

Coordinated Search with Unmanned Aerial Vehicle Teams



Paul Ward
Kellogg College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Michaelmas 2013

Abstract

Advances in mobile robot technology allow an increasing variety of applications to be imagined, including: search and rescue, exploration of unknown areas and working with hazardous materials. State of the art robots are able to behave autonomously and without direct human control, using on-board devices to perceive, navigate and reason about the world. Unmanned Aerial Vehicles (UAVs) are particularly well suited to performing advanced sensing tasks by moving rapidly through the environment irrespective of the terrain. Deploying groups of mobile robots offers advantages, such as robustness to individual failures and a reduction in task completion time. However, to operate efficiently these teams require specific approaches to enable the individual agents to cooperate.

This thesis proposes coordinated approaches to search scenarios for teams of UAVs. The primary application considered is Wilderness Search and Rescue (WiSaR), although the techniques developed are applicable elsewhere. A novel *frontier-based* search approach is developed for rotor-craft UAVs, taking advantage of available terrain information to minimise altitude changes during flight. This is accompanied by a lightweight coordination mechanism to enable cooperative behaviour with minimal additional overhead. The concept of a *team rendezvous* is introduced, at which all team members attend to exchange data. This also provides an ideal opportunity to create a comprehensive team solution to relay newly gathered data to a base station. Furthermore, the delay between sensing and the acquired data becoming available to mission commanders is analysed and a technique proposed for adapting the team to meet a latency requirement.

These approaches are evaluated and characterised experimentally through simulation. Coordinated frontier search is shown to outperform greedy walk methods, reducing redundant sensing coverage using only a minimal coordination protocol. Combining the search, rendezvous and relay techniques provides a holistic approach to the deployment of UAV teams, meeting mission objectives without extensive pre-configuration.

Acknowledgements

Firstly, I am extremely grateful to my supervisor Stephen Cameron. His guiding approach to supervision has made my time as a research student all the more enjoyable, supporting and encouraging me through topic changes and non-academic turbulence.

My DPhil experience in the Computing Laboratory has been enhanced by my lab-mates, including Muzammil Hussain, Andrew Symington, Sonia Waharte, Sarfraz Nawaz, Julian de Hoog and Helen Flynn. In particular, Julian and Helen have assisted my work by sharing their knowledge of rescue robots and sensing respectively.

I am fortunate to have been supported in this work by the EPSRC funded project SUAAVE (Sensing Unmanned Autonomous Aerial VEHicles - grant reference EP/F064217/1). Beyond financial assistance, this has provided me with opportunities to learn about related topics and gain a wider perspective on the field. In particular, I would like to thank Chunbo Luo for his efforts during our collaboration.

During the production of a thesis, often those close to the author will have a significant hand in the final outcome. My wife Karen has been alongside me for every step of the way, showing extreme patience, providing encouragement and above all, being a light in dark places. It is no exaggeration to say that this thesis would not have been completed without her support and I would gladly place her name on the front page if it were possible. Thank you to our beautiful baby daughter Scarlett, for allowing me to have just enough sleep over the past few months to complete this work and always brightening up my day with her smiles.

My family have always encouraged me, especially in academic endeavours. They have given support in difficult times and their confidence in me has always provided an extra motivation. Special thanks to my sister Emma, who stepped in at short notice with some expert proof reading.

Contents

1	Introduction	1
1.1	Motivation	2
1.2	Proposed Approach	4
1.3	Progression of Research	5
1.4	Thesis Structure	6
2	Literature Review	7
2.1	Robot Beginnings	8
2.2	Multi-Robot Systems	9
2.2.1	Definition and Distribution of Control	9
2.3	Mobile-Robot Applications	10
2.3.1	Search and Rescue	11
2.3.2	Remote Exploration	12
2.3.3	Surveillance and Tracking	13
2.3.4	Hazardous Area and Environmental Monitoring	14
2.3.5	Object and Environment Manipulation	15
2.4	Ground-Based Multi-robot Systems	16
2.4.1	Coordination	17
2.4.1.1	Explicit Coordination	17
2.4.1.2	Implicit Coordination	19
2.4.1.3	Communication for Coordination	20
2.4.2	Localisation and Mapping	21
2.4.2.1	Occupancy Grid Representations	22
2.4.2.2	Simultaneous Localisation and Mapping	22
2.5	Aerial Multi Robot Systems	23
2.5.1	Rotor-craft Vehicles	24
2.5.2	Coordinated Deployments	24
2.5.2.1	Communication	25

2.5.2.2	Shared Probabilistic Maps	26
2.5.2.3	Localisation	26
2.6	Summary	27
3	Problem Description	29
3.1	Wilderness Search and Rescue	30
3.1.1	Robot Assistance	32
3.2	Goals	33
3.3	Assumptions and Scope	34
3.4	Performance Metrics	35
3.4.1	Search Coverage	35
3.4.2	Team Performance	36
3.4.3	Data Latency	37
3.5	Summary	38
4	Individual Robot Behaviour	39
4.1	Tasks	40
4.2	Area Sensing	40
4.2.1	Representing Spatial Information	41
4.2.2	Frontier Search	44
4.2.3	Frontier Generation	45
4.2.4	Frontier Selection	50
4.2.4.1	Altitude Change Weighting	52
4.2.4.2	Within the Team Environment	54
4.2.5	Frontier Coverage	55
4.3	Relaying	55
4.3.1	Opportunistic Altruistic Relay	56
4.4	Behaviour Transitions	57
5	Team Organisation	59
5.1	Search Team Organisation	60
5.1.1	Information Flow	60
5.2	Coordinated Frontier Search	61
5.3	Rendezvous	64
5.3.1	Rendezvous Timing	66
5.3.2	Selecting Rendezvous Location	67
5.4	Team Relay	68

5.4.1	Designated Single Relay	69
5.4.2	Multi-Robot Relay	70
5.4.2.1	Relay Configuration	71
5.5	Adaptive Data Latency Driven	
	Multi-Robot Relay	73
5.5.1	Adaptation Mechanism	76
6	Simulation Experiments	81
6.1	Simulation Usage	82
6.2	Simulation Detail	83
6.3	Approaches Under Consideration	86
6.4	Methodology	88
6.4.1	Environments	90
6.5	Benchmark	92
6.6	Performance Analysis Experiments	96
6.6.1	Frontier Search	96
6.6.1.1	Basic Search Rate	96
6.6.1.2	Search Rate for Team Deployments	97
6.6.2	Altitude Change Preference	100
6.6.3	Rendezvous and Opportunistic Relay	101
6.6.3.1	Altruistic Relay	102
6.6.3.2	Search Team Rendezvous	104
6.6.3.3	Team Relay Effort	108
6.6.4	Rendezvous with Designated Relays	109
6.6.4.1	Comparison with Altruistic Relaying	109
6.6.4.2	Frontier Selection Heuristics for Rendezvous	112
6.6.5	Data Latency	113
6.6.5.1	Multi-Link Relay	114
6.6.5.2	Adaptive Relay	118
6.7	Characterisation Experiments	124
6.7.1	Scalability	125
6.7.2	Communication	127
6.8	Summary of Results	130

7	Discussion	131
7.1	Examining Coordinated Frontier Search	132
7.1.1	Frontier Approximation Methods	133
7.1.2	Communication of Agent Location and Intention	135
7.2	Further Analysis of Data Latency	137
7.3	Usage	142
7.3.1	Team Size and Restricted Regions	142
7.3.1.1	Multiple Independent Teams	144
7.3.2	Robotic Platforms	145
7.3.2.1	Sensing	146
7.3.2.2	Motion Capability	147
7.4	Failure Modes and Response	148
7.5	Spatial Progression of Search	150
7.6	Towards Practical Deployments	153
7.6.1	Integration with Search and Rescue Teams	153
7.6.2	Deployment Example	154
8	Conclusions	157
8.1	Summary of Contributions	158
8.1.1	Comparison with Previous Work	160
8.2	Limitations	161
8.3	Directions for Future Work	162
8.4	Summary	165
A	Multi-robot Simulation Software	167
B	Additional Experimental Results	173
C	Coordination in Multi-Tiered Robotic Search	177
C.1	Introduction	178
C.1.1	Robot Tiers	179
C.2	Related Work	180
C.3	Simulation Experiments	183
C.3.1	Experimental Configuration	184
C.4	Results	186
C.4.1	Ground to UAV Communication	186
C.4.2	Number of Ground Robots	187

C.4.3 Robot Velocity	188
C.5 Appendix Conclusions	189
Bibliography	191

List of Figures

1.1	Example of complex terrain	3
4.1	UAV sensing configuration	41
4.2	Spatial information layers used during search	43
4.3	Frontier exploration for ground-based robots	46
4.4	Data sources for frontier generation	47
4.5	The frontier generation process	49
4.6	Sample results from the frontier generation process	51
4.7	UAV behaviour transitions	58
5.1	Team information flow	61
5.2	Coordinated frontier search	62
5.3	Team rendezvous	65
5.4	Example rendezvous locations	67
5.5	Single designated relay	69
5.6	Multi-link relay chain configuration	71
5.7	Information flows for rendezvous and multi-relay operation	73
5.8	Adding a relay to the multi-link chain	78
5.9	Subtracting a relay from a multi-link chain	79
6.1	Approaches under consideration	86
6.2	Sample randomly located priority regions	89
6.3	Terrains used during simulation experiments	91
6.4	Information known at base during a search	93
6.5	Search progress for different terrains	94
6.6	Rate of information gain during search	95
6.7	Comparison of search methods using information gain	97
6.8	Mission completion times for a range of search methods	99
6.9	The effect of the altitude change weighting parameter	102
6.10	Information gain rate of self-relay and altruistic relay strategies	103

6.11	Mission completion with combinations of team rendezvous and relay approaches	106
6.12	Information shared within the team	107
6.13	Time spent on relaying activities	108
6.14	The effect of search distance to base on completion times	111
6.15	Frontier selection heuristics	113
6.16	Average data latency for multi-link relay chains	115
6.17	Maximum data latency for multi-link relay chains	116
6.18	Minimum data latency for multi-link relay chains	116
6.19	Mission completion for teams using multi-link relay	118
6.20	Average data latencies for designated and adaptive relay configurations	120
6.21	Completion times for designated and adaptive relay configurations . .	120
6.22	Number of relayers used by the adaptive relay approach	122
6.23	A comparison of relays used and information gain with the adaptive relay method	123
6.24	Effects of team size on mission completion time	126
6.25	Effects of team size on data latency	127
6.26	Effects of communication range on mission completion times	129
6.27	Effects of communication range on data latencies	129
7.1	Search recoverage using different search techniques	132
7.2	Frontier approximation methods	134
7.3	The effect of frontier approximation on mission completion time . . .	134
7.4	The effect of multi-hop communication on completion time	136
7.5	Data latencies by environment location	139
7.6	Data latency profiles for different relay configurations	141
7.7	Information gain per UAV for a range of team sizes	145
7.8	Spatial progression of search	151
7.9	An example practical deployment	155
A.1	Mission commander planning interface	169
A.2	Simulation overview interface	170
A.3	Agent knowledge maps during simulation	171
B.1	Large mountainous terrain	173
B.2	Rate of information gain while searching the plains terrain	174
B.3	Rate of information gain while searching the undulating terrain . . .	174

B.4	Data latencies by sensing location	175
C.1	Three tier multi-robot search	184
C.2	Simulation in progress	185
C.3	Time required to complete search tasks	187
C.4	Analysis of task auctions	188
C.5	Effects of team size and search velocity	189

List of Tables

6.1	Ground height ranges for sample terrains	90
6.2	UAV behaviours during a search mission	95
A.1	Typical Simulation Parameters	171
C.1	Multi-Tier Simulation Parameters	186

List of Algorithms

4.1	Generation of Search Frontiers	48
5.1	Frontier Selection	63
5.2	Rendezvous and Multi-robot Relay Configuration	72

Chapter 1

Introduction

As technology improves the opportunity for deploying autonomous mobile robots into a range of real-world situations greatly increases. There is a wide variety of vehicle types with wheels, wings, rotor-blades and propellers, operating in an equally large set of environments, on land, in the air, on and under water. From their static beginnings in industry, mobile robots are becoming commonplace in many walks of life; for example: cleaning the home [115], exploring the surfaces of the Moon [9] and Mars [99] and urban aerial surveillance [126]. Furthermore, in many of these applications, devices have made the leap from the academic and industrial world to successful, and sometimes affordable, commercial products, such as: vacuum cleaners [43], lawn mowers [118] and gutter clearers [72].

All current state-of-the-art mobile robots share a common set of components, including: on-board computing, sensing devices, wireless communication radios and data storage. Many of these robots are currently controlled by human operators; however there is a growing trend for automation. This will enable robots to operate in environments that are more difficult to access, unsuitable or unsafe for humans, or beyond the reach of direct communication.

The principle aim of this research is to investigate effective methods for the deployment of groups of autonomous mobile robots in scenarios where spatial positioning is paramount to achieving the goal. A common example is assisting search and rescue

operations; robots must move through the environment using sensing devices in an attempt to locate missing persons (or other assets). Emphasis is placed on devising suitable methods for coordination to maximise the performance of the group as a whole.

1.1 Motivation

There are a range of application areas now being considered for the use of groups of mobile robots, or *multi-robot systems*, many of which can be described as problems of search, exploration or monitoring. Typically this requires mobile robot platforms to move through an environment and gather data with on-board sensors for the purpose of locating a target, building a map or simply collecting information. In this work the primary motivation is drawn from the Wilderness Search and Rescue (WiSaR) scenario; the DPhil is funded by the EPSRC sponsored project SUAAVE (Sensing Unmanned Autonomous Aerial Vehicles) [131]. However, many of the approaches taken in mobile robotics research are not strictly dependent on the application and can be applied more widely.

Wilderness Search and Rescue (WiSaR) refers to the task of locating and assisting missing person(s) in uninhabited natural environments such as national parks, and has traditionally been carried out by human responders. The key objective is to find missing persons as quickly as possible to reduce the delay in providing medical attention and extraction from the scene. A significant challenge is posed by the environment, which is often difficult to traverse on foot and finding an injured person in a concealed location may only be possible at short range. For example, see the complex unstructured terrain in Figure 1.1.

A new tool that has become available is the small Unmanned Aerial Vehicle (UAV), with the ability to carry sensing and computing equipment on-board. Small UAVs are generally from one of two categories, fixed wing (aeroplane) and rotor-craft (helicopter), with digital cameras as the primary sensing approach. Such aerial robots



Figure 1.1: Complex terrain in Portstewart, Northern Ireland, United Kingdom

are particularly suited to this application, combining a high degree of mobility with an alternative viewpoint on the terrain; especially rotor-craft UAVs with their ability to hover while sensing¹. Currently, the choice between these two vehicle types is a trade-off between the high sensing accuracy provided by low-altitude battery powered rotor-craft and the much longer flight times of fixed-wing platforms. In addition, it is possible to consider deploying both categories of aerial robot simultaneously in a heterogeneous team.

This type of equipment could play a much larger role in WiSaR if a sufficiently high level of autonomy can be achieved, reducing the requirement of human operators. In live emergency situations the proposed robot team will become an additional resource and hence should be usable by non-experts and provide suitable information to the human rescuers.

¹Traditional manned helicopters are very effective in this role, but are extremely expensive to use.

1.2 Proposed Approach

Search scenarios vary widely in their details, but here the focus is on the key issues of coordination for a WiSaR framework. Accordingly, this thesis considers a proposed solution to the multi-robot search problem; aiming to meet the dual objectives of sensing at required locations and providing regular updates to a human mission commander situated at a base station. One assumption made here is that long-range wireless communication is unreliable. Although possibly surprising, this assumption was borne out by physical experiments within SUAAVE, as terrain interferes and antennae become misaligned. Also, an abstraction made is that the exact nature of the sensing task is ignored, instead, we assume that the UAV must pass over (“cover”) all if its search area in order to detect a missing person. The approach taken in this thesis attempts to meet these goals by designating *search* and *relay* as separate robot activities. These behaviours are framed within a multi-robot setting, striving for efficiency through cooperation.

A novel *frontier-based* search approach is developed for rotor-craft UAVs, taking advantage of terrain information obtained from topographical maps to minimise altitude changes. This is accompanied by a lightweight coordination mechanism to enable cooperative behaviour with minimal additional overhead. Cooperation between robots is promoted by the new concept of an arranged meeting, or *rendezvous*, for all team members. This also provides an opportunity for a remotely deployed team to consolidate relay activity; a single robot can be responsible for relaying new information for the whole team. Furthermore, the inherent delay in new data becoming available at base is addressed with a method that adapts the behaviour of a homogeneous team.

Operating mobile robots brings considerable inherent uncertainty in many aspects, including: motion, sensing, communication and environmental conditions. Rather than attempting to mathematically model this uncertainty, the approaches in this thesis try to mitigate its effects by designing systems that are adaptable. Robot

behaviour is therefore informed by information gained and progress towards the particular scenario goals; for instance, updates concerning areas that have been visited as part of a search task. This methodology is also aligned with the goal of allowing a rapid practical deployment for emergency situations.

The novel work in this thesis includes:

1. Coordinated frontier search for UAVs
2. Rendezvous meetings for the entire search team
3. Adaptation of relay configurations during a deployment
4. Use of real terrain information from topographical maps to inform behaviour

1.3 Progression of Research

Preliminary work was performed to examine the use of structure in multi-robot teams. This demonstrated the use of coordination approaches between different classes of robots to complete sensing tasks in an abstract environment. In particular, the use of task allocation methods was investigated. As a result, the importance of taking into account the individual robot behaviours while constructing teams was highlighted; especially when the problem contains multiple objectives. A report on this work is provided in Appendix C.

The primary research effort was placed on the investigation of coordinated robot teams for the realistic WiSaR application, yielding the approaches described in this thesis. This work required the use of real terrain information and the development of suitable simulation software.

The multi-robot simulator enables the deployment of teams of robots to be examined in great detail and promotes rapid development of approaches. The importing of real terrain data is supported and used in creating realistic environments for simulated search missions. All aspects of robot behaviour can be logged for subsequent

analysis and a graphical interface enables the observation of deployments in progress. Further detail of the simulation software is provided in Chapter 6 and Appendix A.

1.4 Thesis Structure

A review of relevant literature is presented in Chapter 2, including an introduction to mobile robotics and their applications, followed by a description of previous approaches taken in deploying multi-robot systems. Chapter 3 describes the problem domain of wilderness search and rescue. Full details of the goals, assumptions, scope and performance metrics are given. In Chapters 4 and 5, the proposed solution is described, beginning with the behaviour of individual robots and followed by approaches that consider actions in a team context. Experimental work is presented in Chapter 6, covering the methodology and results. Chapter 7 provides further analysis and a discussion of issues arising from the experiments. The thesis concludes with Chapter 8, which summarises the contributions made and potential for further work in this area.

Chapter 2

Literature Review

This literature review begins with an introduction to the field of mobile robotics, with particular focus on Multi-Robot Systems (MRS), followed by a description of applications with particular relevance to this thesis. A review of MRS research is organised by type of vehicle used and discusses the main issues faced and approaches taken.

2.1 Robot Beginnings

Robots are being considered for many tasks that could otherwise be carried out by humans, offering a number of advantages outlined below.

- Robots can withstand extreme environments, including temperature, chemical toxicity and lack of breathable atmosphere.
- Robots offer increased sensory capability by carrying such equipment as infrared cameras and chemical sensors.
- Robots are highly mobile allowing them to be positioned in locations challenging for humans, to both perceive and affect the world.
- Robots can be extremely precise, for example in a repeatable industrial setting.
- Robots can be cheaper, especially for long term tasks such as continuous surveillance.
- Robots are expendable.

This thesis considers the use of mobile robots, typified by a mobile vehicle platform that carries on-board sensing and computing capability. Early work in robotics and artificial intelligence was primarily focused on attempting to construct models of the world. These would then be used to reason and plan complete optimal solutions such as a specific control sequence. A major shift occurred towards building robots that use minimal internal state to produce control inputs from sensor information [16], known as reactive [3] or behaviour-based robotics [17]. This introduced to robotics the concepts of *embodiment*, having a specific physical body, and *situatedness*, being located in a physical environment that can be perceived and affected. Throughout this work, particularly during the simulation experiments of Chapter 6, each robot is considered as an individual autonomous actor. The modelling of a physical body restricts robot movement suitably for the class of vehicle considered. All robots

are situated in the simulated environment, this impacts on their sensing capability, perceiving only the immediate area, and limiting their communication with others. Furthermore, an agent's motion can be affected by the topology; it must ensure it does not collide with any obstacles.

2.2 Multi-Robot Systems

Due to vehicle improvements, in both capability and reduced size, using groups of robots has become a realistic proposition. Although this can be considered as a specialisation of general Multi-Agent Systems (MAS), in practice the additions of situatedness and embodiment greatly negate the effectiveness of many MAS approaches (robots should not interfere with each other). Utilising multiple robots commonly aims to reduce the time taken to achieve goals and improve resilience to failures. The latter implies the need for the system to adapt during execution to ensure all necessary functions are performed. Furthermore, some tasks cannot be completed without multiple actors, a classic example is pushing a heavy box. In this case a single robot alone would not be able to complete the task; hence multiple cooperating robots are required.

2.2.1 Definition and Distribution of Control

An important aspect of any MRS is the definition of overall goals and how the various robots are controlled to achieve them. A common approach that suits many applications is to select a series of sub-goals, which when completed indicate the end of the overall problem. Therefore by dividing the work into distinct packages, focus shifts to an efficient assignment of robots to sub-goals. Although generally this takes the form of a one to one relationship (one robot, one sub-goal at a time), it is possible that certain goals require multiple robots or a single robot can simultaneously complete multiple goals [50].

A simple robot-task pairing defines both the goal and chosen agent but says noth-

ing about how the goal may be accomplished. In order for a mobile robot to complete the required task a suitable control scheme must be defined. At the most complete level, a task itself can represent a fully planned activity for the robot, achieved with a sequence of control inputs. This implies minimal decision making during execution as the robot attempts to navigate a fixed trajectory. Conversely, a robot can select a defined high-level behaviour to match the task, for example measure air temperature in an area. The robot then behaves in the required manner, such as moving through the environment using on-board sensors and processing to avoid obstacles and gather data. The general trend in this area is toward a higher level of autonomy for individual robots, facilitated by improved technology and miniaturization.

Early work in MRS used centralised control, such as a single base-station, to simultaneously direct multiple robots; this simplification enables solutions that target optimal MRS behaviour. However, the communication reliability and range required to achieve this is seldom available in real-world applications. Therefore much of the current research concentrates on distributed approaches, with robots taking decisions based on locally available information. Although this often results in sub-optimal performance, it is more suited to uncertain and complex real-world scenarios.

2.3 Mobile-Robot Applications

There is a group of applications relevant to this thesis that have significant similarity and consequently approaches can often be applied more widely than the original problem domain. The following list describes some of the common characteristics typically found in such application areas.

- Operate in real-world environments that can be both unstructured and dynamic
- Requirement to perceive/sense at challenging physical locations using a high degree of mobility
- Inherent communication restrictions; due to distance, obstacles or topography

- New information must be transferred to a base-station or fixed infrastructure
- Uncertainty and variability in the task during execution

Mobile multi-robot systems are being considered for a range of applications, the majority of which fall into the following main categories.

2.3.1 Search and Rescue

One of the most common applications considered by the research community is search and rescue; locating injured or lost humans in challenging environments. These range from vast outdoor wildernesses, such as a missing hill-walker lost in a national park, to unstable enclosed urban areas, for example a collapsed building following an earthquake. One of the first live deployments of rescue robots was during the response to the World Trade Centre collapse in 2001 [22]. The type of robot deployed in various scenarios varies widely, see [87] for a review. Robots are used primarily to locate missing persons and notify human responders of the areas searched and what (if anything) has been found. Any rescue procedure is likely to be carried out by highly skilled humans for the near future, however robots could deliver life-saving supplies such as oxygen before rescuers are able to reach a casualty. A key characteristic of search and rescue scenarios is the restriction placed on inter-robot and human-robot communication by the environment; this can be due to large distances involved and/or attenuation caused by obstacles, such as walls.

In the Wilderness Search and Rescue (WiSaR) problem, robots are tasked with locating a missing person in an unstructured wilderness type environment [51]. One of the key criteria is the time required to find and attend to the missing person, as expected survivability decreases over time. UAVs are particularly suited to this task, using a high degree of mobility to visit sites of interest over challenging terrain. For instance the sFly project [127] aims to develop small UAVs for operations in obstacle-ridden environments with particular focus on localisation challenges, [152] for example. Cameras are often used to attempt automated detection of missing

persons using computer vision techniques [36, 133, 148]. Manned helicopters are very effective in this domain, but are vastly more expensive to purchase and operate than UAVs.

The Robocup rescue scenario is a testbed for the search and rescue area; robots are deployed into a post-disaster situation where they may be people trapped in a partially collapsed building [80]. Custom developed robots make their way through the building to search for survivors. As part of this task robots often attempt to build a map of the environment, which may have changed substantially, as they go; a common method uses laser range finders in a SLAM (Simultaneous Localisation And Mapping) scheme. The HELIOS system is a team of heterogenous robots for urban search and rescue operations [53]. The highly specialised tracked platforms are able to move through complex environments, including climbing stairs and self-righting [54]. These vehicles are able to be equipped with sensing and/or effector devices, such as an arm capable of grasping and opening a door. This allows accurate mapping of indoor environments prior to entry by human rescuers.

2.3.2 Remote Exploration

The key goal of exploration type applications is to traverse an unknown environment to gather information about that specific physical environment. Generally the aim is to build an accurate map representation of the area using one or more robots. Hence the key supporting technologies include localisation, accurate sensing and perception and merging maps. Operation of a multi-robot group can be optimised for required scenario goals including speed, accuracy and area covered. Typical environments for this type of robot deployment are generally difficult for humans to work in and include: underwater [94], underground (such as mines [137] or caves [68]) and on other planets [25].

2.3.3 Surveillance and Tracking

In surveillance and tracking applications it is specific targets or areas that robots are deployed to perceive using sensor devices. One of the early robotic problems is the positioning of a surveillance robot in order to provide maximal sensor coverage in an indoor environment. This is extended using mobile robots to perform security sweeps; research efforts have examined the calculation of trajectories to ensure optimal time-varying coverage and identify intruders with minimal delay. Using robots in place of humans for surveillance can lead offer benefits such as reduced cost and greater coverage. For example, consider a large border region between countries, whether on land or at sea, that must be monitored for smuggling and illegal immigration. Currently, sufficient surveillance can be extremely difficult to achieve without substantial human resources due to the large area of terrain. Due to their high mobility researchers are beginning to consider mobile robots for this task [55]; especially UAVs. A high degree of autonomy is required along with the ability to quickly raise the alarm if necessary.

Robotic platforms can be ideal for gaining alternative or enhanced views in changing situations. For instance UAVs under manual control can be used as an “eye in the sky” in civilian applications such as crowd monitoring at major events such as the Olympic Games. Several UK Police forces have already trialled UAVs under manual control, using on-board thermal cameras to track suspects and make arrests [10]. The military have used robotic surveillance platforms, especially airborne ones, for quite some time, with recent advances allowing smaller, lower cost and lower altitude vehicles. These can be deployed for reconnaissance gaining key information about building layouts, enemy positions and IEDs (Improvised Explosive Devices).

The ability to not only acquire but track moving targets is applicable to both civilian and military contexts. The difficulty of this task is highly related to the characteristics of the tracked target; passive or predictable targets, such as a car travelling on a road, being markedly easier than adversarial or unpredictable ones, for instance criminal suspects attempting to evade capture. The Oxford Aerial Tracking System [59] tracks a moving vehicle in an outdoor environment using a single helicopter

equipped with an active gimballed camera. Using a computer vision technique called mean-shift tracking [26] and neural network learning, the combined movement of the platform and gimbal attempts to maintain a continuous view of the target.

Researchers have built teams of UAVs to track moving targets, attempting to improve tracking capability and robustness. In [13] multiple cooperating quadrotor UAVs are used to track a moving car in a restricted indoor environment. By maintaining multiple continuous views of the target the accuracy of the tracked location was significantly improved. A similar task involves a group of fixed wing UAVs tracking a single target in an outdoor urban setting [77]. Flying at a maximal altitude to gain the widest camera field of view, circular paths are planned using current estimation of the target's location to provide sensor coverage.

2.3.4 Hazardous Area and Environmental Monitoring

Some environments are particularly hazardous for humans and hence provide an ideal opportunity to deploy robots which can withstand such harsh conditions and are ultimately expendable. For instance, nuclear material remaining from power generation is kept in storage pools filled with water. Underwater robots are being considered for the crucial monitoring that must be carried out [100]; with localisation in such a cluttered environment one of the key research goals [69]. Similarly, long distance power lines must be periodically inspected for faults [2]. Currently this dangerous work is performed by humans using specially equipped helicopters and automating this task with aerial robots offers a much lower risk alternative. Robots must be able to carry highly specialised equipment, for instance both visual and contact sensors may be required [67]. This provides additional engineering challenges as many current robot platforms have limited weight capacity and operation time.

The offshore environment is particularly suitable for the use of robots, being difficult to access due to the distances involved and potential danger in bad weather conditions. Tasks in this area include: the inspection of offshore wind turbines [101], monitoring the thickness of sea ice [58] and inspection of offshore oil platforms [11].

Deploying robots for such operations could both increase safety by removing humans from harm and allow continued operation in poor conditions.

An uncontrolled release of hazardous substances (eg: liquids, gases or radiation) can result from major incidents such as fires, industrial leaks or use of weaponry. Although extremely challenging it is vital to monitor the spread of these contaminants in order to protect human health, wildlife and gain control of the incident. UAVs have been proposed as a suitable technology to aid in this difficult task [29]. The authors note that minimising the time to assess the situation by measuring the presence and spread of harmful substances is crucial in reducing unwanted consequences. This type of task often requires very specialised sensor equipment including air sampling chemical and radiation sensors, placing an additional payload burden on the robots. For example, the gases produced at landfill refuse sites are harmful to the environment; mobile robots are being equipped with devices such as a laser absorption spectrometer to map the presence of these pollutants [60]. In the ongoing cleanup operation at the Fukushima nuclear reaction in Japan, robots are being used in locations that humans would not safely be able to work in. The Quince rescue robot has been modified specifically for this task with the addition of new sensing devices to detect radiation and water contamination levels [98].

Robots have been successfully deployed for handling explosive devices, including landmine detection, de-mining and bomb disposal. Automating the search for landmines requires the ability to traverse challenging terrain, positioning sensors such as metal detectors near to ground level. Such capability has been demonstrated by both a four-wheeled buggy [142] and six-legged walking [34] robots.

2.3.5 Object and Environment Manipulation

Whilst the majority of current applications consider robots chiefly as mobile sensor platforms, they can also affect the environment. Early robot research examined object manipulation behaviours such as pushing, grasping and lifting. Recent research in multi-robot systems has attempted to coordinate a group to perform similar tasks

cooperatively, including pushing boxes [109] with constant orientation and lifting an object using multiple UAVs [97].

Certain applications contain a higher degree of complexity, with a greater understanding of the environment required in order to successfully manipulate it. In the site clearance task [111], robots must first examine an area before clearing away obstacles. This is seen as a key capability for future space travel to other celestial bodies; ensuring a clear landing zone for both manned and unmanned missions.

The Robocup soccer tournament [79, 103] pits teams of robotic players against each other. They attempt to both score and prevent goals by affecting the ball through various means, including kicking by a legged robot or simply blocking. Although this may seem trivial, it is surprisingly complex to develop behaviours such as passing the ball to team-mate and shooting for the goal. This requires an awareness of the robot's location on the pitch, relative position to its team-mates and the ball. It is hoped that developments in this field can be transferred to other more productive applications [5].

2.4 Ground-Based Multi-robot Systems

Early work in MRS considered using ground-based robots to improve performance over single robot solutions or perform tasks that could not be carried out by an individual vehicle. Some of the first research examined the control of a multi-robot team through combined motion planning; producing a solution of planned paths for the group of robots. This has been applied to exploration/mapping [46, 12] and cooperative object transportation [156]. Such approaches do not scale well with group size and environmental complexity, and also have difficulty with dynamic environments. An alternative research trend shifted away from approaches that generate plans for each agent to those that produce more dynamic coordinated behaviour, utilising the behaviour-based paradigm for individual robot control. For instance in order to achieve coordinated motion, inspiration has been taken from biological

systems, especially insects and birds, to find local control rules that produce an overall group effect [92]. For example a group of robots moving in formation whilst avoiding obstacles in the environment [7]. Similarly, the majority of the emphasis in multi-robot research is placed on dynamic cooperation and coordination techniques to achieve more complex application goals and respond to new information gained during operation.

2.4.1 Coordination

A key aspect of a multi-robot system is the manner in which the behaviour of individual robots is coordinated. Several existing surveys characterise coordination methods in various ways including: by coordination type [41], by communication and computational effort [38], and by intentionality [70].

Often the general strategy is determined by the application scenario; in many cases this results in a division of work, creating either sub-goals to be accomplished or roles to played by individual robots. Examples include segmenting the physical space in a sensing deployment and robots playing different positions in robot soccer [83]. Less common are applications where the goal cannot be reduced to individual robot actions, instead requiring strongly coordinated behaviour, for instance cooperatively pushing a box with high precision. This type of control problem is outside of the scope of this thesis.

Once a strategy is chosen, producing a set of potential individual robot tasks or behaviours, methods are needed to enable robots to perform suitable actions in a coordinated manner. The two types of approaches are explicit coordination (a direct assignment of robots to tasks) and implicit coordination with robots selecting a task based on currently available information.

2.4.1.1 Explicit Coordination

Explicit coordination approaches aim to assign tasks or roles to robots to produce a maximal overall outcome; achieving primary objectives with minimal unwanted

redundancy. Such task allocation has been exercised in a wide range of applications including: area exploration [120], collaborative box pushing [49] and security sweeping [74].

A common approach that has been transferred from general multi-agent research is that of auction or market based allocation. Robots place bids for tasks that they are able to complete based on the expected cost of completion, and the robot with the best bid wins. With mobile robots often the major component of this cost is moving to the required location. The Traderbots system [35] uses a weakly centralised approach with one of the robots acting as auctioneer when a task is required. Robots that have won tasks are also able to sub-contract them to others, if for instance the task can be completed for a lower cost. In this way the robots attempt to minimise the overall cost to complete all tasks. A broadly similar approach is taken in [49], with weakly centralised auctions for tasks demonstrated on a physical robotic system. A publish/subscribe broadcast communication system is used to reduce messaging requirements. The auctioneer is responsible for monitoring task progress by querying the assigned robot, allowing the task to be reassigned if insufficient progress has been made or the robot has failed. Market based allocation has also been demonstrated for area exploration where candidate unexplored regions are assigned to robots based on the estimated utility [159]. In [129] a distributed implementation allows individuals to bid for the task of exploring available areas using their travelling distance, expected information gain and proximity to others as their offer. This balances having the robots spread out when allocated to explore with staying close to others to improve coordination. The result is that the group splits into sub-networks that tend to search the region together.

Role assignment is commonly used in RoboCup soccer [79] to coordinate which robots perform required behaviours, for example goalkeeper, defender, attacker. In [83] roles are assigned in a centralised manner by a coach agent that is in communication with all the robot players. Information from all of the robots sensors is shared, allowing roles to be assigned based on the current position of robots in relation to

the play area and ball. To provide robustness against robot malfunction, thereby reducing the team size, roles are assigned in priority order with the striker first.

Certain goals cannot be achieved by a single robot alone; for example: box pushing, area clearing, strongly cooperative search, and localisation [110]. The use of multi-robot strategies is combined with a market based allocation approach in [73]. This type of approach can also be combined with simpler forms of coordination in a hybrid approach, for example to coordinate multiple robots in a surveillance sweep application [74]. [134] combines a market-based approach with schema mapping of sensor-actuator information to permit both strong and weak cooperative behaviour in a site clearing scenario. Robot coalitions are formed using pre-determined schemas that describe how robot sensors, actuators and communication is used in the team [109]. Task allocation is performed using a bidding system with centralised control. Attempts have also been made to modify coalition formation methods from MAS for the multi-robot domain [147].

2.4.1.2 Implicit Coordination

Multi-robot coordination can also be achieved without explicit agreement or assignment between robots. Implicit coordination relies on the sharing of information to allow each robot to choose a task, role or behaviour using a utility function. Typical data types used in making this decision include maps of the environment, known locations of team-mates and progress made towards objectives.

This type of coordination has been demonstrated in multi-robot exploration using the frontier approach [157]. Each robot maintains a map of explored areas, periodically sharing and merging this information with others when in communication range [18, 33]. Individual robots designate frontiers as constrained sections of the boundary between explored and unexplored areas, each of which is valued with a utility function and the most rewarding selected. This enables a coordinated exploration of the environment without direct assignment of robots to frontiers.

Implicit coordination has been used in RoboCup soccer to assign roles without

a centralised controller [146]. All robots share information, such as their estimated position of the ball and themselves, by periodic broadcast. Each then calculates the suitability of itself and all of its team mates for each of the roles, again using a priority order. As each robot is using the same information they should all reach the same role allocation without any directed communication. The broadcasted information is used additionally to avoid collisions and to aid robots who are unable to currently sense the ball directly. This work also uses a potential fields approach [75] to determine the player's role dependent on target location, pulling an attacking player towards the ball for instance. Another distributed approach, using an emergent methodology for football role allocation, is presented in [153]. The authors note the increased cost of switching tasks often, as roles tend to require the robot to be in different areas of the field, with robots becoming specialised by preferring roles they have performed recently. One disadvantage of this approach is that a player may make a poor role decision in new circumstances because it has recently been performing that role.

2.4.1.3 Communication for Coordination

The ability to communicate in multi-robot scenarios is vital to support coordination and information sharing activities. However despite improvements in wireless radio technologies the operating environment has a significant impact and can severely restrict both basic connectivity and reliable bandwidth. This occurs across the spectrum of potential environments, from attenuation caused by walls inside buildings to the distances involved in large scale outdoor operations.

Researchers have attempted to mitigate this issue by restricting the movement or location placement of robots. For example by maintaining a line-of-sight between wirelessly connected neighbours in order to preserve a continuous robot network [4, 66]. Similarly in outdoor environments placing limits on the inter-robot distance allows communications to be constantly maintained [119, 135]. Although these methods undoubtedly provide certain advantages, such as increased information sharing and consequently greater coordination efficiency, by limiting mobility overall progress

towards the goal can be hampered. This is especially the case in large scale sensing tasks like wilderness search and rescue, where a significant physical space must be searched by a limited number of available robots.

An alternative method of combating communication problems is to organise pairs of robots to meet at predefined locations and times known as a rendezvous [122]. This has been demonstrated in search and rescue robotics, in particular for disaster response with the frontier exploration algorithm [33]. In this role-based example each rendezvous allows a *searcher* and *relay* to exchange data. The period between each wirelessly connected meeting enables the searcher to explore new areas and the relay robot to transport new information to a fixed base station. This approach takes advantage of the high degree of mobility offered by contemporary vehicles by removing any restriction on their movement.

2.4.2 Localisation and Mapping

One of the key capabilities required for the use of situated mobile robots is localisation; determining the position and pose of the robot in some defined frame of reference. This is vital to enable autonomy [28] and supports such operations as path planning, sensing tasks, multi-robot coordination and combining multiple sensor readings. Localisation approaches rely on sensors to collect odometry information and/or make observations of the local environment. Some early work modified the environment using beacons in order to find a robot's position, however this is of limited use in many applications. Probabilistic localisation relative to a known starting position can be achieved by tracking the position of a moving robot, using techniques including Markov localisation and the Extended Kalman Filter (EKF) [139]. The ability of a robot to localise itself from an unknown starting position is particularly useful in many applications. This can be achieved using methods such as grid-based localisation [130] and the particle-filter based monte-carlo localisation [44].

In some applications prior maps of the environment are either unavailable or cannot be relied upon. Robots can attempt to construct a new map either as a

primary goal or to support other activities, such as path planning. Mapping involves collating sensor data into a useful representation from either a single robot, with multiple sensor readings from various locations, or a team of robots. Typical sensor devices include ranging, using ultrasonic or laser distance measurements, and visual, using digital cameras with computer vision processing.

2.4.2.1 Occupancy Grid Representations

Grids are used in robotics as both a convenient method of representing spatial data and to reduce complexity through abstraction. In particular, occupancy grids detail a robot's current belief about what occupies each grid location, allowing sensor measurements to be integrated over time [40]. Using probabilistic techniques, uncertain and noisy measurements can be included and used to improve the map. Occupancy grids have been used for collision avoidance [14], mapping indoor environments with range sensors, and representing the belief in the location of a missing person [148]. Comparing robot generated occupancy grids to a prior environment map is an additional method of localisation [125].

Occupancy grids also allow information gathered by multiple robots to be combined. Approaches to *map merging* include: computing a set of transformations and assessing the fit [21], and robots checking relative positions via observation when they meet to exchange information [45].

2.4.2.2 Simultaneous Localisation and Mapping

The Simultaneous Localisation and Mapping (SLAM) problem applies when a robot has neither knowledge of its position or a map of the environment. It is a fundamental area of mobile robotics research, where new approaches are continually proposed and refined. When performing SLAM a robot attempts to both map the environment and find its pose relative to this map using just observations of the world and the control inputs applied as it moves. A survey of the main approaches can be found in [138]. An extension to the standard problem is to allow multiple robots to share information

and hence improve both their maps and localisation performance. Approaches to multi-robot SLAM include, particle filters [65], information filters [140] and EKF's [158].

2.5 Aerial Multi Robot Systems

Recent technological advances in aerial robotics have enabled researchers to consider the possibility of deploying them in groups. One of the principal characteristics of aerial robots is the ability to travel at high speed irrespective of the ground terrain, thus greatly increasing the potential range. This makes them especially suited to tasks in large open outdoor environments, as found in several application areas, such as search and rescue and environmental monitoring.

These vehicles can be divided into two classes, fixed-wing and rotor-craft. Fixed-wing craft generate lift efficiently by continuously moving through the air at speed, enabling extended operating time. However, these kinematic restrictions require paths to be planned in advance, for example [78, 141] and the limited manoeuvrability reduces the suitability for some applications, for example positioning a sensor in a specified location for a continuous period would not be possible. In addition, there is a hybrid class of UAV that behaves like a powered glider, requiring forward movement to maintain altitude, but with a slow stall speed. Although manoeuvrability is restricted, a long flight time and large payload capacity makes for a good aerial photography and photogrammetry platform [136].

Although safety and ethical concerns exist about the use of this type of platform, its usage is becoming increasingly common. Regulations and guidelines regarding the operation of UAVs are being developed to enable their use in civilian applications; for example a recently published road-map by the Federal Aviation Administration in the USA [42].

2.5.1 Rotor-craft Vehicles

Rotor-craft vehicles have been increasingly miniaturised and improved in recent years, with multi-rotor configurations emerging as the most popular. Lift is generated by all of the rotors with changes in attitude (orientation), and hence movement, controlled via the associated motor currents. For example in a quad-rotor configuration (4 rotors arranged in a cross shape), tilting the craft can be performed by increasing the speed of one rotor and a corresponding decrease of the opposing rotor. Autonomous stabilisation is possible due to on-board sensors including accelerometers and gyroscopes. These enable the attitude to be continuously estimated and high frequency control procedures to apply the required rotor speeds to maintain stable flight.

Advantages of rotor-craft UAVs over fixed-wing include the ability to hover, operate near to ground level and even fly inside buildings. This enables sensors to be positioned and maintained in locations that would be impossible for fixed-wing vehicles. This capability is especially important when a camera is fixed in position on the UAV, enabling the required images to be captured, for example when tracking a moving target. With current technology the main disadvantages of rotor-craft UAVs are flight time and lift capacity, with both being dependent on the available on-board battery power. A full discussion of the engineering of these UAVs is provided in [81] covering the trade-offs in scale, weight capacity and acceleration. The authors suggest a total weight of approximately 1kg is required to enable full autonomy, with batteries contributing a significant proportion. Usable flight times of current electric rotorcraft are in the range of only 20 minutes, but this is increasing as new lightweight components and, more importantly, new battery technologies become available.

2.5.2 Coordinated Deployments

Teams of UAVs can in general be coordinated in a similar manner to ground-based robots, with many of the same methods being applied. In [84] a type of market-based task allocation is demonstrated, with auctions determining which robot is assigned

each task. This approach is further enhanced by adding a balancing coefficient, attempting to spread the workload equally. Similarly, negotiation based task allocation can be used with UAVs in a battlefield search task [132]. Improved performance is demonstrated when additional target information is shared between UAVs, especially with limited range sensor devices. Some of the particularly significant characteristics of coordinating multi-UAV systems are described in the remainder of this section.

2.5.2.1 Communication

In common with ground-based robot teams, communication is key to supporting the information sharing needed for coordinated deployments. Assuming reliable equipment, disruption in the wireless connectivity from UAV to UAV or UAV to base-station can be caused by either distance or obstacles in the transmission path. Due to the very highly mobile nature of these vehicles, especially in outdoor environments, constant communications are not possible unless particular attention is given to maintaining connections. However many of the approaches for UAV teams require consistent transfer of information to support coordination approaches.

Many authors attempt to mitigate communication difficulties by limiting the relative movement of UAVs, thereby maintaining a connected network. Teacy *et al* propose a method that learns the radio propagation characteristics of the particular environment during operation and use this information in selecting flight paths[135]. This allows UAVs to be moved to locations that will preserve network connectivity and routing capability. Similarly, a method for determining UAV trajectories to avoid breaking network connections is described in [30]. This method promotes clustering of UAVs to form cooperative teams whilst allowing mobility and separate clusters to merge. In [56] a prior map of the environment is used to calculate expected communication capabilities for different locations. The Voronoi method partitions the area around known obstacles allowing each region to be rated on connectivity quality. An alternative approach is to deploy UAVs primarily for the purpose of creating a communication link. Such *relay chains* enable a continuous high-bandwidth link to

be established, for example between a distant UAV providing surveillance cover and a base-station [102].

2.5.2.2 Shared Probabilistic Maps

A recent trend in multi-UAV teams is the use of shared probabilistic maps to guide behaviour. For instance, in search and rescue scenarios a prior map can be built using an assumed last known location and a model of potential travelling velocities of the missing person. An extension of this approach uses terrain gradients to further infer movement [86]. From such a representation robot motions can be planned to sense those areas of highest probability, for example with a single fixed-wing aircraft [85]. Similarly, paths for a small group of fixed-wing aircraft can be planned based on a shared probabilistic map, generated using the movement of a target drifting at sea [15]. This type of representation has also been used with a small number of rotor-craft UAVs in a flat ground environment with artificial obstacles [148]. A further example uses iso-probability curves that surround the last known location of a missing person to describe the belief in the current location of a missing person and deploy search robots [91]. These approaches in particular require careful consideration of the communication requirements, as the amount of data to be shared between team members could be significant.

2.5.2.3 Localisation

Spatial coordination of UAVs requires robust methods for determining robot location in three-dimensions. In outdoor environments, GPS is typically relied upon for localisation, often in common with barometric pressure sensors to improve altitude accuracy. This is certainly the case in open environments, where signals from many satellites will reach a UAV without obstruction. However in obstacle ridden or urban areas, GPS coverage cannot be assumed to be sufficient for acceptable positioning. Vision based localisation approaches have been applied to UAVs to handle such scenarios. Various visual SLAM approaches have been developed to both localise and

build maps during flight [150, 152]. Where GPS is intermittently available, it can be combined with vision based localisation to achieve better accuracy in locations that have poor satellite coverage [154]. Computer vision techniques can also be used to complete specialised maneuvers that otherwise may not be possible, such as automated landing using a dedicated optical flow sensor [82].

Indoor environments pose an altogether different challenge; GPS is almost certainly unavailable. Vision approaches have also been applied to this problem, monocular SLAM for example [23]. Other types of sensors may also be used, for instance, laser ranging is compared with stereo vision in [1]. These approaches tend to be specific to an indoor environment, taking advantage of the the structure provided by common visual features, such as the corners of doors. A recent trend in indoor aerial robotics uses motion capture systems to enable high precision (and impressive) UAV behaviours, in [37] for example. The University of Pennsylvania's GRASP lab has used this technology to demonstrate aggressive maneuvers with quadrotor helicopters [95] and subsequently extend these techniques to highly coordinated collaborative lifting [96]. While these results certainly are dramatic, they require a very high localisation accuracy (to within a few millimetres) often achieved with a full instrumentation of the environment, changing the research emphasis to primarily a control problem. In addition, these approaches are severely limited in their application, as very few real world application scenarios enable such an intensive preparation of the operating environment.

2.6 Summary

This review covers applications and approaches from the robotics literature, with particular focus on the use of multi-robot systems. Deploying groups of robots is being considered for a wide range of scenarios, especially to assist the work of humans or remove them from potential harm. Such systems aim to provide robustness, cooperation and increased performance by using multiple platforms. These goals can only be

achieved when appropriate coordination mechanisms are developed to take advantage of the available vehicles.

The primary focus of this thesis is addressing the use of multi-UAV teams for the WiSaR application. A previous approach to this scenario has concentrated on manually controlled UAVs, such as [51], with the major disadvantage of requiring at least one human controller per UAV used [27]. Autonomous aerial robots have also been considered using techniques to plan search paths through the environment, for example [85, 148]. In contrast, the approach taken here is for each autonomous UAV to define its own search behaviour using local information by selecting a region to search. Another common feature of the existing work is restricting the movement of robots to enable continuous communication links, examples include [135, 66]. In this work, this limitation is relaxed to allow robots the freedom to search outside of wireless communication range and hence take advantage of the high degree of mobility offered by UAV platforms.

Techniques developed in this thesis are inspired by work in the ground-based rescue robotics field, particularly frontier search [157], rendezvous meetings [122, 33] and robot relay configurations [32]. In the following chapters, novel approaches are described that build upon these ideas for multi-UAV teams in the WiSaR domain.

Chapter 3

Problem Description

The principal problem considered by this thesis is that of assisting wilderness search and rescue operations with groups of autonomous robots. This chapter describes the general problem domain and defines the goals, assumptions and performance metrics that will be used throughout.

3.1 Wilderness Search and Rescue

Wilderness environments can vary greatly depending on the local terrain and climate; even within a single region there can be any combination of sub-environments including mountains, rolling hills, wide open plains and lakes. Further variation occurs due to season, time of day and weather conditions. Their remote location can make access challenging, with many square kilometres of terrain not reachable via road. People often visit these areas for leisure activities including hillwalking, mountaineering/climbing and other extreme sports. In addition they may unintentionally come to be there, for example as the result of an emergency aircraft landing. Irrespective of the activity, there exists the possibility for people to become lost and disorientated or stranded due to injury or a medical episode.

Traditionally search and rescue operations have been carried out by teams of human rescuers. The main method of search is by human search teams walking across the terrain looking for any signs of the missing person such as discarded belongings (eg: clothing, food wrappings) or indications of human activity (eg: footprints, evidence of fires). In addition further specialised resources may be used including, off-road vehicles, trained search dogs and manned helicopters. Planning the search will use any available information about the missing person such as a general profile, planned activities and the last known location and time. For example: 35 year old male hiker, car left in a parking area, last seen at 5pm the previous day. Detailed local maps, UK Ordnance Survey for instance, provide a reference of the terrain, describing the relief (terrain heights), features (roads, tracks, rivers, etc) and type of ground (fields, forest, beach). In addition, human search teams have significant local knowledge gained from many years of experience operating in the specific environment. This can be applied to prioritise areas to visit or discount potential walking routes (for both the missing and searchers) due to impassible terrain.

A typical scenario following a reported missing person is as follows.

1. Report of missing person is received by emergency service handler
2. Local search team is notified, mission commander identified
3. Local search team convene at agreed meeting place; accessible by road but as near as possible to suspected missing person location
4. Mission commander makes a search plan based on available information and resources
5. Search begins.

Finding missing people in wilderness regions can be a very difficult task, especially when the main method is searching on foot. In addition to the large physical area to be covered, several factors can make it difficult to spot a prone person even at a short distance. These include,

- Vegetation (tall grass, bushes or trees)
- Variations in ground height (ditches or hidden ground depressions)
- Poor visibility (rain or fog)
- Line of sight obstructions (walls or farm buildings)

Furthermore a lack of information about the missing person (eg: last known position, planned walking route) increases the priority search area and hence search time. Human search efforts may need to be cancelled or postponed, usually due to increased dangers for the search team. Common causes include, darkness/nightfall, adverse weather conditions and dangerous terrain, for example deep snow, swamps or an overflowing river.

Studies have been performed to examine the behaviour of people who become lost in a wilderness environment. For example, in [61] the behaviour of missing people in

Alberta, Canada was examined; the distance traveled by these people was found to depend on their state of mind. Similarly, [113] collates data from many years of search and rescue situations in the UK. Information is provided on the category of missing person, their activity before becoming lost, any medical issues and where they were located. Using this knowledge some attempts have been made to predict the likely behaviour of lost people [90, 86]; the resultant models can then be used to inform a robotic search strategy [91].

3.1.1 Robot Assistance

Due to the large and unstructured nature of wilderness environments, the use of UAVs in particular has gained significant attention. The first attempts used small platforms carrying groundward facing cameras under direct and constant manual control [51]. The camera feed is wirelessly transmitted to a portable display unit, enabling human responders to gain an alternative view of the landscape via an “eye in the sky”. This provides an additional resource to the search commander, providing a view into difficult to reach areas and beyond line-of-sight obstacles such as walls, vegetation, hilltops and ridges. However, each of such vehicles require at least one human controller [19, 27]; some schemes use multiple personnel, for example a sensor operator to control and monitor an active camera, in addition to the pilot.

In order for UAVs to have an impact on search and rescue missions, the sensor data produced (mainly still images and video), must be automatically processed to detect missing people. The research in this area has mostly focused on applying people detection techniques to visual images from the on-board cameras of aerial robot platforms, for example [107, 133]. These approaches use various types of computer vision techniques to assess images for the presence of humans, for instance in [117] the shadows cast by upright people on the ground are used to classify images on the basis of whether they contain people or not. Furthermore, researchers have examined the use non-visual light sensing. In [48] imagery from a fixed wing UAV at a 60m flight altitude is used to identify both people and vehicles using both visual light and

infrared sensors. Lightweight thermal imaging sensor devices have recently become available, enabling a new direction for people detection research [123].

More recently, overall approaches have concentrated on increasing the level of robot autonomy and consequently reduce the human requirements for UAV assisted search, for example [148, 20]. Hence enabling multiple robots be launched by as little as one human operator and removing the requirement for constant manual control of each vehicle. This allows more advanced search strategies to be considered, for example allowing UAVs to operate out of communication range and even out of sight. Such strategies are not currently possible in the UK due to operating regulations, but such restrictions may not be permanent; authorities in the USA are currently investigating this issue [42].

3.2 Goals

The primary goals of the work in this thesis are:

1. Search designated areas. Robots must traverse the physical environment and using sensor devices gather information. The search should continue until the required area has been covered unless otherwise directed.
2. Transfer newly gathered information to a base-station, including search results and overall progress updates. Human responders located at the base-station will use this information to direct their own search effort and perform rescue operations.
3. Efficiently use the available robots. In order to complete the mission as quickly as possible the available resources should be used efficiently when meeting goals 1 and 2. In practice these goals may conflict and providing a suitable balance is a key issue for this thesis.

3.3 Assumptions and Scope

There are many facets to consider when designing, building and deploying practical multi-robot systems, many of which remain active research areas in their own right. This work considers each robot as an operational system with a set of certain capabilities. The primary attention of this research is on determining robot behaviour, including strategy selection, both overall and local, and providing organisational constructs such as assigning roles.

Below is a list of the main assumptions, including robot capability, the operating environment and type of task to be performed.

1. Robots are physically situated in the environment.
2. Robots are kinematically constrained, consistent with the vehicle type.
3. All robots are able to localise to a common frame of reference. This permits the merging of maps and other spatial data.
4. Robots have the capability to autonomously control their own motion relative to the local environment. This enables path planning and obstacle avoidance; a UAV can avoid the ground for example.
5. Communication is limited, both inter-robot and robot to base-station, due to physical separation in the environment. This is modelled by a range limitation and line of sight constraint.
6. UAVs are able to perceive the ground level while in flight. Sensing is performed from a restricted height above the ground and at a limited velocity.
7. Robots begin the scenario at the manned base station.
8. Where prior information about the environment is available, all robots will have this on-board at the start of operations.

9. Robots have an internal clock mechanism and are able to effectively plan to reach a location by a required time.

3.4 Performance Metrics

A set of definitions and metrics are needed in order to evaluate robot deployment strategies. These formulations describe the previously stated goals quantitatively and enable the comparison of approaches. Typical metrics used in multi-robot tasks can include: total area sensed, time to task completion, distance traveled and number of communication messages sent. The following metrics are specifically developed for the work of this thesis and include measures to both assess overall performance and the usage of available robots.

Overall, the metrics chosen follow the trends of the multi-robot search and exploration research field. Search progress is measured by the area that has been covered and is hence now known to the team; such metrics are used in [31] for example. Similarly, analysis of overall team performance is frequently examined through the time taken to search a given area, [124] for instance. In addition to these common metrics the concept of data latency, the delay between sensing and resultant data being available at base, is introduced. This metric is often used in the related field of Wireless Sensors Networks (WSN) where data is to be retrieved from physically remote deployed sensors, such as by a mobile vehicle [76].

3.4.1 Search Coverage

The primary goal requires robots to gather information by examining the scene at ground level with sensor devices. Furthermore this information can be shared between members of the team subject to communication availability. The team's combined information can be described in terms of the ground area explored.

$$A_{TEAM} = \bigcup_{i=1}^n A_i \quad (3.1)$$

where A_i denotes the ground area that robot i has information about and the team contains n robots. This measures the overall search progress and allows comparisons in the search rate of various approaches.

Any information gathered has maximum value when available to the human responders at the base station. The quantity of information about the target search area available at base at time t is written A_{BASE}^t . This metric is vital in assessing the overall effectiveness of the robot search team and its potential impact on the human-led search and rescue operation.

3.4.2 Team Performance

A simple method to evaluate search progression is to calculate a basic rate of information gain both by the team and at the base station.

$$InfoGainRate_{TEAM} = \frac{A_{TEAM}}{t} \quad (3.2)$$

$$InfoGainRate_{BASE} = \frac{A_{BASE}}{t} \quad (3.3)$$

where t is the current time step and the deployment began at $t = 0$. Similarly the search rate can be assessed for a particular time window, for example to contrast different approaches during the final stage of a search.

One of the goals of the deployment is to fully search a specified area; achieved by at least one robot passing a sensor over the required area. The time taken to complete this task can be written as follows.

$$CompletionTime = t \quad , \quad A_{TARGET} \subset A_{BASE}^t \quad (3.4)$$

where A_{BASE}^t denotes the area information known at the base station at time t and A_{TARGET} is the required search area.

Beyond the metrics that examine overall performance in the search task it is useful to measure the behaviour of the robot team. This allows increased understanding and fine-tuning of strategies. A typical set of basic measurements about each robot can include distance traveled and time spent on different behaviours (such as moving to a search area, returning to base, sensing and being idle). In addition, it is particularly important to look at how new information gained by sensing activity permeates through the team and to the base station.

Autonomous robots rely on information to guide future behaviour; in the multi-robot scenario data acquired through sensing has the greatest benefit when shared among the team. The information sharing metric measures the proportion of total team information that is known by each robot, at a particular time. It is defined,

$$InfoShareAvg = \frac{1}{n} \sum_{i=1}^n \frac{A_i}{A_{TEAM}} \times 100\% \quad (3.5)$$

for a team of n robots. This enables different team configurations operating on the same task to be compared in terms of sharing efficiency. For an efficient team $InfoShareAvg$ will track A_{TEAM} closely over time, demonstrating that new data is dispersing quickly between robots.

3.4.3 Data Latency

In time critical applications the latency from data gathering to availability, at a base station for instance, can be crucial in achieving a successful outcome. This delay for the information gained about a particular area, $A_{INTEREST}$ is described by,

$$SenseToBaseDelay(A_{INTEREST}) = t_{BASE} - t_{SENSE} \quad (3.6)$$

$$t_{BASE} = \min(t) \quad , \quad A_{INTEREST} \subset A_{BASE}^t \quad (3.7)$$

$$t_{SENSE} = \min(t) \quad , \quad A_{INTEREST} \subset A_{TEAM}^t \quad (3.8)$$

As the delays in delivering information are measured piecewise an overall picture of the team delay characteristics can be constructed. These include simple measures such as calculating average and maximum delays, and more complex analysis, for instance comparing the latency for target areas that are close to and distant from the base station.

3.5 Summary

This chapter sets out the problem of searching a wilderness area using a team of UAVs to achieve the goals of full search coverage and supplying newly gathered information to base. In addition, the robot team should operate efficiently to make use of the limited number of vehicles. The majority of assumptions made relate to the capabilities of the robot platforms; these are consistent with current and near-future technologies. Perhaps the most important of these is limited wireless communication, a very real issue that is not always considered in multi-robot systems research.

The described performance metrics are designed to capture elements of the multi-robot system directly related to the overall goals. Search progress is monitored by the area-information metrics, the transfer of new data to base by the data latency measurement and the efficiency of the team by the completion time and rate of information gain.

Chapter 4

Individual Robot Behaviour

Each robot must be capable of behaving autonomously and without continuous external control. It should use the available data to determine how to operate and maximise performance of both itself and the team as a whole. This chapter describes the individual robot behaviours and selection processes used during execution for the distributed sensing scenario. The novel frontier search approach for UAVs is introduced, which takes advantage of available terrain information to designate regions for search.

4.1 Tasks

To realise the goal of autonomous robotics each robot must, when disconnected from any base station or human controller, select its own behaviour based upon the best available information. This is especially important in the WiSaR scenario due to the large operating environment where communication links between vehicles and vehicle to base cannot be assumed to be reliable. In addition this requires each robot to regularly re-evaluate its actions, especially when new data becomes available. A common technique in the robotics literature is the representation of robot behaviour as selected, or assigned, tasks.

In this work each robot can perform one of two main functions at a time; either *search* or *relay*. Searching the environment for signs of missing persons is required to meet the primary objective of assisting search and rescue operations. This is the default behaviour for each UAV, to increase the information gathered by visiting more of the terrain. In this type of application, the human element is crucial, and hence robots should transport new information to the base station. This includes both general search progress and specific data such as captured images or video for inspection. The relay task requires UAVs to move physically to within communication range of base, this can take significant time depending on the terrain.

When selecting a new behaviour to perform, each robot defines a specific task with a full set of parameters. For example a search task is primarily defined by the area to be examined. Tasks are not necessarily always continued until completion; if new information causes a reexamination of the situation a robot can abandon and switch to a new task.

4.2 Area Sensing

The primary objective of the search task is to sense the environment by positioning mobile robots with sensing devices at various locations. The WiSaR scenario in particular contains defined large continuous regions of the ground to be searched. An

important assumption in this work states that these spaces can be discretised; data stored for a single location represents a real world dimensioned space. For convenience this work, as with many others, uses a regular grid, but all of the approaches discussed can be implemented with other types of spatial representation.

Using a suitable scale, each grid location can be completely perceived by a UAV with a downward facing camera from a low altitude. Figure 4.1 describes this scenario for a small quad-rotor with a low weight camera. Operating from a low altitude with a stable platform allows a high quality of sensory information to be gathered. In addition, the maneuverability of such vehicles enables a view near to and even underneath obstacles such as trees and buildings. The trade-off is a reduced amount of ground coverage achieved from each single location sensor capture.




Figure 4.1: UAV sensing configuration. Each UAV carries a downward facing camera in order to sense the environment on the ground surface. A spatial grid is constructed such that a single location can be fully perceived from a fixed *search height* above ground level.

4.2.1 Representing Spatial Information

In this work the environment is represented as a regular grid, with the dimensions of each square location defining the resolution of information stored. Using multiple

grids allows different types of geographical data to be stored and aligned. This relies heavily on the perfect robot localisation assumption specified in Chapter 3 to correctly combine information from varying sources with correctly geolocated data sets. Using layered grids to represent spatial information provides a convenient method for sharing, combining and referencing multiple data sources. This is particularly suited to a cooperative multi-robot scenario and especially important for the frontier search methodology. An overview of the layers of spatial data is shown in Figure 4.2 for a multi-UAV team.

All robots are provided with information about the terrain at initialisation. The most important feature of this data is that it describes the ground height throughout the operating environment and is a vital input for the frontier search algorithm described in Section 4.2.2. This type of data is typically presented in survey maps used by those undertaking activities in wilderness environments, such as national parks or nature reserves. For example in the UK the Ordnance Survey publishes maps at different scales for the entire country. Furthermore, electronic versions of this data has become easily available. Ground height data is typically produced via surveys carried out by low flying aircraft equipped with distance ranging devices.

Through pre-processing this complex data is transformed into a base terrain data set for use with the UAV based sensing technique described in Section 4.2 and the frontier search approach in Section 4.2.2. Firstly, if required, the resolution is altered to align the dimensions of each grid location with the sensing coverage of a single UAV at minimum search height. Depending on the source data format this may require sampling at an appropriate spacing or interpolation to generate the required grid. As a simplification step the ground height values are discretised to reduce the complexity of planning the UAV search. The granularity of this process should be tailored to the flight dynamics of the vehicle type, promoting significant altitude changes that will not become so large as to reduce sensor coverage beyond acceptable levels. As an example, the experimental work in this thesis targets a 10m resolution grid (each square 10m by 10m) with discretised ground height data of a 20m step size.

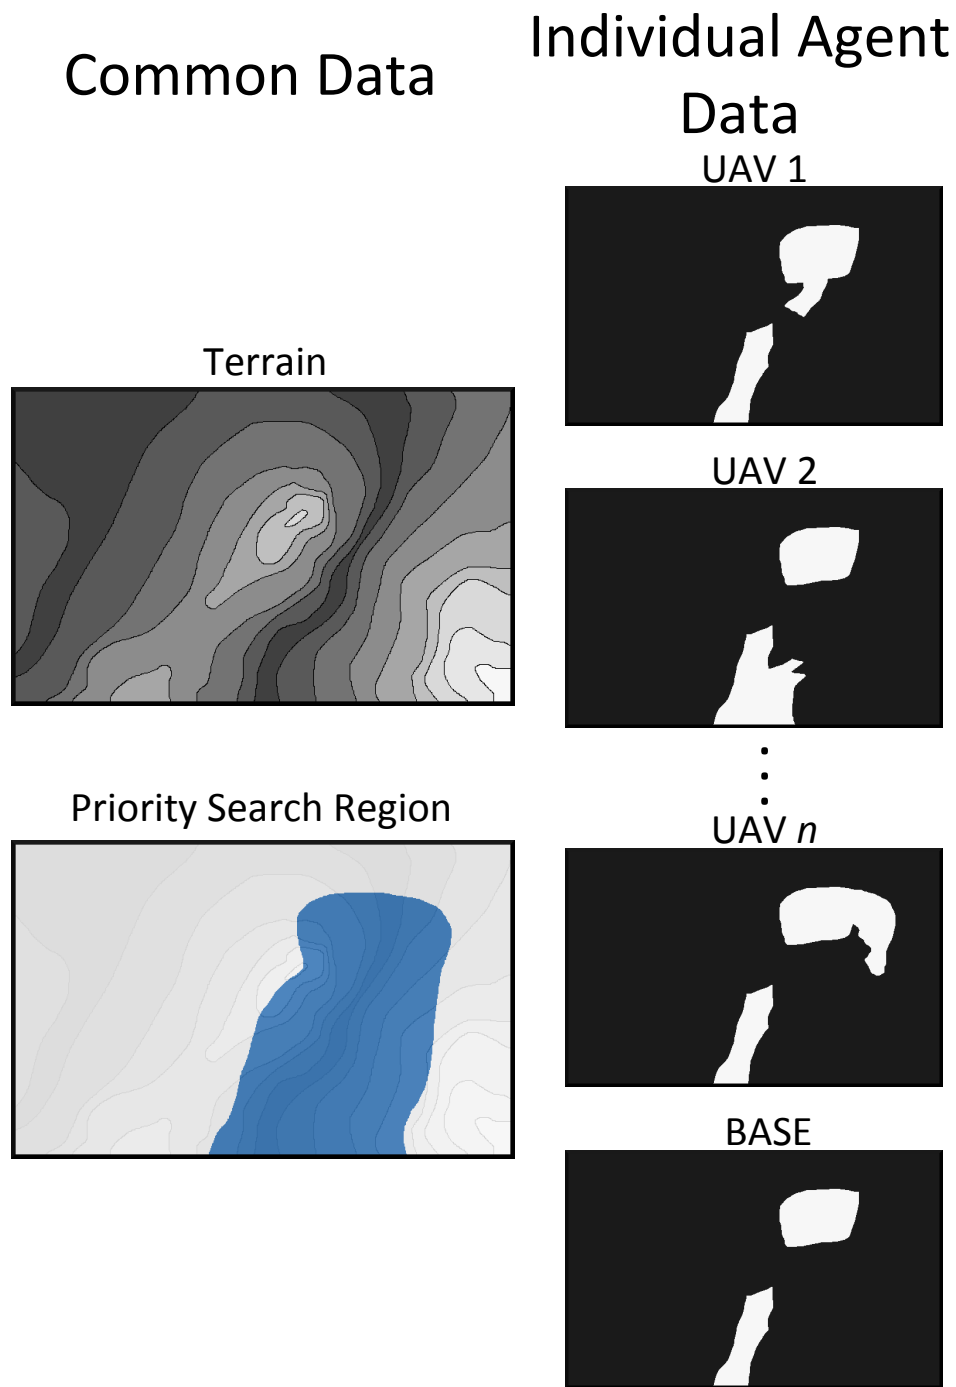


Figure 4.2: A snapshot of spatial information layers during search. Terrain data describes the ground height of the operating environment, lighter shading indicates higher ground. The priority search region provided by the mission commander is shown in blue. Each active agent maintains an individual map of known searched (white) and unsearched (black) areas.

Capturing additional data can be achieved by adding further layers to the spatial data set; particularly useful for applications that require human-machine interactions. In a search and rescue task the mission commander uses various information sources to formulate a search plan. This can include information including a missing person's last known location, type of activity or a known walking route. A key part of the search plan is the deployment of different resources to different parts of the landscape. Priority search regions selected by the mission commander for the UAV team can be captured by a single layer. This may describe a target region that may be considered the most likely location of the missing person or be especially challenging for other actors (human searchers) to reach. The form this takes can be as simple as a binary grid (priority/normal) or multiple distinct regions labelled along with an associated importance score.

Each UAV maintains a current view of which parts of the area have been searched. This is updated whilst the UAV is actively searching and when other information is available via sharing. Similarly, the base station must keep a general map of visited and unvisited areas, updated only when robots with new data have a direct communication link. This grid takes the basic binary form; each location denoting *searched* or *unsearched*. Greater complexity can be added to preserve the identity of the data's origin, either this UAV or another agent, and to store the time at which an area was searched. Furthermore, probabilistic modelling of the sensing task can be used without altering the overall approach; the frontier search mechanism can be updated to include such a notion.

4.2.2 Frontier Search

The frontier exploration method uses a map of known visited space to select new *frontiers* to explore. Each potential frontier is a defined area on the edge of a previously visited region. A mobile robot can select a frontier to visit using a suitable decision system such as a simple heuristic.

This technique originates from the use of ground-based robots using distance

sensors to explore, map and search unknown areas [157]. Figure 4.3 shows an example progression for a ground-based robot performing the frontier exploration method. Frontiers are generated by examining the border of known territory and segmenting the edge according to obstacles encountered. For example, a frontier would exist in an open doorway that has yet to be visited. A robot is then tasked to visit the central location of the frontier and gather data using a sweep of its ranging sensor; new information is then incorporated into the robot's map. This cycle repeats until the area is fully explored or other goal criteria reached.

This section describes the frontier approach adapted for low altitude aerial robots searching in outdoor environments. It relies heavily on the availability of prior information about the terrain, typically ground height data, in order to generate potential frontiers. Each UAV uses its currently available information to select new frontiers to search.

4.2.3 Frontier Generation

For this application area potential search frontiers are created as *regions* instead of the edge sections of boundary between visited and unvisited space. Each UAV generates frontiers using its currently held information regarding the searched area, in combination with terrain data provided prior to the mission start, see Figure 4.4. Spatial data is overlaid and combined to describe a previously unsearched space with terrain data.

A potential frontier region should,

1. Cover previously unvisited space; according to the UAV's current knowledge.
2. Border the known visited area.
3. Be searchable by a single UAV from a constant altitude; chosen to enable a sufficient perception quality and resolution for the task using the on-board sensor device.

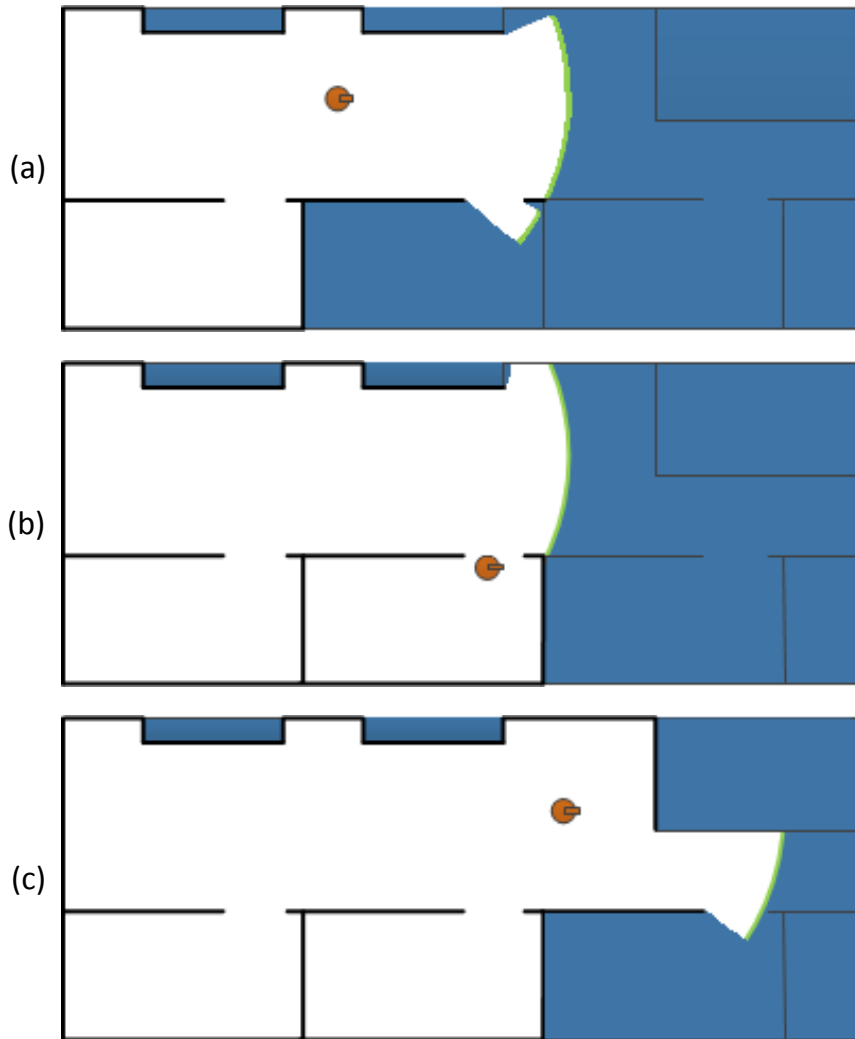


Figure 4.3: The frontier exploration approach for ground-based robots equipped with 360 degree ranging sensors. The robot is shown in orange, explored space in white, unexplored in blue, obstacles with a black line (thicker if discovered) and frontiers in green. (a) After sensing, two separate frontiers are available, the nearest one is chosen. (b) Sensing provides coverage for the side room, revealing no new frontiers, the only remaining frontier is selected. (c) Through sensing more of the environment is understood, a new area is now available for exploration.

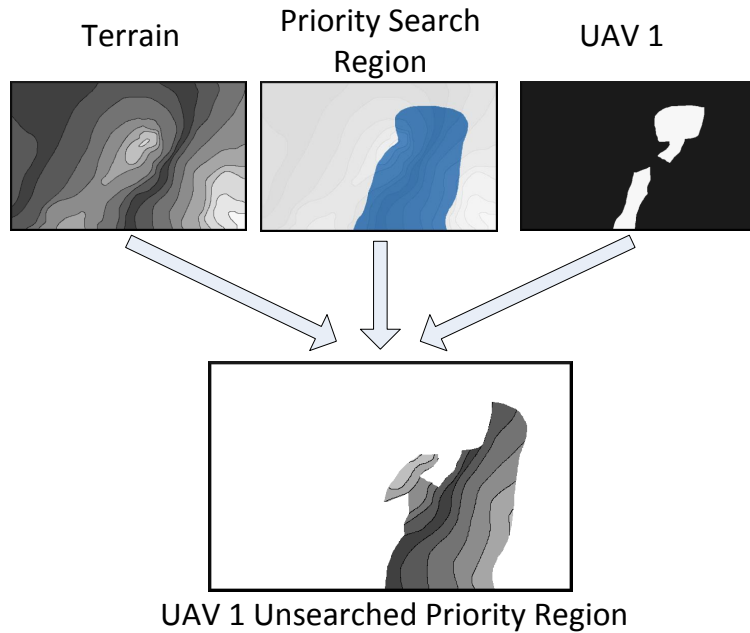


Figure 4.4: Data sources used during frontier generation to create a search region. A base terrain map is masked by the priority search region (include) and the UAV’s individual search record (exclude).

The generation of potential frontiers is required under two different sets of constraints; search within a set priority region and once the priority search area has been completely covered. Figure 4.5 illustrates both scenarios from the perspective of a single UAV. The procedure is also fully detailed in Algorithm 4.1.

When generating frontiers within a priority region the aim is to cover the unvisited space with new searchable frontiers. Starting in an arbitrary location, a region growing procedure groups cells (grid locations) according to the terrain at this position. All traversable adjacent cells with the same discretised ground height are eligible to become part of the same frontier. These neighbouring locations are added to the candidate queue, from which a new location is selected for inclusion during each iteration of the region growing procedure - see Algorithm 4.1. This process continues until the maximum size for a single frontier is reached or no further suitable candidate cells are available. The order in which cells are considered, either as frontier start points or for inclusion in a growing frontier, is randomised by shuffling the relevant

Algorithm 4.1 Generation of Search Frontiers

```
function GENERATEPOSSIBLEFRONTIERS(UAV uav, Region r, Integer MaxSize)
  Map map = getVehicleMap(uav)
  List unsearchedLocations
  if map.regionFullySearched(r) then
    unsearchedLocations = map.getBoundaryLocations()
  else
    unsearchedLocations = map.getUnsearchedLocationsInRegion(r)
  end if
  unsearchedLocations.shuffle()
  for all loc  $\in$  unsearchedLocations do
    if !map.partOfFrontier(loc) then
      Frontier newFrontier = ExpandNewFrontier(loc, map, MaxSize)
      possibleFrontiers.add(newFrontier)
    end if
  end for
end function
```

```
function EXPANDNEWFRONTIER(Location start, Map map, Integer MaxSize)
  List frontierCells
  Queue candidateCells
  Float frontierTerrainHeight = map.getTerrainHeight(start)
  Integer frontierSize = 0
  Location locationPtr = start
  while frontierSize  $\leq$  MaxSize & candidateCells.notEmpty do
    if map.getTerrainHeight(locationPtr) = frontierTerrainHeight then
      if !map.partOfFrontier(locationPtr) then
        frontierCells.add(start)
        frontierSize = frontierSize + 1
        map.markAsFrontier(locationPtr)
        List neighbours = map.getNeighbourLocations(locationPtr)
        neighbours.shuffle()
        candidateCells.add(neighbours)
      end if
    end if
    locationPtr = candidateCells.next
  end while
  return frontierCells
end function
```

lists. Therefore, the process is non-deterministic and no two executions will produce the same frontiers; ensuring there is no hidden bias for creating a particular type of search region.

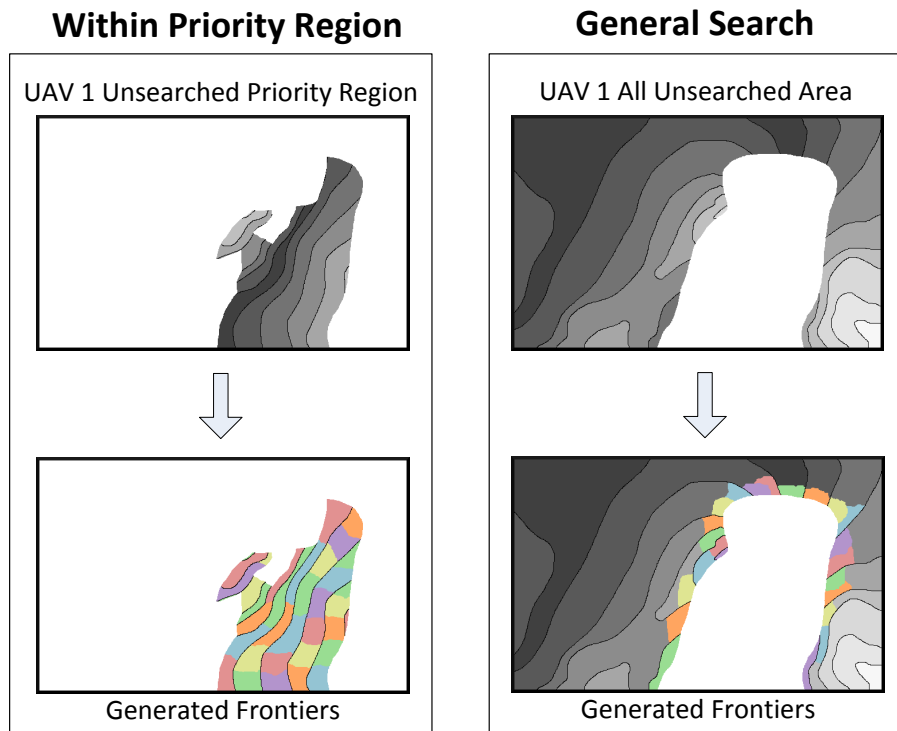


Figure 4.5: Frontier generation when searching within a priority region and in a general setting. Generated frontiers are shaded randomly to illustrate size and position. Each frontier is limited by area to restrict maximum size.

Without the limitations of a priority region, frontiers are generated adjacent to previously visited space that meet the required criteria. Firstly, the boundary between searched and unsearched space is detected, creating an edge. This provides the starting point for the region growing procedure. Beginning with any grid location that falls upon the boundary, a frontier region is grown according to the ground height description on the prior terrain map. This process continues until all boundary locations are part of a frontier.

Under particular circumstances some of the frontiers produced may have undesirable properties. For example very large frontiers would take a considerable amount

of time to search and not be as amenable to a coordinated team approach. This would occur when the terrain is particularly flat, placing little constraint on region growth. Selecting a suitable threshold for the maximum frontier size addresses this issue. It can be determined through experimentation and varied to suit specific scenarios. Hence a large continuous plateau in the terrain is covered by multiple frontiers instead of being considered as a single one. This type of constraint is present in the sample frontiers shown in Figure 4.6. This figure shows a range of examples of frontier generation results using the procedure described, both inside a priority region and in a general setting using different terrains. The shape of the created frontiers is affected by the underlying terrain. On a flat surface, the regions tend to be more regular as the region growing process will often be unrestricted by terrain and cease at the maximum size limit. Over steeper terrain, the frontiers generated will typically be constrained by changes in the ground height and hence their shape will reflect the underlying landscape.

Unrestricted, the frontier generation procedure can result in a significant number of potential frontiers. Despite providing more options to the UAV it can create a computational challenge; especially with limited on-board capability. The region growing process can be streamlined using knowledge of the robot's preferences for frontier selection. For example, with a greedy selection process using a simple heuristic estimating the distance to the frontier, frontier generation should begin with the creation of frontiers near to the current UAV location and progressively examine areas further away.

4.2.4 Frontier Selection

Upon becoming idle, on completing a previous task for example, each autonomous UAV must choose a new frontier to search. Using the output of the frontier generation process, each option must be assessed and compared. Operating individually a standard technique called greedy search allows a robot to select its next task according to a rating system. A heuristic scoring mechanism can reduce the complexity of




Figure 4.6: Sample results from the frontier generation process: (a) and (b) within an elliptical priority region; (c) and (d) on a relatively flat terrain; and (e) and (f), on a steep and highly variable landscape. Searched space is shown in white, unsearched in black and frontiers in bright colours. The white area in the upper left corner is a zone surrounding the base station that is not required to be searched.

performing a full analysis of each competing option by taking simple measurements. The applicable heuristics for a single UAV acting alone include:

- **Distance from current location to centre of frontier region**

A general estimation of how distant an entire frontier is to the current robot position; minimised to completely search a frontier in the shortest time.

- **Distance from current location to nearest frontier point**

A measurement of the separation between the robot and the nearest location at which new information can be gained; minimised to promote rapid exploration by minimising time to collect new data.

- **Size of frontier**

The area that a UAV can search without excessive maneuvering, such as changing altitude. This can also be considered an estimation of the amount of new information to be gained from a frontier.

- **Distance from frontier to base station**

Frontiers close to the base station will allow the new information to be relayed quickly after acquisition.

Furthermore, such heuristics can be combined to influence robot behaviour. For example, preferring large frontiers that are near to the base station, to enable the quick relay of new data once the search is complete. These combinations can be implemented simply by taking a weighted sum of multiple estimations.

4.2.4.1 Altitude Change Weighting

Currently the most common type of UAV platform being considered for multi-robot sensing operations is the multi-rotor (mostly quadrotor) helicopter. For rotorcraft, unlike fixed wing aircraft, increasing altitude is always accomplished using additional thrust, and hence energy [63]. However for any practical deployment the exact power consumption characteristics will depend on the vehicle used and in particular payload

carried. The addition of sensing devices and computing hardware can also reduce the expected flight time on these vehicles. The impact of power consumption and flight time will depend on the application area; greater for long term deployments such as surveillance and less for relatively short emergency operations.

The frontier search approach relies on the premise that areas need to be searched from a certain height range above ground level. Along with energy, making changes in altitude also uses valuable flight and searching time, so should be limited to necessary occurrences only. Therefore it is desirable for UAVs to have a preference for maintaining current altitude when possible.

For this work, the preference for not changing altitude will be modeled by a weighting parameter, ACW , and is applied as follows:

$$MovementCost(A, B) = Cost(A_x, B_x) + Cost(A_y, B_y) + ACW * Cost(A_z, B_z)$$

$$ACW \geq 1 \tag{4.1}$$

for given locations A and B , each with coordinates $\langle x, y, z \rangle$, and where $Cost(p, q)$ is the cost of moving between points p and q along a common axis. This defines the relationship between the agent's preference for moving laterally and vertically. Although descending does not incur the same energy cost as ascending, the weighting is symmetrical, due to the expected future cost of both energy and time when later regaining height. This weighting is used in planning future behaviour, principally in the selection of frontiers. For instance when calculating a distance measurement heuristic any required change in flight altitude will be penalised if $ACW > 1$. Values of the parameter are evaluated experimentally via simulation in Section 6.6.2. For practical deployments where energy consumption is a critical factor, measurement and modeling of the power required to increase altitude for the specific vehicle would allow this relationship to be tailored to the situation.

4.2.4.2 Within the Team Environment

All team members, robots and base station, are only able to communicate when permitted by wireless radio capability. Typically in this outdoor scenario the primary limitation is the distance between participants, although ground topography can also be significant. When a connection is available agents periodically share their current maps and hence update their respective knowledge describing the previously searched area. As each UAV reaches a point where a new frontier is required this updated information reduces the potential search area and helps to lesson instances of redundant multiple coverage (locations searched by several UAVs). Through this information sharing the team is able to operate in a basic cooperative manner without explicit organisation of their behaviour.

This straightforward approach has two obvious issues that would hinder overall team performance,

- **Undetermined time period between data sharing**

This will reduce the overall effectiveness and efficiency of the team in covering the required search area; especially in large open environments where UAVs may operate for long periods without a random encounter.

- **No coordination on current search behaviour**

Without any further cooperation UAVs may needlessly search the same areas in between data sharing events. In the extreme case, two nearby agents may hold the same knowledge base after merging maps and both select (almost) identical search frontiers.

These issues need to be addressed to create an efficient search team; approaches to improve on this basic information sharing approach are discussed in Chapter 5.

4.2.5 Frontier Coverage

Once a UAV has selected a frontier, the area it covers must be searched. The precise method used by an agent to perform this sensing task is of low importance for the overall approach. As frontier regions can be of an arbitrary shape and size, it is unlikely that one search trajectory technique will be optimal for covering all potential regions¹. In this work, including the simulation implementation in Chapter 6, frontiers are searched starting with the outside edge and working towards the centre, in a spiral-like pattern.

Firstly, the UAV calculates the point within the selected frontier that is nearest to its current location; planning a flight path to reach this point at the appropriate sensing altitude. The UAV then traces the edge of the region, moving to an adjacent unsearched location that lies on the frontier boundary. If no suitable location exists, it simply finds the closest unsearched location within the frontier region, moves there and continues the process. Once the frontier is completed, the UAV will therefore often be well inside the boundary of the region (perhaps near the centre of a approximately circular region for example); at this point a new frontier will be selected for search.

4.3 Relaying

To meet the objective of providing information to a fixed base station, UAVs perform a specific relay behaviour. Due to the large physical space of the operating environment, this involves agents moving sufficiently close to base to enable a direct wireless link to be established. An autonomous agent must decide when it appropriate to select a relay task instead of continuing the default search behaviour with a new frontier.

An initial approach to consider is imposing a hard limit on data gathering performed by each UAV before returning to base. This can take different forms including,

¹Investigating search trajectories for arbitrary regions is beyond the scope of this thesis and an interesting direction for future work

- **Total new information limit**

Each agent should monitor the amount of new information it has individually collected and relay once a threshold is achieved. In the grid-based WiSaR scenario this involves simply counting the number of new grid locations searched between relay tasks.

- **Limited search time**

A UAV should return to base immediately once a set amount of time has passed since last being in contact with base.

While this relay method does ensure regular updates are received at the base station, depending on threshold levels, it takes no advantage of the multi-robot team. A significant duplication of relay effort is to be expected, as UAVs having shared their information among the team will then all proceed to relay this same data.

4.3.1 Opportunistic Altruistic Relay

It is desirable to reduce such duplication of relay effort caused by UAVs operating fully independently. A simple approach to improve the team's performance is to enable agents to pass responsibility for new data when sharing takes place. This alters the agent's behaviour from purely self-interest, searching and relaying their own data exclusively, towards a more team centred approach. In this *altruistic relay* technique a UAV will accept the burden of responsibility for transporting a certain quantity of data to base in order to free a different UAV to gain new information more quickly.

After nearby UAVs have shared information, a short negotiation protocol can allow one to accept the responsibility for the combined new data. Again heuristics are used in selecting an agent for this task; for example the nearest UAV to base station assumes responsibility for ferrying the new information to base. The non-responsible UAV is then free to continue searching without the restriction of relaying previously acquired data. Both UAVs will continue to store all information on-board to mitigate against the failure of either platform. As the event of map sharing is

unplanned, this relay technique also only occurs opportunistically when two UAV happen to come within communication range. Therefore the impact on overall team performance may be limited; this is assessed through simulation in Chapter 6.

4.4 Behaviour Transitions

The UAV's default behaviour is to gain new information by searching frontiers, only selecting a relay task when required. Figure 4.7 captures the agent's overall behaviour and this transition between search and relay activities. Searching a frontier requires the UAV to firstly move to a location at which sensing can be performed, including the required height above ground level, and secondly, perceive the environment using sensors. This is repeated until the entire frontier has been completely covered. This cycle is broken when the conditions for beginning a relay activity are met; for example when the threshold for storing new information has been reached. A relay task then takes full priority until it has been completed; with the new data having been passed onto another agent, either a UAV or the base station depending on relay method used. UAV behaviour is also influenced by other team members. For example, if new information received from another agent indicates that the currently selected frontier has already been searched, the generation of new frontiers will be triggered. The mission is complete when the agent's knowledge base shows that full coverage of the target area has been achieved; the UAV will then return to base.

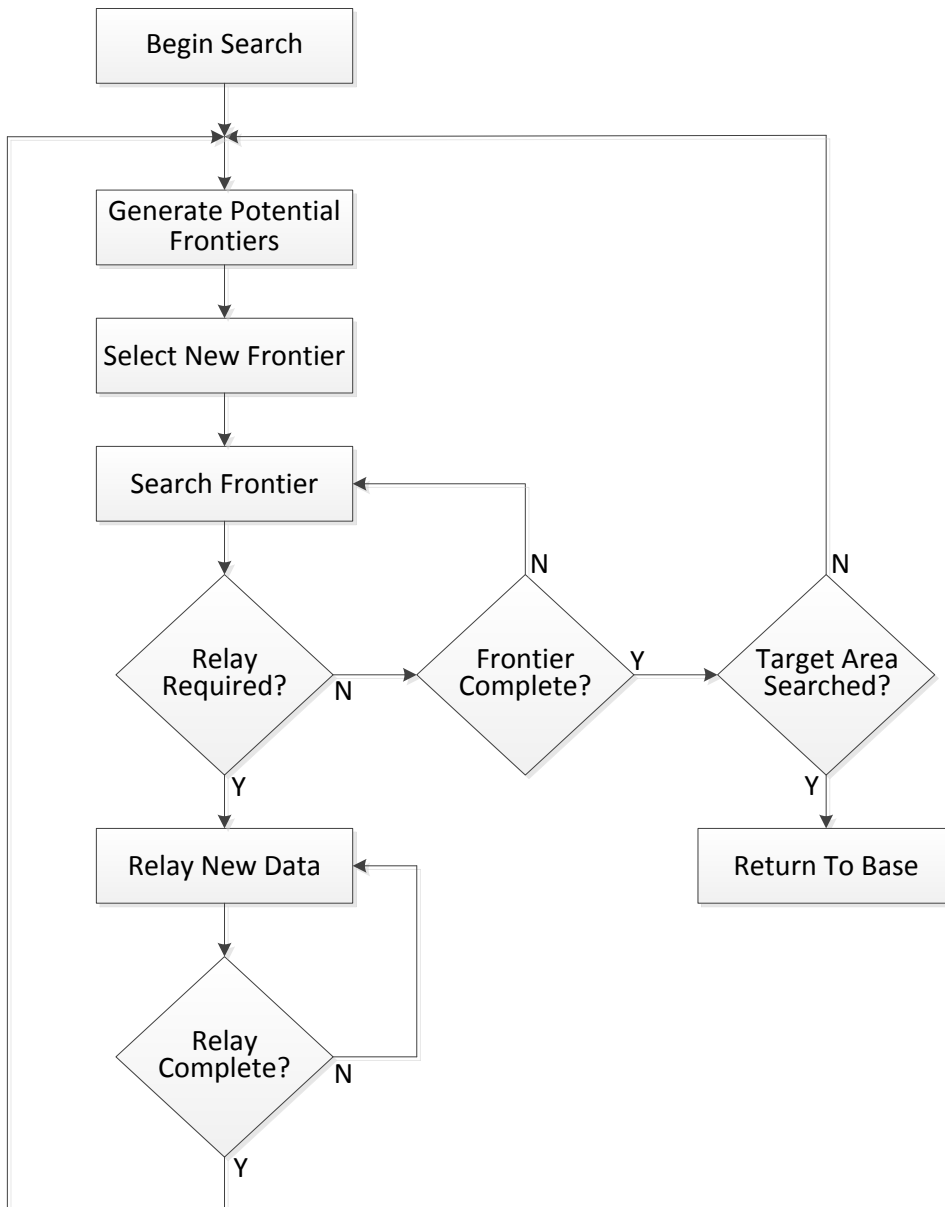


Figure 4.7: The behaviour and transitions between tasks for a UAV using frontier-search.

Chapter 5

Team Organisation

The robot team must operate efficiently to provide maximum benefit to the search effort. This chapter describes techniques to organise the robot team to achieve cooperation on both search and relay goals. The novel approaches presented include: coordinated frontier search, team rendezvous and adaptive data latency driven relay. The methods use a combination of coordinated activities and individual autonomy to create balanced deployments.

5.1 Search Team Organisation

With only a limited number of robots available for the deployment it is important to extract the maximum benefit from the autonomous search team. In order for the team to operate efficiently the general approach taken is to maximise cooperation towards the main goals of search and relay. The methods described here aim to achieve this through reducing the duplication of robot effort that can occur in multi-robot systems.

Although cooperative, these approaches still seek to preserve the autonomy of deployed robots without applying fully centralised control or long term planning. This should enable sufficient flexibility to respond to the changing task during operation, making decisions based on the latest available information. For instance, the search methodology does not impose restrictions on UAV movement for continuous communication links, instead prioritising search freedom.

In this chapter the behaviour of a single team is discussed, however in general a large available group of robots may well be split into smaller teams. For example, in such a scenario, each team being assigned a distinct region on which their efforts should be prioritised. The scalability of team size and consequent effects on efficiency are examined through simulation in Chapter 6.

5.1.1 Information Flow

The approaches in this section rely on information available to each autonomous robot at the point a decision is made. An overview of the general pattern of information flow is shown in Figure 5.1. Newly acquired information is transferred between searching UAVs when wireless communication permits. A relay robot, when used, can gather this new data from multiple searchers. Relaying robots will then physically move to within range of the fixed base station's location and provide a full payload of fresh information from the deployed robot team. This figure only shows the primary information flows of new information, implementations will include multiple instances of less significant data transfer.




Figure 5.1: An overview of the primary information flows between deployed robots engaged in searching (S) and relaying (R) and the base station (B). The size of arrows gives a general indication of the data quantity.

5.2 Coordinated Frontier Search

The frontier search technique described in Section 4.2.2 detailed a method for both generating and selecting areas to search. Heuristics were proposed for a UAV to choose where the search should be continued for a single lone agent; with no consideration given to other team members. In a team environment a significant amount of redundant repeat coverage would be expected as multiple agents will often simultaneously select similar areas for search; hence reducing overall team efficiency. This is despite a basic level of cooperative behaviour via the sharing of coverage maps between agents. Further coordination methods can be applied to the team, attempting to reduce redundant coverage.

Many coordination methods use a negotiation or bidding process to assign tasks to agents, market-based task allocation for example. This requires all participants to communicate with each other or a central entity to, for instance, place bids or accept a task. Due to the large operating environment, communication between UAVs cannot be assumed with sufficient regularity to enable this type of coordinated approach. Instead of restricting UAV movement to ensure a continuously connected network, the approach taken here uses a minimal amount of inter-robot communication.

The implicitly coordinated frontier approach is illustrated in Figure 5.2. Each

UAV continues to behave with complete autonomy, generating potential search frontiers from its individual spatial data. However, when selecting the next search task the presence and intention of nearby team members is taken into account. This requires additional information beyond the spatial data each agent maintains. Each UAV should periodically broadcast both a current location and search intention. Rather than transmitting a full description of a targeted search region, a simple bounding box uses a very small data packet to convey an approximation of the same information. Only team members within communication range will receive this broadcast and hence be able to utilise the information. Whilst this does limit the opportunity for UAVs to take other’s behaviour into account, spatial coordination between those in close proximity should have a greater effect.




Figure 5.2: An overlay of known locations and operating extents of team members on potential search frontiers. The UAV in question (shown larger) uses this combined data to select its next search task. Note, the bounding boxes of other UAVs need not be aligned with the displayed frontiers, as each robot calculates and select its own potential search frontiers, using an individual map.

The additional information regarding nearby team members is utilised during the frontier selection process described in Algorithm 5.1. Firstly, potential frontiers are generated using the procedure described in Section 4.2.2. These are then pruned to remove any that intersect with the declared operating areas of nearby searchers. The

remaining frontiers are evaluated and sorted using the chosen scoring method, such as a heuristic cost function. As an additional step the UAV considers the highest rated frontier from the perspective of other agents. Using the same scoring function the frontier will only be selected if this UAV achieves a score greater or equal to all other known team members. Although rare, it is possible that this will lead to no frontier being selected; for example, when a UAV is completely surrounded by other nearby team members. If this occurs the highest scoring non-overlapping frontier is selected, ensuring the agent does not remain idle. By simply moving, the result of the agent’s frontier process will likely provide a different result at its next execution.

Algorithm 5.1 Frontier Selection

```

function SELECTNEXTFRONTIER(UAV uav, Region r, TeamMembers team)
  Map map = getVehicleMap(uav)
  List possibleFrontiers = generatePossibleFrontiers(map, r)
  for all pf  $\in$  possibleFrontiers do
    if overlap(pf, team) then
      possibleFrontiers.remove(pf)
    else
      pf.score = calculateScore(pf)
    end if
  end for
  List sortedFrontiers = sortMaxScore(possibleFrontiers)
  Frontier highestScoringFrontier = sortedFrontiers.firstItem
  for all sf  $\in$  sortedFrontiers do
    if sf.score < maxTeamMemberScore(sf, team) then
      uav.assignFrontier(sf)
      return
    end if
  end for
  if uav.noFrontierAssigned then
    uav.assignFrontier(highestScoringFrontier)
  end if
end function

```

This approach is designed to avoid redundant multiple coverage of the same area and in particular discourage clustering of search robots. By assessing frontiers from the perspective of nearby team members, a UAV can approach an estimation of the likely behaviour of other autonomous agents. This relies on the assumption that

robots in the same vicinity will have similar, but not necessarily identical, knowledge of the searched and unsearched regions via data sharing activities.

All team members should continuously listen for and keep record of received updates from the team; making this information available whenever a new frontier is required. However it is important to only rely upon recent updates; older data may now be incorrect. This can be implemented simply via a freshness threshold, where data received too long ago is simply ignored. Furthermore, an update received from a team member can cause replanning. Any incoming message that states another UAV is intending to search a region that intersects with the currently selected frontier will require a new frontier to be selected. The full generation and selection procedure is performed to ensure the most recent spatial and team data is used.

When scoring potential frontiers all of the heuristics described in Section 4.2.4 are still applicable. In addition, data from nearby team members can be used to create both new and variant grading criteria, for instance:

- **Average distance from centre of frontier region to known UAV locations**

An estimation of how distant an entire frontier is from other team members; maximise to encourage team dispersal and/or limit to reduce team spread.

- **Shortest distance from frontier boundary point to other UAV search extents**

A measurement of the separation between the considered frontier and the nearest stated search area among the team; maximise to encourage search in differing parts of the the environment.

5.3 Rendezvous

For any group of robots to behave as a team, some method of communication must be present; in this case the only option is directly between agents¹. Due to the limitations of the predominant communication technology, wireless radio devices, the

¹In other applications, communication can take differing forms, *stigmergy* for example uses communication via the environment [64]

relative location of agents within the environment greatly impacts their ability to communicate. In this particular scenario the primary concern is signal attenuation, typically caused by inter-robot distance and obstacles present in the path. Reliable communication sufficient for data sharing can only be assumed over relatively short distances with a clear line of sight.

An organised meeting, or *rendezvous*, between UAVs provide an opportunity to share data and hence cooperate. Each rendezvous specification consists of both a predetermined location and time (or time remaining until the meeting). All members of the team should attempt to reach within a short distance of the rendezvous point by the specified time; sufficiently near to communicate directly with all other robots present.

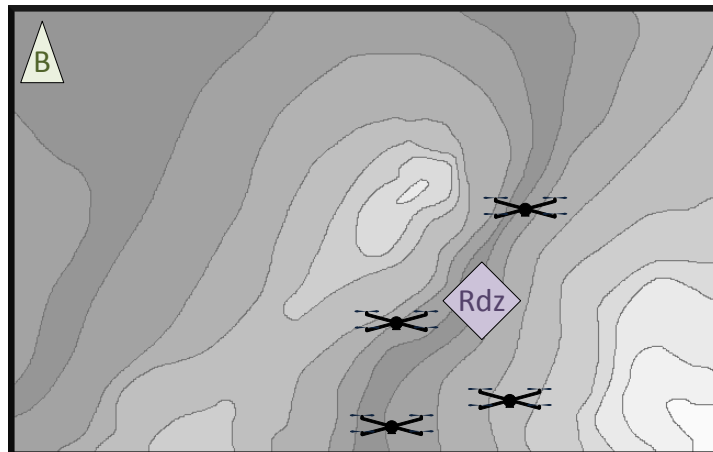


Figure 5.3: An overview of a search operation in progress, showing the base station B , next rendezvous location Rdz and current UAV positions.

This technique allows UAVs to utilise their mobility to rapidly and flexibly search between rendezvous, while ensuring regular contact to promote cooperation. A common alternative method, especially for aerial vehicles, is limiting movement to maintain continuous communication links. Although this does have the advantage of enabling new information to quickly disperse, it fails to fully take advantage of the excellent range of motion capability available with recent UAVs.

The rendezvous approach has been previously demonstrated with ground-based

robots during exploration type tasks, see for example [122, 33]. However, these works consider each meeting to be between a paired explorer and relay, for the single purpose of retrieving new data to base. Here, the concept is applied to a team of multiple robots, all visiting the specified meeting. This enables data sharing and coordination of both search and relay activities. It is an ideal opportunity for a specific relay robot to collect data from a deployed search team - as discussed further in Section 5.4.

Continuing the overall approach of making decisions during operation, the details of each rendezvous are chosen on-the-fly rather than pre-planned before deployment. The specification of a rendezvous is agreed only at the previous meeting. Meetings can therefore be tailored to the current situation using the most recent information. Both the timing and location of the following rendezvous are selected by the UAV team without direction from a base station.

When used in combination with frontier search, the rendezvous location influences the areas searched and hence the search effort overall. For instance, immediately following a rendezvous the frontiers selected are likely to be those in close proximity, requiring the shortest travel distances. The rendezvous method actually adds a further reference point which can be included by a UAV when choosing from the potential frontiers. Additional heuristics become available using this data including:

- **Distance between frontier and next rendezvous location**

As a heuristic for the travel time required to reach the next rendezvous location when required.

- **Estimation of time needed to complete frontier**

Determines if a frontier is likely to be completely searched before the UAV needs to travel toward the next rendezvous.

5.3.1 Rendezvous Timing

The time period between team rendezvous essentially determines the balance of the UAVs freedom to search and need to cooperate. Allowing a larger amount of time

before the next rendezvous provides more time for searching and promotes rapid information gathering. However, the disadvantage is that during this period any type of communication between UAVs is not guaranteed. The most important consequence is the lack of data sharing which impacts on team search efficiency and would likely result in undesirable multiple coverage.

A starting approach to determining the inter-rendezvous period is considering it as a tunable parameter of the deployment. Either set and fixed per search operation or variable on-line during execution. This assumes the purpose of the rendezvous is only to share data and hence coordinate search behaviour. However, a rendezvous also provides an opportunity to cooperate on relay tasks, which has a further impact on the time period before the next meetings. The effect of team relay techniques on rendezvous timings is described in 5.4.

5.3.2 Selecting Rendezvous Location

Working in tandem with the frontier search approach, rendezvous location selection occurs under two different conditions; inside a set priority region and without, the general case. Figure 5.4 shows examples for both situations.




Figure 5.4: Sample rendezvous location placements under different search conditions. The priority region (left) shown in blue is yet to be searched, whereas in the general case the green shading indicates the area already covered by the team.

When the team has an assigned, and still unsearched, priority region the ren-

rendezvous point can simply be set to a central location within this area. This is calculated as the centre of mass of the priority region and at a safe height above the ground level at this position for UAVs. This takes into account the shape of the region, placing the meeting point relatively closer to larger search areas. Potentially this can place the rendezvous location outside of the arbitrarily shaped priority zone but this should not cause any particular problem.

In the general setting, including continuing the search beyond a requested region, the rendezvous location guides the search progress by ensuring UAVs will visit this approximate part of the environment. To advance the search beyond the currently visited area, the next rendezvous is arranged for a location along the boundary between searched and unsearched space. There are various options for further refining where along the boundary is chosen, such as:

- **Nearest to current location**

Continue the search with minimal delay.

- **Shortest distance to base**

Encourage search near to the base station for short relay.

- **Farthest distance from base**

Promote search at distant location from base, examining areas that human responders may take longer to reach.

5.4 Team Relay

At each rendezvous the entire team aims to meet, sharing newly gathered information including both search coverage maps and actual images captured. This provides an ideal opportunity for the team's combined information to be relayed back to base by one or more robots. Furthermore, this enables a simple method for the base station to disseminate new instruction and otherwise acquired data to the remotely deployed team. For instance, the mission commander may wish to provide a new

priority region or instruct UAVs to land safely or return to base. This approach also enables multiple independent teams, with a common base station, to share data. Any information received in this way may not have been recently updated, but it will help inform the search should the separate teams area's of operation overlap.

5.4.1 Designated Single Relay

A common approach used in multi-robot systems is using a robot as a designated relay or data ferry. This provides a delay tolerant link between deployed robots, often involved in sensing tasks, and a fixed base station. Applied to this scenario, an assigned UAV travels between the base station and successive rendezvous points, as shown in Figure 5.5. New information is gathered from all members of the team during each rendezvous and ferried back to base. The designated relay must travel to within communication range of the base station to deposit the new data.

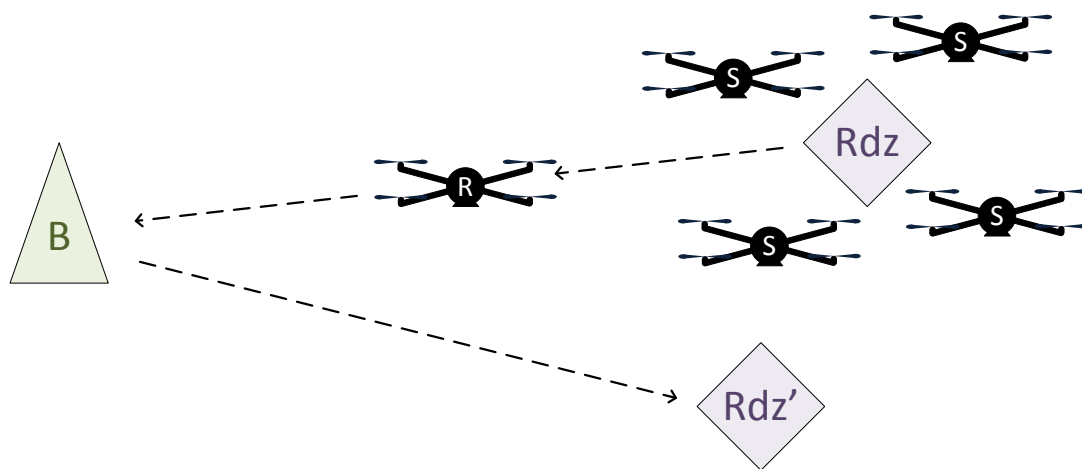


Figure 5.5: The assigned UAV R relays data from the current rendezvous Rdz to base before navigating toward the next relay Rdz' .

Although this removes a UAV completely from searching it fulfills completely the secondary goal of transferring new information to base. Using a single robot for this purpose allows the remaining team members to concentrate solely on search tasks. This configuration also supports the use of a heterogeneous team, with a specialised

relay robot. Such a vehicle may carry enhanced communication equipment in lieu of sensors to improve performance on this role. A potential approach for configuring the communication of a specific relay robot while servicing a multi-robot team at a rendezvous is presented in [89].

The relay time, the period required to successfully transport data from a rendezvous to base, is a combination of travel and data transfer durations. Under most circumstances in large operating environments the travel component outweighs information transfer. In turn, time spent by the relay traversing the route will vary with both the base to rendezvous distance and environmental features. This causes relay times to vary throughout a search deployment as the rendezvous location changes.

With the addition of the designated relay the inter-rendezvous period will also be affected. Assuming each rendezvous must include the relay, the minimal time between meetings is equal to the total time taken for this UAV to complete the relay task and reach to the subsequent rendezvous location. As this varies with the relative base-rendezvous positions the team cannot be assured of regular meetings. For instance, when searching at a large distance from base, there may be a significant delay in both providing new data to base and receiving new instruction. Conversely, when the team is operating a short distance from base the relay robot may complete this behaviour rapidly and be available to collect new data before any significant further searching has been achieved. Therefore a basic threshold on the minimum time between rendezvous should be enabled, ensuring searching UAVs have the freedom to cover new ground in between team meetings.

5.4.2 Multi-Robot Relay

To reduce the time between relay-enabled rendezvous, especially for meetings distant from base, multiple robots can be deployed as relays. Forming a chain with agreed meeting points along the route between base and team rendezvous location, each robot travels only a fraction of the complete relay distance. The inter-rendezvous period is therefore determined by the time taken for the final relay robot to execute

meeting, the entire subsequent multi-relay chain must also be configured. This allows each relay robot to be notified of its role in relaying data from Rdz' as the relay from Rdz occurs. In theory as a rendezvous completes all robots present will have identical maps of search coverage, so any could take responsibility for computing the following rendezvous location. In multi-robot relay operation, the relay robot R_r must have full knowledge of the relay configuration for the subsequent rendezvous at Rdz' ; this specification is passed along the chain to notify all relays. Therefore it follows that the responsibility of selecting the new rendezvous location and producing the relay configuration should be carried out by the relay present at the team rendezvous.

At the conclusion of rendezvous Rdz , the relay present, R_r , executes Algorithm 5.2 to full configure the subsequent rendezvous-relay stage. Once the location of Rdz' is selected, a full relay path from this point to base can be computed. This is then split into r segments, one for each active relay robot, with meeting locations M' at the path boundaries. These meeting locations are transmitted downstream along with the relayed data payload, providing all relayers a meeting location to navigate towards once the current cycle is complete. Finally the next rendezvous specification is broadcast to the search team. This marks the end of the current rendezvous, search operations resume and the current relay chain (Rdz to base via M) begins.

Algorithm 5.2 Rendezvous and Multi-robot Relay Configuration

```

function CONFIGURENEXTRENDEZVOUS(UAV uav, Region region, Integer numrelays, Location base, Time fixedRdzPeriod, Location rdz)
    Map map = getVehicleMap(uav)
    Location rdz'location = calculateNextRendezvousLocation(map, region)
    Path rdz'fullrelaypath = planPath(rdz'location, base)
    List meetings' = endPoints(split(rdz'relaypath, numrelays))
    uav.setNextRendezvous(rdz'location, fixedRdzPeriod)
    uav.broadcastNextRendezvous(rdz'location, fixedRdzPeriod)
    uav.beginRelay(rdz, meetings[r])
end function

```

This process is further described by Figure 5.7 in terms of the information flow between participants. During a rendezvous new data is transferred from searchers to the relayer present at the meeting, R_3 in this example. This relayer then performs

the algorithm to calculate the next rendezvous details and full relay configuration. The specification of the subsequent rendezvous is immediately communicated to the searchers present. Throughout the current relay cycle, both new data and the next cycle’s relay configuration is passed along to all relay UAVs in the chain. Finally, the base station receives the new information. All UAVs will now know of the next rendezvous or relay configuration and hence can alter their behaviour to make sure they fulfill the required commitments.




Figure 5.7: Information flows between participants during rendezvous and multi-relay enabled search.

5.5 Adaptive Data Latency Driven Multi-Robot Relay

Whilst employing multiple relay robots reduces the inter-rendezvous period, it requires more of the available vehicles and the overall relay time is still highly dependent on the base-rendezvous distance. Expected data latency, sensed to availability at base², is therefore variable throughout the search as rendezvous locations change. From the base station’s perspective the relay operation also determines the regularity of updates received; this is also varied by base-rendezvous distance for a fixed relay approach. It is desirable to obtain more control over data latency and the period between base updates to ensure a certain degree of freshness. This helps to guarantee

²For a full definition of data latency, see the performance metrics in Section 3.4.

that the data collected by the autonomous search team will be useful for the overall operation.

The proposed adaptive approach seeks to alter the number of robots deployed in a multi-robot relay chain to meet a data latency demand. The targeted data delay threshold can be fixed for the entire period or altered during operation.

In the general case for an organised team, the maximum possible data latency occurs when new data is sensed immediately following the end of a rendezvous. This new information is then collected by a relay at the following rendezvous before being relayed back to base. An estimate of the maximum possible time between sensing and data availability at base is

$$\begin{aligned}
 EstMaxLatency(Rdz', BASE, r) &= (Rdz'_{TIME} - t) + Rdz'_{DURATION} \\
 &\quad + EstMultiRelayTime(Rdz', BASE, r)
 \end{aligned}
 \tag{5.1}$$

where Rdz'_{TIME} is the time of the next planned rendezvous, t is the current time, $Rdz'_{DURATION}$ is the duration of the next rendezvous, and $EstMultiRelayTime(a, b, r)$ estimates the time required to relay data between a and b with a relay chain of r robots.

The major unknown component of this estimation is calculating the relay time and hence the potential maximal data latency. This includes the total period from the end of a rendezvous to data collected at that meeting becoming available at base. All robots are able to navigate by planning paths through the environment, this includes estimating the time required to traverse a chosen route. Therefore a predicted time required to complete relay paths is available to UAVs during operation. The estimated time required for the multi-robot relay can be written

$$\begin{aligned}
EstMultiRelayTime(Rdz', BASE, r) &= EstRelay(Rdz', M'_{r-1}) \\
&+ EstRelay(M'_{r-1}, M'_{r-2}) \\
&+ \dots \\
&+ EstRelay(M'_2, M'_1) \\
&+ EstRelay(M'_1, BASE) \quad (5.2)
\end{aligned}$$

$$EstRelay(a, b) = PathTime(a, b) + DataTransfer(I) \quad (5.3)$$

where M'_i denotes a meeting point between relays R_{i-1} and R_i , $PathTime(a, b)$ is the time taken to traverse a path between a and b , and $DataTransfer(I)$ is the time required to transfer data packet I between two robots within mutual communication range. Furthermore, the minimum time before the final relay link could reach the next rendezvous can be calculated

$$Min(Rdz'_{TIME}) = t + EstRelay(Rdz, M_{r-1}) + PathTime(M_{r-1}, Rdz') \quad (5.4)$$

Note the distinction between the current and subsequent rendezvous, Rdz and Rdz' respectively. The relay present at the meetings, travels from the current location to a relay meeting point and then on to the next rendezvous, which in general will be at a different location in the environment.

This provides a method to determine the timing of the next rendezvous based on the current relay chain. If the number of relays is increased each link becomes shorter, allowing the final relay robot R_r to reach the next rendezvous sooner. This reduction in minimum inter-rendezvous period in turn reduces the overall data latency. In practice, it is desirable to place a limit on the shortest time before another rendezvous can occur. This ensures that the search team is able to perform a reasonable amount of sensing activity and achieve a balanced usage of time. Without such a limitation

rendezvous could occur with a very high frequency, with only a small amount of new information shared at each meeting. Hence, the actual time for the next team meeting is

$$Rdz'_{TIME} = t + \text{Min}(\text{EstRelay}(Rdz, M_{r-1}) + \text{PathTime}(M_{r-1}, Rdz'), \text{MIN_INTER_RDZ}) \quad (5.5)$$

where *MIN_INTER_RDZ* is the parameter that limits rendezvous frequency.

An earlier assumption states that the time spent by robots moving through the environment far exceeds that required for communicating, including rendezvous, relay meetings and sharing with base. Excluding these data transfer terms demonstrates how the physical movement of UAVs completely determines data latency. Also note that the path costs used include the full route between meeting locations. In practice, robots need only travel sufficiently close to these points to establish communication with the other participants.

5.5.1 Adaptation Mechanism

In order to alter the relay configuration during operation a mechanism is needed that can be carried out by the deployed autonomous team. The approach used here allows the number of relays used to be increased or decreased step-wise; changing by one between successive relay cycles. It would also be possible to running an optimisation routine to select the number of relayers as required following each rendezvous. The suitability of the chosen approach is discussed with relevant simulation results in 6.6.5.2.

At the end of each rendezvous, once all data sharing has been completed, the relay present is able to assess the rendezvous and relay situation for the following cycle, as described in Section 5.4.2.1. Once a relay configuration for the next rendezvous is

calculated using Equation 5.1 with the current number of relay robots r , it can be assessed against the target latency criteria.

$$r' = \begin{cases} r + 1 & \text{if } EstMaxLatency(Rdz', BASE, r) > LATENCY_THRES \quad \& \\ & EstMaxLatency(Rdz', BASE, r + 1) \\ & < EstMaxLatency(Rdz', BASE, r) \\ r & \text{otherwise} \end{cases} \quad (5.6)$$

where r' is the number of relays to be used in the next cycle and $LATENCY_THRES$ is the preferred maximum data latency. When the current number of relays r is insufficient to meet the latency threshold, the addition of an extra relay is considered. However, another is only recruited if it would be expected to reduce the data latency. This will only occur if the additional relay agent enables the rendezvous to happen sooner. This is not guaranteed to be the case, as the time until each subsequent rendezvous may already be below the minimum threshold. Therefore, further increasing the relay chain will not actually reduce latency. In addition, a new relay is only added if the team will still contain at least one searcher. The mechanism for affecting this change is demonstrated in Figure 5.8. The relay robot present at the rendezvous will request that the searcher currently nearest to base joins the relay effort, becoming the final link for the subsequent rendezvous. The full relay chain configuration for the next cycle is calculated and passed to other agents along the chain during the current relay operation. Therefore, each agent knows the location of a meeting point to head towards following the completion of its immediate task. For the newly recruited relay UAV, this will be the location of the next rendezvous Rdz' .

As the team's rendezvous location changes the relay distance may decrease, leading to shorter relay times and reduced data latency. It is therefore important to reduce the number of relays and correspondingly increase the amount of searching UAVs. An additional rule describes how this is assessed.




Figure 5.8: Recruiting another relay to the multi-link chain. Following the current rendezvous Rdz the nearest searcher to base changes role to relaying, becoming R_3 . The arrows indicate the movements of relayers to prepare for the next cycle.

$$r' = \begin{cases} r - 1 & \text{if } EstMaxLatency(Rdz', BASE, r - 1) < LATENCY_THRES \\ r & \text{otherwise} \end{cases} \quad (5.7)$$

Simply, if fewer relays would result in an acceptable data latency, below the set threshold, then reduce the length of the relay chain. To implement this change, the final relay link R_r completes the immediate relay as previously configured before beginning a search task, as shown in Figure 5.9. For the next cycle relay R_{r-1} becomes the final link in a reduced chain, providing service at the next rendezvous. Hence, its position must be considered when calculating the time of the subsequent team meeting. It must first complete its role in the current relay activity, before heading to the rendezvous. The calculation of the soonest time this meeting can occur is therefore modified from Equation 5.4 to,

$$\begin{aligned}
Min(Rdz'_{TIME}) &= t + EstRelay(Rdz, M_{r-1}) \\
&+ EstRelay(M_{r-1}, M_{r-2}) \\
&+ PathTime(M_{r-2}, Rdz')
\end{aligned} \tag{5.8}$$

This takes into account the time required to complete two relay stages, immediately adjacent to the current rendezvous Rdz , and for the new team relayer to reach the next meeting. In addition, this is subject to the limitation of the parameterised minimum time period between rendezvous'.

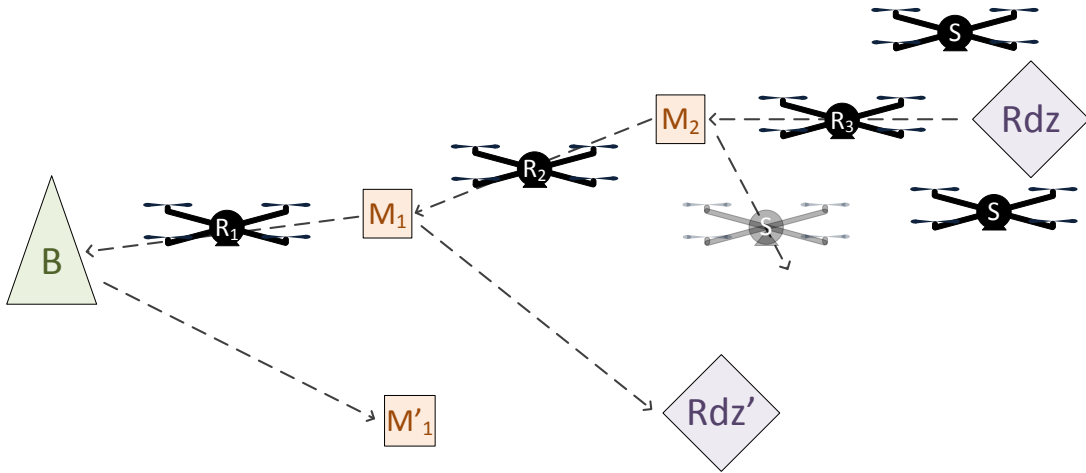


Figure 5.9: Reducing the number of relayers. The final relay R_3 is relieved from relay duty and takes up the role of a searcher. Agent R_2 now becomes the relay responsible for collecting data at the next rendezvous Rdz' and will move to that location after completing its current relay stage.

As with the numbers of searchers, the number of active relays is subject to a minimum of one, ensuring a certain minimal data relay service for the team. When the number of relays decreases, the expected search rate of the team will improve. This is due to both a longer inter-rendezvous period and one additional active searching robot.

Chapter 6

Simulation Experiments

This chapter describes experiments used to evaluate robot deployment approaches for wilderness search and rescue. These are performed in simulation using a range of real terrain data to model the search environment. Firstly the usage and implementation of the simulation software is described, followed by a set of benchmark results. Each of the approaches is then examined with suitable experiments, followed by a characterisation of the best performing strategies.

6.1 Simulation Usage

Simulation is a common tool in robotics research, enabling development and analysis to occur without or in combination with field work. This is especially the case for multi-robot systems, where the difficulties of performing representative practical experiments are magnified. Experimental work performed in simulation uses models of objects and the interactions of a physical world, including robots, the operating environment and other entities. Such models are often purposefully abstracted from their real-world counterparts to reduce both the implementation effort and the computational resources needed to execute the simulation.

Advantages of using simulation instead of, or alongside, real robotic experiments include: reduced costs and time required; the ease of comparing approaches and deployment parameters; and unrestricted data capture. Performing practical experiments can involve considerable financial costs to commission, maintain and operate autonomous robots. In addition to the material cost of physical vehicles and equipment, significant investment is required for both personnel and to use a representative environment. For example, a multi-UAV WiSaR experiment requires a minimum of one adequately qualified safety pilot per robot and all of the equipment may need to be deployed into a challenging remote location. In contrast, simulation can be performed using modest computing resources. Using models also provides increased opportunity to investigate potential approaches in environments where it would be extremely difficult to or even impossible to run practical experiments. For example, in congested urban locations, nuclear waste storage pools, the ocean depths and other planets. When carrying out multi-robot simulations any data required is immediately available with minimal effort. In addition, mission and robot parameters can easily be varied and tuned to examine their effect on system performance.

The main disadvantage of relying on simulation is a reduction in relevance to the targeted real-world application¹. Inaccuracies can occur in almost all areas of

¹Even physical robot experiments are not automatically relevant to the final *live* scenario. Results can be easily biased by the experimental configuration if sufficient care is not taken.

simulation, with the abstract notions only approximating the actual systems they model. However, the extent to which simulation is relevant varies with the nature of the system under consideration. For instance, robot control in a dynamic environment is notoriously difficult to model and successfully translate to a physical system; consider, for example, navigating a full 6 degree of freedom UAV in a windy obstacle ridden environment. Similarly with perception; computer vision systems require significant real world data for testing, development and classifier training. For these reasons, approaches aiming to fully plan significant periods of future behaviour or tightly coupled perception/action systems should not be expected to translate well from simulation to physical robots. Furthermore, the potential for unexpected robot behaviour or failure is not always well described by simulation.

The scope of this work makes the robot deployment approaches particularly well suited to examination by simulation. Each individual robot has defined basic behaviours, see the global assumptions in Section 3.3, which do not require detailed modelling. Furthermore the period of time each behaviour is performed before re-planning is very short relative to the total typical operation length. Therefore if similar approaches were to be transferred to a physical robot system, the potential discrepancy between simulation and physical systems should be limited. For instance any increase in the variability of navigation times would have a limited impact. Similarly, the performance of a sensing task requires only an outline of the time taken to sense a given area and the associated altitude required.

6.2 Simulation Detail

A purpose built Java-based software package is used to conduct the simulation experiments described in this chapter. This enables simulation of robot team deployments in three-dimensional environments. Terrain maps from the Ordnance Survey are imported, processed using Matlab and accessed using the matlab-control Java library [93]. The software features a graphical user interface to observe simulation progres-

sion, examine the current maps of all the agents (UAVs and base) and view an overall summary of their locations. Priority search regions of arbitrary position and shape can be also be user-defined.

Each agent is fully independent in the simulated world, with complete autonomy to take decisions such as selecting which behaviour to enact. UAVs have individual spatial data storage; with a copy of the known terrain information provided at intialisation. These are represented on regular grids as described in Section 4.2.1. The simulation is run with discrete time steps, however each agent cannot access the global clock. Agents do not have access to a common clock and must maintain their own. Events that rely on timing, such as a rendezvous, are agreed after a specific duration, for example the next rendezvous will occur at 250 time steps following its announcement.

Robots can move through the environment laterally on the 4-connected grid (north, south, east, west) or change altitude (up or down) only. Hence the minimal distance moved is equal to the side of a grid square (10m) as determined by the grid resolution. UAVs are subject to a limited maximum velocity per time step and must maintain a minimum altitude of 20m above the ground at all times. Navigation is aided by an A* path planner [57] to select routes to required locations, such as a start point for frontier search or when relaying information to base. Whilst fully described optimal paths can be found using this method, it is computationally expensive, especially over long distances for larger maps. Therefore, UAVs use the A* algorithm to plan a route across randomly placed nodes on the map, with each agent creating an individual set of nodes at initialisation. Segmenting the entire environment (the ground surface with expected maximum altitude) into three-dimensional boxes, a node is randomly placed in each segment to ensure a reasonable spread. When a new route is required, the nearest nodes to the start and goal locations are calculated and a path planned between them via the navigation nodes. Shorter paths still rely on an optimal A* approach, such as from the current location to the start node and between nodes.

Localisation is a key aspect of mobile-robotics with errors having a great impact, especially in practical deployments. As detailed in Section 3.3 robots are modelled with perfect localisation for both lateral (x/y) and height (z) dimensions. This is consistent with the overall theme of the thesis, which excludes an examination of any vehicle dynamics and control issues. For most outdoor robotics, GPS technology provides a convenient relatively low-error estimation of longitude and latitude, but poor height data. Therefore, it is often fused with a barometric pressure or groundward facing ultrasonic distance sensors to achieve acceptable accuracy.

Sensing is performed by a UAV at a set height above ground level (20m) and whilst not exceeding a predetermined search velocity (significantly less than maximum). The entire grid location is assumed fully sensed by the UAV provided these conditions are met. In these simulations moving between adjacent grid locations (10m in distance) and perceiving a grid location requires a single time step.

Agents are modelled with wireless communication capability which is limited by distance and due to environmental interference. The range limitation is a simulation parameter that can be varied to assess the importance of communication on the deployment. Typically in physical devices transmission power and antennae will be the significant factors in determining effective range in clear space. Hence, the base station is assumed able to achieve an increased range over UAVs due to larger antennae and increased power supply, both of which are limited on-board mobile robots due to size and weight constraints. Transmission is assumed to be only permitted between agents if a clear line-of-sight exists between their locations. This is implemented by ray-tracing between the positions and consulting the terrain data to establish if a straight line path exists. Attenuation of wireless signals due to obstacles can be modelled in more detailed manner by calculating the reduction in signal strength. However, as the only obstructions in this environment consist of the ground (a high density obstacle), the pessimistic approach is taken to not permit any communications where a clear line-of-sight does not exist. A messaging system is used to organise communication between agents. This allows both broadcasting, a simple

alive message for example, and directed communication, such as a request to join the relay effort.

A simulated deployment begins with the definition of the total operating area, associated terrain data, base station location and (if used) a priority search region. UAVs are initialised at random locations at a limited maximum distance away from base. Many aspects of the simulation are logged continuously over time to aid analysis, these include: each UAV’s behaviour, the portion of map searched by UAV, the information available at the base, and the data latency between sensing and availability at base.

6.3 Approaches Under Consideration

The simulation experiments in this chapter will use multiple approaches for each of the main aspects of the problem; search, relay and team coordination. A deployment configuration consists of a technique for each aspect, chosen to examine a specific part of the team’s performance. This enables comparison and analysis of the relative merits of various deployment configurations. The approaches under consideration are shown in Figure 6.1. This illustration is used throughout this chapter to highlight the methods being investigated and compared at that point.

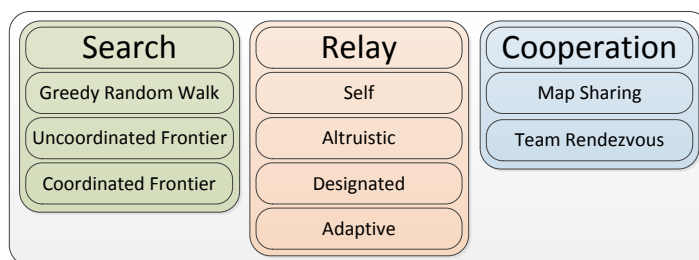


Figure 6.1: Approaches to be evaluated through simulation experimentation.

The main approach for UAV-based search is the frontier method described in Section 4.2.2 but to provide comparison a *greedy-random-walk* algorithm will also be used. Each UAV acting individually will search by randomly moving through the

environment to visit unsearched locations; behaving greedily by selecting the nearest first. The UAV randomly selects any neighbouring unsearched grid square that can be searched without changing altitude, moves there and performs sensing activity. If no suitable locations exist, using its individual map of searched/unsearched areas, the nearest unsearched point is chosen and navigated towards. This continues until the required area has been covered. Where multiple UAVs are deployed, they will each attempt to periodically share coverage maps with those in wireless communication range; basing search decisions on currently held data at that instant. Examination of the frontier search method is particularly focused on the comparison of various heuristics available for frontier selection.

A range of tactics for relaying fresh data to base are considered in Chapters 4 and 5. The basic approach, *self-relay*, relies on each UAV to transport and communicate any new data it has acquired to the base-station. This behaviour is triggered upon reaching a predetermined limit for the quantity of new data held in its cache. Varying the threshold allows more or less searching to be carried out before relaying must occur. Under this scheme, each UAV only considers data acquired via sensing and not shared by other team members. By introducing the concept of information responsibility the *altruistic-relay* approach aims to reduce redundant relaying. When nearby UAVs share their gathered information, the agent nearest to the base station assumes responsibility for all of the new combined information, acquired since the previous relay activity.

When using the team rendezvous approach (see Section 5.3) specifically designated relay robots enable data to be collected from a team of multiple searchers. This *designated-relay* configuration may contain one or more UAVs acting as a chain to ferry data to base from the remotely deployed search team. The number of robots in the chain remains static, however their meeting points (between adjacent relay robots) vary as the rendezvous location changes. The most complex relay technique described enables UAVs to switch roles between searching and relaying to meet a data

latency demand. Using this *adaptive-relay* approach, the number of robots involved in relaying and searching changes depending on relative rendezvous and base locations.

All simulated deployments that involve at least 2 UAVs will contain a minimal element of cooperation. Opportunistic *map-sharing* provides a basic method for data to pass between agents, either directly when nearby or via the base station as successive UAVs visit during relay operations. By sharing information team members use data gathered by others to inform their own decisions; an implicit form of coordination. The basic frontier search methodology is enhanced via the broadcasting of selected search regions. These notifications enable *coordinated-frontier* search using a team of UAVs; each considering other agent's operation when taking action. With the addition of *team-rendezvous* these elements of coordination can be encouraged. The enforced contact between UAVs aims to increase search efficiency by promoting data sharing.

6.4 Methodology

In order to adequately compare different approaches, multiple simulation runs are used to eliminate bias created by the scenario configuration. Random elements in UAV positioning and navigation prevent simulations with identical overall parameters becoming deterministic. For example by initialising robots at random locations and with navigation nodes being randomly placed, both within limitations, different frontiers will be selected across numerous trials. Hence, unless stated all results presented are a combination of multiple simulations, usually calculated as the mean².

Team deployment approaches are assessed by experimenting with the key parameters and examining the effect on the metrics presented in Section 3.4. These variations include elements of the robotic deployment, the application scenario spec-

²During the course of producing these experimental results the data has been analysed with further measures, including the variance. This was found to be very limited and not to offer additional insight so is not shown on the plots for clarity.

ification and mission demands. In Section 6.5 a set of results is presented for some benchmark simulations, these provide a starting point for the experimental analysis.

Experiments are conducted by deploying a team of 1 or more UAVs on a search mission, the principle goal of which is to search either a defined priority region only or the entire map beginning with the region. Priority regions are initially selected with a fixed area but at random locations a specified distance from the base station. This allows multiple simulation runs of an identical configuration except for differing search region locations, eliminating any bias from region position. Where a comparison is to be made between approaches, an identical set of missions is used, consisting of a map and a set of priority regions, one for each deployment. In particular this work targets search operations at locations distant to the base station, this is reflected in the selection of priority search regions within the environment. Figure 6.2 shows an example of 10 priority regions of identical size, but randomly located with their centres at a fixed distance from the base station. This is a typical configuration used during these experiments; a total map size of 4km^2 , with a priority region of approximately 0.25km^2 at a distance of 1.5km from the base station.




Figure 6.2: Randomly located priority search region ellipses. Each region is shown by its outline in red, with the base station in orange at the upper left corner.

Many of the metrics use a measurement of the amount of information gathered

or held by an agent as a proxy for the search coverage. In the simulation results these quantities are expressed as a percentage of the total area to be searched in that particular mission; either the entire map or the specified region.

6.4.1 Environments

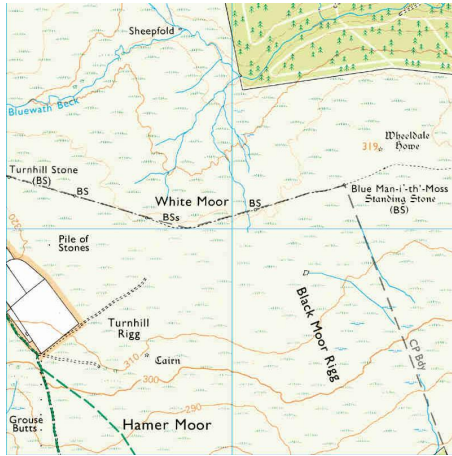
There exists a wide range of potential operating environments for the deployment of a UAV-based search team. One of the most significant variations is the terrain relief; how the height of the ground varies across the region. Throughout these simulations performance metrics are evaluated over three environment types: *plains*, *undulating* and *mountainous*. The differing terrain relief will impact many aspects including the constraints of frontier search, inter-agent communication and UAV navigation.

Figure 6.3 shows the terrain samples used for these experiments along with the associated discretised ground height data that is available to agents. Each search area covers 2km by 2km, providing an operating space of 4km². This data is provided by Ordnance Survey [145] through the MasterMap [105] and Land-Form PROFILE Digital Terrain Map products [104]. All data describes real-world locations in the United Kingdom, with the plains, undulating and mountainous terrains representing parts of North Yorkshire Moors, Dartmoor and Caingorm National Parks respectively. The range of ground height varies substantially between different environments, see Table 6.1. Moreover, the rate of height change is significant, for instance mountainous terrain contains incredibly steep gradients.



Table 6.1: Ground height ranges of sample terrain maps.

Terrain	Ground Height Range (m)	Relative Difference (m)
Plains	257.9 - 320.9	63.0
Undulating	358.7 - 546.2	187.5
Mountainous	728.7 - 1190.9	462.2

Plains



Undulating



Mountainous



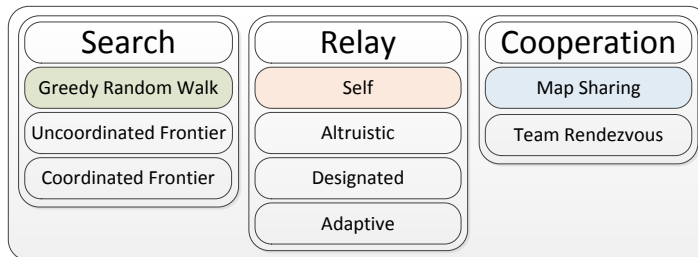


Figure 6.3: Terrain maps and associated discretised ground height data used for simulated deployments. All maps and data are ©Crown Copyright/database right 2013. An Ordnance Survey/EDINA supplied service [145, 39].

6.5 Benchmark

A series of basic benchmark experiments are presented to establish a baseline performance across several metrics using an uncomplicated approach. The deployment uses the greedy random-walk search algorithm, self-relay of new data by each UAV and basic cooperation via map-sharing when possible. In these trials each agent will attempt to share its current map with any other agents in communication range every 50 time steps. Figure 6.4 shows the progress of a single UAV searching a priority region, with both the data known by the UAV and at the base. The elliptical search region has maximal dimensions of 0.8km by 0.4km for a total area of approximately 0.25km² and the center is located at a distance of 1.5km from the base station. The UAV is configured to relay once 0.03km² of new ground has been covered. When actively searching, the UAV's data information increases at an approximately constant rate, whereas the base station receives an update only after a relay has occurred, during which time no new search data is acquired.



Similarly, the search behaviour of a team can be examined by looking at the information gained. Figure 6.5 shows the information available at base for a team of 5 identically configured UAVs searching the entire map of each terrain type. The search environment affects search progress with the increased ground height range of mountainous terrain resulting in a reduced search rate and hence longer completion time. The highly variable terrain poses additional difficulty to the team, potentially blocking communications, increasing navigation complexity and limiting search area before changing altitude. As only a limited amount of communication is used in this




Figure 6.4: A single simulation run of an individual UAV searching a elliptical region in the undulating terrain map using a greedy random walk algorithm. Information gathered by the UAV is shown along with data available at base over time.

benchmark trial, it is the latter two issues that are likely the cause for a slower search rate over the mountainous terrain. Each UAV spends an increased amount of time simply moving around the environment, with each height change costing time steps that otherwise would be spent searching.

The team’s behaviour can be examined more thoroughly by examining the rate at which new information is gathered. Figure 6.6 uses the InfoGainRate metric over a sliding window of 500 time steps; each data point being the mean rate for the preceding period. During the initial phase, the restriction of searching a fixed area, the priority ellipse, results in a lower maximum rate of information gain, which further reduces as the size of the remaining unsearched area decreases. As all UAVs are attempting to search in the limited area, it is inevitable that some undesirable interference will occur; both by physically obstructing one another and by sensing the same locations multiple times. Immediately following the completion of this phase the team achieves its maximum search rate, as UAVs are unrestricted in their movement, sensing in any unvisited location. Again, as the available search space reduces over




Figure 6.5: Search information available at the base station during a team search deployment. The measurements are a result of 10 simulation runs for each terrain type, using a predefined set of priority regions to begin the trials.

the remainder of the operation less information is acquired per unit of time. A further dramatic reduction is observed in the final period before full coverage is achieved. During this period agents must visit and sense at any remaining unvisited parts of the map, which may be widely distributed. This is exacerbated by the greedy walk strategy, which takes no account of the spatial distribution of unsearched area when selecting next sensing locations. Rates of information gain across the three terrain types all follow the same overall trends, with slightly reduced rates for the undulating and more so the mountainous environment. This is a result of agents expending more effort on navigating via additional altitude changes.

Table 6.2 shows the individual behaviour of UAVs during one deployment. Using map sharing as the only method of coordination, each agent is essentially searching and relaying without any regard for other team members. Hence, each robot uses approximately the same amount of time for searching - both sensing and moving to a required location - and relaying by returning to base. Similarly, search coverage for each UAV is very closely matched. However in combination the team has sensed




Figure 6.6: Rate of information gain for the search team of 5 UAVs for each terrain, calculated over a sliding window of 500 time steps.

an additional area equivalent to 33.2% of the environment. The lack of coordination, apart from opportunistic map sharing, results in a considerable amount of recoverage and hence wasted effort. Even for nearby agents, within communication range, the period between map sharing events (50 time steps) provides ample time to re-search the same ground.

Table 6.2: Individual UAV behaviour during a simulated deployment on the undulating terrain.

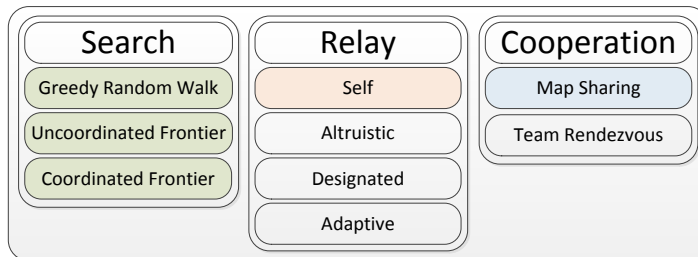
Agent	Time Searching	Time Relaying To Base	Area Sensed (% of total)
UAV-Search-1	13999	2390	26.47
UAV-Search-2	14099	2290	26.44
UAV-Search-3	13995	2394	26.67
UAV-Search-4	14130	2259	26.74
UAV-Search-5	14134	2255	26.92
		Total	133.24

6.6 Performance Analysis Experiments

The main set of experiments presented here attempts to analyse the performance of various deployment configurations. Each aspect is addressed separately, comparing the benchmark results to the frontier, rendezvous and relay approaches proposed in this thesis. Additional experimental results to the main narrative can be found in Appendix B.

6.6.1 Frontier Search

The frontier search method is used individually by each UAV to designate the next part of the environment for searching. A list of potential frontiers are generated using a region growing approach on the agent's search progress map, one of which is then selected for immediate search.



6.6.1.1 Basic Search Rate

A basic comparison of search rates with the greedy random walk approach is shown in Figure 6.7. For these simulations a single UAV is deployed to search a elliptical priority region placed on the mountainous terrain. To enable a direct comparison of information gathering in isolation, the agent is not required to relay data to base until full coverage is achieved. Whilst the greedy walk approach reaches a higher maximum rate early in the search, the frontier method completes the mission in a shorter time by maintaining a more consistent rate throughout. Greedy walk deployments also exhibit a greater reduction in search rate as full coverage is approached. This variability is

caused by the nature of random walk, fragmenting the space into smaller searched and unsearched sections. However the frontier approach manages improved consistency by completing each chosen continuous frontier region before moving on the next. A similar pattern is repeated for the plains and undulating terrains, these additional results can be found in Appendix B, Figures B.2 and B.3.




Figure 6.7: Information gain rate for a single UAV searching an elliptical priority region only, using both the greedy walk and frontier search strategy. For this experiment, no relaying occurs and the mission is complete once the UAV has covered the required area. Data is the result of 10 averaged simulation runs.

6.6.1.2 Search Rate for Team Deployments

Frontier search can be operated in two modes: uncoordinated, using only the available map data, and coordinated, using information about other team members to inform selection. The coordinated method requires all UAVs to periodically broadcast their location and a representation of the size of the frontier currently being investigated, for which we use a bounding box surrounding their current search frontier. For these experiments this relatively short message will be transmitted every 10 time steps to those in communication range. Each agent uses this data to reject frontiers that

either intersect other operating areas or would be scored higher for another searcher. Only information received within a set previous time period (100 time steps) is used.

Whether coordinated or uncoordinated, a scoring method is required for frontier selection, with several possibilities, as described in Section 4.2.4. To compare the frontier selection methods a series of identical deployments were simulated, using the various selection methods. Figure 6.8 shows the completion times to search the entire map for a team of 5 UAVs, across the 3 terrain types, for a range of selection heuristics. As with the benchmark tests, the mission dictates that a priority region be searched before the entire map is explored. Each UAV acts as its own relay, transporting acquired data to base once a threshold is met. The heuristics demonstrated include: *minDist*, measuring the minimum distance from the current location to a position from which the frontier can be searched; *minDistCentre*, calculating the minimum distance to the frontier's centre of mass; and *MinDist*Size*, which uses the size of the frontier to modify the minimum distance measure, encapsulating a preference for larger frontiers. In all distance measurements, an altitude change preference weight *ACW* of 5 is used to calculate the cost of changing flight altitude, see Section 4.2.4.1.

The completion times shown indicate, as before, that increased variability in ground height lengthens the search operation consistently across all methods. Considering the approaches as three groups, the order of increasing performance is uncoordinated frontier, greedy walk and coordinated frontier. In terms of completion time the uncoordinated frontier techniques perform poorly, taking longer to finish the search than a greedy random walk strategy. All deployments using coordinated frontiers manage to search the required area in a shorter time.




Figure 6.8: Completion times for various search methods by separate terrain type.

Despite the uncoordinated frontier approach outperforming greedy walk in terms of its basic information gathering rate for a single UAV, see Section 6.6.1.1, when applied to a team the opposite is true. As the frontier selection policy is an individual decision there is no impediment to stop multiple UAVs simultaneously selecting overlapping regions to search. This can occur even though the individual frontier generation process is randomised, meaning that identical maps will only produce the same frontiers very rarely. Consider two agents located in a similar location, for example near the base station. If they have shared maps recently, incorporating each others data, they will both generate frontiers that cover similar unsearched areas of the environment. The frontiers selected by each agent using the same heuristic may share common ground with a reasonable likelihood. This can result in significant recoverage, with agents often searching parts of the terrain only recently covered by another. Furthermore the issue will be amplified for larger team sizes.

The coordinated frontier approach mitigates the issue of inter-agent inefficiency by each UAV taking into account other team members when selecting frontiers. Through the use of a simple broadcast mechanism the search completion times are greatly improved. Each of the heuristics shown offer a very similar level of performance, with *minDistCentre* achieving the quickest searches. By providing a measure derived from the entire frontier size and shape *minDistCentre* marginally outperforms *minDist*, which just uses the single nearest point to the current agent location. The final heuristic, *MinDist*SizeRatio*, combines spatial separation with the frontier size to express a preference for larger frontiers near to the current location. However, this does not seem to offer any performance benefit. This could be due to the frontier generation process producing many frontiers at or very near to the maximum size.

6.6.2 Altitude Change Preference

For a rotorcraft vehicle to increase flight altitude additional energy and time is required compared to maintaining a constant height velocity. To allow agents to specify a preference for maintaining level flight, the parameter Altitude Change Weight

(ACW) was introduced in Section 4.2.4.1. The effect of this parameter can be assessed by performing a series of simulations for a range of ACW values. The minimum value it can take is 1, meaning any planned travel distance towards a frontier will be given the same cost for the UAV irrespective of the direction of motion. Increasing the value causes any altitude changes to be penalised with a higher cost when comparing potential actions.

Taking the best performing search approach demonstrated so far, coordinated frontier search with the *minDistCentre* heuristic, the effect of ACW on completion times is shown in Figure 6.9. For operations over the mountainous terrain the parameter has a measurable effect on overall search time, whereas no impact is observed on the undulating or plains environments. This is due to the minimal amount of variation in ground height producing a much lower requirement for UAVs to alter their altitude. With values of ACW less than or equal to 5 a minimal increase in completion time is seen even for the mountainous terrain; beyond which the gradient of the plot is marginally steeper. From these results it appears that there is a benefit from each UAV being greedy in its approach, selecting nearby frontiers to search without regard for power usage. However for many, if not all, practical deployments power consumption for this type of vehicle is a concern and potentially the main limiting factor in some scenarios. For the remainder of the simulations ACW will be assigned a value of 5, giving some consideration to power consumption and sacrificing a small amount of overall performance. In a specific practical deployment this value can be set using recorded power data from actual vehicles combined with knowledge of the specific mission priorities, for example time critical or a long term deployment.

6.6.3 Rendezvous and Opportunistic Relay

The dual concepts of a team rendezvous and relay strategy attempt to improve overall performance by addressing information flow within the team and towards the base respectively. This set of experiments examines the use of team rendezvous and a

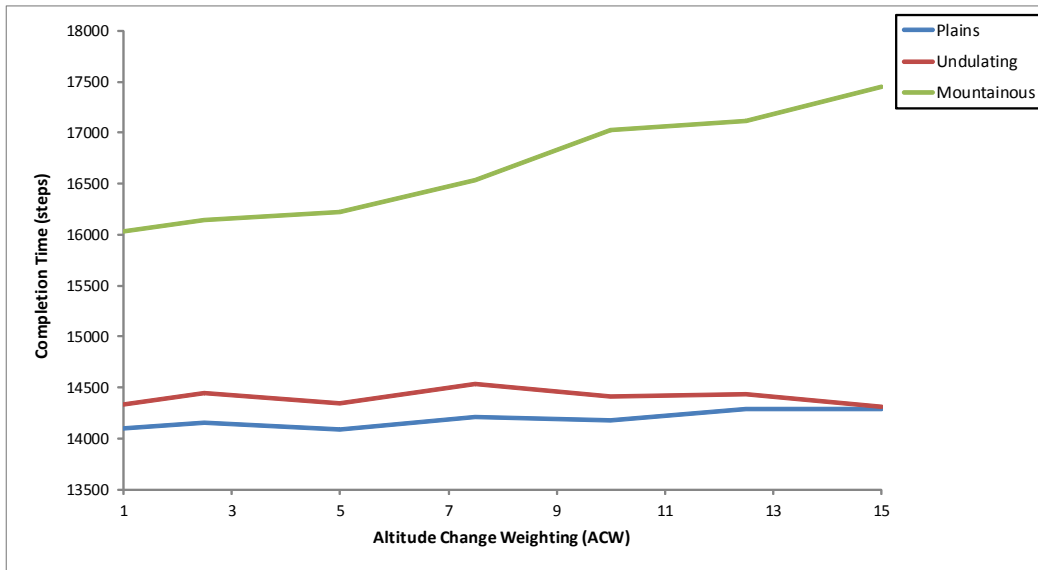


Figure 6.9: Completion times for a 5 UAV team using coordinated frontier search for a range of Altitude Change Preference Weight values. Each data point is the mean of 10 simulations.


basic opportunistic relay approach to reduce the collective time spent transporting new data to the base station.

6.6.3.1 Altruistic Relay

In the approaches examined thus far a significant amount of effort has been expended by each agent to convey newly gathered information to the base. The altruistic relay approach allows agents to pass the responsibility for relaying data to other agents whenever map sharing occurs. As such events are unplanned this approach is purely opportunistic and the impact is therefore difficult to predict. When agents meet³, exchanging new sensory data and current locations, they both calculate their respective distances to the base station, with the closer one taking responsibility for the task. This is confirmed with a simple message protocol.

The altruistic relay strategy can be applied to both frontier and greedy random

³Here, an agent meeting, refers to two agents coming within direct communication range of each other and communicating.



walk search methods; the information gain rate for these deployments is shown in Figure 6.10. The addition of altruistic relay offers a performance improvement for teams using greedy walk or frontier search; deployments using both of these search methods are able to complete the search in a shorter time with altruistic rather than self-relaying agents. The overall pattern of search rate remains, the maximum being achieved shortly after the priority region is complete, slowly trending downwards as the remaining unsearched space reduces in size. These results further illustrate how frontier approach enables the team to maintain a more consistent search rate. This is especially noticeable as the priority region becomes almost complete, where the effectiveness of greedy walk to gather new data is drastically reduced.




Figure 6.10: Rate of information gain over a sliding window of 500 time steps for self and altruistic relay strategies. A team of 5 UAVs was tasked to search the mountainous terrain.

6.6.3.2 Search Team Rendezvous

All UAVs will try to meet, moving to within communication range, at each arranged rendezvous to share gathered data and their records of search progress. The meetings are arranged successively, being defined by a location and an amount of time before the next occurrence. This approach can be used with or without an accompanying team relay strategy. When the team at rendezvous does not contain a designated relayer, as in the simulations in this section, the UAV that is nearest to base will select the subsequent location. Whilst searching is still required within the assigned priority region the rendezvous is always at its centre. Beyond this, the next rendezvous location will be placed at the nearest point to the current meeting that lies along the boundary of searched and unsearched space (see Section 5.3.2). This enables searching UAVs to continue exploring with minimal delay following a meeting's conclusion. With this configuration rendezvous occur at fixed intervals, every 300 time steps for these simulations.




Figure 6.11 shows the completion times for team deployments using a combination of search, relay and rendezvous techniques. Looking at the left-hand-side of the figure, the addition of team rendezvous does not provide any reduction in overall search times whilst UAVs are self-relaying acquired data. Although this ensures that opportunities to share information occur regularly, it diverts active time away from searching. Further inefficiencies are possible with this combination, for example agents heading towards and participating in rendezvous' when their search cache threshold is almost reached. There is also an element of duplication of effort in sharing

data as self-relay agents will receive updates on teammate's search progress indirectly via the base station. Applying the rendezvous method to a greedy walk search team performs particularly poorly; forcing agents to be periodically close together results in significant recoverage, impacting overall performance.

However, examining the right hand side of Figure 6.11, using the altruistic relay approach alone (without team rendezvous enabled) does offer a marginal benefit; taking advantage of the chance meetings between searching UAVs to collate responsibility for relaying data. This can be seen for both the greedy random walk and frontier strategies.

The maximal performance for coordinated frontier search is reached by employing both team rendezvous and altruistic relay strategies, as seen at the extreme right of the figure. This combination facilitates regular information sharing between agents and compresses the required relay effort into fewer vehicles (as described in Section 6.6.3.3). Despite the additional effort expended by each UAV to travel to, wait at and participate in the rendezvous and then travel to the next task, a clear increase in performance is present. Coordination of search is vital in the success of this approach, as agents are brought very close together at each meeting. When they return to sensing tasks, agents must be careful to avoid interference with each other. The coordinated frontier method achieves this by each UAV taking into account the position, operating region and likely behaviour of others. Conversely the uncoordinated greedy walk search method is not compatible with any approach that forces or encloses agents into a limited space.




Figure 6.11: Completion times for deployments with or without team rendezvous' and using self or altruistic relay strategies.

The effect of rendezvous' on the information flow within the team is shown in Figure 6.12 for deployments using altruistic relay and with or without rendezvous. As the team meets, the average proportion of information known by each agent increases to 100% due to the data exchange that takes place. This reduces as agents spread out during the search phase, until they must head towards the next meeting. The frequency of rendezvous' is the result of the fixed interval (300 time steps) used in this example; in a general setting rendezvous timing is shown to be variable, see Section 6.6.4. Information sharing when using rendezvous' is almost always higher than without, through the entire search period. This allows the team to search more efficiently, to reduce re-searching of the same locations, and to complete the mission in a shorter time.




Figure 6.12: Average information known by each UAV as a percentage of total team information, shown for deployment with and without periodic team rendezvous, both using altruistic relay. These results are averaged over all three terrain types.

6.6.3.3 Team Relay Effort

From the simulations described above it is clear that altruistic relay offers a benefit in search rate and therefore completion time. The effect it has on time spent by the team on relay activities is demonstrated by Figure 6.13. The time spent relaying is shown for each of the 5 UAVs and team as a whole. Rendezvous' are used in both self-relay and altruistic relay configurations. The figure is compiled by sorting the UAVs in terms of the time spent relaying. Hence, *UAV_1* shows the average time cost of relaying, for the agent in each deployment that used the minimal amount of time on this behaviour.

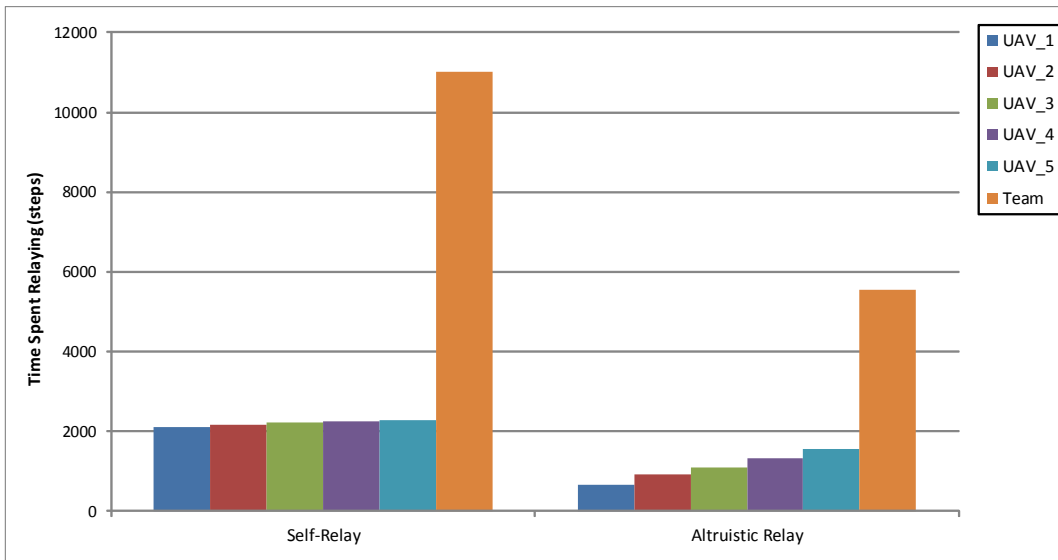


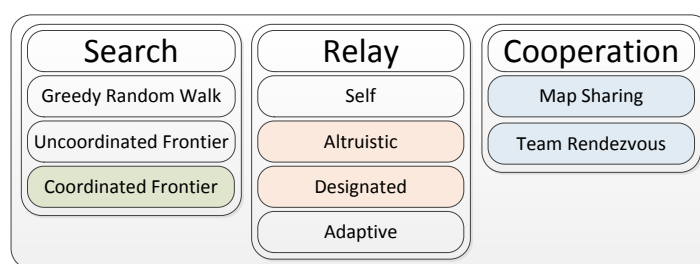
Figure 6.13: Time spent on relay activities for altruistic and self relay methods. Data is averaged over 10 simulations on each terrain type for 30 total trials. Both configurations use team rendezvous. Individual UAV data is sorted by time spent relaying to demonstrate the range between greatest and smallest relay effort per deployment.

Altruistic relay reduces the overall time required by a team to ferry its data to base and therefore enables a larger proportion of effort to be directed to search. This results in shorter completion times for the mission. The distribution of relay effort within the team is significantly different for the two strategies. When self relaying all UAVs have

a very similar time usage, matching very closely how much searching and relaying are performed. Altruistic relaying concentrates the relaying responsibility into fewer of the agents, effectively splitting the roles between team members. This pattern can be easily seen when observing a graphical representation of a deployment in progress. As the majority of the team searches near a remote rendezvous location, usually a single UAV emerges as a de facto relay, taking new data from the team at a meeting before transporting it to the base station. During some missions, particularly those further from base, the relaying agent will return shortly before the next rendezvous. Following the data exchange it is once again the nearest agent to base and therefore is reselected to relay for the team.

6.6.4 Rendezvous with Designated Relays

An alternative approach to meet relay demands is assigning specific vehicles to act solely as data relays; also known as mules or ferries. In rendezvous enabled deployments, this creates a periodic delay-tolerant link between the base station and remotely operating team. Designated relays also have a specific role in organising each subsequent rendezvous and communicating the required information to the team.



6.6.4.1 Comparison with Altruistic Relaying

To compare the use of a designated relay to the previously described altruistic relaying approach a range of scenarios is required. The primary factor affecting relay behaviour is the time taken to complete each relay task. This in turn is influenced by the

distance from the deployed team to the fixed base station at the time of relay activity. When fully searching a large area, this distance varies significantly as the rendezvous location changes. Therefore, this series of simulations requires the team to search an elliptical region, randomly located at a defined Manhattan distance from base⁴. Each priority search region has maximal dimensions of 1.1km and 0.6km for a total area of approximately 0.5km² and are oriented with the long axis vertically or horizontally, chosen at random. Search missions are performed on the mountainous terrain using a larger map size to enable searching at a greater distance from base; see Figure B.1 for the larger terrain map. For each given rendezvous to base distance, an identical set of 10 regions are used for each combination of other variables.

The difficulty in evaluating these different relay approaches is due to the non-trivial relationship between rendezvous meetings and relay methods, both of which are dependent on associated parameters. When a fixed relay is used, the period between rendezvous is set by the relay as the estimated time for its immediate relay task. This is subject to a minimum time, 300 time steps, to enable a reasonable amount of searching to be completed, even if the relay will be ready for the next meeting in a shorter amount of time. Each searcher agent is free to search unhindered until the next rendezvous. For the altruistic method, the inter-rendezvous period is fixed, also 300 time steps, but this approach requires an individual data cache threshold be set. The combination of these determines if an agent will spend the entire inter-rendezvous period searching, or need to relay before the next team meeting. In addition, UAVs under this scheme pass responsibility for relaying data to others, often collating new data on fewer agents and causing some agents in particular to carry out relay behaviours sooner. In the following experiments a range of data cache thresholds is used with the altruistic relay approach to examine this dynamic.

A comparison of the completion times for several configurations of altruistic relay and a single designated relay are shown in Figure 6.14. Each deployment contains

⁴The Manhattan distance measure more accurately indicates the time required for performing a relay task than a Euclidean calculation. This is due to the limited movement of agents in the grid based simulation.

a team of 5 homogeneous UAVs in total, reducing the number of searchers to 4 where a fixed relay is used. Despite using a fewer number of potential searching UAVs the designated relay approach completes the search task more quickly for all regions except those nearest to base. The completion time varies little as the region moves further away from base; the 4 searchers cover the required space similarly, irrespective of its location within the environment. Completion times for the altruistic method all increase with the distance from base, reflecting additional time spent relaying. Where the search region is near to base and the data cache large, this approach can outperform a fixed relay deployment. All UAVs use the entire time between rendezvous to search, simply storing all new data which is then shared at the rendezvous. Only at this point does any relay hold sufficient data to trigger a relay task. If the relaying UAV returns to the search area before the next rendezvous it will use any available time to carry out sensing tasks. In this way the altruistic method can perform more efficiently in these limited circumstances.




Figure 6.14: Various altruistic and fixed relay configurations for teams searching regions at a range of distances to base. The data cache threshold is specified for the altruistic relay deployments. It describes the number of grid locations (10m by 10m) searched before the limit is met. Every data point is the average of 10 missions over a set of defined regions for each of the base-rendezvous distances.

The designated relay approach cements the relationship between rendezvous and relay, binding them together. This enables a more predictable rate of search irrespective of the relative distance to base. Whilst the altruistic approach accomplishes the maximal performance, its opportunistic nature is particularly sensitive to the specific situation. Furthermore it is less predictable, taking advantage of chance meetings to share data and pass relaying responsibility between successive rendezvous’.

6.6.4.2 Frontier Selection Heuristics for Rendezvous

Deploying the coordinated frontier search method with team rendezvous offers an additional reference point when selecting frontiers, the location of the next rendezvous. New heuristics can be calculated using this extra piece of data: including *minDistRdz*, the distance from the centre of the frontier to the rendezvous location; and *minDistRdz+Centre*, the sum of *minDistRdz* and the previously described *minDistCentre*.

The overall search behaviour for teams using a range of frontier selection heuristics is shown in Figure 6.15. These deployments use a single designated relay, with the remaining 4 UAVs carrying out search behaviour only. Using the extra knowledge of the next rendezvous offers no benefit over the best known performing heuristic, *minDistCentre*; conversely, it actually causes an increase in completion times. As agents are already being brought together periodically around a rendezvous point, they are unlikely to move any significant distance apart by simply using a greedy selection method. Furthermore, as the following rendezvous will be placed at the nearest searched/unsearched boundary, each agent will have limited navigation to reach the subsequent meeting. Using only *minDistRdz* - a distance measure of the frontier to the rendezvous location - alone performs particularly poorly, as agents will repeatedly move large distances to their next frontier, simply because it is slightly nearer to the upcoming rendezvous.

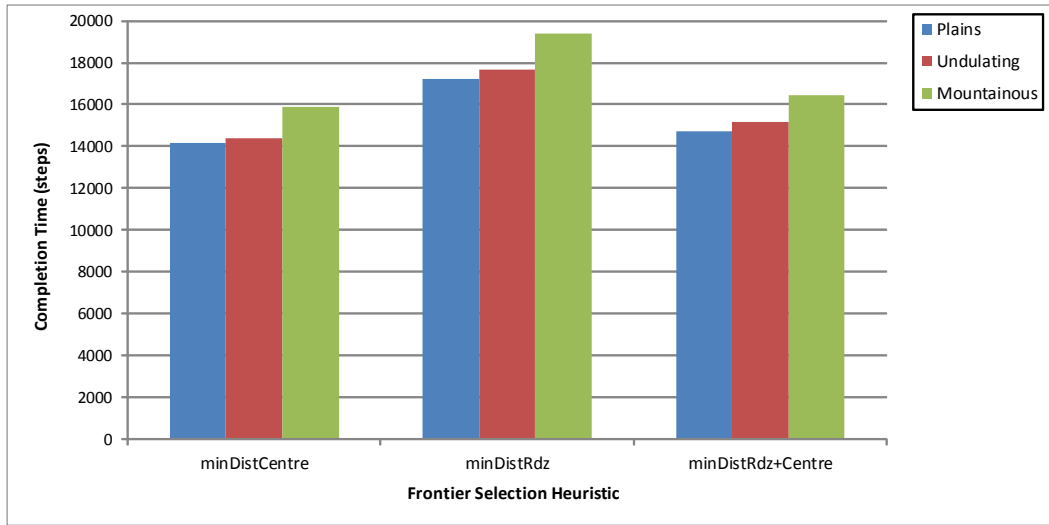


Figure 6.15: Completion times for search mission using various frontier selection heuristics.

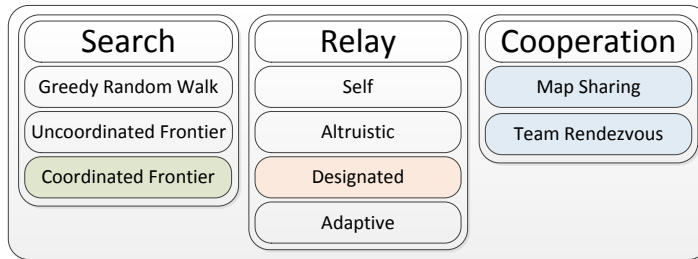
6.6.5 Data Latency

The majority of results presented thus far in this chapter focus on search performance and how information is disseminated within the team. However, the delay between sensing a location in the field and the associated data (including still images and video) becoming available at base affects the utility of this data to the mission. The data latency parameter described in Section 3.4 captures this delay per grid location searched. In these simulations the delay is recorded the first time each location is covered by a UAV only; counting the time from sensing until the relevant data is relayed to base.

In addition to providing a sink for acquired data, the relay behaviour provides a delay tolerant link to the team for human operators situated at base. This facility can be used during operations to issue new commands, such as to search a different region, or bring the deployment to an end, for instance *return to base* or *land immediately* if there is a potential safety issue. Therefore, there is additional motivation to limit the length of time required for communication between base and the remotely deployed team.

6.6.5.1 Multi-Link Relay

The aim of deploying multiple agents as relays is to reduce the time taken for new data to reach base. Selected agents are configured into a chain that operates along a planned path between base and the team rendezvous. The route is divided into segments with agent meeting points in between where data exchange occurs. Therefore each UAV travels only a fraction of the entire relay distance, resulting in a shorter time between rendezvous' for the search team as the relay servicing the team can return sooner. The inter-rendezvous period is also subject to a lower bound, ensuring searchers have an opportunity to gather new data, set to 200 time steps for these experiments.



Any number of available UAVs could potentially be used in the linked relay chain for the mission to be completed; provided at least one is available for sensing. To evaluate the impact this has on team performance, experiments must consider search missions at different distances from base. These simulations use a series of 10 randomly placed elliptical regions (approximately 0.5km^2 in area) at each 500m interval from 1500m to 5000m in Manhattan distance from base. Agents must completely search the given region of the mountainous terrain, with the simulation completing once all the required data is available at base. If a searching agent believes full coverage has been achieved, it proceeds to return immediately to base; it will not be necessary to attend any upcoming rendezvous. In these trials the maximum velocity of the vehicles is set as 20m per time step, double the maximum velocity permitted while searching but lower than the limit of 30m per time step used in previous exper-

iments. This ensures total relay journey times are sufficiently long enough to assess the effect of the various relay configurations.

The data latency achieved by teams containing 1, 2 and 3 relayers is demonstrated by Figures 6.16, 6.17, and 6.18, showing the average, maximum and minimum delays respectively. For search regions that are a relatively short distance from base, up to 2000m, the delay in sensing to base availability is similar irrespective of the relay chain length. At such limited distances from base, even a single UAV can complete the relay activity and return before the next rendezvous occurs, at the minimum period after the previous meeting. Therefore, additional relays offer very little advantage in terms of timely delivery of data to base.




Figure 6.16: Average data latency for a range designated relay chain lengths. Teams are tasked to search an elliptical region at a range of set distances from base. Each data point is the mean of 10 simulations.

As the distance between base and the remote search activity increases, the inter-rendezvous period becomes determined by the relay configuration. Each rendezvous must allow all team members to participate, including the assigned relay, and hence is planned based on the expected relay length. Where this exceeds the lower bound on rendezvous periodicity, for example a single relay requiring a round trip of 200 time




Figure 6.17: Maximum data latency for a range of designated relay chain lengths.




Figure 6.18: Minimum data latency for a range of designated relay chain lengths.

steps, data latency will increase. At this point a step change occurs in the relative delays of systems using different numbers of relayers; using one additional relay now provides lower latencies. This is the result of the lower number of relays not being able to return to rendezvous within the minimum time. This effect is observed in the plots at around 2000m from base for 1 relay and 4000m from base for a 2 relay chain.

Maximum data latency occurs when a location is searched immediately following a rendezvous, and this data is only delivered to base after a full rendezvous and relay cycle. This metric can be highly variable for an individual simulation, being subject to outliers where for instance a UAV fails to attend a rendezvous and therefore data is not relayed until after the subsequent meeting. This also affects the results presented here, although to a lesser extent as they are the average of multiple trials.

The minimum data latency is found when a UAV searches a location very shortly before a rendezvous. This new information packet begins transfer towards base almost immediately by the relay. A full relay path, rendezvous to base, is covered in a very similar amount of time independent of the number of relayers. Hence, in all configurations the minimal delay is very similar for each defined search region.

By assigning multiple UAVs to purely relay duty the number remaining to perform search tasks is consequently reduced. The effect this has on overall search rate is shown by the mission completion times in Figure 6.19. For each configuration the overall time required is determined by the number of UAVs active in sensing tasks, the relay chain configuration has little effect. It is unaffected by the distance from base that operations take place, as the relayers fully handle the responsibility for ferrying new data. Any minor increase in the overall mission time comes as a result of increased time to reach the search region after initialisation and potentially return once full coverage is achieved.





Figure 6.19: Completion times for teams using a fixed multi-link relay approach when searching a priority region at a range of distances from the base station.

6.6.5.2 Adaptive Relay

The adaptive relay approach attempts to balance team performance and data latency by enabling UAVs to switch between the roles of searcher and relayer. Hence, the composition of the multi-link relay chain varies during the mission. Additional agents are recruited for relaying when searching at a greater distance from base, and then returned to the sensing team when operations move closer to the command centre. The experiments in this section will evaluate this approach against teams with a fixed number of relayers.



To compare these varied strategies, the mission used involves searching a large area of 3.5km by 2km (7km²) to provide full coverage, along with relaying data to base in a timely manner. Agents must perform sensing tasks at a range of relative distances to base, from as little as 250m to over 4km away. In these trials the adaptive relay team is deployed with a range of values for its maximum data latency target threshold; which is used by the team to predict if the current configuration is satisfactory or requires modification. As the search area is particularly large, a team of 6 UAVs will be employed along with a minimum inter-rendezvous period of 200 time steps. Therefore, a higher basic search rate is expected compared to other experiments, due to both the increased number of agents and the longer period of time in between meetings to carry out sensing tasks. This parameter is selected in combination with the search area dimensions to allow the adaptive relay approach to use 1, 2, 3 or even 4 relays, if required by the circumstances.

The primary results from these experiments are the average data latencies, Figure 6.20, and completion times, Figure 6.21. Data from fixed relay teams is also provided for comparison. The adaptive approach is able to achieve latencies that match or improve upon fixed configurations with 1 or 2 relays and closely approach that offered using 3 or 4 permanent relayers. Using a lower maximum latency target, the adaptive approach will use additional relays whenever it provides a reduction in predicted data delay. However, this penalises the overall search rate and therefore mission completion time is extended. This trade-off is observed across the range of adaptive approaches trialled.

With only a single fixed relay, long delays can be expected between sensing and availability at base. The single UAV must travel the entire distance from rendezvous to base and back before the subsequent meeting can occur. The additional time provides freedom for searching agents to rapidly gather new data and the search is completed in the shortest time. The search area chosen is actually well suited to an approach using 2 designated relays, resulting in good performance on both metrics; a large portion of the search area allows the 2 relay chain to function efficiently. It will




Figure 6.20: Average data latencies for teams with a range of configurations, including various numbers of fixed relayers and adaptive relay methods.




Figure 6.21: Completion times for teams using a range of relay methods whilst searching a 7km² area.

operate less efficiently when either very near to base, where the second relay becomes redundant, or at a great distance, as data latency becomes unacceptably high. Therefore, the 2 relay approach would likely suffer if the search operations were required at only short distances to base or at greater distance from base. Deploying 3 or 4 full-time relays produces the minimum latencies but at a large cost in longer completion time; with only 3 or 2 searchers remaining respectively, the rate of information gain is significantly reduced.

The behaviour of the adaptive relay approach is influenced by both the minimum inter-rendezvous period and target maximum data latency parameters. Whilst the former does affect data delay, it should in general be set to enable an acceptable search rate for the team. The latency threshold chosen must be greater than or equal to the minimum rendezvous period. At its minimal value, a new relay will be recruited if the predicted delay with an additional relay is less than the estimate under the current configuration. This situation occurs when the final relay in the chain cannot return for the subsequent rendezvous before the minimum period has elapsed.

Increasing the maximum latency target allows the team to continue to use fewer relays for rendezvous' at greater distances from base; enabling more UAVs to participate in search operations. This is demonstrated by Figure 6.22, showing the average number of agents involved in relaying for a range of data latency threshold values. For example, using a value of 500 results in only one relay being utilised for the majority of the operation; hence, the latency and completion achieved is similar to a fixed single relay approach. Conversely, if the scenario requires data be delivered within lower delay, setting the latency target close to the minimum rendezvous period will enable this, but at the expense of increased relay role usage and longer search time.

The changing roles within the team affect search performance; this aspect of a single deployment is examined in Figure 6.23. The number of UAVs acting as relays over time is shown along with the rate of information gain for the team, measured by averaging over a sliding window of the preceding 1000 time steps to smooth short term variations. In general, when additional agents take up the role of relayer, the




Figure 6.22: Average number of agents operating as relays during search missions for adaptive relay teams with a range of target maximum latency threshold values.

search rate reduces as less sensing activity occurs. However, this is combined with some additional dynamics resulting from searching in a bounded area.

At initialisation, a single relay is used as the team searches the specified priority region, giving a limited search rate that sharply reduces as full coverage of the restricted region approaches. Once this stage is complete, the team approaches its maximum search rate as the entire remaining map is now available to search. At around $t = 2000$ the number of relayers increases to 2 as the result of a rendezvous location selected further from base; the information gain rate correspondingly decreases. This is observed as a delayed change to the team reconfiguration as the metric averages over 1000 time steps. Similarly this effect is present at $t = 10500$, relayers increasing to 2, and at $t = 26000$, beyond which 3 relayers are used. The reverse is also demonstrated, when the number of relay agents reduces from 2 to 1, at $t = 5500$, the team's search rate jumps back to near its maximal as 5 searchers are again active.

However, sudden step changes in search rate can also be result of the spatial dynamics of searching in restricted regions. For example, the sharp fall and rise that




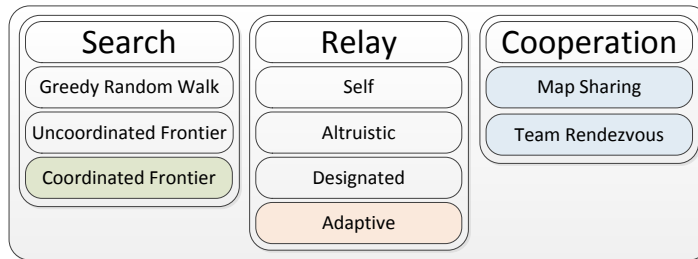
Figure 6.23: The number of relays used by an adaptive team over time and corresponding team information gain, calculated with a sliding window of 1000 time steps. Data taken from a single simulation for a team of 6 UAVs searching the undulating terrain, with a target latency threshold of 250.

occurs due to the priority restriction and subsequent removal of the limit. This phenomena is also observed when the selected rendezvous location has a limited amount of unsearched space nearby. Team members either spread widely, each attempting to find a unique frontier, or interfere with one another, both resulting in reduced sensing activity. This effect can be seen in this simulation at $t = 23000$, when the information gain rate falls suddenly, remains low for a period and then recovers back to its nominal rate for a 2 relayer/4 searcher configuration. The situation is resolved when a rendezvous location is selected that offers greater space for exploration, typically in a significantly different part of the environment.

Using this adaptive relay approach the number of relays is only permitted to change by one, increase or decrease, between successive relay cycles. An alternative approach would be to run an optimisation process when planning each multi-robot relay to select the number of UAVs to include. This would enable the number of relays used to vary arbitrarily by any amount to suit the context. However, the evidence from these simulations, both anecdotal and in Figure 6.23, suggest that the number of relays varies only infrequently. This is caused by successive rendezvous often occurring in very similar locations, hence the relay configuration need not be altered. Therefore the optimisation approach may have very limited opportunity to impact the relay usage.

6.7 Characterisation Experiments

The simulation results have demonstrated the use and performance of the considered approaches. In this section, deployments using the best performing techniques are further characterised in terms of scalability and sensitivity to inter-agent communication. The team is configured with coordinated frontier search, team rendezvous and adaptive relaying.



6.7.1 Scalability

All of the approaches presented in this work can be applied with any amount of agents, subject to some basic limitations. A minimum of one searching UAV is required at all times in order for the mission to progress and if designated relays are to be used, the number of agents must exceed those required for the relay configuration by one additional vehicle. Furthermore, from the simulation results a general guideline could be taken to keep at least 2 active searchers in order to achieve an effective search rate and to benefit from cooperative behaviour.

Teams using the adaptive relay method, demonstrated in the previous section, are able to change the usage of vehicles throughout the operation as the focus of attention moves in the environment. To evaluate how this approach scales with respect to the number of agents deployed, these experiments compare teams with between 4 and 14 UAVs on varied terrains. Figure 6.24 shows the completion times achieved by these search teams. An identical trend is observed for all three terrain types, as increasing the number of UAVs reduces search time. The improvement is quite significant for team sizes below 10 agents, for example deploying just a single additional UAV to grow the team from 5 to 6 can provide a reduction in completion time in excess of 15%. However, this improvement is not proportional to the number of agents and the return for adding more vehicles diminishes.




Figure 6.24: Completion times for range of team sizes, using an adaptive multi-link relay approach with rendezvous' and coordinated frontier search.

Once the team reaches a certain size, around 10 to 12 members, the search capability of the team is near its peak. There is insufficient time between rendezvous' for all of the UAVs to spread out, each find a unique area to search, and to gather a reasonable quantity of new information before returning for the next meeting. This saturation also causes a marked increase in interference between agents, especially at rendezvous'. Under this simulation, and also a typical safety practice with real aerial robots, a safe distance is maintained between nearby vehicles. With so many attempting to get close enough to the rendezvous location, in order to make contact with others present, planned paths often become blocked; causing additional maneuvering, costing time and energy. This result is particularly relevant to the size of search area tested ($4km^2$), when operating in larger areas the team may be able to use the additional robots more effectively.

Examining the average data latency for these teams, see Figure 6.25, data delivery to base is relatively unaffected by team size. Only the mountainous terrain exhibits any noticeable trend, with a slight increase as the number of UAVs grows. The

adaptive relay method only uses a sufficient number of agents as relays to meet the supplied latency target, this results in a stable data latency irrespective of the team composition.




Figure 6.25: Average data latency over various sizes of search team for each terrain type.

6.7.2 Communication

Communication plays a vital part in enabling the team to operate successfully; indeed it is key to the central requirement to transfer new information to base. Sharing information within the deployed team enables cooperation and a reduction in search recoverage. The free-space wireless communication range experienced by physical platforms is determined largely by the communication hardware, including transmission power generated and antennae. Such design choices are always compromises for rotorcraft vehicles. For example, increasing range by fitting a larger antennae and providing more power comes with a weight and energy consumption penalty, reducing flight time. To evaluate the importance of communication range, its effect on the team is investigated in this section.

Using the adaptive relay method a team of 5 UAVs is assigned to search a full standard map of each terrain type. The deployment is configured with a communication range between 100m and 700m; all other experiments thus far have used 200m. This distance applies when a clear line of sight exists between two UAVs. When communicating with base a minimum range of 300m is available, as the base station is assumed to be equipped with greater capability, when both transmitting and receiving, than the mobile robots. The time required to complete the search using these configurations is shown in Figure 6.26. Increasing the communication range reduces the time required to fully search the area for all terrain. However a significantly more capable communication system is needed to find large overall improvements. This result provides additional evidence that the rendezvous and coordinated frontier search approaches are able to offer a good level of performance in communication limited scenarios. Communication range influences the information flow in the team; a greater range allows UAVs to share data when further apart and hence more frequently. In turn this provides each searcher with an enhanced knowledge base from which to select frontiers, improving coordination and reducing recoverage.

The communication range also affects the delay between gathering new information through sensing and availability at base, see Figure 6.27. With increased communication range, average data latencies are reduced as information arrives in less time through two separate mechanisms. Firstly, relayers taking new data to base are able to transfer this information over a greater distance, shortening the relay journey time. Secondly, when searching in areas that are relatively near to base, agents can directly transfer sensory data even before the rendezvous occurs. This effect can be observed in the figure for ranges above 300m; this being the approximate shortest relative distance to base that UAVs are required to search.




Figure 6.26: A comparison of the mission completion times for teams with different communication capabilities.




Figure 6.27: The effect of communication range on data latency for adaptive relay search teams.

6.8 Summary of Results

This chapter has systematically examined the performance of several approaches to the three main aspects of team configuration, search, relay and cooperation. Coordinated frontier search was shown to offer increased search rates over both greedy random walk and uncoordinated frontier methods. Applying the team rendezvous concept further improves the team's efficiency, through enhanced data exchange and an opportunity to deploy a common relay solution for the entire team. Data latency, between sensing and availability at the base station, was shown to be significantly affected by the distance between base and search operations. A multi-link relay configuration enables data latencies to be reduced when searching at remote locations. With the addition of adaptive relay the team is able to change the number of vehicles involved in search and relay during the mission, balancing the compromise between search rate and data latency.

Chapter 7

Discussion

Further analysis and conclusions resulting from the experimental work are presented in this chapter. Firstly, the behaviour of coordinated frontier search is examined and potential improvements suggested. Data latency is investigated and in particular the performance of the adaptive relay approach discussed. The practical usage of these techniques is considered, especially regarding the suitability, and number, of potential robot platforms. Qualitative assessments and data visualisations that highlight aspects not captured by the simulation results are used to illustrate the arguments made.

7.1 Examining Coordinated Frontier Search

During experimental analysis coordinated frontier search was shown to outperform both greedy random walk and uncoordinated frontier search, enabling mission completion in a shorter time (see Section 6.6.1). This technique uses a small periodic broadcast of each agent’s location and bounding box approximating their current search region to enable implicit coordination of search behaviour.

Figure 7.1 shows the recoverage that occurs using the considered search approaches. In this context, recoverage is simply defined as occurring when a portion of the search area is sensed by more than one UAV. Using coordinated frontier search the recoverage is significantly lower than both greedy random walk and the uncoordinated frontier method. This considerable difference demonstrates how teams taking the coordinated approach are behaving much more efficiently; expending less effort on sensing the same space multiple times. For the team overall, this enables a higher rate of search and hence missions are completed more rapidly.




Figure 7.1: Recoverage during sensing using different search methods.

7.1.1 Frontier Approximation Methods

The coordinated frontier approach presented thus far uses an axis aligned rectangular bounding box as a basic approximate representation of frontier regions. This enables very short messages to be used to achieve the coordinated search behaviour¹. However, this type of representation will often not describe an arbitrarily shaped region well, covering a much larger area than the frontier itself. A range of other standard geometric representations exist and are commonly used in simplifying geometrical computation problems; for example, when testing for collisions in the physics engines used in computer games.

Figure 7.2 shows a selection of these methods applied to an arbitrarily shaped frontier. They range in initial computation complexity, the size of resultant representation and complexity of comparing between geometries. The axis aligned and minimum bounding rectangle are described by two and four vertices respectively. A convex hull is the smallest convex region that contains all of the required points (the frontier). Hence, the resultant representation can include any number of vertices up to the full boundary description. In terms of computation effort, finding the axis aligned bounding box is trivial. Convex hull algorithms can complete in $O(n \log n)$ time [108] and are actually required before a minimum bounding rectangle can be computed. The rotating calipers technique [143] can then quickly find this oriented box from the convex polygon.

The effect of these approximation techniques on search behaviour is demonstrated in Figure 7.3; showing the simulation results of a 5 UAV team with 1 designated relay using coordinated frontier search over the standard terrains. Where a more descriptive method is used, a small reduction in completion time is observed, most noticeably with the full boundary descriptor. Looking at the different terrain types, the mountainous environment sees the most significant improvement when replacing the basic bounding box approximation with other techniques. This is due to the general shape of frontiers created in the different terrains; generally, the more variation in ground height, the

¹For detail of how this information is used to coordinate the search, see Algorithm 4.1




Figure 7.2: Frontier approximation methods. The frontier region is shown in orange with its approximation outlined in blue.

more arbitrary the frontier regions, see Figure 4.6 for examples. For such irregular shaped regions, the basic approximation methods are a particularly poor descriptor, and in the context of the frontier selection algorithm, will exclude many more potential frontiers. This results in an agent having to travel further to find search regions that do not intersect with the known operating areas of other team members, and hence reduces the rate of information gain.




Figure 7.3: A comparison of the effect of frontier approximation methods on mission completion times. A 5 UAV team was tasked to search a 4km² area over the terrain types indicated. Data is the averaged result of 10 simulation runs.

The disadvantage of using these more accurate representations is the extra computation and, more significantly, communication requirements. Additional computation is required for both the generation of the approximation and when computing the intersections with similar representations of other regions². For each participating UAV the most frequent communication task is to broadcast the “alive” message, including details of its current location and searching area. Hence, the additional message length required, especially for the full boundary descriptor, may degrade communication performance elsewhere. Furthermore, the longer message may have an increased likelihood of erroneous transmission, particularly at the limits of the wireless communication channel.

Selecting the most appropriate frontier approximation technique for a given deployment will depend on the computation and communication capability available on the specific vehicles chosen. Where communication is the primary limitation, either the axis aligned bounding box or minimum bounding rectangle should be chosen. Computation concerns indicate the basic bounding box should be used, without sacrificing much performance. Convex hull offers a middle ground and utilising a full boundary description will increase the team’s search efficiency when neither on-board processing or communication limitations are a concern.

7.1.2 Communication of Agent Location and Intention

Coordinated frontier search requires that each agent periodically broadcasts their current location and an approximation of their currently selected frontier. Therefore, this information will only be received by those UAVs close enough to be in direct contact with the agent in question. The deployed robot group, subject to communication constraints, actually forms a multi-hop communication network and so potentially could be used to disseminate this information more widely. In this section, the utility of re-broadcasting these short messages is examined.

²Although the more complex methods will require additional processing, for a typical UAV this will be of little impact when compared to intensive tasks such as sensing and map merging

A team of 5 UAVs with 1 designated relay is tasked to search over the standard terrains using coordinated frontier search and team rendezvous. When receiving an “alive” message from another searcher, agents will consider re-broadcasting this information to others in communication range. Messages contain an extra value to count the number of hops it has travelled along. A new simulation parameter determines the maximum number of hops that a message will accrue before no longer being forwarded. When the maximum number of hops is 1, only direct communication is permitted.

Figure 7.4 shows the effect that varying the maximum number of permissible hops per message has on mission completion times. For this configuration, no appreciable impact is observed on the team’s search performance. This demonstrates that direct communication only to those in range provides sufficient information for coordinated the frontier search. In this case, taking only local decisions about which frontier to select cannot be improved upon using the multi-hop communication method.




Figure 7.4: The effect of multi-hop communication on completion times for a 5 UAV team using coordinated frontier search. The parameter is varied between 1, direct communication only, and 4, potentially passing the information to all team members where communication constraints permit.

The use of team rendezvous has a significant influence on this behaviour, by (almost) guaranteeing periodic data sharing for all team members. Following a team meeting, all agents present will have a similar knowledge base from which to generate and select frontiers; with all of this information gained via a single communication hop (directly). Hence, the use of multi-hop communication in this setting does not offer any advantage in team performance and uses the often limited communication availability.

The study of sensitivity to overall communication range in Section 6.7.2 demonstrated teams using coordinated frontier search were able to operate successfully with limited inter-agent communication. Reducing the basic wireless range down to 100m showed only a small reduction in overall performance; this is only 10 times the size of an individual grid location in our simulation. This investigation examined the effect of communication range for all types of connectivity, both short coordination messages and the larger exchanges required for map sharing. Combining this study with the above result suggests that deploying coordinated frontier search with team rendezvous may be applicable in other communication limited scenarios. For example in congested urban areas or for disaster search and rescue.

These conclusions are applicable to the configurations used here; namely a search team of about 5 UAVs with regular rendezvous meetings (at a minimum of every 300 time steps). For larger teams, or where there are extended periods between rendezvous, using multi-hop communication to disseminate this coordination data may offer an advantage. Further investigation is required to examine this issue under a wider range of conditions.

7.2 Further Analysis of Data Latency

Through simulation experiments the approximate relationship between a sensing location's distance from base and delay in that data becoming available at base has been established. This enables relay configurations, whether designated or adaptive,

to be planned against the scenario requirements. Although calculating the average latency for a specific deployment provides a useful metric to compare approaches, it does not fully capture the data latency profile.

Maximum data latency occurs when a previously unsearched area is sensed immediately after a rendezvous. Therefore, the newly acquired information will not reach base until an entire rendezvous-relay cycle is complete. This is the critical factor in ensuring maximal utility for the gathered data; if it does not arrive at base in a timely manner it becomes less meaningful and less useful to the mission objectives. Hence, approaches that limit the instances of large latencies may be preferable to those with a wide spread of latency values, including very short and very long, despite being comparable using the average metric. This is especially important in applications with mobile or even adversarial targets, where definite specific action is to be taken once a target is identified.

When the mission dictates that the important search regions are spread over a wide range of distances from the base station, it is difficult to select a designated relay configuration to provide adequate latency for all search areas without sacrificing significant search capability. Using too few relayers will result in unacceptably long delays in data delivery, whereas too many compromises search rate and overall completion time. The adaptive relay approach attempts to lessen the impact of this trade-off, by altering during operation how many of the limited number of UAVs are used for relaying. Simulations investigating this issue are presented in Section 6.6.5.2. Using this same data, Figure 7.5 shows the data latencies achieved by the location at which sensing occurred for 1 and 2 designated relays as well as the adaptive approach. For completeness the results for fixed relay configurations with 3 and 4 relays are shown in Figure B.4.

All of the methods are able to quickly relay data that is gathered near to base, but vary significantly for search operations at the opposite extreme. Data from approximately half of the map can be relayed with a reasonable latency (up to 300 time steps) using just 1 relay. However, for the more distant areas of the environment, this




Figure 7.5: Data latencies for data delivered to base shown at the original sensing location. Each data point is the average of 10 simulation trials on the undulating terrain using the specified relay configuration. Latency values are in time steps. The base station is located at the upper left corner, searching of the immediate surrounding area is not required. The total area is approximately 7km^2 with dimensions 3.5km by 2km .

configuration is insufficient and latencies increase rapidly. Deploying a configuration of 2 relays enables acceptable latency for the majority of this size of environment. Finally, the adaptive relay approach is able to offer further improvement by adding additional agents to the relay chain when searching the most distant locations. In this example, from the team of 6 UAVs, a maximum of 4 relays are used and an average of 2.3 throughout the entire operation. For the environment dimensions shown, the adaptive approach appears to offer a minimal improvement over 2 designated relays. However, for any further extension of the search area the latency performance of the fixed approach would continue to deteriorate significantly, as observed on a smaller scale with the single fixed relay. The adaptive approach however, continues to limit data latency by employing additional relayers until the team size limit is reached.

The data latency profile for these approaches can be further examined with a frequency plot, see Figure 7.6. The single relay configuration yields a multi-peak result; likely a consequence of the minimum inter-rendezvous period, set at 200 for these experiments. While the team is relatively near to base, the single relay can return to each subsequent rendezvous before this limit and must wait for the searching agents to arrive. The expected data latencies would therefore likely follow a somewhat normal distribution; as sensing occurs sometime during the constant inter-rendezvous period. However, when this barrier is breached, with successive rendezvous' becoming further apart, an increase in all latencies will be experienced. The magnitude of this increase depends on the additional inter-rendezvous time, in turn a consequence of the rendezvous to base distance. For a 2 relay configuration, the beginning of the same pattern can be observed, as the most distant rendezvous locations stretch the inter-rendezvous duration beyond its lower bound. Hence, for search regions at even greater distance the average latency for this approach would likely increase significantly.

The adaptive relay approach modifies the team to successfully reduce the occurrence of high latency data delivery, as demonstrated by the steep reduction in frequency above approximately 300 time steps. This technique is expected to be able to continue to manage data latency until the maximum number of relays is reached.

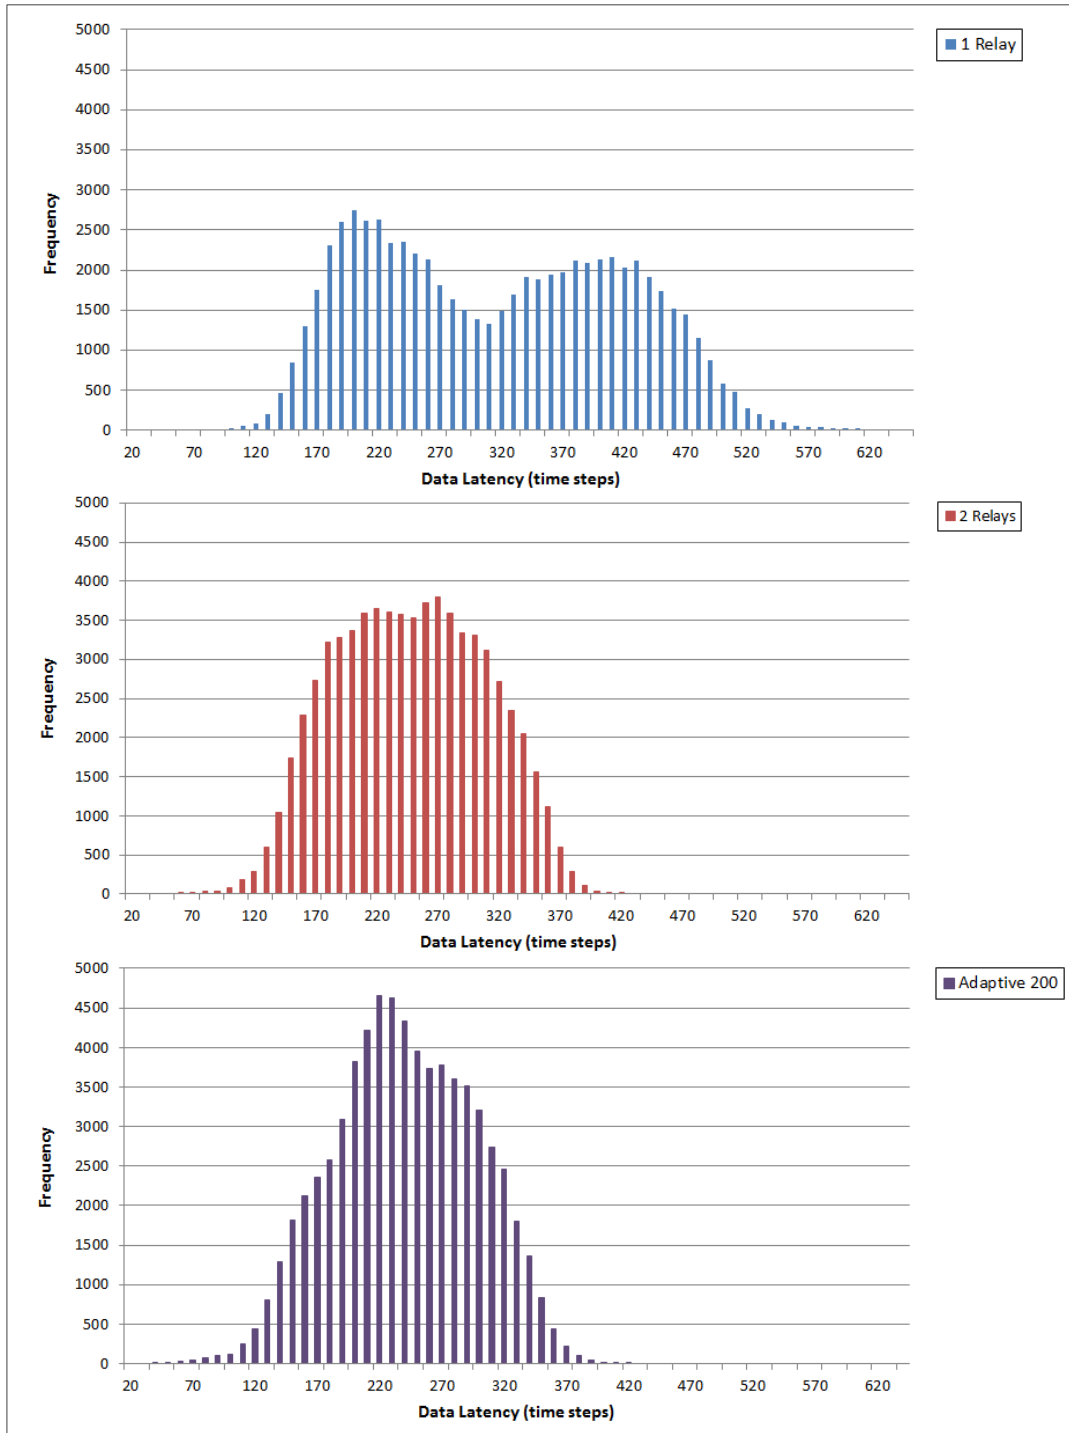


Figure 7.6: Data latency profiles for different relay configurations. The range and frequency of data latencies is shown for each approach.

Beyond such a point, it will exhibit the same behaviour as a fixed relay team and the overall number of UAVs deployed may be insufficient to realise an acceptable performance level. The adaptive technique is therefore particularly well suited to missions that require searching to be carried out at locations that cover a large range of distances from base; either a significant continuous space or multiple distinct regions of importance. It also supports the rapid deployment required for emergency scenarios, not requiring any prior calculation to select an appropriate number of relays for the given situation.

7.3 Usage

The simulation experiments presented in the previous chapter evaluate various approaches to multi-robot deployment through a range of mission parameters and terrains. In particular, the results aim to characterise the frontier search, team rendezvous for coordination and adaptive relay methods. This section discusses how, when and where these techniques, and others with similar qualities, should be used.

Although the simulations focus on the particular methods developed in this thesis, they are qualitatively similar to other approaches. For instance, frontier search is an example of a region segmentation approach, dividing the total area into smaller sections; this is in contrast to walk or path based methods, where agents will choose a path through the environment, searching on route. As a result the overriding conclusions may be applicable more widely to multi-robot system design.

7.3.1 Team Size and Restricted Regions

An evaluation of the scalability of frontier-based search with rendezvous was presented in Section 6.7.1. Increasing the number of UAVs was found to positively affect search progress, to a diminishing degree, particularly beyond around 10 agents in a single team. In general this is caused by a lack of nearby space that requires sensing and physical interference where UAVs disrupt each others movements. These issues are

a result of the rendezvous forcing all team members to periodically converge on a small area. For practical deployments, a large number of agents may cause additional issues where potentially large quantities of data must be shared by many participants³. These factors imply that a sensible recommendation for live deployments would be to impose a limit on the number of UAVs used in a single operation.

One of the main drivers for the proposed approaches was to enable a high level of autonomy to take full advantage of the excellent maneuverability of new robot platforms. However, a significant issue arises in trying to keep them from adversely affecting each other, especially when operating in restricted regions. This is best demonstrated by examining the rate of information gain achieved by a search team, see for example Figures 6.10 and 6.23. At any point when there is a limited amount of potential unsearched area for the team to explore, a sharp drop in the rate of new data accumulation is observed; particularly as a priority region or indeed the entire map nears fully coverage. Furthermore, this can occur at any point when the team rendezvous draws all agents to an area of the map that is almost completely searched. This will be exacerbated by larger team sizes, as squeezing more agents into a smaller space would not be an efficient usage of the available robots. Typical metrics do not allow a particularly close inspection of this issue; however, it can be observed through a graphical interface of the mission in progress. When limited to searching in small areas, agents can be viewed as replanning frequently and hence wasting possible time searching, as they receive many updates from other team members in quick succession.

Whilst coordinating UAV search behaviour (as with the coordinated frontier method) will reduce recoverage and inter-agent interference, such approaches make no allowance for limitations on the total unvisited area remaining for the number of active searchers. For example, under the currently implemented frontier approach a UAV will simply select the highest scoring frontier in the case that all of the potential regions are discounted due to conflicts with other known agents. This practical

³A potential approach for a single relay UAV to collect data from multiple transmitting vehicles via wireless communication is explored in [89]

approach ensures no UAV idles when it could be searching, but does not take into account the total unsearched area left to be visited by the team.

7.3.1.1 Multiple Independent Teams

The reduction in search efficiency experienced for both large team sizes and when operating in restricted regions are not specific to the individual techniques used here. Any approach in which many agents may be active in a limited region, even if this is not an intended team behaviour, will suffer with similar issues to some degree. This section considers whether dividing the total available number of UAVs into separate teams is likely to enable an overall improved performance.

Scalability of the number of agents is assessed in section 6.7.1, but the analysis looks at the team as a whole. Figure 7.7 shows the overall rate of information gain on a per agent basis, for a range of teams sizes. This is an alternative view of the data used previously and is generated from adaptive relay teams, with the team size indicating the total number of UAVs deployed, including any operating as relays. These results suggest that team sizes of between approximately 4 and 8 will offer the best search rates per agent deployed. When only 3 UAVs are present the team is essentially fixed, having 1 relay and 2 searchers, and cannot take advantage of the adaptive relay approach, limiting the search rate. With more than 8 UAVs in a team there is significant decline in search efficiency as the additional agents are under utilised.

The data suggests that when sufficient robots are available, splitting these between multiple teams may be a useful approach. The frontier, rendezvous and relay techniques do not contain any elements that preclude more than one team from operating in the same area and sharing the same base station (or indeed using multiple base stations). New information gathered by one team would actually be available to the others indirectly via the base station, with relaying agents able to update their knowledge whilst depositing new data. Take for example a deployment of 12 UAVs; rather than dispatch them as one team, creating two teams of 6 should enable an over-

all higher rate of search and hence shorten completion time. In addition, each team is able to use the adaptive relay approach to manage data latency at the required level. To minimise any inter-team interference, each team can be assigned a distinct region to search. Depending on the scenario, this could take the form of a priority region per team or simply dividing the entire map into equally sized segments. In a WiSaR scenario the mission commander may naturally assign such priorities based on their experience. As in the existing approach, searchers are only permitted to exceed their assigned search area once it is fully covered. Hence, for a two team scenario, a team's UAVs will not idle once their assigned region is complete, but continue to operate until the total mission objective is reached.




Figure 7.7: The average rate of information gain per UAV over the entire deployment. Data shown for a range of team sizes for all terrain types.

7.3.2 Robotic Platforms

The approaches examined in this thesis have primarily targeted the current (and near future) commonly available multi-rotor helicopter platforms. However, the capability of all facets of autonomous vehicles increases continuously as new technologies

emerge and reduce in cost. Whichever basic platform is chosen, a large degree of configuration is necessary to select equipment, such as, sensors, on-board computing, communication devices and battery size. With the simulation results known, the suitability of platforms can now be assessed for practical deployments of featured techniques. Furthermore, alternative aerial robots can be included for consideration, such as fixed-wing or hybrid vehicles.

7.3.2.1 Sensing

One of the major assumptions used in this work is that the UAV must be within a small range of heights above ground level in order to detect missing people. With the advancement of both digital vision sensors and accompanying detection techniques a wider range may be acceptable. In this case the discretisation of terrain height data can be simply adjusted, allowing the search approaches to be used unchanged. For example, if sensing is acceptable between 30m to 60m above ground level, adjusting the discretisation to a 30m step size (minus an amount for potential inaccuracy of the data) will enable generated frontiers to be covered without flight altitude changes despite the ground height variability within a specified area. Continuing with this example, UAVs would need to change altitude less frequently during operation, reducing navigation effort and consequently both search and relay times. For instance, the plains terrain contains a maximum difference in ground height data of just 63m and hence, using this updated sensor configuration, it could be completely searched from an altitude range covering just 40m.

Selection of a sensing paradigm forms an integral part of the robot specification when constructing or specifying vehicles for this type of system. Comprising of both sensor devices and associated processing techniques, it impacts the platform both mechanically and computationally. Using heavier sensors, such as cameras with glass or zoom lenses, or even gimbals, negatively effects overall weight, airborne time and flight dynamics. Any additional data produced, for instance from higher resolution multiple vision sensors, places extra burden on both data storage, processing and

exchange with other agents. Therefore, the ease with which the frontier search approach can be adapted to sensing requirements can influence this important decision at the design stage. This flexibility may allow a lighter lower resolution sensor to be selected, with the frontier method limiting the required altitude changes (and associated costs) during sensing. The alternative may be a heavier sensor that enables sensing from further above ground level but severely restricts flight time and requires additional on-board computation capacity.

7.3.2.2 Motion Capability

Simulation experiments have shown that, rather than being simply a secondary concern to searching, relaying has a significant impact on overall mission performance. While the motion capability of a UAV does play a role in searching, primarily when moving to the required location, it will have greater effect on relay activities due to the larger distance to be covered. Where a group of homogeneous vehicles is used, generally increasing the maximum velocity will reduce expected data latency. A positive effect would also be experienced for completion times if the adaptive relay approach is used; quicker relays would enable more agents to remain as searchers for a given latency threshold.

The majority of the approaches featured can also be implemented using a heterogeneous team, with a range of capabilities. In particular team rendezvous' are compatible with the use of purpose built relay vehicles, with a higher maximum velocity and enhanced wireless communication configuration to simultaneously collect data from multiple searching agents. Such specialised platforms can save weight by foregoing sensor devices and the computing abilities to process acquired data. This potential extends beyond rotorcraft and can also be implemented with unmanned fixed-wing aerial vehicles acting as relays and data collectors. These are particularly suited to the task, capable of high speeds but having restricted maneuverability which limits their potential for low-altitude sensing tasks. The team rendezvous provides a planned point that such a relay can collect data from a search team efficiently. Fur-

thermore, specialised relays are compatible with the multiple team approach; planning a route to visit more than one team's rendezvous in a single trip away from the base station.

7.4 Failure Modes and Response

One of the advantages of distributed multi-robot systems is a robustness to failures of any single participant. These systems should be sufficiently flexible to recover from a variety of potential problems during execution. In this domain there is a range of potential failure modes, from catastrophic to a minor error; the causes of which can originate from the environment or within vehicle hardware and software. These occurrences can cause complete loss of one or more robots, the impact of which may adversely affect the entire team. This section examines the specific failure modes and potential recovery for the relay/rendezvous approach and how the consequence varies depending on the current role of the failed agent. While the consideration of potential failure scenarios does not form a principal area of this work, it is a vital element of any practical usage of multi-robot systems. As part of the process of developing these approaches towards real-world usage, the impact of these failures would need to be assessed, firstly in simulation and then in practical trials.

Failure of a Searching UAV

At any given time a searcher is performing a task that it has individually selected, with no interdependence on other agents. Therefore, any failures during search activity will have only a limited impact. Immediately, any new data gathered by this UAV since the previous rendezvous, or recent data sharing opportunity, will be lost. Over the longer term, the overall performance of the team will suffer in line with the reduction in total team size. The bounded period between team meetings ensures that this type of failure has limited impact; any areas searched by the failed UAV will simply be revisited by others via the normal frontier process.

Failure of a Relaying UAV

The failure of an agent while involved in relaying will require a greater response to ensure the team is able to return to full functionality. Furthermore, the actions taken by other agents depend on the particular circumstances. A relaying agent never holds more information than the combined rest of the team, so whenever it is lost the team as a whole will not lose any information. However, additional data latency will almost certainly occur as the planned relay configuration is interrupted.

Where a multi-link relay chain is in operation, a single failure results in the failure of the whole chain to function effectively. Therefore, appropriate action is required to return the team to a full functional state. Other relay agents in the chain should wait at their predefined meeting points for a set period beyond the expected meeting time, in case the other participant is simply delayed. Once this timeout has been reached they will then become free to abandon their specific relay responsibilities and rejoin the main team. As the full relay configuration is generated at the rendezvous, and then passed along the chain toward base, all relayers have knowledge of the complete setup. Therefore, in the circumstance of a broken chain, they are able to head towards the known rendezvous location to locate the rest of the team. Such UAVs revert to searching as their default behaviour, selecting frontiers as normal. If they are not able to reach a known planned rendezvous, they will also behave as any other agent that has missed the team meeting, as described in the following section.

Searcher Misses a Rendezvous

A searching agent can fail to be present at a team rendezvous if it underestimates the actual length of time required to reach the location from its searching position. This can be caused by poor path planning accuracy or unexpected interference along its chosen route. The latter occurs often, even in simulation, especially when a large number of UAVs are operating in a small area. The only negative impact is increased data latency for any new information held exclusively by the UAV and potentially a small reduction in search efficiency as a data sharing opportunity is missed. To

mitigate against these consequences, all team members will share, if known, the details of the next rendezvous with others on request. Therefore, an agent can recover by learning the location for the next team meeting. If an agent becomes completely lost and separated from the team, it can relay periodically its own data to base, still taking any opportunity to learn about a rendezvous and fully rejoin the team. This procedure is implemented in the simulation software, enabling recovery from this type of event.

Relay Misses a Rendezvous

From the perspective of searchers waiting at a rendezvous, if the relay does not arrive an alternative method of organisation is required; usually the relay agent present specifies the subsequent meeting. The searchers present will continue to wait until a timeout duration has passed. At this point, the UAV with the shortest distance to base will select the next rendezvous, notify the team and switch roles to become a relayer. Therefore, the team continues to operate as normal almost immediately. Additional data latency may be expected, especially if a multi-link relay was previously in place. However, the team will continue to adapt adding more relays when necessary.

7.5 Spatial Progression of Search

As an aside to the main goals of improving search and latency performance, the spatial progression of multi-robot sensing is rarely discussed. Each different search technique will cover the required area differently, depending on the movement and sensing pattern of the agents. Figure 7.8 shows coverage maps at various stages of completion during a single mission simulation using both the greedy random walk and coordinated frontier search methods. The plains terrain type is used here as limited altitude changes are needed, hence allowing the basic search behaviour in a large open space to be more visible.

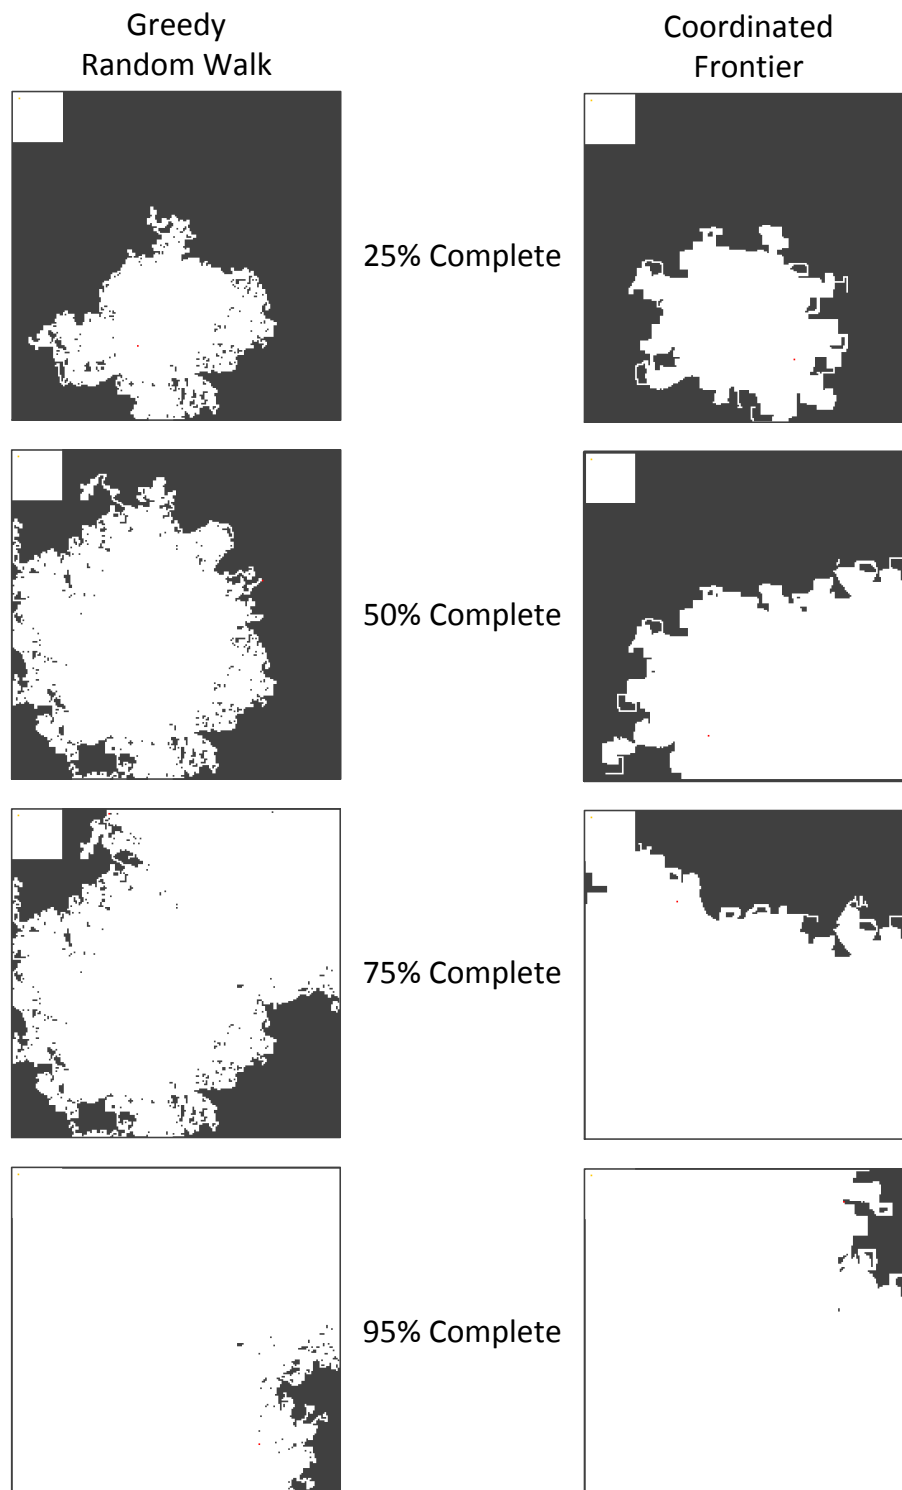


Figure 7.8: Coverage maps at set points during a mission for greedy random walk and coordinated frontier search approaches. Searched areas are shown in white and unsearched in black. The search occurs on the undulating terrain, beginning with an elliptical priority region just to the south of the map's centre.

After the completion of the initial priority region, the overall progression of search is heavily influenced by the random element involved in rendezvous location selection. Agents are forced to visit the team rendezvous' and hence will move to that part of the environment. This overall pattern changes with every deployment but is of little significance; for instance in the provided frontier sample the upper right corner is the final area to be searched, whereas the greedy walk example finishes with the lower right corner. Examining the coverage maps more closely reveals a key difference in how the space is covered. Both methods create a main searched area with a ragged searched/unsearched boundary at the periphery; the result of search activity that was interrupted by the need to attend a rendezvous or when new information caused replanning. However, for the frontier approach this main space is large and continuous, whereas the greedy walk approach will leave numerous small separate patches unvisited. These areas will later require searching and hence will consume valuable flight time.

In addition to the extra effort expended to cover any unsearched locations, this issue may impact the usefulness of the search to real-world applications. For instance, the WiSaR scenario will always be a collaboration between the human rescuers and robot search team. Therefore, at any given point during the operation the human commander may assess the situation or consider reallocating resources, the large continuous searched area provides highly valuable information. If some random locations within the space are not yet unvisited, this area will certainly require continued search effort, tying up resources that could potentially be redeployed.

The frontier method can be termed a *region-based* approach; search behaviour is carried out one region at a time, which is generally completed before moving on. In comparison greedy random walk is a *path-based* technique; each agent follows a path, generated according to a rule set, through the environment, searching at the locations visited. There are many other approaches that also use paths, pre-planned or otherwise, that may also exhibit this type of behaviour; hence this issue is worthy of further investigation.

7.6 Towards Practical Deployments

The simulation experiments presented in Chapter 6 highlight the complexity involved in preparing a group of robots for deployment. The sheer range of options can be overwhelming - even with the small selection of approaches trialled here. Indeed, many combinations of “good” approaches may not perform satisfactorily when applied together. For example, the application of team rendezvous when agents are responsible for self-relaying their own gathered data, actually reduces overall team performance (see Section 6.6.3.2). However, when combined with even a simple team relay approach, it is advantageous. This type of situation should be expected to arise readily when deploying complex multi-robot teams. Therefore it is vital to assess the deployment as a whole, rather than selecting which techniques to use individually. During the development of the coordinated frontier and adaptive relay approaches, the simulation software has been invaluable in this regard.

7.6.1 Integration with Search and Rescue Teams

For multi-robot search to truly impact real-world distributed sensing tasks they must be usable by non-experts. Furthermore, many scenarios, search and rescue included, require a rapid deployment process to minimise delay in beginning operations. To support this usage, we propose the creation of template team configurations that can be chosen from to suit the specific mission priorities. This will require preparation work to be carried out, both with the physical robot platforms and in simulation. For the WiSaR scenario, such efforts will form part of their regular training exercises.

Beyond investigative research work, simulation tools can be used for actual live systems to trial sample team configurations and select parameter values. During these preparation activities, the mission commander can use their experience to specify, and deploy, the UAV search group on typical mission types. Beyond gaining practical experience with this new resource, the data gathered can be used to tune and improve the robot search team. For example, real data about actual achieved

flight velocities and time required for sensing will customise simulation models to the particular terrain, vehicle type and usage required. Automated simulation tools can then be used to further optimise team configurations and parameters, creating sample deployments that are ready at short notice. For example: teams should contain at least 4 and no more than 8 UAVs; and configurations for maximum search rate and when low latency updates are required.

When deploying in a real scenario the mission commander must first gather all relevant data and form an overall plan. As with current practice, this should describe the roles of all available resources, including an overall objective for the UAV search team, such as a specific region (or regions) to be covered. A pre-configured search team template can then be selected to suit the particular mission. A well developed base station solution should be able to provide a suitable interface to the mission commander, enabling the team configuration to be fully described. The entire configuration can then be wirelessly communicated to the aerial robots at initialisation.

7.6.2 Deployment Example

This section describes an example simulated search deployment, using 3 different UAV team configurations. The scenario begins with a report of a missing person, with a last known location at significant distance from the nearest accessible location. The mission commander identifies a region that will be very difficult and time consuming for the human searchers to cover on foot. He uses the base station interface to select this region, as shown in Figure 7.9. For this example, three different team configurations will be examined; under real conditions, just a single one is chosen. Human responders place the 6 available UAVs on the ground and power them up. The base station wirelessly communicates the required information to all robots, including: the relevant discretised terrain data; the base station's location; the priority region; and their starting role (as searcher or relay). The UAVs take off and the search mission begins.

Figure 7.9 shows an overview of the terrain, with base station and selected priority

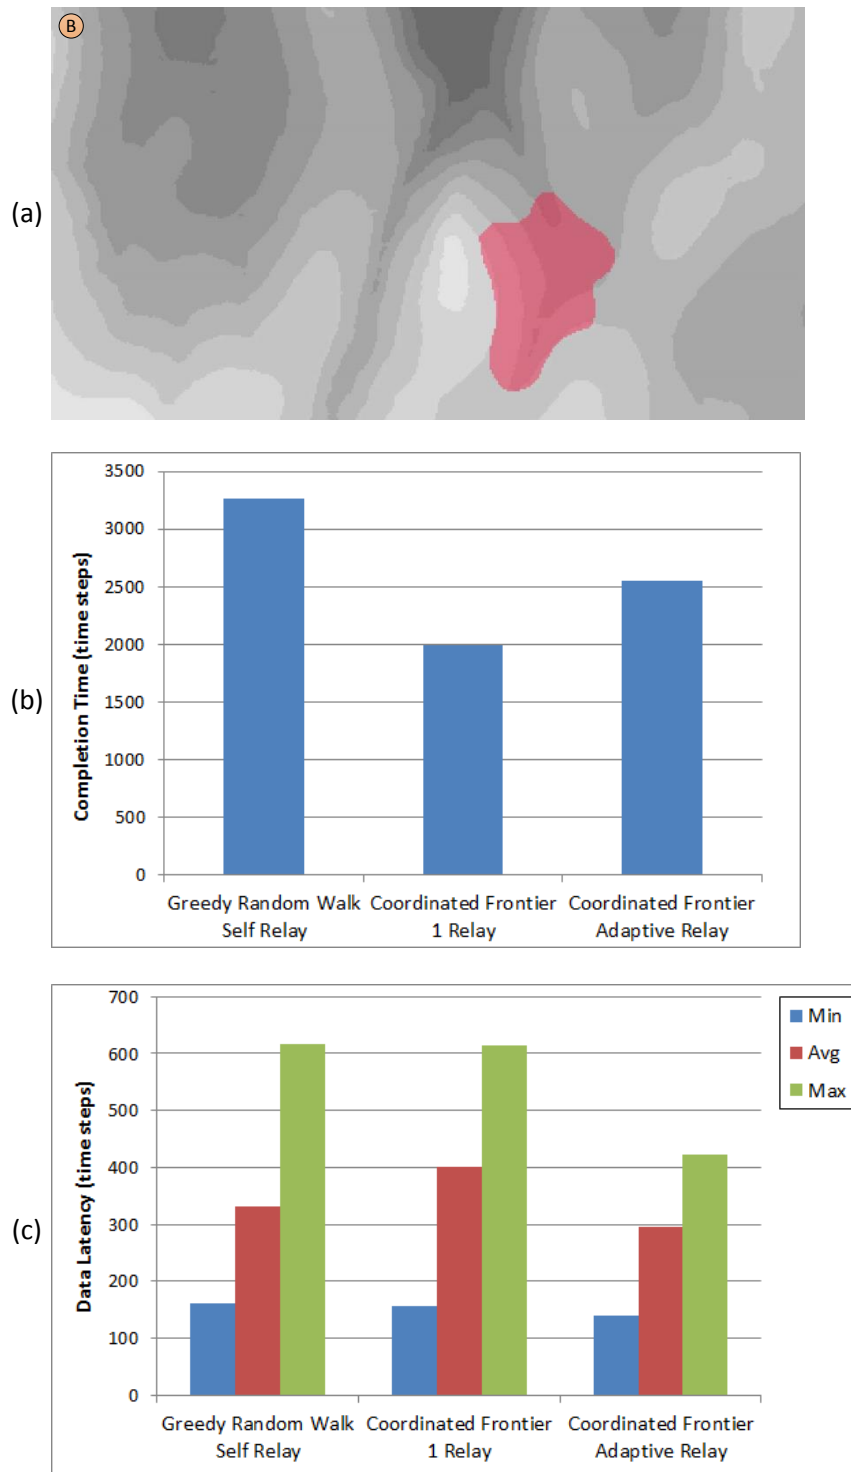


Figure 7.9: An example practical deployment using three team configurations: greedy random walk search with self-relay, coordinated frontier search with team rendezvous using both a single designated relay and the adaptive relay method. (a) An arbitrary priority region on the undulating terrain, (b) Mission completion times, and (c) Data latency performance.

region, along with the simulated search results for three team configurations. To complete the search in the shortest possible time, the team should be configured to use coordinated frontier search; with either a single designated relay, or none at all, depending on whether any updates are required before full coverage is achieved. If low data latency is desired, the adaptive relay approach is able to offer this at the cost of a slightly longer search completion time. Results using the greedy random walk method are provided as a comparison with the benchmark approach; it is not able to match either method for completion time and is worse than adaptive relay for data latency. In addition, this example shows that the combination of coordinated frontier search with adaptive relaying provides good performance on both search rate and data latency without extensive pre-configuration. This makes it particularly well suited to rapid deployment scenarios.

Chapter 8

Conclusions

This chapter summarises the overall body of work, highlighting the most important contributions and discussing the limitations of the proposed techniques. Potential directions for future work are described, including further developments to the approaches presented and related open questions.

8.1 Summary of Contributions

This thesis presents practical approaches to the deployment and organisation of a team of robots for distributed sensing operations. The principal developments are coordinated frontier search, team rendezvous and adaptive relay. In addition, real terrain information from common sources is used by the autonomous robots to select their own behaviour.

Coordinated Frontier Search

Frontier-based search has been adapted and developed for use in open outdoor terrain by aerial robots. Previously this technique has only been applied to ground-based robots, usually in restricted indoor environments. The developed frontier approach uses available terrain information to generate potential search regions that can be sensed without changing flight altitude. In this new method, terrain data is analogous to known obstacles in the original frontier search technique, with both helping to define edges to frontier regions. Each autonomous agent generates and then selects a frontier to begin searching immediately. The selection process can be implemented using a variety of heuristics to grade potential regions. Experimental analysis has shown that simply preferring frontiers whose centres are close to the robot's current location provides a good level of performance.

A coordinated version of the frontier approach yields efficiency improvements for a multi-robot search team. Agents successfully cooperate using a lightweight implicit coordination mechanism that requires only a small amount of information to be passed between team members. Each robot must, when communication permits, periodically notify others of its current location and an abstract representation of its current search area. This is sufficient information for a modified frontier selection process to take into account the behaviour of other agents.

Team Rendezvous

The introduced team rendezvous concept provides an opportunity for data exchange, coordination and shared relay of new data. This extends the previously used concept of a rendezvous as an agreed meeting between two cooperating agents. When combined with a designated relay agent, this technique allows newly acquired data from multiple robots to be collected by a single robot for transfer to a base station. Together this system of relay and rendezvous creates a delay tolerant link between a team of remotely operating robots and a fixed base station, facilitating both the transfer of new data and control commands during operation.

Adaptive Relay

The proposed adaptive relay approach is designed to change the team's configuration during operation to reduce the delay in new data becoming available at base, with a minimal negative impact on the search rate. It uses multiple agents in a relay chain and facilitates the change of roles between searcher and relayer depending on the current context. In particular, it was shown that the distance between base and the search operation, represented by the next rendezvous location, is the primary factor that determines data latency. The adaptive technique predicts the expected delay for the following rendezvous/relay cycle and adjusts the team as necessary by recruiting or releasing agents from relay duty. Hence when near to base the team uses fewer of the limited number of agents as relayers, enabling more of the team to conduct search tasks. The reverse occurs when search must be carried out at areas distant to base. This process is controlled by parameters that determine a target maximum data latency and the minimum period between team rendezvous.

A significant advantage of this approach is that it removes the need to preselect the number of relay robots used. This not only potentially speeds up deployment in emergency situations but also enables efficient use of the available resources over a wide range of situations. For example, if the scenario requires search to take place both very near to and at a great distance from the base station. The adaptive

method can adjust the team accordingly, whereas a non-adaptive relay configuration will perform less efficiently for a significant part of the mission.

Using Terrain Information

Information describing the search region's terrain is acquired from topographical maps and used in determining robot behaviour. In coordinated frontier search it enables the generation of frontier regions that can be searched by a UAV without having to change altitude. Furthermore, this data allows the estimation of more accurate relay paths, as required when planning the timings for team rendezvous meetings and especially in predicting data latency in the adaptive relay technique. Without this additional information source, the estimations of time required to navigate these routes would likely be more error prone, particularly over highly variable terrain.

8.1.1 Comparison with Previous Work

The contributions in this work attempt to maximise the autonomy of UAVs in WiSaR. In contrast, many of the previous approaches use full manual control [27] or only partial autonomous behaviour [51]. Coordinated frontier search extends previous work with ground-based robots [157] to enable this technique to be applied to aerial robots. This area-based search method varies from the typical path-based techniques that have previously been applied to UAVs, examples include [85, 148].

Previous approaches to multi-UAV systems have tried to mitigate limited wireless communication range by restricting UAVs movement in order to keep close together and ensure continuous connection [135, 66]. Team rendezvous, as deployed here, extends a previous idea from rescue robotics [122, 33] to provide periodic data sharing opportunities. This ensures the highly mobile vehicles have sufficient freedom for sensing activities between team meetings.

8.2 Limitations

These approaches target the particular problem presented in Chapter 3 and therefore have been designed to operate under a specific set of circumstances. There are certain limitations of the work that should be remembered when considering it for deployment. The most important of these are the ability of robots to limit uncertainty in sensing through accurate localisation and the avoidance on unknown obstacles in the environment.

1. Localisation

Robots in all of the approaches must have the ability to localise themselves. Multi-robot sensing and merging spatial data from multiple UAVs, in particular, depend heavily on localisation. Errors in this area would also cause issues in team rendezvous and relaying planning. While GPS can generally be relied upon for WiSaR scenarios, it is likely to be ineffective in certain locations, such as canyons and below the tree line. Other solutions, computer vision-based for example, can be used to supplement it for improved localisation, see Section 2.5.2.3.

2. Robot Maneuverability

Robots must be highly maneuverable, able to maintain position and have sufficient control to operate near other vehicles. This is required to ensure that multiple robots attending a team rendezvous can all move to within communication range of each other and avoid collisions. Hence, the use of certain types of UAV platform is precluded in a homogeneous team, including fixed wing aircraft (unless highly specialised navigation expertise is available). A high degree of maneuverability is also required in order to search the arbitrarily shaped frontier regions.

3. Availability of Prior Terrain Information

Coordinated frontier search relies on the availability of prior terrain information

to generate frontiers. Although it has been shown to perform well across a range of terrain types, unexpected changes in the ground height would likely reduce its effectiveness. Any unplanned alterations in flight altitude during the search of a frontier will negatively impact on performance.

4. Navigation Estimation

The team rendezvous and relaying approaches require an estimation of the time required to complete potentially long distance navigation. If the accuracy of these predictions is poor it could result in relaying agents regularly being late to attend team meetings. Where possible, a pessimistic approach to path estimation should be taken, as overestimation of actual time required is definitely preferable.

8.3 Directions for Future Work

Multi-robot systems is still a relatively immature research field and there remain many avenues for future work. This section describes some potential future directions for this topic, some of which may also apply more generally to the field.

Frontiers in Complex Spatial Data

In this work the frontier search method is adapted for use with ground height data, generating regions that can be sensed from a constant altitude. However, there is theoretically no restriction to the same technique being applied to other data types. For example, terrain classification information or probabilistic belief maps. Rules concerning how grid location similarity is measured must be developed for each type of data used. For instance, terrain classification data describes the ground conditions and how the land is used, such as farming, marsh, pasture, growing crops. Hence, neighbouring grid positions would only be added to the frontier if they were of identical type. Frontiers can then be searched with different techniques according to the ground class, for example a quick search over open ground but slower, more detailed

sensing when there is significant vegetation cover. Acting on probabilistic maps, the process would attempt to find continuous regions, where all locations fall within a limited belief range. This enables agents to search by region and use the frontier coordination mechanism in a multi-robot setting. Furthermore, more complex multi-layer frontier approaches are possible, where neighbouring locations are eligible to be part of the same frontier if they match on several types of data.

Coordinated Search Patterns

While the coordinated frontier approach reduces recoverage of the same ground, it does not explicitly coordinate the simultaneous behaviour of multiple agents. There is a body of research work that examines the use of coordinated search patterns such as flocking or formation flying with aerial robots, for example [116, 24]. Under such approaches, sensing can be performed by the team while moving in the coordinated manner. These methods may not be particularly suitable for long term operation, as they offer no long term strategic plan. However, the short term periods between rendezvous' may offer an ideal setting for highly coordinated behaviour, efficiently covering the search space.

Furthermore, these search patterns may have additional practical use if they can be successfully applied depending on the relevant terrain information. Particular patterns may perform better on specific terrain features. For instance, with rotorcraft UAVs with fixed cameras, given the relative camera to ground surface angle, what is the best method to search a steep slope?

Automated Terrain Analysis

The current approach enables a mission commander to direct the search effort by specifying a priority search region for the team. However, given the full range of terrain information available (ground height data, land use, paths, tracks, rivers, etc) such directions could be generated automatically. There are at least two ways to approach this issue, from the perspective of the missing person and by examining

the available resources. Some work on predicting the movements of missing people has already been carried out, for instance [86, 90]. However, it may be possible to automatically process maps to identify areas that would be especially challenging for human searchers to visit, assigning them to robot teams. In addition, any specific terrain features that are suitable for highly coordinated search patterns, such as formation sensing, can be identified.

Multiple Independent Teams

As discussed in Section 7.3.1, a significant reduction in search rate is observed when either a larger number of agents are in a single team or are forced to operate in a restricted space. Therefore, further investigation into deploying multiple independent teams is warranted, with the aim of maintaining an overall higher rate of search. The main issues to be examined are the division of work between teams, assignment of resources (how many agents in each team) and potential negative effects when different teams are in close proximity.

Spatial Progression of Search

The illustrations in Figure 7.8, and accompanying discussion in section 7.5, demonstrate how various search methods, in this case greedy walk and frontier, cover the space differently during search operations. This aspect is not captured at all by the typical metrics, which examine the overall performance of the team and its members. However, in practical applications it may be preferable for the space to be searched in a manner more typical of human search teams, for instance by region or along a particular landscape feature, for instance a path or river. The additional utility of knowing a particular continuous space has been completely covered makes the robot search team's work more useful to the overall combined search and rescue effort.

It may be possible to develop advanced metrics to evaluate performance on the basis of spatial progression. For instance, measuring the largest continuous searched

region covered by the team throughout the mission. This way of assessing team performance may also be applicable to other distributed sensing applications.

8.4 Summary

Many challenges remain in multi-robot system research before practical systems can be deployed in the field by non-experts. Beyond the maturation of vehicle technology, robust techniques are required to enable the deployment of truly autonomous mobile robot teams without the need for humans in the loop. This thesis attempts to progress the current state of the art towards this ultimate goal by developing coordinated approaches that maximise the use of autonomous behaviour. New application scenarios will continue to be proposed and the flying robots of the future will need to operate over the hill and far away from their human masters.

Appendix A

Multi-robot Simulation Software

All of the experimental work for this thesis was performed using custom-built simulation software. It features the import of real terrain data, fully independent simulation of multiple robots and a graphical user interface. This appendix describes some of its attributes and implementation details.

Implementation Detail

The simulator is primarily written in Java, with Matlab used during the import of terrain data (via the matlab control library [93]). Additional geometric calculations are performed using the Java Topology Suite library [144]; including implementations of the Graham Scan algorithm [52] for computing convex hulls, and the rotating caliper method [143] for calculating minimum bounding rectangles. Each simulation requires a significant number of parameters to fully specify the experiments. This is implemented using XML configuration files, fully describing a simulation, including for example, the terrain map, number of agents, robot capabilities and mission objectives.

Simulations are run in discrete time steps, with all agents having an opportunity to act during each iteration. The time resolution chosen allows each UAV to move at a maximum velocity (see following parameters list), and if not traveling too quickly perform sensing. This process is implemented in a single thread; all agent

behaviours are therefore calculated sequentially. Hence, the time required to execute these simulations can be considerable. A significant proportion of this computation involves planning navigation paths, especially when using larger map sizes. The standard A* algorithm is used for finding optimal paths over short distances. For longer routes, each agent will plan an overall path via randomly placed nodes; each step of which can then be fully calculated. Furthermore, the implementation uses a storage heavy approach to reduce memory allocation operations¹; path planning nodes are reused rather than re-instantiated. In combination these elements increase the speed at which experiments can be performed and also provide a random element that more closely represents the movement of real platforms.

Importing Terrain Data

The terrain information used in this work is from the Ordnance Survey product Landform-PROFILE [104] and provided through the Edina service [39]. The data is available as tiles of 5km by 5km; the UK Ordnance Survey tile system covers the country with labeled tiles, for example the plains, undulating and mountainous terrains are tiles *SE79NE*, *SX68NW* and *NO09NW*, respectively. Two versions of the data are available, contours (describing lines of identical height across the area) and, as used here, a Digital Terrain Map (DTM) (a grid of height values). This information was originally generated by low altitude aircraft and photogrammetric techniques; the data has a vertical resolution of 5m or 10m depending on area, with a claimed accuracy of $\pm 2.5m$ and $\pm 5m$, respectively.

For this work, DTM terrain tiles are acquired in *.tif* format and loaded into Matlab. The *geotiffread* function from Matlab's *Mapping* toolbox is able to import this data; with this particular data source the result is a vector containing the terrain data grid, a reference vector and a bounding box. The reference vector allows this data to be referenced correctly against other geographical data sources if required. A

¹While Java's garbage collector is convenient, creating and destroying many temporary objects was found to significantly burden simulation processing

newer product is now available, TERRAIN 5 [106], which features a higher resolution grid, improved accuracy and is available in a delimited ASCII file; trivial to import into any programming environment.

User Interface

The Graphical User Interface (GUI) provides several useful functions: an overall view of the locations of all agents during a mission, a current knowledge map for each agent and a planning interface. When required the commander interface allows the definition of an arbitrary priority search region for the team, see Figure A.1. This proved highly useful during development and testing, enabling quick variation of the mission to test the implementation under differing conditions.




Figure A.1: Mission commander planning interface. An arbitrary priority region has been defined, shown in red.

An overall view of a search mission in progress is shown in Figure A.2. A team of 5 UAVs is shown, with 4 searchers operating at a remote location to the base station. The single relay is shown approximately equidistant between the team and the base; during a relay task. The communication ranges shown demonstrate the maximum distance at which agents are able to communicate if at the same altitude.

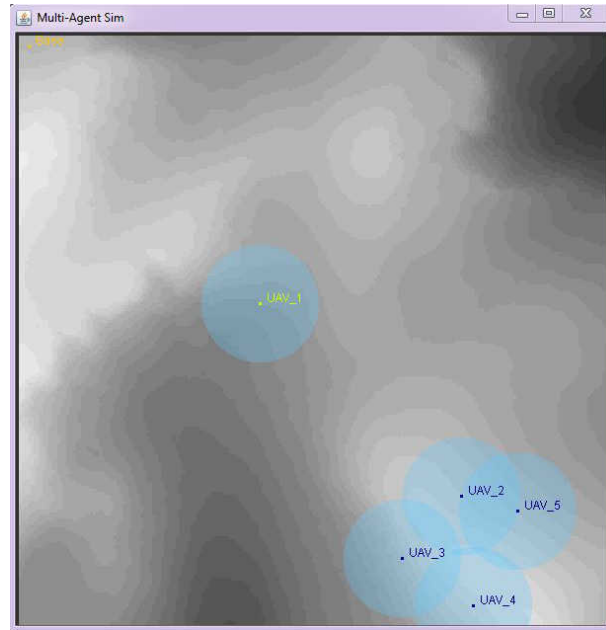


Figure A.2: Overview of a simulation in progress, including the locations of all agents. Searcher UAVs are shown in dark blue, a relay in yellow and base station at the upper left in orange. Blue circles denote a UAV's communication range in plane.

In addition, the current knowledge map for each agent is available to be viewed by the user, see Figure A.3 for examples. These views show the agent's current information about search progress; which spaces have been searched and those yet to be visited. Also shown are the most recently generated frontiers, one of which will have been selected for the agent's current sensing task.

Simulation Parameters

A wide range of parameters are required to configure the simulation software. Table A.1 shows the typical values used during the experiments of Chapter 6; any differences are mentioned alongside the associated results.




Figure A.3: Sample knowledge maps for a UAV at different points during a deployment, showing search inside and outside an elliptical priority region on the left and right respectively. Searched space is shown in white, unsearched in black and generated frontiers in bright colours.

Parameter	Value
Total Search Area	2km x 2km (4km ²)
Elliptical Priority Region	0.8km x 0.4km (0.25km ²)
Number of UAVs	5
UAV Maximum Velocity	30m per time step
UAV Sensing Velocity	10m per time step
Altitude Change Weighting	5
UAV Communication Range	20m
Base Station Communication Range	40m
Map Share Period	50 time steps
UAV Data Cache Threshold	300 cells (0.03km ²)
Maximum Frontier Size	100 cells (0.01km ²)
Coordinated Frontier Alive Message Broadcast Period	10 time steps
Minimum Inter-rendezvous period	300 time steps
Target Data Latency Threshold	250 cells (0.025km ²)

Table A.1: A set of typical simulation parameters. Distances are in metres (m) or kilometres (km). Time is measured in discrete time steps. Each cell (grid location) is 10m by 10m (100m² = 0.0001km²)

Appendix B

Additional Experimental Results

This appendix contains extra figures that are referred to in the main text.

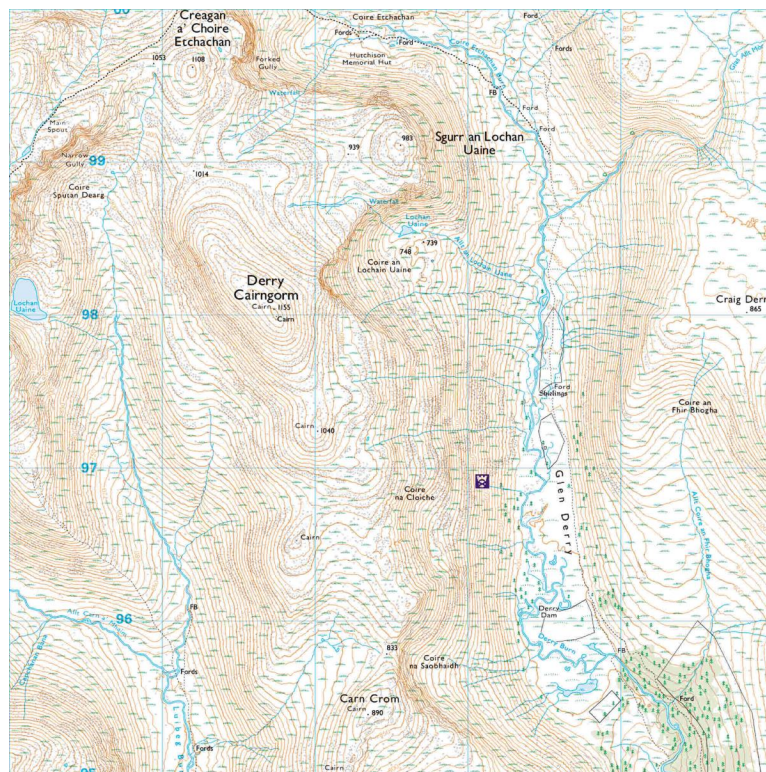


Figure B.1: Mountainous terrain used for experiments where agents search at greater distance from the base station. Total area covers 25km^2 (5km by 5km). Contours show isoheight lines at 10m intervals. ©Crown Copyright/database right 2013. An Ordnance Survey/EDINA supplied service[145, 39]

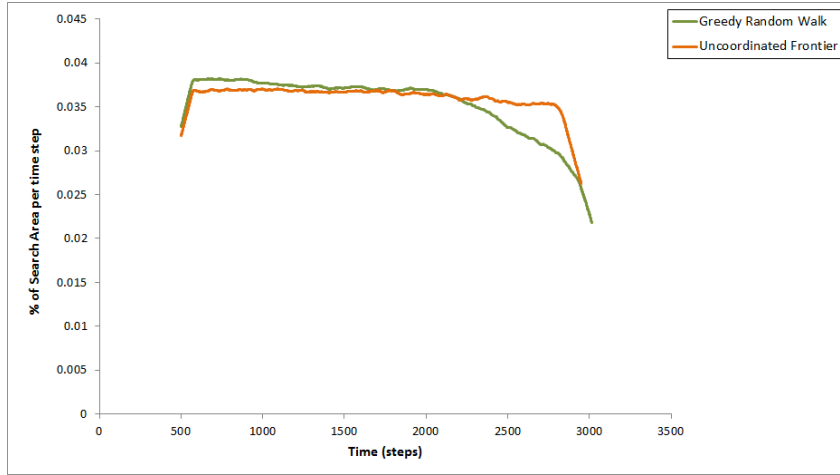


Figure B.2: Information gain rate for a single UAV searching an elliptical priority region only on the plains terrain, using both the greedy walk and frontier search strategy. For this experiment, no relaying occurs and the mission is complete once the UAV has covered the required area.




Figure B.3: Information gain rate for a single UAV searching an elliptical priority region only on the undulating terrain, using both the greedy walk and frontier search strategy. For this experiment, no relaying occurs and the mission is complete once the UAV has covered the required area.




Figure B.4: Data latencies for data delivered to base shown at the original sensing location. Each data point is the average of 10 simulation trials on the undulating terrain using the specified relay configuration. Latency values are in time steps.

Appendix C

Coordination in Multi-Tiered Robotic Search

This appendix contains a report on some preliminary work performed on the organisation of multi-robot teams. In particular, it focuses on structuring the team such that the coordination problem can be simplified. An abridged version of this work is published in [151].

Abstract

There are many problems which might be advantageously tackled with the assistance of heterogeneous teams of robots. For example, in search-and rescue one might deploy small helicopters to search from the air, and use these to direct ground-based robots towards potential targets, while using fixed-wing aircraft as communication bridges above them all. With such a varied set of resources the issues of tasking the robotic agents becomes non-trivial. In this study we define the notion of a robotic *tier* to capture the sorts of information connectivities that we are interested in, and consider a simulation of such a configuration in which we emphasise the information flow. For this we use a market-based approach to allocate resources, although the techniques considered will be applicable to other allocation strategies. The results confirm the importance of effective communication in such situations.

C.1 Introduction

Over recent years robot vehicle technology has matured to allow groups of mobile robots to be simultaneously deployed as *Multi-Robot Systems* (MRS), whose advantages include robustness to failures, reduced completion time and the ability to tackle more complex goals. A major design consideration is effective coordination, such as division of work, especially for heterogeneous groups. Such variety of robots is often handled using a generalised task allocation formulation to allocate tasks to individuals with suitable capability.

An important target application area for MRS is assisting in search and rescue operations, for example locating missing persons in the wilderness or post-disaster areas. The main goal is to firstly locate potential areas of interest and then gather data using on-board sensors of sufficient quality that the presence of people can be detected either automatically, by computer vision techniques for example, or by human responders. Such scenarios involve challenging environments where detailed information about the physical world is either not available or cannot be relied upon due to recent events. This puts additional emphasis both on the autonomy of the robot group and on its ability to operate independently of fixed infrastructure — such as a base station — and then to react dynamically. Such increased autonomy also reduces the number of human operators required for larger deployments. Furthermore typical wireless communication between robots can be greatly affected by the environment and may be only reliable over short distances or with clear line of sight.

This work is part of the SUAAVE project; SUAAVE (Sensing, Unmanned, Autonomous, Aerial VEHicles) [131] targets the investigation of robotic assisted search and rescue using quad-rotor helicopter UAVs (Unmanned Aerial Vehicles) as an experimental platform. The base platforms are out-fitted with computing hardware, imaging sensors and wireless communication equipment to form a networked group of autonomous robots. Therefore a heterogeneous group can be created with individual robots specification matched closely to their purpose.

This paper introduces a methodology for organising heterogeneous MRS in order to aid coordination. A review of some related work is given before an example deployment is presented. The simulation of this search and rescue scenario is described along with a discussion of the results.

C.1.1 Robot Tiers

A major aspect of any MRS is the number and variety of robots available. To simplify the vast possible variation, they can be classified into broad groupings using their type and purpose. This is also supported by the fact that robot vehicles tend to be designed to perform certain functions; the general purpose robot does not yet exist for real world environments. Furthermore, for the purpose of discussing coordination, a system can be described by the number of different agent classes, or *tiers*, it contains. For clarity, this way of categorising systems is not the same as general agent heterogeneity, which implies an arbitrary variation amongst the population. This designation looks from a system design perspective, where the participants are designed to perform one of a set of specific functions within a multi-agent scenario, each of which will be labeled a tier. The number of tiers in a system influences the suitability of coordination techniques.

Consider the simplest system, containing a single type of agent all with identical purpose, for example a group of identical ground vehicles all performing the same search behaviour. At this level, coordination mechanisms tend to relate simply to the division of work amongst agents, for example assigning separate areas to each robot for exploration. A common two tier system consists of a search or exploration tier and an exploitation tier, for example the site preparation problem [111] requires firstly the detection of obstacles within an area and subsequent clearance; two tasks that require markedly different capability.

In the search and rescue domain, the two tier system includes search agents able to rapidly sense an area to find potential target sites before a more specialised vehicle is tasked to investigate these points. We may then wish to add further tiers to

include communication infrastructure or even human responders as a resource to be coordinated with (as in Figure C.1). Adding both of these (i.e., communication and responders) generates a system with four tiers; some system (e.g., critical military command-and-control) may have even more tiers. A specific example of a search-based 3 tier system, comprising a ground vehicle, a low flying rotor-craft UAV and a high flying fixed wing UAV, is presented in [149]. However in this instance autonomy is limited, with each tier containing only a single agent and all activity is coordinated from a base station.

Outside of the robotics domain, there is a parallel with the organisation of emergency disaster scenarios, with services including police, fire, ambulance, specialist rescue crews and mission control. Considering each of these groups as a tier, behaviour is coordinated using members of each tier to cope with the specific requirements of many separate challenges.

C.2 Related Work

Task allocation for general multi-agent systems (MAS) typically concentrates on the optimal assignment of tasks to agents, without considering constraints on communication or energy usage. These non-embodied agent based systems are not restricted by the physical configuration of the computing hardware and thus can transfer resources between agents. Communication between agents is assumed to be lossless and uninterrupted. The traditional paradigm of MAS is that each agent will act to maximise its own utility, with the area of cooperative MAS, or the similar field Distributed Artificial Intelligence, aiming to increase the performance of the entire system. The latter is closer to the typical MRS as they are deployed to achieve global goals, for instance mapping an environment.

One approach used in MAS is to form coalitions or teams of agents in response to currently unassigned tasks. This assumes that the tasks are able to be performed efficiently by a team, which in general can be achieved with informational agents

due to the transferability of resources. Coalitions can be formed from the agent population in a distributed fashion [128]. In this work all possible coalitions are considered; this number is reduced by heuristics and finally the lowest cost solution is selected. Smaller coalitions are preferred due to the extra communication and coordination cost required to operate larger ones. Agents are able to be part of multiple coalitions, provided their role in each does not overlap. A further step is to permit only a limited number of base inter-agent connections, reducing the possible teams that can form [47]. The agent network is adapted by rewiring connections, based on past performance of that connection. For example, connections that have been used in successful coalitions will remain and those that have been unused will be disconnected.

Task allocation in MRS has been studied for some time; the work on area exploration is of particular relevance, including constraints such as line-of-sight [4, 66] and utility heuristics such as frontier size [18, 119] or with market bids [159, 129]. The scenario found in the RoboCup Rescue Virtual Robots Competition is particularly relevant to our own work; here using a 3D physics simulation allows virtual ground and aerial robots to cooperate in helping to detect casualties after some make-believe disaster. Including real-time scanning into these models has allowed these to be incorporated into this setting [80] with good results [114, 32]. One ‘grand challenge’ project was recently staged in the UK in which the participants used the type of communication structure explored in this paper. The Ministry of Defense (MoD) required investigation of the use of robotic technology in the problem of locating armed threats in an urban setting. Using an army training village teams typically used tiers of aerial vehicles for communications and surveillance, directing ground vehicles to observe potential threats (i.e., guns and bombs!) [8]. Although tested over a small scale – the training village extended over a few hundred metres only – communications were discovered to be a real issue, and a limiting factor in the ability of the teams to perform their tasks.

In [50] multi-robot task allocation is considered by the type of tasks present in

the problem being tackled, using three major properties: (a) can the robots concurrently perform single or multiple tasks; (b) can the task be performed by a single robot or only by cooperating robots; and (c) are forthcoming tasks predictable. For robots that are only capable of single-tasking then the issue is of task allocation. An approach that has been carried forward from MAS research is that of auction or market based allocation; robots place bids for tasks that they are able to achieve based on the expected cost of completion, for example [35, 50, 121]. The RoboCup programme provides a testbed for several robotics research directions, with task allocation being relevant in both the soccer [79] and rescue [80] scenarios. A common approach in RoboCup soccer is to assign roles to the team members, with reallocation occurring throughout the game [83, 146, 153]. (In our setting the classic soccer game is comparable to a 1-tier system with each player acting individually, although often implemented as a 2-tier system with a common sensing module sending positional data to the players. Later games also incorporate a special ‘coach’ agent, which can be thought of adding an additional tier to the situation during time-out periods only.)

In the multi-robot exploration problem, a group of robots must explore the environment, often building a map and/or searching for targets. A common approach is to assign separate robots to the currently unexplored areas to improve efficiency [155, 120]. [119] takes into account the constraints of wireless networking. In contrast, coordinated multi-robot exploration is possible with a distributed setup [129], which uses the popular frontier approach [157] and a distributed form of market based task allocation.

Certain goals cannot be achieved by a single robot alone; for example, box pushing, area clearing, multi-UAV lifting [97], strongly cooperative search, and localisation [110]. The use of multi-robot strategies is combined with a market based allocation approach in [73]. This type of approach can also be combined with simpler forms of coordination in a hybrid approach, for example to coordinate multiple robots in a surveillance sweep application [74]. [134] combines a market-based approach with schema mapping of sensor-actuator information to permit both strong and weak cooperative behaviour

in a site clearing scenario. Robot coalitions are formed using pre-determined schemas that describe how robot sensors, actuators and communication is used in the team [109]. Task allocation is performed using a bidding system with centralised control. Attempts have also been made to modify coalition formation methods from MAS for the multi-robot domain [147].

C.3 Simulation Experiments

In order to investigate a multi-tier robotic system a search and rescue scenario was devised, with the goal of searching an unknown area for the presence of targets and gathering data through close proximity sensing. A three tier system was chosen, see Figure C.1, with each covering one of the key roles, namely data collection, search, and communication infrastructure; containing the robot types *GroundGatherer*, *UAVSearcher* and *UAVCommBridge* respectively. To achieve the goal, targets must first be located by a *UAVSearcher* and assigned to a *GroundGatherer* using a task allocation procedure, who will then approach the target. We make the assumption of limited communication at near ground level due to the environment, hence the deployment of *UAVCommBridges* specifically to increase connectivity via multi-hop networking. This forms the platform on which distributed task allocation can be performed. The UAVs are modeled after quad-rotor helicopters with the ability to hover in position and downward facing target sensors where applicable. Each type of UAV operates at a different fixed altitude; *UAVSearcher* close to the ground for detecting targets and *UAVCommBridge* at a suitable altitude to reduce wireless interference by maintaining line of sight to other aerial vehicles.

This 3-tier scenario was chosen as a careful balance of generality and simplicity. Having more than two tiers means that we have to consider more than one type of communication transfer, whereas the task and communication simplifications allow us to focus on the inherent complexity introduced by the communication layers. Particular focus is placed on coordination in the tiered system and how this is affected




Figure C.1: Three Tier Multi-Robot Search

by the system configuration. The actual task performed once the target location is reached is of no interest here, and the assumption that communication is simple at the highest tier will be true for many real systems (but not all).

C.3.1 Experimental Configuration

The multi-tier scenario was implemented with a discrete time simulation in a 100×100 grid-based environment with a minimum resolution of 1 square. All robots are modelled as occupying one square, with velocity limited holonomic mobility and localisation assumed to be in a common frame of reference.

Three tiers are present in this robotic search scenario each containing a different vehicle type, *GroundGatherer*, *UAVSearcher* and *UAVCommBridge*. Ground robots operate in the ground plane, navigating around obstacles using a type of Bug algorithm [88]. When allocated a task they move to the target location to gather data at close proximity, before returning to a default dormant state (to conserve energy). The middle tier contains *UAVSearchers* that actively seek out targets using a spiral pattern to ensure area coverage with minimal overlap. The sensing footprint – the ground area in which targets are visible – is a square whose side is simulation parameter. Finally, vehicles in the communication tier attempt to remain in communication with *UAVSearchers* by keeping spatially close. Using positional information gained from regular alive messages, each *UAVCommBridge* moves toward the average posi-

tion of nearby *UAVSearchers*. In addition they move away from each other to avoid clustering and increase communication coverage. This demonstrates how the tiered approach can simplify coordination; with communication bridges using information from both UAV tiers.




Figure C.2: Graphical representation of simulation, with 3 types of robots, obstacles in black, targets in red, communication ranges as blue circles and UAV sensing footprints as red squares

Wireless communication is modelled simply by a hard distance limit, with range determined by the vehicle types. Connectivity between *UAVCommsBridge* robots is assumed to be unrestricted due to the lack of interference at higher altitudes. An important parameter is the communication distance limit between ground and aerial-based robots as this models the effect of the environment. Multi-hop networking is used with topology restricted by robot-tier, see Figure C.1, and connectivity assumed subject to the range limitations. In this configuration there is a maximum of 5 hops between any two vehicles (*GroundGatherer* - *UAVSearcher* - *UAVCommsBridge* - *UAVCommsBridge* - *UAVSearcher* - *GroundGatherer*).

For simplicity distributed task allocation is handled by an auction based mechanism, with the UAV search tier coordinating the ground tier. When a *UAVSearcher* locates a target, it pauses search activity and begins an auction by broadcasting an announcement. Any *GroundGatherer* robots that receive this message and do not currently have a task will reply with their bid, calculated with a simple metric of distance to the target location. After a short timeout period the winner, with the lowest bid, is notified and begins performing the task. The overall goal of the multi-robot

search is to locate and gather information about a number (20) of static targets. At any time there are 10 targets present in the world, each being replaced by a new randomly positioned one when visited by a ground robot.

Parameter	Value
World Dimensions	100 x 100
Number of GroundGatherers	5
Number of UAVSearchers	5
Number of UAVCommBridges	3
Number of Concurrent Targets	10
Max. Velocity GroundGatherer	1
Max. Velocity UAVSearcher	2
Max. Velocity UAVCommBridge	2
UAVSearcher Sensing Footprint	5 x 5
GroundGatherer to UAVSeach Communication Range	15
UAVSearcher to UAVCommBridge Communication Range	25

Table C.1: Simulation Parameters (units are % of side of world)

C.4 Results

The described simulation was run with a range of different parameters to examine the behaviour of the multi-tier system. A full set of default values are given in Table C.1, with results presented as deviations from this base. The data presented is the result of averaging 50 independent runs, each of which is initialised with robot and target locations drawn from a uniform random distribution. (A small number of experiments were also tried with a much larger world size, and a corresponding increase in the number of agents; as expected, the results were very similar.) The following sub-sections highlight the more interesting effects that were discovered.

C.4.1 Ground to UAV Communication

In this scenario communication between the ground-based and aerial robots is vital to allow distributed task allocation. Limitations due to the environment can be

roughly modeled by manipulating the communication range. For example, a cluttered environment provides a greater degree of interference and thus lower inter-robot connectivity than a open area with free space. The effect on overall system performance from the communication link between these tiers is shown in Figure C.3, where the time required to complete 20 tasks reduces as communication range increases. This is largely due to an increased number of participants in each task allocation round, see Figure C.4(a), which in turn gives a higher probability that once a target is located it can be allocated with the first task allocation procedure. An additional result of increased communication at ground level is the reduction in message hops between auctioneer and the allocated robot, Figure C.4(b). Where this connectivity is more limited, multi-hop routes through communication bridges are relied upon to include robots in task allocation; as demonstrated by a higher number of average hops from auctioneer to task winner. Conversely, larger communication distances result in a greater proportion of tasks assigned to robots within direct communication with the auctioneer.




Figure C.3: Average time to complete 20 tasks for a variety of robot group sizes with varying communication ranges

C.4.2 Number of Ground Robots

The number of ground robots represents the availability of the critical capability to achieve the goal. Increasing this number boosts both the total strength of the robot group and reduces the expected distance between any located target and an




Figure C.4: (a) Average number of bids placed per task auction, (b) Average Number of communication hops between the auctioneer and successful bidder

individual capable of completing the task. Figure C.3 demonstrates the intuitive reduction in time required to meet the goal when larger numbers of ground robots are deployed and similarly to communication range, saturation is achieved where further increases do not yield improvement. Furthermore increased availability leads to more participants for each task allocation, Figure C.4(a), due to reduced average distance between robots and hence a higher number of contactable robots. Similarly there is an increased likelihood of a ground robot being within a single hop of the auctioneer, resulting in fewer hops between these robots in Figure C.4(b). In some application areas the energy consumption, largely consisting of power for mobility, is a critical factor. Figure C.5(a) shows the relationship between the number of task capable ground robots and the average distance traveled per individual. This shows a roughly inverse relationship, indicating good utilisation of the ground agents (in this relatively crowded environment). In a situation in which energy consumption is critical more global planning of paths would be expected to become important, both to conserve energy and to move in order to replace it.

C.4.3 Robot Velocity

Robot movement has a direct and intuitive impact on their performance, for example the time required for ground robots to reach target locations and UAVs to search an area. The combination of velocity and sensing footprint gives the area

searched by each UAV per time step and hence influences the rate at which targets are found. However an additional result is the effect on network connectivity as higher velocity travel breaks communication links by moving out of range. Figure C.5(b) demonstrates this occurring as the velocity of *UAVSearchers* increases the number of participants per task allocation round decreases heavily. This is a particular problem for distributed algorithms that require a period of consistent connectivity for several message exchanges, such as the implemented auction system. (An example of an alternative strategy that might work better in this situation would be to have each aerial vehicle keep track of each ground vehicle, and to try tasking a specific vehicle rather than using a multi-pass allocation procedure. In the presence of uncertainty about the status of vehicles this might not succeed every time, but even with retries it could work better for moderate numbers of vehicles.)




Figure C.5: (a) Distance traveled per ground robot, (b) Average number of bids placed per task auction for a range of *UAVSearcher* velocities

C.5 Appendix Conclusions

This paper presents a method for organising heterogeneous multi-robot systems into tiers, according to the roles performed. The main advantage of this approach is to classify and hence simplify the coordination problem for heterogeneous systems. To demonstrate this proposal a robotic search and rescue scenario was implemented in

simulation, containing a variety of robot vehicle types in 3 tiers, this being regarded as interesting, non-trivial, and potentially useful.

The results generated from this experimental work highlight the important variables of the multi-tier system. The ability to communicate was shown to have a significant effect on the overall performance measure of time to meet the goal criteria. When connectivity to some members of the robot group is limited, typically due to their spatial relationship and the environment, key capabilities cannot be accessed on demand and hence the full potential of the group is not realised. We demonstrated that deploying robots specifically to aid communication by connecting multi-hop routes can counteract this problem; in this case by allowing more participants in task allocation protocols. An additional benefit of increasing connectivity potential is improved ability to report data quickly to nearby external computing infrastructure, such as the sinks used in sensor network applications.

A further outcome was the demonstration of the dual roles played by all robots in a dynamic system, highlighted in this instance by the searching UAVs. Whilst the current state of the art vehicles are able to travel at high speed, hence completing operations such as sensing quickly, we demonstrated that this has a detrimental affect on the vital communication links. This is particularly important for implementations that require periods of consistent connectivity, such as data streaming and the task allocation auctions deployed here. There is a clear trade-off to be investigated for highly mobile networked robotic systems and successful implementations need to match carefully the constraints of connectivity against the advantages of mobility.

The future direction of this work is toward further investigation of the potential for the multi-tier organisation paradigm as a means of simplifying the coordination problem. We intend to use characterisation data from the available physical UAVs to inform more detailed simulations. Particularly we are interested in combination of complex spatial relationships due to unstructured real-world environments and robotic systems with a high degree of autonomy; for example, the need to take into account complex path-planning constraints and communication pathways.

Bibliography

- [1] Markus Achtelik, Abraham Bachrach, Ruijie He, Samuel Prentice, and Nicholas Roy. Stereo Vision and Laser Odometry for Autonomous Helicopters in GPS-denied Indoor Environments. In *Proceedings of SPIE, Unmanned Systems Technology XI*, page 7332. SPIE, 2010.
- [2] Jean-Francois Allan. Robotics for distribution power lines: Overview of the last decade. In *2nd International Conference on Applied Robotics for the Power Industry (CARPI)*, pages 96–101, 2012.
- [3] Ronald C. Arkin. Integrating behavioral, perceptual, and world knowledge in reactive navigation. *Robotics and Autonomous Systems*, 6(1-2):105–122, June 1990.
- [4] Ronald C. Arkin and Jonathan Diaz. Line-of-sight constrained exploration for reactive multiagent robotic teams. In *AMC 7th International Workshop on Advanced Motion Control*, page 455461, 2002.
- [5] Minoru Asada, Manuela Veloso, GerhardK. Kraetzschmar, and Hiroaki Kitano. A review of robot world cup soccer research issues robocup: Today and tomorrow. In *Experimental Robotics VI*, volume 250 of *Lecture Notes in Control and Information Sciences*, pages 369–378. Springer London, 2000.
- [6] Ascending Technologies. <http://www.asctec.de>. Accessed January 2014.
- [7] Tucker Balch and Ronald C. Arkin. Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*, 14:926–939, 1999.
- [8] Joseph Barnard. The MOD grand challenge 2008: the Barnard team perspective. In *Unmanned Air Vehicle Systems*, pages 12.1 – 12.5, Bristol, 2008.

- [9] BBC News. China's jade rabbit moon rover sends back first photos. <http://www.bbc.co.uk/news/world-asia-25393826>. Accessed January 2014.
- [10] BBC News. Merseyside police drone tracks car theft suspects. <http://news.bbc.co.uk/1/hi/england/merseyside/8510370.stm>. Accessed January 2014.
- [11] M. Bengel, K. Pfeiffer, B. Graf, A. Bubeck, and A. Verl. Mobile robots for offshore inspection and manipulation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. (IROS 2009)*, pages 3317–3322, 2009.
- [12] M. Bennewitz and W. Burgard. An experimental comparison of path planning techniques for teams of mobile robots. In *Tagungsband des 16. Fachgesprächs Autonome Mobile Systeme (AMS)*, 2000.
- [13] Brett Bethke, Mario Valenti, and Jonathan How. Cooperative vision based estimation and tracking using multiple uavs. In *Advances in Cooperative Control and Optimization*, volume 369 of *Lecture Notes in Control and Information Sciences*, pages 179–189. Springer Berlin Heidelberg, 2007.
- [14] J. Borenstein, Y. Koren, and Senior Member. The vector field histogram - fast obstacle avoidance for mobile robots. *IEEE Journal of Robotics and Automation*, 7:278–288, 1991.
- [15] Frederic Bourgault. Coordinated decentralized search for a lost target in a Bayesian world. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2003. (IROS 2003)*, pages 48–53, 2003.
- [16] Valentino Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press, 1986.
- [17] R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, 2(1):14–23, March 1986.
- [18] W. Burgard, M. Moors, C. Stachniss, and F. Schneider. Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3):376 – 378, 2005.
- [19] J.L. Burke and R.R. Murphy. Human-robot interaction in usar technical search: two heads are better than one. In *13th IEEE International Workshop on Robot*

- and Human Interactive Communication, 2004. (ROMAN 2004)*, pages 307–312, 2004.
- [20] S. Cameron, S. Hailes, S. Julier, S. McClean, G. Parr, N. Trigoni, M. Ahmed, G. McPhillips, R. de Nardi, J. Nie, A. Symington, L. Teacy, and S. Waharte. Suaave: Combining aerial robots and wireless networking. In *25th Bristol International UAV Systems Conference*, 2010.
- [21] Stefano Carpin. Fast and accurate map merging for multi-robot systems. *Autonomous Robots*, 25(3):305–316, October 2008.
- [22] J. Casper and R.R. Murphy. Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics.*, 33(3):367–385, 2003.
- [23] K. Celik, Soon-Jo Chung, M. Clausman, and A.K. Somani. Monocular vision slam for indoor aerial vehicles. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. (IROS 2009).*, pages 1566–1573, 2009.
- [24] A. Chapman and M. Mesbahi. Uav flocking with wind gusts: Adaptive topology and model reduction. In *American Control Conference (ACC), 2011*, pages 1045–1050, 2011.
- [25] Raja Chatila, Simeon Lacroix, Thierry Simeon, and Matthieu Herrb. Planetary exploration by a mobile robot: Mission teleprogramming and autonomous navigation. *Autonomous Robots*, 2(4):333–344, 1995.
- [26] Dorin Comaniciu, Visvanathan Ramesh, and Peter Meer. Real-time tracking of non-rigid objects using mean shift. In *IEEE Conference on Computer Vision and Pattern Recognition, 2000.*, volume 2, pages 142–149. IEEE, 2000.
- [27] Joseph Cooper and Michael A. Goodrich. Towards combining UAV and sensor operator roles in UAV-enabled visual search. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction, HRI '08*, pages 351–358, New York, NY, USA, 2008. ACM.

- [28] Ingemar J. Cox. Blanche-an experiment in guidance and navigation of an autonomous robot vehicle. *Robotics and Automation, IEEE Transactions on*, 7(2):193–204, 1991.
- [29] K. Daniel, B. Dusza, A. Lewandowski, and C. Wietfeld. Airshield: A system-of-systems muav remote sensing architecture for disaster response. In *3rd Annual IEEE Systems Conference, 2009.*, pages 196–200, 2009.
- [30] K. Daniel, S. Rohde, N. Goddemeier, and C. Wietfeld. A communication aware steering strategy avoiding self-separation of flying robot swarms. In *5th IEEE International Conference on Intelligent Systems (IS), 2010.*, pages 254–259, July 2010.
- [31] Julian de Hoog, Stephen Cameron, and Arnoud Visser. Role-based autonomous multi-robot exploration. In *International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE)*, November 2009.
- [32] Julian de Hoog, Stephen Cameron, and Arnoud Visser. Dynamic team hierarchies in communication-limited multi-robot exploration. In *IEEE International Workshop on Safety, Security & Rescue Robotics (SSRR 2010)*, Bremen, July 2010.
- [33] Julian de Hoog, Stephen Cameron, and Arnoud Visser. Selection of rendezvous points for multi-robot exploration in dynamic environments. In *Workshop on Agents in Realtime and Dynamic Environments, International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, May 2010.
- [34] P. Gonzalez de Santos, J.A. Cobano, E. Garcia, J. Estremera, and M.A. Armada. A six-legged robot-based system for humanitarian demining missions. *Mechatronics*, 17(8):417 – 430, 2007.
- [35] M B. Dias, R. Zlot, M. Zinck, , J.P. Gonzalez, and Anthony Stentz. A versatile implementation of the traderbots approach for multirobot coordination. In *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS-8)*, 2004.
- [36] Patrick Doherty and Piotr Rudol. A UAV search and rescue scenario with human body detection and geolocalization. In *AI 2007: Advances in Artificial*

- Intelligence*, volume 4830 of *Lecture Notes in Computer Science*, pages 1–13. Springer Berlin / Heidelberg, 2007.
- [37] G. Ducard and R. D’Andrea. Autonomous quadrotor flight using a vision system and accommodating frames misalignment. In *IEEE International Symposium on Industrial Embedded Systems, 2009. (SIES ’09).*, pages 261–264, 2009.
- [38] Gregory Dudek, Michael R. M. Jenkin, Evangelos Milios, and David Wilkes. A taxonomy for multi-agent robotics. *Autonomous Robots*, 3:375–397, 1996.
- [39] Edina Digimap. Edina digimap service. <http://digimap.edina.ac.uk/>. Accessed January 2014.
- [40] A. Elfes. Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, June 1989.
- [41] Alessandro Farinelli, Luca Iocchi, and Daniele Nardi. Multirobot systems: a classification focused on coordination. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 34(5):2015–2028, 2004.
- [42] Federal Aviation Administration (FAA). Integration of Civil Unmanned Aircraft Systems (UAS) in the National Airspace System (NAS) Roadmap. Available at http://www.faa.gov/about/initiatives/uas/media/UAS_Roadmap_2013.pdf, 2013.
- [43] Jodi Forlizzi and Carl DiSalvo. Service robots in the domestic environment: A study of the roomba vacuum in the home. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-robot Interaction, HRI ’06*, pages 258–265, New York, NY, USA, 2006. ACM.
- [44] Dieter Fox, Wolfram Burgard, Frank Dellaert, and Sebastian Thrun. Monte carlo localization: Efficient position estimation for mobile robots. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI ’99)*, pages 343–349, 1999.
- [45] Dieter Fox, Jonathan Ko, Kurt Konolige, Benson Limketkai, Dirk Schulz, and Benjamin Stewart. Distributed multi-robot exploration and mapping. In *Proceedings of the IEEE*, volume 94, pages 1325–1339, 2006.

- [46] Kikuo Fujimura and Karansher Singh. Planning cooperative motion for distributed mobile agents. *Journal of Robotics and Mechatronics*, 1996.
- [47] Matthew E. Gaston and Marie desJardins. Agent-organized networks for dynamic team formation. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 230–237, New York, NY, USA, 2005. ACM.
- [48] A. Gaszczak, T.P. Breckon, and J. Han. Real-time people and vehicle detection from UAV imagery. In *Proc. SPIE Conference Intelligent Robots and Computer Vision XXVIII: Algorithms and Techniques*, volume 7878, 2011.
- [49] Brian P. Gerkey and Maja J Mataric. Sold!: Auction methods for multi-robot coordination. *IEEE Transactions on Robotics and Automation, Special Issue on Multi-robot Systems*, 2001.
- [50] Brian P. Gerkey and Maja J. Mataric. A formal analysis and taxonomy of task allocation in multi-robot systems. *The International Journal of Robotics Research*, 23(9):939–954, 2004.
- [51] M.A. Goodrich, J.L. Cooper, J.A. Adams, C. Humphrey, R. Zeeman, and B.G. Buss. Using a mini-uav to support wilderness search and rescue: Practices for human-robot teaming. In *IEEE International Workshop on Safety, Security and Rescue Robotics, 2007. (SSRR 2007)*, pages 1 –6, sept. 2007.
- [52] Ronald L. Graham. An efficient algorithm for determining the convex hull of a finite planar set. *Information Processing Letters*, 1(4):132–133, 1972.
- [53] M. Guarnieri, R. Kurazume, H. Masuda, T. Inoh, K. Takita, P. Debenest, R. Hodoshima, E. Fukushima, and S. Hirose. Helios system: A team of tracked robots for special urban search and rescue operations. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. (IROS 2009)*, pages 2795–2800, Oct. 2009.
- [54] M. Guarnieri, I. Takao, P. Debenest, K. Takita, E. Fukushima, and S. Hirose. Helios ix tracked vehicle for urban search and rescue operations: Mechanical design and first tests. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2008. (IROS 2008)*, pages 1612–1617, 2008.

- [55] Chad Haddal and Jeremiah Gertler. Homeland security: Unmanned aerial vehicles and border surveillance. Congressional Research Service, July 2010. Available at <https://www.fas.org/sgp/crs/homsec/RS21698.pdf>. Accessed January 2014.
- [56] O.M. Hammouri and Mustafa M. Matalgah. Voronoi path planning technique for recovering communication in uavs. In *IEEE/ACS International Conference on Computer Systems and Applications, 2008. (AICCSA 2008)*, pages 403–406, 2008.
- [57] P.E. Hart, N.J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *Systems Science and Cybernetics, IEEE Transactions on*, 4(2):100–107, 1968.
- [58] J. Haugen, E. I. Grøtli, and L. Imsland. State estimation of ice-thickness using mobile sensors. In *Proc. of the IEEE Multi-Conference on Systems and Control*, pages 336–343, Dubrovnik, Croatia, October 3-5 2012.
- [59] Heiko Helble and Stephen Cameron. OATS: Oxford aerial tracking system. *Robotics and Autonomous Systems*, 55(9):661 – 666, 2007.
- [60] V.M. Hernandez Bennetts, A.J. Lilienthal, A.A. Khaliq, V. Pomareda Sese, and M. Trincavelli. Towards real-world gas distribution mapping and leak localization using a mobile robot with 3d and remote gas sensing capabilities. In *2013 IEEE International Conference on Robotics and Automation (ICRA 2013)*, pages 2335–2340, 2013.
- [61] C.Donald Heth and Edward H. Cornell. Characteristics of travel by persons lost in albertan wilderness areas. *Journal of Environmental Psychology*, 18(3):223 – 235, 1998.
- [62] Gabe Hoffmann, Dev Gorur Rajnarayan, Steven L. Wasl, Phd C, Claire J. Tomlin, and Assistant Professor. The stanford testbed of autonomous rotorcraft for multi agent control (STARMAC). In *Proceedings of the 23rd Digital Avionics Systems Conference*, 2004.

- [63] Gabriel M. Hoffmann, Haomiao Huang, Steven L. Wasl, and Er Claire J. Tomlin. Quadrotor helicopter flight dynamics and control: Theory and experiment. In *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2007.
- [64] Owen Holland, Chris Melhuish, and Coldharbour Lane. Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, pages 173–202, 1999.
- [65] Andrew Howard. Multi-robot simultaneous localization and mapping using particle filters. *International Journal of Robotics Research*, 25(12):1243–1256, December 2006.
- [66] Andrew Howard, Maja J. Mataric, and Gaurav S. Sukhatme. An incremental deployment algorithm for mobile robot teams. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2002)*, pages 2849–2854, 2002.
- [67] C. Huerzeler, G. Caprari, E. Zwicker, and L. Marconi. Applying aerial robotics for inspections of power and petrochemical facilities. In *Applied Robotics for the Power Industry (CARPI), 2012 2nd International Conference on*, pages 167–172, 2012.
- [68] Ammar Husain, Heather Jones, Balajee Kannan, Uland Wong, Tiago Pimentel, Sarah Tang, Shreyansh Daftry, Steven Huber, and William L. Whittaker. Mapping planetary caves with an autonomous, heterogeneous robot team. In *2013 IEEE Aerospace Conference*, pages 1–13, 2013.
- [69] Muzammil Hussain and Niki Trigoni. Adaptive node placement for improving localization accuracy in clutter-prone environments. In *IEEE Wireless Communication and Networking Conference (WCNC)*, 2013.
- [70] Luca Iocchi, Daniele Nardi, and Massimiliano Salerno. Reactivity and deliberation: A survey on multi-robot systems. In *Balancing Reactivity and Social Deliberation in Multi-Agent Systems*, pages 9–34, 2000.
- [71] iRobot. iRobot Warrior. <http://www.irobot.com/us/learn/defense/warrior.aspx>. Accessed January 2014.

- [72] iRobot. Looj Gutter Cleaning Robot. <http://www.irobot.com/us/learn/home/looj.aspx>. Accessed January 2014.
- [73] Edward Gil Jones, Brett Browning, M. Bernardine Dias, Brenna Argall, Manuela M. Veloso, and Anthony Stentz. Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks. In *Proceedings of the 2006 IEEE International Conference on Robotics and Automation (ICRA 2006)*, pages 570–575, 2006.
- [74] Nidhi Kalra, Dave Ferguson, and Anthony Stentz. Hoplitest: A market-based framework for planned tight coordination in multirobot teams. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA 2005)*, pages 1170–1177, 2005.
- [75] O Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *International Journal of Robotics Research*, 5(1):90–98, 1986.
- [76] Donghyun Kim, B.H. Abay, R. N. Uma, Weili Wu, Wei Wang, and A.O. Tokuta. Minimizing data collection latency in wireless sensor network with multiple mobile elements. In *INFOCOM, 2012 Proceedings IEEE*, pages 504–512, March 2012.
- [77] J. Kim and Y. Kim. Optimal circular flight of multiple uavs for target tracking in urban areas. In A. Lazinica, editor, *Intelligent aerial vehicles*. IN-TECH, Vienna, Austria, January 2009.
- [78] Derek Kingston, Randal Beard, Timothy McLain, Michael Larsen, and Wei Ren. Autonomous vehicle technologies for small fixed wing uavs. In *AIAA Journal of Aerospace Computing, Information, and Communication*, pages 2003–6559, 2003.
- [79] Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, Eiichi Osawai, and Hitoshi Matsubara. Robocup: A challenge problem for ai and robotics. In *RoboCup-97: Robot Soccer World Cup I*, volume 1395 of *Lecture Notes in Computer Science*, pages 1–19. Springer Berlin Heidelberg, 1998.
- [80] Hiroaki Kitano and Satoshi Tadokoro. Robocup rescue: A grand challenge for multiagent and intelligent systems. *AI Magazine*, 22(1):39–52, 2001.

- [81] Vijay Kumar and Nathan Michael. Opportunities and challenges with autonomous micro aerial vehicles. *International Journal of Robotics Research*, 31(11):1279–1291, September 2012.
- [82] S. Lange, N. Sunderhauf, and P. Protzel. A vision based onboard approach for landing and position control of an autonomous multirotor uav in gps-denied environments. In *International Conference on Advanced Robotics, 2009. (ICAR 2009)*, pages 1–6, 2009.
- [83] N. Lau, L.S. Lopes, G. Corrente, and N. Filipe. Multi-robot team coordination through roles, positionings and coordinated procedures. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. (IROS 2009)*, pages 5841–5848, Oct. 2009.
- [84] T. Lemaire, R. Alami, and S. Lacroix. A distributed tasks allocation scheme in multi-uav context. In *2004 IEEE International Conference on Robotics and Automation, 2004. Proceedings. (ICRA '04)*, volume 4, pages 3622–3627 Vol.4, 2004.
- [85] Lanny Lin and Michael A. Goodrich. UAV intelligent path planning for wilderness search and rescue. In *Proceedings of the 2009 IEEE/RSJ international conference on Intelligent robots and systems, IROS'09*, pages 709–714. IEEE Press, 2009.
- [86] Lanny Lin and Michael A. Goodrich. A Bayesian approach to modeling lost person behaviors based on terrain features in wilderness search and rescue. *Computational and Mathematical Organization Theory*, 16(3):300–323, September 2010.
- [87] Jinguo Liu, Yuechao Wang, Bin Li, and Shugen Ma. Current research, key performances and future development of search and rescue robots. *Frontiers of Mechanical Engineering in China*, 2(4):404–416, 2007.
- [88] V. Lumelsky and A. Stepanov. Dynamic path planning for a mobile automaton with limited information on the environment. *IEEE Transactions on Automatic Control*, 31(11):1058 – 1063, nov 1986.

- [89] Chunbo Luo, P. Ward, S. Cameron, G. Parr, and S. McClean. Communication provision for a team of remotely searching uavs: A mobile relay approach. In *2012 IEEE Globecom Workshops (GC Wkshps)*, pages 1544–1549, 2012.
- [90] A. Macwan, G. Nejat, and B. Benhabib. Target-motion prediction for robotic search and rescue in wilderness environments. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 41(5):1287–1298, 2011.
- [91] Ashish Macwan, Goldie Nejat, and Beno Benhabib. Optimal deployment of robotic teams for autonomous wilderness search and rescue. In *Proceedings of the 2011 IEEE/RSJ international conference on Intelligent robots and systems, IROS’11*, pages 4544–4549, 2011.
- [92] Maja J. Mataric. Designing emergent behaviors: From local interactions to collective intelligence. In *In Proceedings of the International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, volume 2, pages 432–441, 1993.
- [93] Matlab-Control. <https://code.google.com/p/matlabcontrol/>. Accessed January 2014.
- [94] T. Matsuda, T. Maki, T. Sakamaki, and T. Ura. State estimation of multiple autonomous underwater vehicles for wide area survey of seafloor. In *OCEANS - Bergen, 2013 MTS/IEEE*, pages 1–9, 2013.
- [95] Daniel Mellinger, Nathan Michael, and Vijay Kumar. Trajectory generation and control for precise aggressive maneuvers with quadrotors. *International Journal of Robotics Research*, 31(5):664–674, April 2012.
- [96] Daniel Mellinger, Michael Shomin, Nathan Michael, and Vijay Kumar. Cooperative grasping and transport using multiple quadrotors. In *Distributed Autonomous Robotic Systems*, volume 83 of *Springer Tracts in Advanced Robotics*, pages 545–558. Springer Berlin Heidelberg, 2013.
- [97] N. Michael, J. Fink, and V. Kumar. Cooperative manipulation and transportation with aerial robots. In *Proceedings of Robotics: Science and Systems*, Seattle, USA, June 2009.

- [98] Keiji Nagatani, S. Kiribayashi, Y. Okada, S. Tadokoro, T. Nishimura, T. Yoshida, E. Koyanagi, and Y. Hada. Redesign of rescue mobile robot quince. In *2011 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 13–18, 2011.
- [99] NASA Jet Propulsion Laboratory. Mars exploration rover facts. <http://marsrovers.jpl.nasa.gov/newsroom/factsheets/pdfs/Mars03Rover041020.pdf>. Accessed January 2014.
- [100] Sarfraz Nawaz, Muzammil Hussain, Simon Watson, Niki Trigoni, and PeterN. Green. An underwater robotic network for monitoring nuclear waste storage pools. In Stephen Hailes, Sabrina Sicari, and George Roussos, editors, *Sensor Systems and Software*, volume 24 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 236–255. Springer Berlin Heidelberg, 2010.
- [101] Oyvind Netland and Amund Skavhaug. Prototyping and evaluation of a telerobot for remote inspection of offshore wind farms. In *2012 2nd International Conference on Applied Robotics for the Power Industry (CARPI)*, pages 187–192, 2012.
- [102] P. Olsson, J. Kvarnstrom, P. Doherty, O. Burdakov, and K. Holmberg. Generating uav communication networks for monitoring and surveillance. In *2010 11th International Conference on Control Automation Robotics Vision (ICARCV)*, pages 1070–1077, 2010.
- [103] Robocup Online. <http://www.robocup.org>. Accessed January 2014.
- [104] Ordnance Survey. Land-Form PROFILE. <http://www.ordnancesurvey.co.uk/business-and-government/products/land-form-profile.html>. Accessed January 2014.
- [105] Ordnance Survey. Master Map. <http://www.ordnancesurvey.co.uk/business-and-government/products/mastermap-products.html>. Accessed January 2014.

- [106] Ordnance Survey. TERRAIN 5. <http://www.ordnancesurvey.co.uk/business-and-government/products/os-terrain-5.html>. Accessed January 2014.
- [107] O. Oreifej, R. Mehran, and M. Shah. Human identity recognition in aerial images. In *2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 709–716, 2010.
- [108] Joseph O’Rourke. *Computational Geometry in C*. Cambridge University Press, New York, NY, USA, 2nd edition, 1998.
- [109] L. E. Parker and F. Tang. Building multirobot coalitions through automated task solution synthesis. *Proceedings of the IEEE*, 94(7):1289–1305, July 2006.
- [110] L.E. Parker, C.M. Reardon, H. Choxi, and C. Bolden. Using critical junctures and environmentally-dependent information for management of tightly-coupled cooperation in heterogeneous robot teams. In *IEEE International Conference on Robotics and Automation, 2009. (ICRA ’09)*, pages 2729–2736, May 2009.
- [111] Lynne E. Parker, Yi Guo, and David Jung. Cooperative robot teams applied to the site preparation task. In *In Proceedings of the 10th International Conference on Advanced Robotics*, pages 71–77, 2001.
- [112] Parrot. Parrot AR Drone. <http://ardrone2.parrot.com/>. Accessed January 2014.
- [113] Dave Perkins and Pete Roberts. The uk missing person behaviour study. Technical report, The Centre for Search Research, 2011. Available at http://www.searchresearch.org.uk/www/ukmpbs/current_report/.
- [114] Max Pfingsthorn, Bayu Slamet, and Arnoud Visser. A Scalable Hybrid Multi-Robot SLAM method for Highly Detailed Maps. In *Proceedings of the 11th RoboCup International Symposium*, July 2007.
- [115] Erwin Prassler, Arno Ritter, Christoph Schaeffer, and Paolo Fiorini. A short history of cleaning robots. *Autonomous Robots*, 9(3):211–226, December 2000.
- [116] Yaohong Qu, Xu Zhu, and YouminM. Zhang. Cooperative control for uav formation flight based on decentralized consensus algorithm. In *Intelligent Robotics*

- and Applications*, volume 7506 of *Lecture Notes in Computer Science*, pages 357–366. Springer Berlin Heidelberg, 2012.
- [117] Vladimir Reilly, Berkan Solmaz, and Mubarak Shah. Geometric constraints for human detection in aerial imagery. In *Proceedings of the 11th European Conference on Computer Vision: Part VI, ECCV'10*, pages 252–265, Berlin, Heidelberg, 2010. Springer-Verlag.
- [118] Robomow. Robomow. <http://www.robomow.com/en-UK/>. Accessed January 2014.
- [119] Martijn N. Rooker and Andreas Birk. Multi-robot exploration under the constraints of wireless networking. *Control Engineering Practice*, 15(4):435 – 445, 2007.
- [120] C. Rossi, L. Aldama, A. Barrientos, A. Valero, and C. Cruz. Negotiation of target points for teams of heterogeneous robots: an application to exploration. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. (IROS 2009)*, pages 5868–5873, Oct. 2009.
- [121] Claudio Rossi, Leyre Aldama, and Antonio Barrientos. Simultaneous task subdivision and allocation for teams of heterogeneous robots. In *2009 IEEE International Conference on Robotics and Automation (ICRA 2009)*, pages 946–951, 2009.
- [122] Nicholas Roy and Gregory Dudek. Collaborative robot exploration and rendezvous: Algorithms, performance bounds and observations. *Autonomous Robots*, 11:117–136, 2000.
- [123] P. Rudol and P. Doherty. Human body detection and geolocalization for uav search and rescue missions using color and thermal imagery. In *Aerospace Conference, 2008 IEEE*, pages 1–8, 2008.
- [124] Rahul Sawhney, K. Madhava Krishna, and Kannan Srinathan. On fast exploration in 2d and 3d terrains with multiple robots. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1, AAMAS '09*, pages 73–80, Richland, SC, 2009.

- [125] Bernt Schiele and James L. Crowley. Comparison of position estimation techniques using occupancy grids. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pages 1628–1634, 1994.
- [126] E. Semsch, M. Jakob, D. Pavlicek, and Michal Pechoucek. Autonomous uav surveillance in complex urban environments. In *IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technologies, 2009. (WI-IAT '09)*, volume 2, pages 82–85, 2009.
- [127] sFly Project. <http://www.sfly.org>. Accessed April 2014.
- [128] Onn Shehory and Sarit Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, 1998.
- [129] Weihua Sheng, Qingyan Yang, Jindong Tan, and Ning Xi. Distributed multi-robot coordination in area exploration. *Robotics and Autonomous Systems*, 54(12):945 – 955, 2006.
- [130] Reid Simmons and Sven Koenig. Probabilistic robot navigation in partially observable environments. In *Proceedings of the 14th international joint conference on Artificial intelligence - Volume 2, IJCAI'95*, pages 1080–1087, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc.
- [131] SUAAVE. Sensing Unmanned Autonomous Aerial VEHicles. <http://www.suaave.org>. Accessed January 2014.
- [132] P. B. Sujit, A. Sinha, and D. Ghose. Multiple uav task allocation using negotiation. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems, AAMAS '06*, pages 471–478, New York, NY, USA, 2006. ACM.
- [133] Andrew Symington, Sonia Waharte, Simon Justin Julier, and Niki Trigoni. Probabilistic target detection by camera-equipped uavs. In *2010 IEEE International Conference on Robotics and Automation (ICRA 2010)*, 2010.
- [134] Fang Tang and Lynne E. Parker. A complete methodology for generating multi-robot task solutions using asymptre-d and market-based task allocation. In *2007 IEEE International Conference on Robotics and Automation (ICRA 2007)*, pages 3351–3358, 2007.

- [135] W.T.L. Teacy, Jing Nie, S. McClean, and G. Parr. Maintaining connectivity in uav swarm sensing. In *2010 IEEE GLOBECOM Workshops (GC Wkshps)*, pages 1771–1776, 2010.
- [136] H. P. Thamm. SUSI 62 A robust and safe parachute uav with long flight time and good payload. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXVIII-1/C22:19–24, 2011.
- [137] S. Thrun, S. Thayer, W. Whittaker, C. Baker, W. Burgard, D. Ferguson, D. Hahnel, D. Montemerlo, A. Morris, Z. Omohundro, C. Reverte, and Whittaker W. Autonomous exploration and mapping of abandoned mines. *Robotics Automation Magazine, IEEE*, 11(4):79–91, 2004.
- [138] Sebastian Thrun. Robotic mapping: A survey. In *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann, 2002.
- [139] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.
- [140] Sebastian Thrun and Yufeng Liu. Multi-robot slam with sparse extended information filters. In *Robotics Research*, volume 15 of *Springer Tracts in Advanced Robotics*, pages 254–266. Springer Berlin Heidelberg, 2005.
- [141] J. Tisdale, ZuWhan Kim, and J.K. Hedrick. Autonomous uav path planning and estimation. *Robotics Automation Magazine, IEEE*, 16(2):35–42, 2009.
- [142] Y. Tojo, P. Debenest, E.F. Fukushima, and S. Hirose. Robotic system for humanitarian demining. In *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, volume 2, pages 2025–2030 Vol.2, 2004.
- [143] Godfried T. Toussaint. Solving geometric problems with the rotating calipers. In *Proceedings of IEEE MELECON*, 1983.
- [144] Tsusiat Software. Java Topology Suite. <http://tsusiatsoftware.net/jts/main.html>. Accessed January 2014.
- [145] Ordnance Survey UK. Ordnance Survey Online. <http://www.ordnancesurvey.co.uk>. Accessed January 2014.

- [146] Douglas Vail and Manuela Veloso. Multi-robot dynamic role assignment and coordination through shared potential fields. In A. Schultz, L. Parker, and F. Schneider, editors, *Multi-Robot Systems*. Kluwer, 2003.
- [147] L. Vig and J.A. Adams. Multi-robot coalition formation. *Robotics, IEEE Transactions on*, 22(4):637–649, Aug. 2006.
- [148] Sonia Waharte and Niki Trigoni. Supporting search and rescue operations with UAVs. In *Proceedings of the 2010 International Conference on Emerging Security Technologies*, EST '10, pages 142–147, Washington, DC, USA, 2010. IEEE Computer Society.
- [149] K. Wahren, I. Cowling, Y. Patel, P. Smith, and T.P. Breckon. Development of a two-tier unmanned air system for the MoD grand challenge. In *24th International Unmanned Air Vehicle Systems Conference (UAVS)*, 2009.
- [150] Chao-Lei Wang, Tian-Miao Wang, Jian-Hong Liang, Yi-Cheng Zhang, and Yi Zhou. Bearing-only visual slam for small unmanned aerial vehicles in gps-denied environments. *International Journal of Automation and Computing*, 10(5):387–396, 2013.
- [151] Paul Ward and Stephen Cameron. Coordination in multi-tiered robotic search. In *Towards Autonomous Robotic Systems (TAROS 2011)*, volume 6856 of *Lecture Notes in Computer Science*, pages 384–385. Springer Berlin Heidelberg, 2011.
- [152] Stephan Weiss, Davide Scaramuzza, and Roland Siegwart. Monocular-slambased navigation for autonomous micro helicopters in gps-denied environments. *Journal of Field Robotics*, 28(6):854–874, 2011.
- [153] Tony White and James P. Helferty. Emergent team formation: Applying division of labour principles to robot soccer. In *Engineering Self-Organising Systems*, pages 180–194, 2004.
- [154] Allen Wu, Eric Johnson, Michael Kaess, Frank Dellaert, and Girish Chowdhary. Autonomous flight in gps-denied environments using monocular vision and inertial sensors. *Journal of Aerospace Computing, Information and Communication*, 2012.

- [155] Kai M. Wurm, Cyrill Stachniss, and Wolfram Burgard. Coordinated multi-robot exploration using a segmentation of the environment. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2008. (IROS 2008)*, pages 1160–1165, 2008.
- [156] Atsushi Yamashita, Masaki Fukuchi, Jun Ota, Tamio Arai, and Hajime Asama. Motion planning for cooperative transportation of a large object by multiple mobile robots in a 3d environment. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA 2000)*, pages 3144–3151, 2000.
- [157] Brian Yamauchi. Frontier-based exploration using multiple robots. In *AGENTS '98: Proceedings of the second international conference on Autonomous agents*, pages 47–53, New York, NY, USA, 1998. ACM.
- [158] X.S. Zhou and S.I. Roumeliotis. Multi-robot slam with unknown initial correspondence: The robot rendezvous case. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems, (IROS 2006)*, pages 1785–1792, 2006.
- [159] R. Zlot, A. Stentz, M.B. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. *IEEE International Conference on Robotics and Automation, 2002. Proceedings. (ICRA '02)*, 3:3016 – 3023, 2002.