVs or by communication with a hovering UAV. Threshold levels of hazard identification, at which non-lethal or lethal weaponry should be deployed, would also need to be researched. While some of the model of UUV used in the studies here can also be applied to ASW, particularly the hybrid UUV/UAV mission model, there are still many elements unique to an ASW mission which warrant extensive further studies in this field.

REFERENCES

[1] Akyildiz, I.F., Pompili, D. and Melodia, T., 2007. State of the Art in Protocol Research for Underwater Acoustic Sensor Networks. In ACM SIGMOBILE Mobile Computing and Communications Review, Vol. 11, Issue 4, 11-22.

[2] Akyildiz, I.F., Pompili, D. and Melodia, T. 2004. Challenges for Efficient Communication in Underwater Acoustic Sensor Networks. ACM SIGBED Review, Jul. 2004. Vol. 1, no. 2.

[3] Baras, J.S., Tan, X. and Hovareshti, P., 2003. Decentralized Control of Autonomous Vehicles. In Proceedings of the 42nd IEEE Conference on Decision and Control (Maui, HI, Dec. 9-12, 2003), Vol. 2, 1532-1537.

[4] Benton, C., Kenney, J., Nitzel, R., Blidberg, R., Chappell, S., Mupparapu, S. Autonomous Undersea Systems Network (AUSNet) – Protocols to Support Ad-hoc AUV Communications. Autonomous Underwater Vehicles 2004 (June 17-18 2004). IEEE Press, Piscataway, NJ, 83-87.

[5] Caliskan, A. 2000. Multiuser Acoustic Wireless Communications for Underwater Vehicles. Thesis. Pennsylvania State University.

[6] Chandler, P., Pachter, M., Rasmussen, S.J. and Schumacher, C. Multiple Task Assignment for a UAV Team. 2003. AIAA Guidance, Navigation, and Control Conference and Exhibit (Monterey, CA, Aug. 2002), 2002-4587.

[7] Chandrasekhar, V., Seah, W., Choo. Y., and Ee., H., 2006. Localization in Underwater Sensor Networks – Survey and Challenges. In Proceedings of the 1st ACM International Workshop on Underwater Networks (Los Angeles, CA, Sept. 25, 2006), 33-40.

[8] Chung, T., Cremean, L., Dunbar, Z., Jin, E., Klavins, D., Moore, A., Tiwari, D. and van Gogh, S. 2002. A Platform for Cooperative and Coordinated Control of Multiple Vehicles. Technical Report. California Institute of Technology.

[9] Correll, N. and Martinoli, A. 2004. Collective Inspection of Regular Structures Using a Swarm of Miniature Robots. In Proceedings of the Ninth International Symposium on Experimental Robots ISER-04 (Singapore, June 18 - 21, 2004). Springer Tracts in Advanced Robotics, New York, NY, 375-385.

[10] De Barros, E.A., Maeda, H., Miyajima, S. and Yamato, H., 1993. Local Path Planning for an AUV Using a Cooperative Control System. In Proceedings of OCEANS '93: Engineering in Harmony with Ocean (Victoria, BC, Oct. 18-21, 1993). IEEE Press, Piscataway, NJ, 175-180.

[11] De Sousa, J.B. and Gollu, A., 1997. A Simulation Environment for the Coordinated Operation of Multiple Autonomous Underwater Vehicles. In Proceedings of the 1997 Winter Simulation Conference (Atlanta, GA, Dec. 7-10, 1997), 1169-1175.

[12] Deshpande, P.D., Sangekar, M.N., Kalyan, B., Chitre, M., Shahabudeen, S., Pallayil, V., and Beng, K.T. Design and Development of AUVs for Cooperative Missions. In Proceedings of OCEANS 2006: The Next Frontier (Singapore, May 16-19, 2006). IEEE Press, Piscataway, NJ.

[13] Dias, M.B., Zlot, R.M., Kalra, N. and Stentz, A. 2005. Market-Based Multirobot Coordination: a Survey and Analysis. Technical Report. Carnegie Mellon University.

[14] El-Abd, M. and Kamel, M., 2005. Factors Governing the Behavior of Multiple Cooperating Swarms. In Proceedings of the 2005 Conference on Genetic and Evolutionary Computation(Washington,

D.C., June 25-29, 2005), 269-270.

[15] Gage, D. 1995. Many-Robot MCM Search Systems. In Proceedings of the Autonomous Vehicles in Mine Countermeasures Symposium (Monterey, CA, April 4 - 7 1995). SSC San Diego, San Diego, CA, 9.56-9.64.

[16] Gerla, M. and Yi, Y., 2004. Team Communications among Autonomous Sensor Swarms. ACM SIGMOD Record, Vol. 33, Issue 1, 20-25.

[17] Gooding, T. 2001. A Framework for Evaluating Advanced Search Concepts for Multiple Autonomous Underwater Vehicle (AUV) Mine Countermeasures (MCM). Thesis. Massachussetts Institute of Technology.

[18] Gupta, S., Ray, A. and Phoha, S. 2009. Generalized Ising Model for Dynamic Adaptation in Autonomous Systems. Europhysics Letters, July 2009.

[19] Haag, M., Agu, E., Komerska, R., Chappell, S. and Bartos, R., 2006. Status Packet Deprecation and Store-Forward Routing in AUSNet. In Proceedings of the 1st ACM International Workshop on Underwater Networks (Los Angeles, CA, Sept. 25, 2006), 86-92.

[20] Hall, M.R. and Benokraitis, V.J., 1988. A Mission Planning Architecture for an Autonomous Vehicle. In Proceedings of the 1st International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems (Tullahoma, TN, 1988), Vol. 1, 582-589.

[21] Hayes, A.T. 2002. How Many Robots? Group Size and Efficiency in Collective Search Tasks. In Proceedings of the 6th International Symposium on Distributed Autonomous Robotic Systems (Fukuoka, Japan, June 25 - 27, 2002). Springer Verlag, Heidelberg, Germany, 289-298.

[22] Hert, S., Tiwari, S. and Lumelsky, V. 1996. A Terrain-Covering Algorithm for an AUV. Autonomous Robots, 1996. Vol. 3, no. 2, 91-119.

[23] Hsieh, C.H., Nguyen, B.Q., Tung, D.J., Chuang, Y., Marthaler, D., Jin, Z., Shi, L., Bertozzi, A., and Murray, R.M. 2004. Joint UCLA-CalTech Report on the Implementation of Multi-Agent Cooperative Control Algorithms. Summer Project Report. California Institute of Technology.

[24] Jadbabaie, J. and Morse, A. 2003. Coordination of Groups of Mobile Autonomous Agents Using Nearest Neighbor Rules. IEEE Transactions on Automatic Control, 2003. Vol. 48, no. 6, 988-1001.

[25] Jin, Y., Minaia, A.A. and Polycarpou, M.M. 2003. Cooperative Real-time Search and Task Allocation in UAV Teams. IEEE Conference on Decision and Control (Maui, HI, Dec. 2003), Vol. 1, 7-12.

[26] Kamrani, F. and Ayani, R., 2007. Simulation-Aided Path Planning of AUV. In Proceedings of the 2007 Winter Simulation Conference (Washington, D.C., Dec. 9-12, 2007), 1306-1314.

[27] Karim, S. and Heinze, C., 2005. Experiences With the Design and Implementation of an Agent-Based Autonomous AUV Controller. In Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems (Netherlands, July 25-29, 2005), 19-26.

[28] Lemarie, T., Alami, R. and Lacroix, S. 2004. A Distributed Task Allocation Scheme in Multi-UAV Context. IEEE International Conference on Robotics and Automation, 2004.

[29] Marques, E.R.B., Pinto, J., Kragelund, S., Dias, P., Madrueira, L., Sousa, A., Correia, M., Ferrera,

H., Goncalves, R., Martins, R., Horner, D., Healey, A.J., Goncalves, G.M., and Sousa, J.B. AUV Control and Communication Using Underwater Acoustic Networks. In Proceedings of OCEANS 2007 (Aberdeen, Scotland, June 18-21 2007). IEEE Press, Piscataway, NJ.

[30] Mataric, M.J., Sukhatme, G.S. and Stergaard, E.H. 2003. Multi-robot Task Allocation in Uncertain Environments. Autonomous Robots 2003, Vol. 14, 255-263.

[31] Mukherjee, K., Gupta, S., Rao, C., Ray, A. and Phoha, S. 2009. Symbolic Pattern Analysis of Sidescan Sonar Images for Underwater Mine Detection. Submitted to IEEE Transactions on Pattern Analysis and Machine Intelligence, 2009.

[32] Partan, J., Kurose, J. and Levine, B.N., 2006. A Survey of Practical Issues in Underwater Networks. In Proceedings of the 1st ACM international workshop on Underwater networks (Los Angeles, CA, Sept. 25, 2006), 17-24.

[33] Prins, R. and Kandemir, M. 2009. Cooperative MCM for Heterogeneous Vehicles Optimized Under Constrained Time. Submitted to International Symposium on Unmanned Untethered Submersible Technology 2009 (Durham, NH, Aug. 23-26, 2009).

[34] Prins, R. and Kandemir, M. 2009. Communication Strategies in Multi-UUV Cooperative MCM Reconnaissance. Submitted to AUVSI's Unmanned Systems 2009 (Washington, DC, Aug. 10-13, 2009).

[35] Prins, R. and Kandemir, M. 2008. Time-Constrained Optimization of Multi-AUV Cooperative Mine Detection. In OCEANS 2008 (Quebec City, QC, Sept. 15-18, 2008).

[36] Proakis, J., Rice, J., Sozer, E. and Stojanovic, M. 2003. Shallow Water Acoustic Networks. In Encyclopedia of Telecommunications. John Wiley and Sons, 2003.

[37] Pugh, J. and Martinoli, A. 2007. Inspiring and Modeling Multi-Robot Search With Particle Swarm Optimization. In Proceedings of the 2007 IEEE Swarm Intelligence Symposium (Honolulu, HI, April 01 - 05, 2007). IEEE Press, Piscataway, NJ, 332-339.

[38] Rao, C., Mukherjee, K., Gupta, S., Ray, A. and Phoha, S. 2009. Underwater Mine Detection Using Symbolic Pattern Analysis of Sidescan Sonar Images. In American Control Conference (ACC 09) (St. Louis, MO, June 10-12, 2009).

[39] Rajala, A., Edwards, D., and O'Rourke, M. 2005. Collaborative Behavior for Vehicle Replacement in AUV Formations. In Proceedings of the 2005 ASME-IMECE (Orlando, FL, Nov. 5-11).

[40] Rajala, A. and Edwards, D. 2006. Allocating AUVs for Mine Map Development in MCM. In OCEANS 2006 (Singapore, May 16 - 19, 2007). IEEE Press, Piscataway, NJ, 1-8.

[41] Richards, M.D., Whitley, D. and Beveridge, J.R., 2005. Evolving Cooperative Strategies for AUV Teams. In Proceedings of the 2005 conference on Genetic and evolutionary computation (Washington, D.C., June 25-29, 2005), 1721-1728.

[42] Roeckel, M., Rivoir, R., Gibson, R. and Linder, S.P., 1999. Simulation Environments for the Design and Test of an Intelligent Controller for Autonomous Underwater Vehicles. In Proceedings of the 31st conference on Winter simulation (Phoenix, AZ, Dec. 5-8, 1999), Vol. 2, 1088-1093.

[43] Ryan, A., Zennaro, M., Howell, A., Sengupta, R. and Hedrick, J.K., 2004. An Overview of

Emerging Results in Cooperative AUV Control. In 43rd IEEE Conference on Decision and Control, Vol. 1, 602-607.

[44] Sariel, S., Balch, T. and Stack, J., 2006. Distributed Multi-AUV Coordination in Naval Mine Countermeasure Missions. Technical Report. Georgia Institute of Technology.

[45] Schoenwald, D.A. 2000. AUVs: in Space, Air, Water, and on the Ground. IEEE Control Systems Magazine, Dec. 2000, Vol. 20, no. 6, 15-18.

[46] Schumacher, C., Chandler, P. and Rasmussen, S.J. 2002. Task Allocation for Wide Area Search Munitions via Iterative Network Flow. AIAA Guidance, Navigation and Control Conference and Exhibit (Monterey, CA, Aug. 2002), 2002-4586.

[47] Seah, W.K.G., Tan, H., Liu, Z. and Ang, M.H., 2005. Multiple-UUV Approach for Enhancing Connectivity in Underwater Ad-hoc Sensor Networks. In Proceedings of OCEANS 2005 (Washington, DC, Sept. 19-23, 2005), 2263-2268.

[48] Simons, D.G., McHugh, R., Snellen, M., McCormick, N.H. and Lawson, E.A. 2001. Analysis of Shallow-Water Experimental Acoustic Data Including a Comparison With a Broad-Band Normal-Mode-Propagation Model. IEEE Journal of Oceanic Engineering, Jul. 2001. Vol. 26, no. 3.

[49] Sujit, P.B., Sinha, A. and Ghose, D., 2006. Multiple AUV Task Allocation Using Negotiation. In Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (Hakodate, Japan, May 8-12, 2006), 471-478.

[50] Turner, R.M. and Blidberg, D.R., 1990. An Investigation of the Impact of Communication Strategies on the Cooperative Behavior of Multiple, Cooperating Autonomous Underwater Vehicles. Final Report, New Hampshire University Durham Marine Systems Engineering Lab, Oct. 1 1990.

[51] Ulam, P., Endo, Y., Wagner, A. and Arkin, R., 2007. Integrated Mission Specification and Task Allocation for Robot Teams – Design and Implementation. In Proceedings of the 2007 IEEE International Conference on Robotics and Automation (Rome, Italy, Apr. 10-14, 2007), 4428-4435.

[52] United States Department of the Navy, 2004. The Navy Unmanned Undersea Vehicle (UUV) Master Plan. Approved for Public Release, 2003 Update.

[53] Vasilescu, I., Kotay, K., Rus, D., Dunbabin, M. and Corke, P., 2005. Data Collection, Storage, and Retrieval with an Underwater Sensor Network. In Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems (San Diego, CA, Nov. 2-4, 2005), 154-165.

[54] Vasilescu, I., Rus, D., Dunbabin, M., and Corke, P., 2008. Experiments with Cooperative Networked Control of Underwater Robots. Experimental Robots, vol. 39., 463-470. Springer Berlin, Heidelberg, Germany.

[55] Vincent, P. and Rubin, I., 2004. A Framework and Analysis for Cooperative Search Using AUV Swarms. In 2004 ACM Symposium on Applied Computing (Nicosia, Cyprus), 79-86.

[56] Wilde, J., DiBiaso, D. and Nervegna, M., 2007. Team Planning for Unmanned Vehicles in the Risk-Aware Mixed-Initiative Dynamic Replanning System. In OCEANS 2007 (Vancouver, BC, Sept. 29-Oct. 4), 1-8.