

**A Game Theoretic Approach to Coordinating
Unmanned Aerial Vehicles with
Communications Payloads**

Philip B. Charlesworth

Submitted in partial fulfilment of the criteria for the
degree of Doctor of Philosophy



School of Computer Science & Informatics
Cardiff University
July 2015

This page is intentionally blank

DECLARATION

This work has not been submitted in substance for any other degree or award at this or any other university or place of learning, nor is being submitted concurrently in candidature for any degree or other award.

Signed(candidate) Date

STATEMENT 1

This thesis is being submitted in partial fulfillment of the requirements for the degree of PhD.

Signed(candidate) Date

STATEMENT 2

This thesis is the result of my own independent work/investigation, except where otherwise stated. Other sources are acknowledged by explicit references. The views expressed are my own.

Signed(candidate) Date

STATEMENT 3

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed(candidate) Date

This page is intentionally blank

Acknowledgements

It would not have been possible to write this doctoral thesis without support.

Firstly, I would like to thank my supervisor, Dr Stuart Allen, for his guidance, advice and encouragement throughout this PhD. Thirty years of professional engineering had taught me how to be an engineer. He taught me how to think like a research student and guided me through the academic process. I never left his office without having learned something.

This research would not have been possible without the support of my employer, Airbus Group Innovations. The use of company computing facilities made it possible to undertake some very long simulations. The support of my colleagues, especially Matt Roberts and Kevin Jones, helped me to get started on L^AT_EX very early in this research. Others, too numerous to mention, helped in many ways such as checking my code, proof reading chapters, providing data backup facilities, and offering encouragement and coffee at appropriate moments.

I must also acknowledge the guiding hand of my late father, Fenwick Charlesworth FInstP. He gave me a love of science and a sense of curiosity, two great gifts that have accompanied me throughout my adult life.

Finally, and most importantly, it would not have been possible to complete this research or write this thesis without the unfailing support of my wife, Ruth. There are no words to adequately describe the depth of my appreciation for the sacrifices she has made while I pursued my PhD.

This page is intentionally blank

Contents

1	Introduction	2
1.1	Introduction	2
1.2	Mobile Communications	3
1.3	Mission Description	5
1.4	Research Challenges	7
1.5	Structure of this Thesis	8
2	Context and Related Work	10
2.1	Terminology	10
2.2	A Brief History of Unmanned Aerial Vehicles	11
2.3	UAVs in non-military applications	13
2.4	Communications UAVs	14
2.5	Choice of UAV Platforms for Communications	16
2.6	UAVs and Autonomy	17
2.7	Path Planning for a single UAV	20
2.8	Path Planning and Coordination for Multiple UAVs	22
2.9	Game Theory in Wireless Networks	23
2.10	Path Planning and Coordination for Multiple Communications UAVs	25
2.11	State of the Art	26
3	Experimental Method and Principles	28
3.1	Chapter Introduction	28
3.2	System Description	29
3.3	Definition of the Standard Models	31
3.3.1	UAV Model and Characteristics	32
3.3.2	Mobile Model and Characteristics	32
3.3.3	Payload Model and Characteristics	34
3.4	Experimental Method	34
3.5	Packing the Payload and Measuring System Performance	37
3.5.1	Payload Packing Method	38

3.5.2	Performance Metrics	39
3.6	Selection of Parameter Values	41
3.6.1	Number and Spatial Distribution of the Mobiles	41
3.6.2	Power, Beamwidth and Altitude	42
3.7	Assumptions	43
3.8	Characterising the System with a Single UAV	45
3.8.1	Sensitivity of the Operating Point to Power and Beamwidth	46
3.8.2	Sensitivity of the Operating Point to Altitude	48
3.8.3	Statement of the Operating Point	48
3.9	UAVs with Relocation	50
3.9.1	Scenario Layout	51
3.9.2	Relocation Periods	52
3.10	Relocation v Continuous Circling	53
3.10.1	Coverage and Power	54
3.10.2	Flightpath for Relocating UAVs	56
3.10.3	Comparison of Results	56
3.11	Chapter Discussion and Conclusions	58
4	Non-Cooperative Games with Two UAVs	60
4.1	Introduction	60
4.2	Game Description	61
4.3	Benchmarks	63
4.3.1	Two Circling UAVs	63
4.3.2	Select Random Strategies	64
4.4	Comparison of Methods for Solving Normal-Form Games	67
4.4.1	Exhaustive Search for a Pure Strategy	70
4.4.2	Fast Search for Best Response	75
4.4.3	Find One Mixed Strategy using Lemke-Howson	79
4.4.4	Find One Mixed Strategy using Chatterjee’s Method	82
4.5	Extensive Form Game	82
4.6	Effects of Initial Conditions	86
4.7	Initial Conditions in which one UAV is “Cornered”	92
4.8	Chapter Conclusions	96
5	Some Implementation Issues	100
5.1	Chapter Introduction	100
5.2	Prioritised Access Schemes	101
5.3	Modifying Behaviour by Biasing Payoffs	103
5.4	Choice of Frequency Band	109
5.5	Games with Larger Payoff Matrices	110
5.5.1	Two-Dimensional Games with Multiple Players	111
5.5.2	Three-Dimensional Games	111
5.6	Execution Times	112
5.7	Sensitivity to the Age of the Location Data	114

5.8	Number of Solutions	116
5.9	Coverage with Multiple UAVs	122
5.10	Chapter Conclusions	123
6	Conclusions	126
6.1	Research Challenges	126
6.1.1	Defining a Suitable Non-Cooperative Game	127
6.1.2	Solving the Game	128
6.1.3	Practical Implementation of the Game	129
6.1.4	Scalability	130
6.2	Application of this Work	130
6.3	Contributions to Knowledge	131
7	Future Research	132
A	Introduction to Game Theory	134
A.1	Key Concepts	135
A.2	Best Response	136
A.3	Nash Equilibrium	136
A.4	Games with Multiple Players	137
A.5	Extensive and Normal Forms	137
A.6	Extensive-Form Games	139
A.7	Normal-form Games	139
A.8	Methods for Solving Games	140
A.8.1	Exhaustive Search	140
A.8.2	Fast Search	142
A.8.3	The Lemke-Howson Algorithm	142
A.8.4	Chatterjee's Method	143
A.8.5	Backward Induction	143
B	Technical Principles and System Models	144
B.1	Section Summary	144
B.2	Communications Overview	144
B.3	Link Budgets	146
B.4	Traffic Models	150
B.5	Kinematic and Navigation Models for UAVs	150
B.5.1	Vertical motion	152
B.5.2	Change of Latitude and Longitude	153
B.5.3	Altitude	155
B.5.4	UAV Attitude and Footprint	157
B.5.5	Navigation Accuracy	158
C	Bibliography	162

This page is intentionally blank

List of Figures

1.1	General Statement of the Problem	6
2.1	Global Hawk UAV carrying the BACN payload (Picture courtesy of Northrop Grumman)	16
2.2	PACT Levels of Autonomy	19
2.3	USDoD Autonomy Trends	21
3.1	Main System Components	29
3.2	Object Hierarchy, Parameters and Constraints	31
3.3	Key Terms for the UAV Payload Definition	34
3.4	Simulation Flowchart	36
3.5	General Shape of the Coverage Characteristic	42
3.6	P_{max} , A_f and Coverage	43
3.7	Coverage and Power for Operating Point Selection	47
3.8	Mean Coverage and Mean Power v Altitude at the Operating Point	49
3.9	UAV Relocation Options	52
3.10	UAV Relocation Manoeuvre	52
3.11	Coverage and Power	55
3.12	Flight Paths and Probability Distributions of Mobiles	57
4.1	The two-UAV System	61
4.2	Example of a Payoff Matrix with Two Players	62
4.3	Coverage Characteristic for Two Circling UAVs	64
4.4	Coverage and Power v Time for Two Circling UAVs	65
4.5	Coverage Characteristic for Random Strategies	66
4.6	Examples of Coverage with Random Strategies	67
4.7	Coverage Characteristic for All Solution Methods	69
4.8	Coverage and Power v Time for Exhaustive Search	72
4.9	Behaviour of Mobile Clusters	73
4.10	Flightpaths and Estimated Payoffs for Exhaustive Search	74
4.11	Coverage and Power v Time for Fast Search	77

4.12	Flightpaths and Estimated Payoffs for Fast Search	78
4.13	Coverage and Power v Time for Lemke-Howson	80
4.14	Flightpaths and Estimated Payoffs for Lemke-Howson	81
4.15	Coverage and Power v Time for Chatterjee's Method	83
4.16	Flightpaths and Estimated Payoffs for Chatterjee's Method	84
4.17	Coverage Characteristic for Extensive Form Game	85
4.18	Coverage and Power v Time for Extensive Form Game	87
4.19	Flightpaths and Estimated Payoffs for Extensive Form Game	88
4.20	Flightpaths Overlaid on Mobile Probabilities: UAVS Start Separated	89
4.21	Flightpaths Overlaid on Mobile Probabilities: UAVS Start Collocated	90
4.22	Coverage v Time: UAVs start in NW and SE Quadrants	91
4.23	UAV Separation and Coverage v time	93
4.24	An Example of Cornering in which the UAV Fails to Recover	94
4.25	An Example of Cornering in which the UAV Recovers	95
4.26	An Example of Countermeasures to Cornering	96
5.1	Coverage v Time for Mobiles with Prioritised Access	103
5.2	PDF for Access Duration	106
5.3	PDF and CDF for the first 300 seconds	107
5.4	Coverage Characteristic at Several Frequency Bands	110
5.5	The Three UAV System in Two Dimensions	111
5.6	Action Set for a Single UAV in Three Dimensions	112
5.7	Timeline for Finding Solutions: $t_{ex} < t_{sol}$	113
5.8	Algorithm Execution Time t_{ex} for Multiple UAVs	114
5.9	Timeline for Evaluating Solution Quality	115
5.10	Coverage v Data Age, Expressed as a Multiple of t_{sol}	115
5.11	Logic for Locating PSNE - 2 Players	118
5.12	Logic for Locating PSNE - 3 Players	118
5.13	Coverage v Time for 3 UAVs	120
5.14	3D Flight Paths for 3 UAVs	121
5.15	Mean Coverage for Multiple UAVs	122
A.1	Game Representations	138
A.2	Exhaustive and Fast Search Methods	141
B.1	Power and Access Management	145
B.2	UAV Attitude	146
B.3	Link Budget Model	147
B.4	Antenna footprints	149
B.5	The Data Rates R_b Used in these Experiments	150
B.6	Banking Turn	151
B.7	Pull-up	153
B.8	Dubins Path Example	154
B.9	Path Segment	155

B.10 Altitude Management	156
B.11 Effective Footprint Area	158
B.12 Change of Altitude	160

This page is intentionally blank

List of Tables

2.1	Gaps in Current Research	26
3.1	MALE UAV Characteristics	32
3.2	Mobile Characteristics	33
3.3	Characteristics of the Operating Points	50
3.4	Cell Separation and Coverage	53
3.5	Performance Metrics for the Default Data Set	58
4.1	Mean Execution Times for Solution Methods (s)	69
4.2	Mean Execution Times for Backward Induction	86
4.3	Comparison of Solution Methods	97
5.1	Effect of Biasing Payoffs	108
5.2	Exhaustive Search: Percentage of Solutions that are PSNE	117

This page is intentionally blank

Abstract

This thesis considers the placement of two or more Unmanned Aerial Vehicles (UAVs) to provide communications to a community of ground mobiles. The locations for the UAVs are decided by the outcome of a non-cooperative game in which the UAVs compete to maximize their coverage of the mobiles. The game allows navigation decisions to be made onboard the UAVs with the effect of increasing coverage, reducing the need for a central planning function, and increasing the autonomy of the UAVs.

A non-cooperative game that includes the key system elements is defined and simulated. The thesis compares methods for solving the game to evaluate their performance. A conflict between the quality of the solution and the time required to obtain that solution is identified and explored. It considers how the payload calculations could be used to modify the behaviour of the UAVs, and the sensitivity of the game to resource limitations such as RF power and radio spectrum. It finishes by addressing how the game could be scaled from two UAVs to many UAVs, and the constraints imposed by current methods for solving games.

This page is intentionally blank

Introduction

1.1 Introduction

Unmanned aerial vehicles (UAVs) are widely perceived as a purely military technology suited for “dull, dirty and dangerous” missions. In recent years there has been a growing realisation that UAVs could be used in a wider range of roles such as border security and disaster management. These civil applications require reliable communications, not only between the UAVs but also with any ground assets deployed in support of the mission.

This thesis addresses the potential use of UAVs as communications nodes in a disaster recovery network comprising UAVs and ground mobile vehicles. The mobiles represent the agents responsible for implementing the incident recovery plan, for example police, fire, ambulance and other rescue resources. They move within the bounded area of the incident but have limited communications due to the terrain. The UAVs provide a backbone network that allows the mobiles to communicate.

This application of UAVs is not yet permitted by airspace regulators, and there is no radio spectrum allocated for this purpose. The trend in the airspace regulatory bodies is towards permitting higher levels of autonomy by UAVs, subject to the development of suitable collision avoidance mechanisms. Similarly, the International Telecommunications Union is working to identify radio spectrum for UAV applications. In the next decade it is probable that UAVs will start to operate as communications nodes.

There has been extensive academic research into the use of wi-fi equipped UAVs

in this role, typically operating at a few hundred feet above the ground. This thesis starts to address the technical and planning issues of providing communications coverage over large areas using UAVs at altitudes up to 20,000 feet.

1.2 Mobile Communications

Terrestrial mobile communications systems use fixed base stations located on elevated sites to maximize their footprint. The shadowing of radio signals by natural and man-made objects reduces the area of the footprint, but this can be partially offset by increasing the number of base stations. The high cost of installation and maintenance of base stations in mountainous, remote or unpopulated areas can result in poor communications coverage of those regions. This can be partially addressed by equipping some mobiles to communicate through satellites.

Satellite on the move (SOTM) terminals enable real-time mobile communications at data rates on hundreds on kbit/s, but introduce a new set of problems. The typical slant range to a geostationary satellite is 40,000 km, resulting in a radio path loss of about 200dB and signal round trip time (RTT) of over 400ms. The high path losses can be partially offset by using highly directional antennae with tracking systems, but such systems tend to be expensive. The long RTT is incompatible with some common data protocols, particularly ones that have timers such as TCP. Satellite footprints are very large, typically 500km to 1000 km diameter, so frequency re-use factors are very low compared to terrestrial systems. Finally, at moderate to high latitudes the elevation of the satellite can be below the radio horizon in hilly or mountainous regions.

One possible approach is to use both terrestrial and satellite communications. This increases the complexity of the communications architecture as mobiles need to be equipped with two radios, each using different frequencies and antennae. Gateways would be required to interconnect the two systems. The difference in RTT between terrestrial and satellite systems increases the complexity of hand-off mechanisms as they need to ensure that data is neither lost nor duplicated. Using both terrestrial and satellite communications systems is feasible, but offer a relatively expensive and complex solution.

It was recognized that a third option, the use of base stations onboard UAVs, could offer a single communications solution that combined the best properties of fixed base stations and satellites. The altitude of a UAV could be higher than

the radio masts of fixed base stations, but low enough to have footprints that are only 50km to 100km in diameter. The smaller footprints allow a significant improvement in frequency re-use over satellites and reduce round trip times for data protocols to under 1 ms. The low altitude improves path losses by over 50 dB when compared to satellites, reducing the need for directional antennae on the mobiles.

One characteristic of UAVs is that they can relocate in response to demand. This allows the capability of a base station to be relocated to where it is needed, regardless of the terrain. This is particularly valuable when responding to incidents where there is no communications infrastructure or, in the case of disaster relief, where the infrastructure has been severely disrupted. Automated relocation is not permitted in current air traffic management systems, however the trend indicates the progressive relaxation of regulations as experience of UAV operations increases. This thesis addresses a time when fully autonomous operation is permitted and UAVs are fully integrated into airspace management systems.

A UAV-based network could be deployed significantly faster than the equivalent number of temporary radio masts. Furthermore the network of UAVs could be rapidly reconfigured in response to the movement of the mobiles. These two properties initiated a programme of research into self-organising UAV networks. It was recognised that research had tended to focus on two problems: the use of single UAVs to interconnect two mobiles and the use of one UAV to relay signals from another UAV to a ground station. The use of UAVs to provide area communications cover was a topic that had yet to be addressed by the research community.

This research focusses on the area cover problem. The specific problem is to maximize the number of mobiles that can be supported by a small number of UAVs. The UAV payloads will have some power constraints:

- The radios will have a hard limit on the power that can be presented to the antenna. This is normally limited by the rating of the power amplifier but can be limited by other factors, for example interference to other onboard systems.
- The UAV will have some limits on the electrical power that can be generated by its engines. The power offtake from the engines tends to decrease with altitude resulting in a reduction of available power for some onboard systems.

Relocating the UAV can reduce the slant ranges to each mobile, hence reducing

the power that is needed to support its radio downlink. When the mobiles cluster there is a possibility of making significant power savings by relocating towards the cluster. The power saved by this relocation can be used to support other mobiles, hence relocation offers the ability to increase the total coverage of the UAV.

Evolutionary algorithms were considered as a method of determining the UAV's next location but it was felt that such methods lacked novelty. Furthermore, cooperative methods require the generation and distribution of a common plan and the implicit need for reliable communications between all agents. The trend in UAV research is towards increasing levels of autonomy, so a method of arriving at a solution without the need for reliable communications was sought.

An analogous problem was the maximisation of market cover by competing companies, a problem typically addressed using game theory. It was decided to address the problem of area coverage by having the UAVs compete to maximise their individual share of the mobiles.

1.3 Mission Description

The mission considered in this research comprises a set of mobile vehicles that are required to communicate while operating within a bounded region. No terrestrial communications network is available, however a set of Medium Altitude Long Endurance (MALE) UAVs, each with identical communications payloads, provides a backbone network through which the mobiles can communicate.

A diagram of the mission can be seen in figure 1.1 with 3 UAVs and 40 mobiles. The UAVs are interconnected by inter UAV Links (IUAVL) forming a tree network.

The signal quality and data rate will be limited by the RF power supportable by each UAV, topography of the region, earth curvature, and other practical issues. Periodically relocating the UAVs in response to the movement of the mobiles can improve the signal quality to any mobiles within range, and increase number of mobiles that can be supported.

The location of the UAVs can be re-planned at regular intervals to maximise coverage of the mobiles. The planning activity uses current and historic data for the locations of the mobiles to generate a set of new locations for the UAVs. The potential benefits of this approach are twofold:

1. It helps to maximise the total number of mobiles that can be supported.

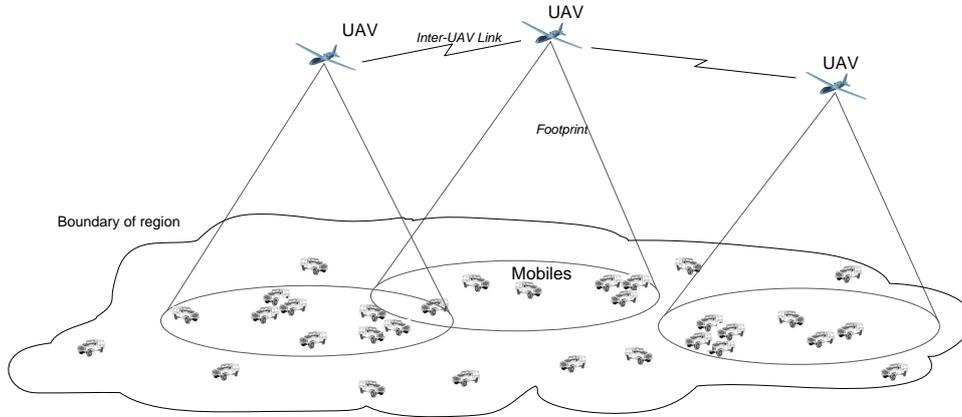


Figure 1.1: General Statement of the Problem

2. It makes the most effective use of the available RF power, thus improving the RF interference environment.

Realising these benefits requires location planning algorithms that can find the best locations for the UAVs. Ideally the algorithms will operate in a distributed manner, allowing the UAVs to operate with significant levels of autonomy.

The problem of optimally locating one or more UAVs is essentially a constrained optimisation problem. A classical approach would be to treat this as problem in multivariable optimisation with constrained solutions. Techniques such as evolutionary algorithms would allow all agents to simultaneously converge on an optimum solution. This collaborative approach is likely to result in good solutions, but the need to distribute a common plan creates a reliance on good communications between all agents. Collaborative approaches are unlikely to be compatible with high levels of autonomy.

MacKenzie commented that game theory “deals primarily with distributed optimization”, noting also that in games “individual, selfish, users make their own decisions instead of being centrally controlled” [59]. This suggests that a game-based algorithm could meet the requirement for autonomous action. This research investigates a heuristic that is based on non-cooperative games as a means of finding optimal locations for the UAVs. Optimal solutions are sought by placing the UAVs in competition to maximise their individual coverage of the mobiles. This is analogous to having companies competing to maximise their share of a market.

The removal of cooperation from the solution significantly reduces the need for

inter-agent communications as only system state needs to be shared; the plan does not need to be distributed. This allows solutions to be obtained with unreliable communications, and implicitly increases each agent's level of autonomy.

1.4 Research Challenges

This thesis addresses two novel ideas. The first is to use UAVs to provide area cover for a community of mobiles. The second idea is to maximise the coverage of the UAVs by having them play a non-cooperative game. The investigation of these ideas posed several challenges.

The first, and most fundamental, challenge was to define a structure for a non-cooperative game that includes all required elements of a game: players, strategies and payoffs. The game structure was required to allow exploration of key variables that affected the game outcomes could be investigated ideally by considering the action set and payoffs of a single player. To quantify any improvements resulting from the game it was necessary to devise a set of performance metrics and to establish a performance baseline.

The solution of the game is clearly a Nash equilibrium. There are several possible methods for finding Nash equilibria for games with two or more players. A significant challenge is to identify a method for solving the game that would be compatible with the processing capability of a UAV and the timeliness requirements of a solution. A perfect method would quickly find an ideal solution. In reality, there is a compromise to be found between the quality of the solution and the time taken to reach that solution. Several methods were investigated during this research.

The game, and solution method, need to be resilient when presented with unusual conditions. This can be tested by setting the initial conditions of the game in an unusual state and establishing how this affects the behaviour of the UAVs. Another, related, challenge is to see if it is possible to deliberately bias the behaviour of the UAVs by changing how payoffs are calculated.

Demonstrating that the game improved performance in a perfect scenario was not enough. For the method to translate into the real world it had to be shown to work at several frequencies. In the real world some messages carrying system state might be lost, corrupted, or only received by some UAVs. There is the possibility that UAVs might be making decisions based on different estimates of system state.

These real-world effects needed to be demonstrated and tested.

Demonstrating the effectiveness and resilience of the game can be achieved with two UAVs. The final set of challenges address the scalability of this method as the number of UAVs, and the size of each UAV's action set, increases. Increasing these parameters causes a change in the size of the payoff matrix, resulting in a consequent increase in the time taken to find a solution. At some values the game will take so long to solve that its solution will be comparable with the "do nothing" option.

1.5 Structure of this Thesis

Chapter 2 places this research in context by summarising the state of the art, identifying topics that are open for research, and identifies the contribution made by this thesis. This includes reference to the author's own contributions to the field of UAV communications published during this research [6] [19] [20] [21] [22] [23] [40].

The method used in this research, together with descriptions of the scenario, metrics for measuring performance, and underlying assumptions are in chapter 3. Chapter 3 then introduces the problem of optimally locating a single communications UAV and uses this to explore the sensitivity of the method to several key parameters.

Chapter 4 develops the ideas described in chapter 3 in the context of a pair of UAVs. It develops a basic model for a two-player non-cooperative game to maximise the joint coverage of the mobiles. In these games the UAVs operate at a fixed altitude, thus they have a limited number of possible locations. This limits the size of the strategy set, and allows the basic characteristics of the games to be explored in a constrained environment. Methods for solving the two-player games are compared to assess their effectiveness at analysing the payoff matrix to find Nash equilibria.

Maximising coverage is only one requirement of the communications network. Other considerations such as the stability of the network topology, making best use of the available RF power, and choice of frequency bands will have a significant impact on coverage. Chapter 5 develops the game to explore how to implement it in a real network.

Chapter 5 starts to lay some groundwork for future research by addressing

scalability issues. It introduces games with larger payoff matrices, for example games in which the UAVs are allowed to move in three dimensions, or the use of more than two UAVs. This chapter leads towards an understanding of the scalability of a general n -UAV problem.

Chapter 6 summarises the key conclusions of this research. Some potentially interesting areas for future research are identified in chapter 7.

Two appendices are included. Appendix 1 provides an introduction to game theory. Appendix 2 summarises the communications, kinematic and navigation models used in this thesis.

Context and Related Work

2.1 Terminology

An Unmanned Aerial Vehicle (UAV) is an aircraft that does not contain any crew. The absence of crew offers many advantages over a manned aircraft, for example the systems required to support life and inform aircrew can be omitted from the design. This saves weight and allows the airframe to be optimised to its specific role. Furthermore, UAVs can execute manoeuvres and missions that would be prejudicial to the health and wellbeing of the crew of a manned aircraft. Freed from these constraints UAVs can be used to conduct dull, dirty and dangerous civil missions. Examples include firefighting [4], climate research [11], traffic monitoring [75], science [78] and contaminant monitoring [96].

Every UAV is constrained by current international airspace treaties, and derived national regulations such as the UK's CAP 722 [16], to have a pilot who is capable of exercising control of the aircraft. The pilot normally exercises this control through a Ground Control Station (GCS). When the UAV is beyond line of sight (BLOS) from the GCS communications are maintained by satellite or intermediate radio relays. UAVs are thus often referred to as Remotely Piloted Air Systems (RPAS) in some literature. The delays and errors inherent to all radio channels mean that the pilot cannot always exercise real-time control so most UAVs have some degree of autonomy.

The combination of the UAV and its GCS are sometimes called an Unmanned Aerial System (UAS).

This thesis addresses a specific class of UAV called a Medium Altitude Long Endurance (MALE) UAV. This class of UAV is designed to operate for 6-24 hours at altitudes up to 25000 ft. While expensive, typically costing \$9m to \$20m, they are beginning to make an impact in civil applications.

There are many other classes of UAVs, details of which can be found in references [5] [68] and [101].

2.2 A Brief History of Unmanned Aerial Vehicles

The earliest unmanned aircraft is commonly believed to be the Curtis-Sperry Automatic Airplane [30]. This first flew in 1916 and, based on Sperry's gyros, had good stability but poor guidance making it unsuitable for any practical purpose. The first practical unmanned aircraft was a modified Tiger Moth known as the Queen Bee which was used by the Royal Navy for gunnery practice.

World War 2 saw unmanned aircraft used as weapons. The German V-1 "buzz bomb" and Henschel Hs 293 stand-off weapon proved to have very effective guidance during their short operational lives.

The Cold War saw the rapid development of unmanned aircraft as platforms for reconnaissance. Spy satellites such as Corona could provide regular coverage of targets, but their predictable orbits allowed targets to be concealed during overflights [32]. High altitude manned aircraft such as the U2 spy plane had significant success until the development of anti-aircraft systems which could reach those altitudes. Furthermore the lengthy and demanding U2 flights showed that there is a limit to the endurance of a pilot on long duration reconnaissance missions [74].

In 1962 the designer of the U2, Clarence "Kelly" Johnson, predicted that such missions would eventually be conducted by "uninhabited aerial vehicles" and developed the Lockheed D-21 reconnaissance drone. The D-21 had a limited operational life but showed that UAV were particularly well suited for long endurance reconnaissance missions. The unmanned Ryan 147 "Lightning Bug" replaced the U2 on over 3500 reconnaissance missions over Cuba, China and Vietnam. The Ryan 147 enjoyed significant success over two decades with 84% of missions completed successfully [92]. Newcome considers that the success of the Ryan 147 signalled

the start of the modern UAV era [66].

Other nations were equally active in developing reconnaissance UAVs. The Tupolev Design Bureau developed a UAV, the Tu-141, which had a range of 1000km and could be launched from a trailer. The Tu-141 and its derivatives were widely employed by Russia throughout the 1970s and 1980s, and were exported to many of Russia's allies. Israel's use of American UAV's during the Yom Kippur war encouraged Israeli Aircraft Industries (IAI) to develop the "Scout" and "Mastiff" UAVs, initiating Israel's leadership of the international UAV industry.

The value and potential of military UAVs was recognised in operation Desert Storm, the international action to liberate Kuwait. The Pioneer UAV, an upgrade of the Mastiff, was used by the US Navy to gather intelligence for amphibious forces. The ability to observe enemy activity remotely, and safely, accelerated the development and introduction of a new generation of American UAVs. The "Predator" and "Global Hawk" are routinely used for missions deemed too dangerous or politically sensitive for manned aircraft.

The development of UAVs after Desert Storm was dominated by the requirements of middle eastern conflicts. The primary role of the present generation of military UAVs is to provide Intelligence, Surveillance, Target Acquisition and Reconnaissance (ISTAR) services. UAVs were used extensively in military operations in Afghanistan and Iraq as they were the only practical way of monitoring activity in such large and inaccessible countries. The attachment of weapons to UAVs added the capability to attack targets immediately after their identification, creating an association between UAVs and controversial military actions in the public mind.

Until recently the dominant applications for UAVs have been in support of military and government activities. The high cost of developing new airborne platforms and the relatively small market have tended to deter potential civil applications. Miniaturisation of electronic navigation and control systems, increased use of composite material in airframes and the development of batteries with high energy density have reduced the size and cost of UAVs making them affordable to non-military customers.

In recent years there have been many studies into new applications for UAVs, for example those by NASA [28] and the European Commission [35]. The prevalent view is that the UAV sector exhibits strong potential for growth and that new civil and commercial applications will continue to be developed.

2.3 UAVs in non-military applications

In 2004 NASA reported on future civilian applications for UAVs [28], and predicted their usefulness in many roles:

- Land management, including forestry and agriculture, vegetation and livestock monitoring
- Commercial, including crop dusting, surveying, and broadcasting
- Earth science, including cloud and aerosol measurements, meteorology, contaminant measurement, glacier and ice-sheet monitoring, extreme weather monitoring, wildlife census
- Homeland security including coastal patrol, forest fire mapping, emergency communications

UAVs have applications in agriculture, and small companies have emerged which will overfly farms to monitor crops and wildlife non-intrusively. This has opened up the possibility of targeting chemicals only at areas of a field that have become infested by pests, the so-called “precision agriculture”. Rotary UAVs have become an essential component of filming as they can fly cameras more cost effectively than a piloted aircraft. UAV cameras are becoming a common site at music festivals, and it is likely that this services will expand into other areas of entertainment.

The research in this thesis addresses the use of UAVs as communications nodes supporting ground mobile users. This application has potential use for “first responders” in large scale incidents or disasters where the communications infrastructure either does not exist or has been rendered inoperable by adverse weather. This is a potentially large market, for example within the European Community floods and storms account for 63% of all natural disasters and 65% of all financial losses [36]. Rapid restoration of communications infrastructure is essential in the management of these disasters.

The growth of UAVs into the civil market is hampered by the regulatory environment. The operation of manned and unmanned aircraft in the same airspace generates significant air traffic management issues. The pilot of a manned aircraft is in close contact with the controls and instruments, can react to verbal instructions from air traffic control, and can see potential hazards through the windscreen.

The UAV has no equivalent systems, and any manual control is exercised remotely over communications links. The UAV is sometimes required to act autonomously in response to delays, errors and outages in the communications channel.

While military air users tend to operate in tightly controlled and segregated airspace, civilian airspace is far more complex. Thus the introduction of UAVs into civil airspace remains a significant regulatory problem. Programmes such as ASTRAEA (Autonomous Systems Technology Related Airborne Evaluation & Assessment) are addressing these issues by establishing equivalence between the procedures used to operate manned and unmanned aircraft. The timescales for the use of UAVs in civil airspace are uncertain but all parties, including regulators, are committed to enable this new capability.

The use of UAVs by governments or non-governmental organisations (NGO) such as police forces and first responders is likely to develop first. Missions such as observation of incidents, searching and rescue or initial disaster response could involve the temporary closure of small volumes of airspace, allowing civil UAVs to operate in similar conditions to military UAVs. Response to large scale disasters such as floods, forest fires and hurricanes could open up opportunities for flying larger UAVs on longer duration missions over a wide area.

2.4 Communications UAVs

The concept of using UAV as communications nodes was first identified by Pinkney in 1996 [71]. A payload architecture for this communications node was a simple transparent transponder at X band and was intended to extend the range of high capacity trunk radios while providing simple data broadcast [70]. It was recognised that communications UAVs could be more than just a replacement for hilltop radio relay systems, but could offer a distinctive capability which fitted somewhere between line-of sight radios and satellites. In 2004 the University of Colorado conducted the “Ad hoc UAV-Ground Network (AUGNet)” experiment with small UAVs equipped with IEEE 802.11b systems [14] [13]. This demonstrated the principle that UAVs could improve connectivity in ad-hoc networks. Experiments by Boeing and the US Naval Research Laboratory [31] emphasised the use of a ground network to support multicast data from a UAV

The doctrine and operational concepts for the use of UAV have been analysed extensively by the world’s armed forces. The need for clear doctrine for UAVs was

first identified following NATO operations in Kosovo [33]. Initial doctrine only considered reconnaissance missions, however doctrine evolved over the following decade to include airborne communications systems [77] [29]. The payload architectures considered in these early papers were basic rebroadcast systems, but this suited the military communications architectures of that period. More recently the US Department of Defense developed new architectures based in IPv6 networks. This concept, known as the Global Information Grid (GIG), includes airborne routers in its communications architecture [90]. This concept of supporting UAVs with onboard routers is one focus of this research as it is relatively unexplored territory.

COST 297 (High Altitude Platforms for Communications and Other Services) (HAPCOS) has generated some interesting concepts for the provision of communications from airborne platforms [84] [85]. While the primary interest of HAPCOS is the application of tethered aerostats, the communications principles are equally applicable to powered UAV. It is not a large step from the tethered aerostats in HAPCOS to the provision of civil communications services from winged UAVs. A recent study sponsored by the European Commission [37] predicted a gradual increase of UAV in civilian applications for the period to 2020, including broadcast and communications service.

HALE UAVs such as the Defense Advanced Research Projects Agency's "Global Observer", Boeing's "Solar Eagle" and Airbus's "Zephyr" have been designed to remain airborne for many days. Multi-day missions at altitudes over 60,000 feet will permit persistent communications coverage of a large area.

A practical implementation of UAV communications appeared in 2011 when Boeing demonstrated a narrowband transponder for small UAVs. This allows a UAV to rebroadcast a single voice or low speed data channel for up to 160 nautical miles [26]. While relatively unsophisticated, it demonstrated that UAVs could offer practical benefits as communications nodes.

In 2012 Northrop Grumman was contracted by US Department of Defense to install "Battlefield Airborne Communications Node" (BACN) payloads onto two Global Hawk UAVs. BACN uses a software radio to convert signals between different waveforms and includes a simple gateway between several dissimilar radio systems. BACN operates in a manner similar to a hilltop radio station as it primarily supports point-to-point transmissions and tactical networks. A picture of BACN can be seen in figure 2.1. It is easily distinguishable from other Global

Hawk airframes by the blade antennae on the underside of the wings.



Figure 2.1: Global Hawk UAV carrying the BACN payload (Picture courtesy of Northrop Grumman)

Other nations have started to consider UAVs for their military communications systems. In 2011 the Chinese Navy started to experiment with communications UAVs [93] for BLOS communications between ships.

Common concepts and technologies for using UAV in communications systems have yet to emerge. Current systems are based on either a UAV which circles a common point, such as BACN, or provision of a communications payload as a secondary system on a surveillance UAV. High level requirements for future military communications payloads are starting to emerge. The US Department of Defence has published its roadmap for developing its UAV capability up to 2030 [68]. It sees the communications payload as needing to offer “... flexibility, adaptability, and cognitive controllability of the bandwidth, frequency and information/data (e.g. differentiated service, separate routing of data based on priority, latency etc)” .

2.5 Choice of UAV Platforms for Communications

The cost of UAVs is a significant factor when considering their suitability as communications relays. The BACN payload can support mobiles up to the geometric edge of its radio footprint, typically up to 100km from the UAV. The provision of continuous cover by a system like BACN requires two or three airframes, allowing

one to be on station while the others are being refuelled and maintained. At \$100m per airframe a HALE UAV can be an expensive way to provide communications cover over a large area.

A few experiments have been conducted to provide short range communications using inexpensive 802.11 communications systems mounted on small UAVs [25] [31] [43] [99]. In theory, area coverage could be provided by using large numbers of small UAVs. In practice, the limited flight time of small UAVs, and the large number of UAVs required, makes this an impractical proposition.

In between these two extremes is the possible use of a small number of MALE UAVs. A MALE UAV will have less available electrical power for payloads, and the radius of a typical footprint will be 30-50km. While it will be less capable than a HALE UAV it has one significant advantage: it is much cheaper.

The provision of continuous coverage requires at least 3 UAVs so that one can be “on station”, one in transit to relieve it, and a third as a maintenance spare and hot standby. In the event of any problems with the “on station” UAV there could be a time lag before its replacement is available, thus a disruption to service across the whole area is inevitable. The RF power required to cover large numbers of mobiles across a large area can be considerable, however the airframe can only provide about 4kW to the communications payload. This limits the number and bandwidth of communications accesses, particularly at the edge of cover. Overall, a single HALE is not seen as a particularly practical solution for area coverage.

It is clear that MALE UAVs will have a smaller footprint than HALE UAVs because of their lower altitude. Maintaining equivalent area cover to a HALE UAV will require 3-5 MALE UAVs to be “on station”. The procurement cost of the MALE airframes will be significantly less than the HALE airframes and provide improved resilience; if one HALE is lost all cover is lost whereas the loss of one MALE only results in partial loss of area cover. The cost and resilience advantages of using several MALE UAVs are obvious. This research is constrained to addressing the use of MALE UAVs as this class of UAV is most likely to be employed in a communications relay role.

2.6 UAVs and Autonomy

The pilot of a UAV is not in intimate contact with the aircraft. Control of a UAV becomes more complex as the pilot’s knowledge of the UAV’s status is limited by

the available bandwidth. Furthermore the delays inherent to radio links, and the variable quality of communications paths, adversely affect the pilot's ability to fly the UAV in real-time. These problems can be overcome if the UAV operates with some degree of autonomy.

The difference between "autonomy" and "automation" is important. An automated system has a defined set of processes which respond to specific inputs and lead to a predetermined set of outputs. If an automated system encounters circumstances for which it is not programmed then it will be unable to make a decision without external assistance or reprogramming. An autonomous system has the ability to make decisions with a reduced level of external assistance. In some circumstances an autonomous system can make decisions based on incomplete information.

NASA discriminates between automatic and autonomous processes in space systems very clearly [86]:

Automated processes simply replace manual processes with hardware/software ones that follow a step-by-step sequence that may still include human participation. Autonomous processes, on the other hand, have the more ambitious goal of emulating human processes rather than simply replacing them.

This suggests that automation and autonomy are alternatives. Such a position is understandable when commanding a rover on Mars or a deep-space probe as, in these circumstances, the lengthy round-trip delays for signals prevent real-time management. Earthbound platforms have shorter round-trip delays which are tolerable for some UAV subsystems. Pilot, navigation and some operational systems require real-time decisions so NASA's view of autonomy and automation is too coarse for UAV operations.

The level of assistance which is offered to the pilot by an autonomous UAV will not be constant but is likely to change between missions, mission phases, platforms and subsystems. One popular taxonomy for UAV autonomy is the Pilot Authority and Control of Tasks (PACT) taxonomy [82]. PACT assigns levels of autonomy from fully commanded to fully autonomous. The PACT levels of autonomy can be seen in figure 2.2.

PACT offers a graduated scale from level 0 (pilot has full authority) to 5 (UAV has full authority). At PACT levels 1 and 2 the UAV offers advice to the pilot,

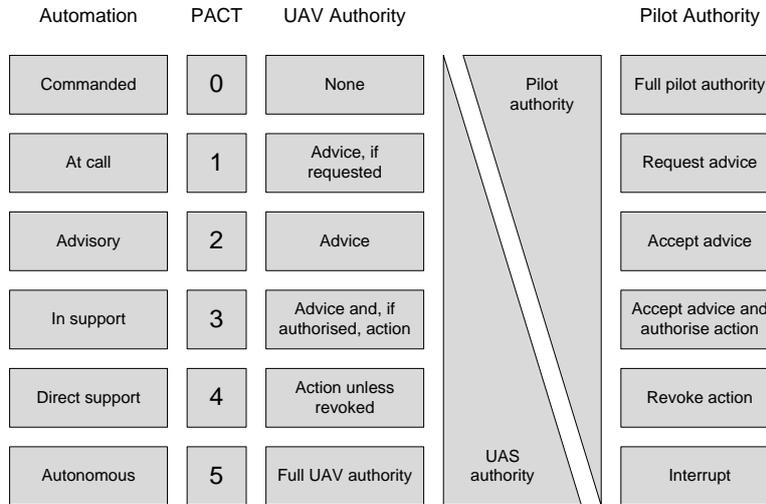


Figure 2.2: PACT Levels of Autonomy

either on request or at its own instigation. The pilot has the choice of accepting this advice. At PACT levels 3 and 4 the roles are reversed as the UAV advises the pilot about actions it intends to take and offers the pilot the opportunity to revoke these actions.

This thesis considers the case where the UAV will implement its resource management and navigation decisions unless overridden by pilot command. The UAV and its payload are thus operating at about PACT level 4. With the integration of sense-and-avoid technologies a communications UAV could operate at PACT level 5, reconciling its mission decisions with the need to recognize hazards such as other air users and terrain, and take appropriate evasive action.

Coordination of the actions and paths of multiple UAVs, all of which are operating with significant levels of autonomy, requires some inter-platform communication. It is recognized [39] that collaboration over radio paths is very sensitive to noisy and unreliable communications, particularly where centralized decision-making is employed. Decentralized decision-making reduces the bandwidth required for inter-UAV communications [50]. It has been shown that the requirements for communications do not scale well in centralized decision-making systems [87]. Furthermore, decentralized decision-making systems can be efficient, particularly if each agent is able to operate with only partial information.

This suggests that a task undertaken by multiple UAVs could be efficiently implemented with distributed decision-making based on incomplete information

and limited communications. Such an approach would be compatible with UAVs operating at PACT level 4.

This is clearly an interesting subject to explore and suggests direction for this research.

2.7 Path Planning for a single UAV

Techniques for planning the paths of unmanned vehicles have existed since the earliest autonomous vehicles. Early path planning systems used gyroscopes, accelerometers and timers and were implemented as electromechanical systems or on analogue computers.

The availability of digital computers and GPS revolutionised the guidance of unmanned vehicles and allowed families of increasingly complex algorithms to evolve, however this subject has seen little original research in recent years.

Jones [52] notes that interest in UAVs diminished during the 1970s and 1980s. This marked a period between the development of effective countermeasures in the 1960s and the deployment of GPS in the mid 1990s. It has been decided to start this assessment of UAV path planning in the mid 1990s as this marks the start of the current era of GPS and digital computers.

In the late 1990s military UAVs were remotely piloted, often over a satellite link. This introduced a time lag of several seconds between the UAV identifying a problem and the pilot taking corrective action, thus UAVs required a low level of autonomy in order to address urgent problems. Flight planning occurred before take-off so the mission followed a pre-defined plan with very few changes being implemented in-flight.

The proliferation of UAVs in the late 1990s made the use of satellite channels increasingly impractical as demand for satellite capacity was increasing for other applications. The combination of pressure on capacity, and smaller flight control systems, initiated the development of algorithms which increased the autonomy of UAVs.

The USDoD's UAV Roadmap 2005-2030 [68] uses the Autonomous Capability Level (ACL) model to illustrate the trends in UAV autonomy. The acceleration towards full autonomy predicted in the Roadmap seen at Fig 2.3 has transpired to be optimistic. This may be due to the inability to choose between the wide range of algorithms rather than any technical limitations.

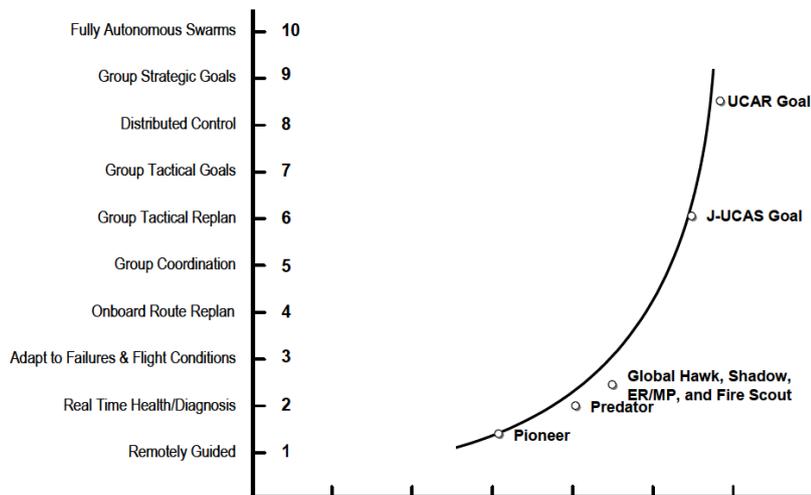


Figure 2.3: USDoD Autonomy Trends

A representative sample of algorithms indicates the diversity of approaches. Bortoff [12] proposed an algorithm for planning the path of a single stealthy UAV that is trying to avoid known radar threats. His algorithm has the UAV follow the edges of a Voronoi diagram. Other path planning algorithms, based on evolutionary algorithms, have been developed for avoiding obstacles [73] or managing performance constraints [51].

Roberge et al showed that genetic algorithms offered better solutions than particle swarm optimization for planning the path of a single UAV [76]. Their analysis considered such factors as fuel consumption and terrain avoidance, demonstrating a new level of complexity in path planning.

Weiss developed an algorithm for obstacle avoidance based on graph theory [95]. Earlier methods generated a fixed flight plan and were unable to adapt the path in response to new information. Weiss' approach planned the path in short segments and introduces the possibility of dynamic path planning, an approach that has been incorporated in this research.

There is little published material on path planning for a single UAV on communications missions.

2.8 Path Planning and Coordination for Multiple UAVs

Coordination of multiple UAVs is an emerging area of research, mainly in the area of coordinated target surveillance and communications relay for range extension of airborne sensors.

A recurring theme for papers in this area is the planning of efficient paths which allow multiple targets to be visited in a region containing forbidden zones such as threats or obstructions. Several path planning techniques have been adopted, for example Dubins paths to provide shortest routes [53] [88], following the edges of a Voronoi diagram [81], carrot following and A* paths [55] and Pythagorean hodographs [80]. While primarily addressing the problem of timeliness of cover, some of these papers recognize that the coordination process requires some communication between the participating UAVs.

Several researchers, notably Holmberg and Burkadov, have addressed the interconnection of UAVs operating beyond line-of-sight with a ground station [49] [15]. The key aspect of this problem is the optimal locating of the relay UAVs. This range extension model has included single or multiple UAVs in the relay chain. More recently there has been research into the path planning and networking of multiple UAVs providing area sensor coverage [100] [98]. Most recently Luo et.al. have considered the use of dedicated communications UAVs within a sensor network, saving communications power and thus increasing the flight duration of the sensor-equipped UAVs [57].

The use of game theory in path planning was suggested by Shen et al [79]. Their approach was to use Markov games as part of an ensemble of tools for the path planning of UAVs on Intelligence Surveillance and Reconnaissance (ISR) missions.

Tortonesi et al recognise the importance, and the complexity, of coordinating UAVs in communications roles [83]. While their paper is aimed at military communications systems, their approach is equally applicable to non-military systems. The paper proposes that any effective coordination system must fulfil three criteria:

1. It must maintain a consistent system state
2. It must enable effective decision making

3. It must be resilient to interference

These criteria have influenced the direction taken in this thesis. All UAVs need to be able to derive similar estimates of the system state, herein defined as the current estimates of the locations for all mobiles and UAVs, by exchanging data. It is permissible for the data to have aged, but it must be possible for the current state to be estimated from that data.

Maintaining a consistent system state, and effective decision making, are addressed by ensuring that all UAVs are capable of sharing information and developing similar estimates of the system state. The state estimates contribute to the calculation of the payoffs and, through that, to the decisions about UAV location and power management. This research partially addresses interference management indirectly by minimising the power required for each link.

Bekmezci et al surveyed the communications implications of multiple UAVs acting together and coined the term “flying ad-hoc networks” (FANETs) to describe the characteristics of multi-UAV communications networks [9]. They recognised that the topology of a FANET changes more rapidly than its terrestrial equivalent, the MANET. Subsequent work in FANETs has tended to concentrate on the allocation of sensor tasks to UAVs rather than networking aspects [8].

It can be seen that research has focussed on range extension, surveillance and sensor missions based on classical optimisation techniques. There is still significant scope for original research in the use of UAVs as communications network nodes. Research into methods for coordinating UAVs on area communications coverage missions is a relatively new topic. This thesis addresses these subjects.

2.9 Game Theory in Wireless Networks

Game theory started to find applications in wireless networking in about 2001. This thesis primarily addresses the physical and MAC layers, particularly the control of radio power, as these subjects have received relatively little attention.

Early papers consider the application of game theory to a single topic such as power management [59] [58], interference avoidance [62], interference management [67], cognitive radio [65] [64] and channel allocation [97].

Mackenzie et al [59] suggests that communications systems based on game theory could make selfish behaviour unprofitable, making the overall system per-

formance optimal. A second paper in the same year [60] recognises that game theory is suited to local decision-making and dynamic topologies and advocates the use of game theory in ad-hoc networks.

There are many papers on the application of game theory to the network and higher layers of the protocol stack. Altman et al reviewed many applications for non-cooperative games in computer networking [3]. They cite many proposed applications at the MAC, network and transport layers for applying game theory to flow control, congestion management, routing and QoS management.

Lim [56] applied game theory to a wireless ad-hoc networks by using “good-put/transmit power” as the utility function. This demonstrated potential applications for game theory at the MAC layer. Shortly after this Pavlidou summarised research into the application of game theory to routing [69].

The interest in applying game theory to communications problems started to widen towards the end of the first decade. Until this time most authors had concentrated on non-cooperative games, seeing these as an alternative to classical optimization techniques. Some authors considered other types of game to address problems that had hitherto been addressed only by non-cooperative games. Menon, for example, applied potential games to interference management [63]. Bhattacharya combined kinematics, communications and differential games to consider how UAVs could avoid an aerial jammer [10]. This started a trend for authors to address the optimization of the whole system rather than a single metric, an approach that influenced the research reported in this thesis.

A survey of game theory in wireless networks by Charilas et al [17] showed that game theory was applicable to all layers of the protocol stack. They also observed that some of the implicit assumptions of game theory, particularly rationality willingness to participate, were often absent in real networks where some processes were inherently selfish, and the human users are often irrational and selfish. They correctly identified that the choice of utility function directly affected the computational resources required to calculate the payoff.

Cong et al advanced the principles from simple normal-form games into auctions by proposing a Stackelberg game for resource allocation [27]. The hierarchical player model they propose was considered for this research but was not pursued, mainly because it sat uncomfortably with the concepts of a self-organising network of peers. Furthermore, the appointment of a leader would create a potential single point of failure and undermine the principle of highly autonomous nodes.

Akkarajitsakul et al [2] continued the exploration of other types of game by proposing Bayesian games for resource allocation.

The fashion for cross layer solutions to networking problems was picked up by Gunasekaran et al. They proposed using game theory to optimise cross layer problems in wireless networks [46].

There have been numerous papers that offer game theoretic solutions to specific communications problems. This thesis takes the more ambitious approach, spearheaded by Bhattacharya, of attempting to address the whole UAV communications system. It links the application of game theory to both communications and path planning as, in this application, the two subjects are interdependent.

2.10 Path Planning and Coordination for Multiple Communications UAVs

Basu et.al [7] consider coordinated flocking of UAVs to provide communications coverage to a community of ground nodes. This paper shows that basic flocking behaviour, and consideration of current UAV attitude, can be used to improve coverage of islands the mobile nodes. After ten years this is one of very few papers that addresses the subject of coordinating multiple communications UAVs.

Agogino et al [1] look to a distant future when large swarms of solar powered HALE UAVs could support ad-hoc networks. This visionary approach suggests that evolutionary algorithms could offer an approach for coordinating communications coverage.

Giagkos and this author have also considered evolutionary algorithms for coordinating communications UAVs [40]. Initial results are encouraging but it is clear that evolutionary algorithms require a higher communications overhead than game-based approaches.

Dac-Tu et al [48] suggest swarms of small, electric, communications UAVs could be coordinated using particle swarm optimisation. The emphasis would be on conserving energy to extend mission duration rather than maximising communications coverage.

The concept of path planning for multiple communications UAVs is a very new topic. After Basu's paper in 2004 the subject was largely ignored for 8 years, possibly because of a surge in interest in UAV sensor networks. The last two years

Subject Area	Key References	Comments
Military use of a single communications UAV	[71] [70] [33] [77] [29] [90] [26] [68]	Extensive research and conceptual papers but few fielded systems
Civil use of a single communications UAV	[84] [85]	Mainly tethered aerostats and small UAVs
Path Planning for a single UAV	[12] [73] [51] [95] [76]	There is relatively little original research into path planning for single UAVs.
Path Planning and Coordination for Multiple UAVs on Missions Other Than Communications	[53] [88] [81] [55] [80] [7] [49] [15] [100] [98] [83]	This is a rich and active research area, particularly for range extension and coordinated reconnaissance & surveillance missions.
Game Theory in Wireless Networks	[59] [58] [62] [67][65] [64] [97] [59] [60] [3] [56] [63] [10] [17] [27] [46]	Extensive activity over the last decade
Path Planning and Coordination for Multiple Communications UAVs	[1] [40] [48]	Very little activity. No activity on game theory
Analysis of the whole system	[10]	Permits study of subsystem interactions

Table 2.1: Gaps in Current Research

have started to see a few authors start to address the problems of coordinating multiple communications UAVs. At the time of writing no-one has considered game theory as an approach to this problem.

2.11 State of the Art

Table 2.1 gives an overview of the state of the art.

The military and civil use of communications UAVs is a relatively new field. There have been many studies into concepts and doctrine for military communica-

tions UAVs but only one system, BACN, has been fielded. Civil communications UAVs have yet to be developed.

The navigation and route planning of a single UAV has been thoroughly researched and little new material has been added to the subject in recent years. The coordination of multiple UAVs has seen many papers, mainly in the context of range extension for UAVs operating BLOS or cooperation in sensor tasks. Few have addressed the use of UAVs as communications nodes and how their coverage could be maximised, indicating that this subject offers significant scope for original research. Similarly, few have addressed their research at the behaviour of the whole systems of UAV, payload and path-planning. No-one has considered game theory as a potential solution to this problem.

This research anticipates the need for a method to coordinate communications UAVs and contributes a potential solution to those problems.

Chapter 3

Experimental Method and Principles

3.1 Chapter Introduction

This chapter establishes the simulation method and experimental principles used in this research.

Experiments were conducted using a simulation environment comprising MATLAB and Systems Toolkit (STK). Simulations were run in MATLAB on a dual core processor with 20 GB of RAM using Windows 7. The operating system was stripped of non-essential processes to improve speed.

The algorithmic aspects of the simulation were coded in MATLAB with the UAVs and mobiles and radio systems represented in STK. A more detailed description of the simulation environment was presented at the IEEE UK Simulation Conference [22], and a technical description of the kinematics and communications models can be seen at Appendix B.

This chapter starts by defining the standard models used in the simulations. It then describes the experimental method and the metrics used to evaluate performance of the system. These metrics are used to explore the parameter space of the system and, from this, to define a set of common parameter values that are used in the experiments.

Using the default parameters, a scheme for relocating a single UAV is defined. This relocation scheme forms the basis of the action sets used in subsequent chapters to define the games. A series of simulations, run with a single UAV, are used to explore the potential benefits of relocating the UAV.

3.2 System Description

The system comprises a set of UAVs, a set of mobiles, and the equipment that allows them to communicate. Radio links can be formed between the UAVs and the mobiles to allow duplex communications. These radio links are implemented using transmitters, receiver and antennae on the UAVs and mobiles.

A diagram of the main system components for the radio link can be seen in figure 3.1. The Communications Controller on the UAVs and mobiles acts as the interface between the navigation and communications subsystems. On the UAVs the Communications Controller implements power and access control, and proposes changes to the UAVs' paths based on its estimates of the future communications demand. On the mobiles it implements power and access control under the direction of the Communications Controller on the UAV. It also periodically broadcasts the mobile's location.

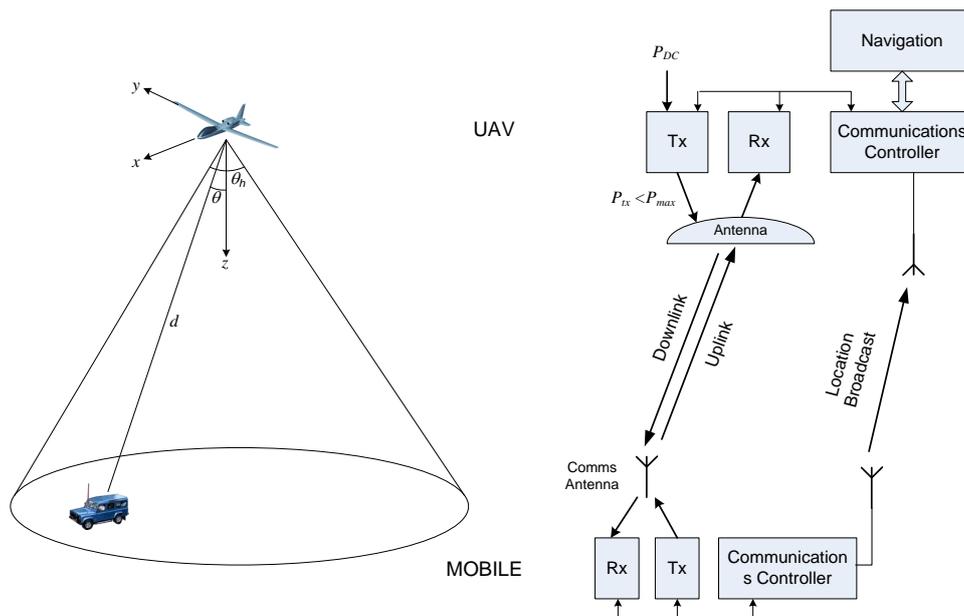


Figure 3.1: Main System Components

Figure 3.1 shows a direct interface between the Communications Controller and the Navigation System. In a real system this interface has security and flight safety implications as it could provide a communications path for potential attackers. A malicious user, or third party, could plant false information into the Communications Controller causing the UAV to relocate away from its supported

mobiles. This false information would be shared with other UAVs over the IUAVL, causing them to reach incorrect decisions. A real implementation of this system needs to address how false messages could be identified to deny attacks through this interface. While outside the scope of this thesis, this is an interesting area for future research.

Consider a single link between a UAV and a mobile. The UAV receives a series of location broadcasts from the mobile informing the UAV of the mobile's latitude, longitude and altitude. The UAV uses these to estimate the future location of the mobile and, by estimating the radio slant range, can calculate the transmitter power P_{tx} required to support that downlink based on its own latitude, longitude and altitude h . It can also calculate the off boresight angle θ and use this to determine whether the mobile is visible to the UAV by testing if θ is less than the half-beamwidth $\frac{\theta_h}{2}$.

The UAV has a maximum power P_{max} . This is usually limited by the transmitter amplifier maximum power output, but at very high altitudes could be limited by the power offtake from the UAV P_{DC} . If $P_{tx} < P_{max}$ then there is sufficient RF power available to support the downlink.

It is assumed in this research that there are no significant power limits on the mobile, therefore the main power constraint on the communications system will be from the UAV payloads's amplifier maximum power, P_{max} . This allows the research to focus purely on the UAV downlink as the limiting factor for communications. This assumption is realistic in systems where one UAV supports symmetric links to many mobiles as the UAV needs to support many downlinks, but the mobile only needs to support one uplink.

The interaction of three parameters, maximum power P_{max} , beamwidth θ_h and altitude h , have a significant effect on the number of mobiles that can be supported from any location. A preliminary investigation of these parameters was conducted to establish a range of values that is of interest. This identified two possible states in which the payload can operate: power limited and footprint limited. These two states are defined in section 3.6 by plotting of a characteristic curve of the trade-off space.

3.3 Definition of the Standard Models

The system elements described in figure 3.1 can be represented as an object hierarchy. This is shown in figure 3.2 along with each object's key parameters and constraints.

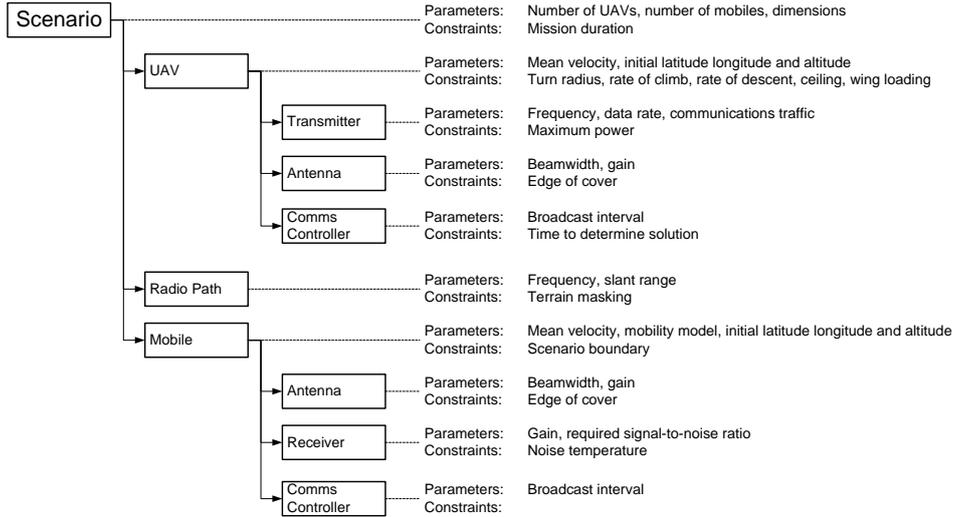


Figure 3.2: Object Hierarchy, Parameters and Constraints

The parameters and constraints are used to derive the other variables in all the simulations. For consistency between experiments these objects have been coded as a set of standard models, representative of real equipment, and some of the parameters have been fixed.

The use of MALE UAVs as communications nodes is a key requirement. It would be easy to choose one MALE UAV from the many platforms that are currently available, however a more generic model was sought. Rather than focus on any specific product the standard UAV was derived from representative performance characteristics of several current MALE UAVs.

Each UAV payload is required to support communications with multiple mobiles. There are currently very few existing communications payloads on which to base a model. It was therefore decided to model the communications payload on a single antenna that could provide area coverage. The obvious choice was a conical horn as this is lightweight, easily manufactured and easily modelled.

The ground mobile is based on an off-road vehicle, for example a Land Rover or Toyota Land Cruiser. The mobile radio uses an omnidirectional antenna as this

removes the need for any tracking systems. This defines the gain and external noise characteristics for the receiver system.

3.3.1 UAV Model and Characteristics

The characteristics of the MALE UAVs used in these experiments are shown at Table 3.1.

Variable	Value	Units
Cruising speed	75	kts
Turn radius	1500	m
Max rate of climb	650	ftmin
Max rate of descent	650	ftmin
Wing load factor	+1.5	g
Maximum altitude	22000	ft
Minimum altitude	1000	ft
Maximum payload power	200	W

Table 3.1: MALE UAV Characteristics

These characteristics are typical of an airframe that is not conducting high energy manoeuvres but is loitering over an area. The maximum altitude is representative of real UAVs. The choice of minimum altitude is less clearly defined but reflects ground clearance requirements in a typical airspace regulatory regime.

The electrical power for the payload is normally generated by power “offtake” from the engine. Some onboard electrical storage is probable on a MALE UAV, normally from batteries or fuel cells. A typical MALE UAV could be expected to provide several hundred Watts of electrical power from these sources to the payload. This will vary according to the UAV’s altitude, energy power used to implement manoeuvres, and the power consumed by other onboard systems.

For the purposes of this thesis the electrical power generation system is ignored. In this research the limiting factor for the payload is the total RF power fed into the antenna, denoted P_{max} , rather than the power offtake from the UAV.

3.3.2 Mobile Model and Characteristics

Ideally, the mobility model for the mobiles should be derived from the movement of real vehicles in a disaster. In practice anonymized data of this type is hard to

obtain for legal and ethical reasons. For this research it was important that the mobiles operate within a defined region and that they periodically form clusters as this is the behaviour of real vehicles in disasters.

Gonzales and others show that actual human mobility tends to be random over short periods but follows reproducible patterns over long timescales [44]. Emergency response vehicles do not follow reproducible patterns as they have no control over the locations of events. They do, however, tend to act within a defined region that could be defined by the geographic limits of a disaster or some governmental limits, for example an administrative boundary.

Mobiles followed a random walk mobility model in which they would change direction with a period of 10 minutes and travel at a mean velocity of 15 m/s, approximately 30 mph. The mean speed of 15 m/s was chosen as it compromise between very slow travel across difficult terrain and high speed travel on metalled roads. The random walk model was modified to give it a reflecting boundary, keeping all mobiles within the scenario area.

Each vehicle was equipped with a radio for communicating with the UAV payload. Its antenna gave hemispheric coverage to avoid the requirement for any tracking systems.

The transmit RF power limit for the radio is unspecified as the uplink is not considered. The receiver specification is defined by the ratio of the receiver system gain G_r to the system noise temperature T_{sys} , measured at the receiver input, $\frac{G_r}{T_{sys}}$. This method of defining a receiver operating in a noisy environment is commonly used in satellite communications [61].

The required signal to noise ratio is normalised to 1 bit/s and 1Hz and expressed as $\frac{E_b}{N_0}$, where E_b is the energy in one bit and N_0 is the noise power in 1 Hz. Normalisation gives the experiments a measure of independence of any particular waveform.

A simplified specification for the mobiles is in Table 3.2.

Variable	Value	Units
Mean speed	15	m/s
Mobility	Random walk	
Receiver specification	-25	dB/k
Desired $\frac{E_b}{N_0}$	10	dB

Table 3.2: Mobile Characteristics

A value of $\frac{E_b}{N_0} = 10dB$ is typical of many low order modulation schemes at useful bit error rates. The choice of $\frac{G_r}{T_{sys}} = -25dB/k$ is representative of an antenna with a hemispheric pattern at the earth's surface.

3.3.3 Payload Model and Characteristics

The UAV's payload supports multiple, simultaneous, connections with mobiles. The payload uses a conical horn antenna that has a circular footprint with its boresight aligned with the z axis of the UAV airframe.

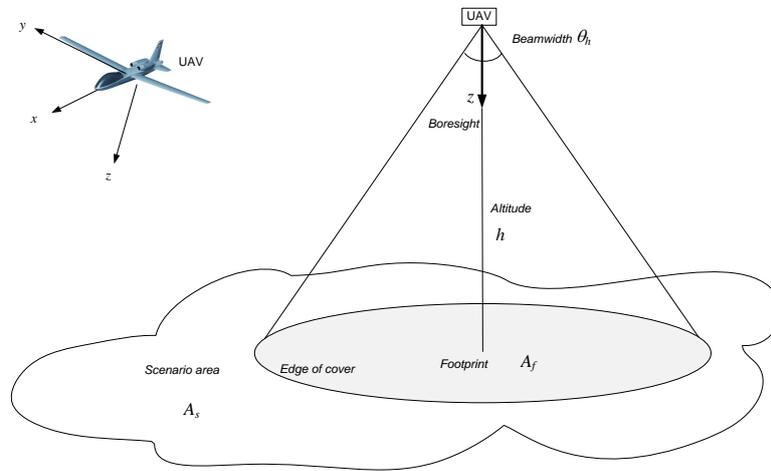


Figure 3.3: Key Terms for the UAV Payload Definition

The edge of cover defines the boundary of the footprint. This is considered to be a hard boundary outside of which no communication is possible.

The antenna is connected to a multi-channel radio that permits individual links to be established between the UAV and each mobile. It is not deemed necessary to specify how this radio is implemented. It could, for example, be multiple parallel radio systems or a single software defined radio that is capable of supporting multiple channels. The latter is preferred as it allows power to be balanced between individual downlinks.

3.4 Experimental Method

This section describes the basic experimental configuration, and the principles that were common to all experiments.

Mobility data sets were created for a set of mobiles $M = \{1, 2, \dots, i, \dots, m\}$. Each mobility data set had different start points for the mobiles and allowed them to move according to the same mobility model. Experiments were repeated with different randomly generated mobility data sets to reduce the possibility of obtaining rogue results from a single data set. Numeric results were calculated from multiple runs of each experiment.

UAVs were started at locations that allowed their footprints to be entirely within the scenario area. In most experiments the footprints did not overlap as this gave each UAV coverage of a unique subset of the m mobiles. A few experiments were conducted to establish the impact of initial conditions on algorithm settling time. In these experiments the UAVs were located so that their footprints completely overlapped as this presented a worst-case start for the algorithm.

It was noted that clusters of mobiles tended to form and disperse in a timescale of tens of minutes, so experiments had a duration of hundreds of minutes. Mindful that a typical MALE UAV mission could last for 6 to 12 hours, a mission duration of 6 hours was used as the default for all experiments. This gave time for algorithms to settle from any initial conditions. Longer periods did not yield any improvement in results.

Unless otherwise stated, the results of any experiments comprise the results of multiple runs of that experiment with different mobility data sets. Results were generated from the mean value of all the individual sets of results.

The flowchart at figure 3.4 shows the basic structure of each simulation.

At the start of each simulation the appropriate mobility data set is loaded and all variables are initialised. This includes the establishment of the initial locations for each UAV.

After initialisation the simulation enters a loop such that each iteration around the loop represents an increment of 1 second in simulation time.

The time taken to complete a single manoeuvre by the UAV(s) is known, thus it is possible to test whether a manoeuvre has been completed. If the manoeuvre is completed then the game, or other algorithm, is run to determine the destination of the UAV on completion of the next manoeuvre. This comprises two stages:

- Estimate the locations of the mobiles at a time that corresponds with the end of the next manoeuvre.
- Use these estimates to select a manoeuvre for the UAV(s) that will give the

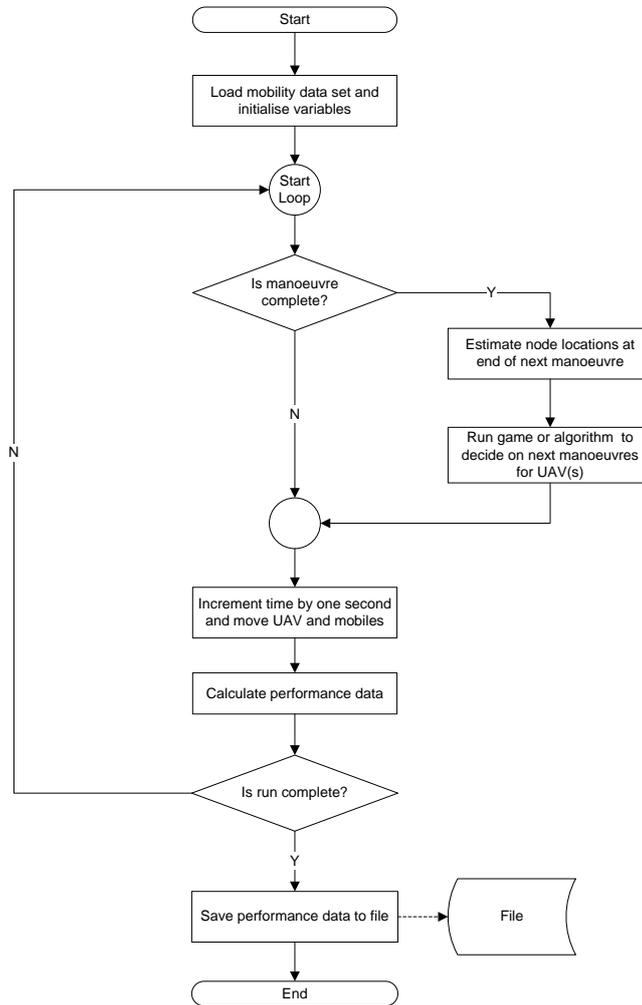


Figure 3.4: Simulation Flowchart

greatest achievable coverage of the mobiles at that time.

The method for estimating the location of the mobiles requires some historic data about the movement of each mobile. Each mobile sends a timestamped location broadcast every few seconds that contains, as a minimum, its latitude, longitude and altitude (LLA). When received by a UAV, the data from the broadcast is committed to memory in the Communications Controller. It is probable that not every UAV will receive every location broadcast; the mobile might be outside another UAV's footprint, obscure by terrain, or subject to interference. The data is retransmitted to the other UAVs over the IUAVL to ensure that all UAVs have similar state data.

If a UAV receives several location broadcasts from one mobile it can estimate the location of the mobile at the time of the next manoeuvre. Early experiments in this research used Kalman filters to estimate the mobile’s location, however it became clear that this was excessively complex. Later experiments showed that an alpha-beta filter gave very similar results but with a significant reduction in processing time, hence Kalman filters were replaced with alpha-beta filters.

The locations of the UAV(s) and mobiles are advanced by one second and the performance data is saved. This data comprises the LLA for each UAV, the power required to support each link from a UAV to a mobile, and the connectivity map for each UAV. This allows the offline calculation of the performance metrics described in section 3.5.2.

The simulation loop is repeated until the simulation is completed, at which time the metrics are saved to file for offline analysis.

3.5 Packing the Payload and Measuring System Performance

There are two important metrics that are used to quantify performance of the system: coverage and power.

Coverage is defined as the total number of mobiles that are connected to any UAV at any time. By implication, if there is more than one UAV, we can define the total coverage as the sum of the unique coverage of all UAVs. The word “*unique*” requires that a mobile can only be allocated to one UAV at a time. Uniqueness ensures that RF power is used efficiently across the network of UAVs, and helps to manage the RF interference environment. One consequence of uniqueness is that it creates a tree structure at the physical layer of the network.

Power is defined as the total RF power required by the UAV payload to support its coverage to a required signal-to-noise ratio. Power is limited to a maximum power, P_{max} , that can be supported by the power amplifier on the UAV payload. The total power is the sum of the power required by all UAVs to support the total coverage.

Coverage and power are related through the method that is used to pack the payload for each UAV.

3.5.1 Payload Packing Method

The packing method can be defined as mapping a set of mobiles M to the set of UAVs N , where $M = \{1, 2, \dots, i, \dots, m\}$ and $N = \{1, 2, \dots, j, \dots, n\}$.

Calculating the RF power required to support multiple mobiles simply requires the summation of the RF power to support individual links. Equation B.2 shows a method for calculating the power P_t required to support the single RF link between two points in order to meet the target signal to noise ratio $\frac{E_b}{N_0}$. This equation was used to calculate the power required to support every link.

If there are n UAVs and m mobiles a $m \times n$ matrix P of the downlink powers required for all UAVs and mobiles can be generated. In matrix P the downlink power from UAV j , to mobile i is denoted p_{ij} .

$$P = \begin{pmatrix} p_{11} & p_{12} & \cdots & p_{1j} & \cdots & p_{1n} \\ p_{21} & p_{22} & \cdots & p_{2j} & \cdots & p_{2n} \\ \vdots & \vdots & & \vdots & & \vdots \\ p_{i1} & \vdots & \cdots & p_{ij} & \cdots & p_{in} \\ \vdots & \vdots & & \vdots & & \vdots \\ p_{m1} & p_{m2} & \cdots & p_{mj} & \cdots & p_{mn} \end{pmatrix} \quad (3.1)$$

The rows of P show the power required by each UAV to support mobile i . The payload packing problem is similar to a 0-1 knapsack problem; each communications links requires a certain power, each mobile can only be assigned to one UAV, and there is a total power constraint for each UAV. The algorithm used to pack the payloads attempts to maximise the number of links that can be supported by each UAV.

The payload packing algorithm searches each row of P to find the UAV that requires the minimum payload power. It allocates mobile i to that UAV and zeros all the other elements in that row. This produces a matrix P' that comprises zeroes where mobiles are unsupported and value p_{ij} if mobile i is assigned to UAV j .

The sum of the column j of P' is a vector that identifies the power required for UAV j to support all the mobiles that could be allocated to it. It is possible that the sum of these powers is $\geq P_{max}$. If this is true then the UAV cannot support all possible mobiles so a second stage of packing is followed.

In this second stage the mobiles that have been allocated to one UAV are

packed in ascending order of power. This allocates the nearest mobiles first as these require the lowest power. Packing continues while there is still unused RF power. New mobiles are only added to the payload if, after their addition, the inequality in equation 3.2 is still satisfied.

$$\sum_{i=1}^m p_i \leq P_{max} \quad (3.2)$$

Allocating each mobile to its nearest UAV, and constraining the total power available to each UAV to P_{max} , results in an allocation matrix X where element $x_{ij} = 1$ if mobile i is allocated to UAV j and $x_{ij} = 0$ if it is not allocated.

The Hadamard product $P' \circ X$ generates a matrix P'' in which the power levels for all unallocated mobiles are zeroed, and the sum of every column satisfies the inequality in equation 3.2.

3.5.2 Performance Metrics

A set of metrics, derived from the matrices X and P'' , are used to quantify the performance of the methods and algorithms investigated in this thesis.

The requirement for unique connectivity implies that a mobile is only allocated to one UAV. The columns of X are the allocation vectors for each UAV, so the allocation vector for UAV j is x_j . The coverage C_j of UAV j can be defined as the sum of column j :

$$C_j = \sum_{i=1}^m x_{ij}$$

The total coverage of all the UAVs, C_{tot} , is the sum of all the cells of X

$$C_{tot} = \sum_{i=1}^n \sum_{j=1}^m x_{ij}$$

The total power required by UAV j to support its allocated mobiles is thus the sum of column j of P'' .

$$P_j'' = \sum_{i=1}^m p_{ij}''$$

Also, the total power P_{tot} required by all UAVs to support all allocated mobiles

is the sum of all the cells in P'' .

$$P_{tot} = \sum_{i=1}^n \sum_{j=1}^m P_{ij}''$$

We normally define the total coverage C_{tot} and total power P_{tot} as the instantaneous values at time t . It is also possible to define metrics over a time interval, for example considering long term values after any transient effects have ended. This is also useful if the data sample is quite large. If we consider the interval δt , delimited by timesteps t_1 and t_2 , and containing n_s samples we can define mean coverage C_{mean} and mean power P_{mean} :

$$C_{mean} = \frac{1}{n_s} \sum_{t_1}^{t_2} C_{tot}$$

$$P_{mean} = \frac{1}{n_s} \sum_{t_1}^{t_2} P_{tot}$$

In some cases coverage and power may be inadequate, for example if the number of mobiles is different between two data sets. In these circumstances a metric that is normalised to a single mobile is more useful. The metric most suited to this comparison is the specific power, that is the power required to support one mobile. The specific power over time interval δt can be calculated as:

$$SP_{mean} = \frac{C_{mean}}{P_{mean}}$$

As a general principle, the time δt is selected so that the transient effects have all terminated. The steady state values of mean coverage and power are more useful indicators of performance than whole mission values. Transient effects are better characterised by their duration than their mean value.

During each simulation the performance of the whole system was quantified by sampling key metrics at 1 second intervals. This permitted offline analysis of the performance change over time. The metrics were:

- The coverage of each UAV C_j
- The power required by each UAV P_j''
- The total coverage of all UAVs C_{tot}

- The total power required by all UAVs P_{tot}
- The latitude, longitude and altitude of all UAVs
- The latitude, longitude and altitude of all mobiles

3.6 Selection of Parameter Values

If the experiments were run with sufficiently high power and beamwidths of 180° then all mobiles would be supported. The only limiting factors would be the curvature of the earth and terrain blocking. At the other extreme, low power and low beamwidths, the coverage would be severely footprint limited so few of the available mobiles would be supported. These extreme cases are uninteresting from both operational and algorithmic viewpoints.

In between these extremes there is a range of parameter values in which relocation algorithms have the potential to improve coverage. Factors such as the spatial density of the mobiles, area of the scenario, UAV altitude, available power and beamwidth influence the effectiveness of algorithms. This takes the research into the realm of real systems in which finite numbers of vehicles operate over large areas.

It is infeasible to evaluate all combinations of parameters and values, thus some boundaries need to be established. Experiments were preceded by a sampling exercise to establish the sensitivity of results to parameter values, particularly at the upper and lower extremes of realistic parameter values. As a result of this the parameter limits were constrained to:

- Mobile spatial densities between 0.005 mobiles/km² and 0.1 mobiles/km².
- Values of P_{max} between $10W$ and $200W$.
- Beamwidths θ_h between 120° and 170°
- Altitudes h in the range $1000ft$ to $22000ft$.

3.6.1 Number and Spatial Distribution of the Mobiles

Consider a set of m mobiles operating in a scenario of area A_s . Assuming a uniform spatial distribution of the mobiles then the average density of the mobiles is $\frac{m}{A_s}$

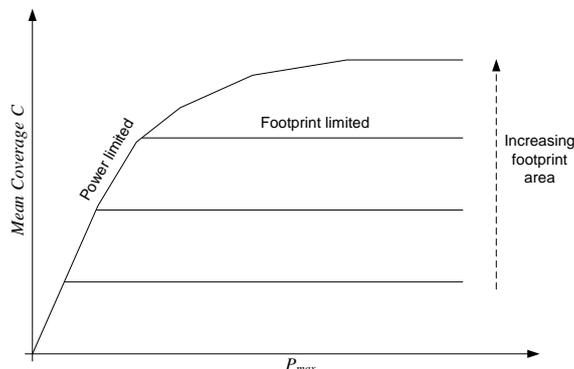


Figure 3.5: General Shape of the Coverage Characteristic

mobiles per square metre. If the footprint of the UAV has an area A_f , and is wholly contained within the scenario as shown in Figure 3.3, then the number of mobiles that would be expected to be found within the footprint would be $m \frac{A_f}{A_s}$.

The geometry of the UAV footprint will clearly affect the number of mobiles that are expected to be found. Increasing altitude h or beamwidth θ_h will increase area A_f and hence the number of mobiles that lie within the footprint.

3.6.2 Power, Beamwidth and Altitude

Parameters such as the altitude h of the UAV, the beamwidth θ_h of the antenna, the spatial density of the mobiles, the maximum power available P_{max} and the dimensions of the scenario all affect coverage. The shape of the set of curves shown in 3.5 shows how the parameters interact.

It can be seen in Figure 3.5 that increasing the footprint area, whether by increasing altitude or beamwidth, allows the mean coverage to increase. Increasing P_{max} with a large footprint area does not generate a linear increase in coverage. As P_{max} increases the rate of increase in coverage is initially high as mobiles close to boresight are supported; these require the lowest power. As P_{max} continues to increase the system starts to support mobiles that are further from boresight and which consequently require higher power. The rate of increase of coverage thus reduces as more power is needed to add the next mobile.

We can thus define the system as being in a **power limited** state if the coverage increases with P_{max} .

There comes a point where all mobiles within the footprint are supported.

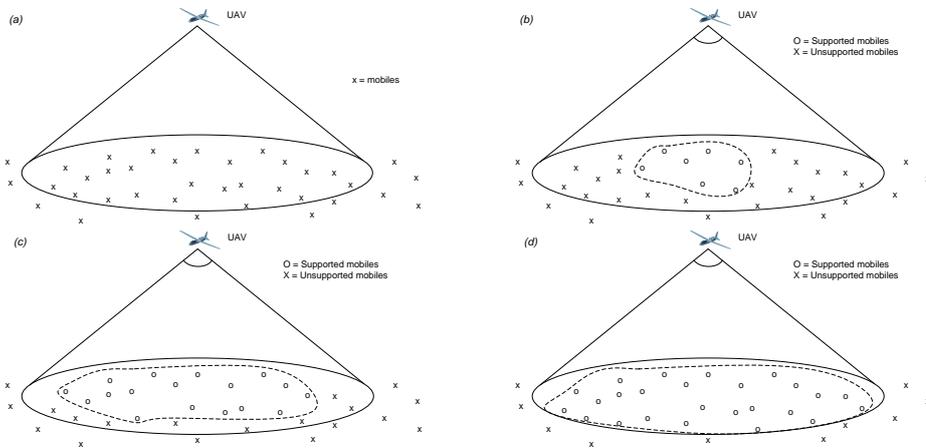


Figure 3.6: P_{max} , A_f and Coverage

Increasing P_{max} beyond this power has no effect because all remaining mobiles are outside the footprint. The system enters a **footprint limited** state where the coverage is limited by the area of the footprint. The power at which this transition occurs represents the maximum power that is required to give that coverage value with that footprint area. If further increases in coverage are required then the only options is to increase the footprint area, either by increasing altitude h or beamwidth θ_h .

It is clear that any graph with the form of Figure 3.5 will be scenario specific, however the characteristics of any system can be adequately explored by evaluating a benchmark scenario and sampling other scenarios to establish general trends.

When P_{max} is very small there is insufficient power to support any mobiles, thus the coverage is zero. As P_{max} increases the mobiles that have the lowest slant range d , and are thus closest to the boresight, are packed first as they require the least power. This is shown in Figures 3.6a and 3.6b.

As P_{max} continues to increase the mobiles that are slightly further from the boresight are supported. Coverage continues to increase with P_{max} until all mobiles within the footprint are supported, as shown in Figures 3.6c and 3.6d.

3.7 Assumptions

The following assumptions have been made to establish a clear boundary around this research:

1. It is assumed that all amplifiers are linear devices so that total power is simply the sum of all individual powers. This simplifies the packing method. Detailed models of power amplifiers are outside the scope of this research.
2. The system operates at an assumed frequency of 5 GHz, placing it in the same band as the current ITU-T allocation for UAV control systems. The use of a common band ensures comparability of results. The same method can be used at any frequency band, and a short example is shown in section 5.4.
3. A spherical earth model has been used rather than a “flat earth” model. This improves the accuracy of slant range calculations and defines the edge of cover. Other earth models, for example WGS84, were not used as the additional calculation complexity was not justified by the marginal improvements in range accuracy.
4. All mobiles broadcast their locations (latitude, longitude and altitude) every few seconds.
5. All UAV payloads are assumed to have a wide coverage receiver antenna that is used to listen to the location broadcasts from all UAVs and mobiles. This gives all UAVs a good knowledge of the locations of mobiles within the radio horizon. Implicitly, the UAVs can use previous broadcasts to establish a history of the locations of each mobile.
6. Information obtained from the location broadcasts is shared with all other UAVs.
7. All UAVs are aware of the locations of all other UAVs.
8. There are no power limits on the mobile, therefore the link performance will be dominated by the power available on the UAV downlink.
9. The maximum power available to the down link, P_{max} , will be limited by the amplifier specification and not by power offtake from the UAV.
10. It is anticipated that the game model will cause the UAVs to maintain a good separation therefore collision avoidance mechanisms are not addressed in this thesis.

3.8 Characterising the System with a Single UAV

In the research described in this thesis, all UAVs have a common payload configuration during experiments. This common configuration comprises a set of values of power limit P_{max} , beamwidth θ_h and altitude h that are common to experiments. It can be considered as a point in the set of possible values of these parameters, and is referred to as the *operating point*.

The use of a standard operating point ensures that results from experiments are comparable. The operating point can be defined as:

The operating point is a point in the parameter space that is common to all experiments.

The ideal operating point would have less than 20% of the mobiles inside the footprint. This ensures that experiments with up to four UAVs could be conducted while maintaining a pool of additional mobiles over which the UAVs could compete. At the operating point, the system would be in a footprint limited state, ensuring that there is significant spare power available for supporting additional mobiles. It is desirable that P_{max} is relatively low as this minimises power offtake from the UAV and helps to control the RF interference environment in real systems.

Mindful of the altitude range of a MALE UAV, and the consequently large horizon ranges, a scenario area of 100km \times 100km was chosen.

Defining an operating point by analysis is feasible. It is desirable that up to four UAV could operate, and that their footprints will not overlap. The geometry of placing four non-overlapping circles inside a square scenario is trivial. The MALE UAVs could operate at altitudes between 1000ft and 22000ft. To allow for manoeuvres in altitude in later chapters it was decided to start the UAV at 10000ft, roughly midway between the upper and lower operating limits.

From this altitude a beamwidth of 165 degrees will ensure that the footprints are as large as possible without overlapping. The area of each footprint will be $A_f = 1.68 \times 10^9 \text{m}^2$, and the area of the scenario $A_s = 10^{10} \text{m}^2$. Assuming that the mobiles are uniformly distributed across the scenario then the footprint should contain $\frac{A_f}{A_s}$ of the 200 mobiles. The expected coverage of the UAV should thus be 33.6 mobiles.

It was shown by experiment that the actual coverage for a uniformly distributed

set of static nodes was 34.3. In this experiment the UAV was circling the centre of the region rather than a fixed point. The small difference between the predicted value of 33.6 and the actual value of 34.3 is because a mobile located at the edge of cover is counted as inside the footprint.

The mean coverage was measured by experiment with mobiles that are subject to random motion, as described in section 3.3.2. It was found to be 38.8 mobiles with a standard deviation of 3.2, a result that is within 2 standard deviations of the predicted value. This apparent increase in coverage is caused by the formation of clusters of mobiles that are tracked by the UAV. These clusters have a short life, typically 20 minutes, but offer the UAV an opportunity to increase its coverage during that period. Algorithms that relocate to exploit the formation and dispersal of clusters can be expected to offer higher coverage than a circling UAV.

3.8.1 Sensitivity of the Operating Point to Power and Beamwidth

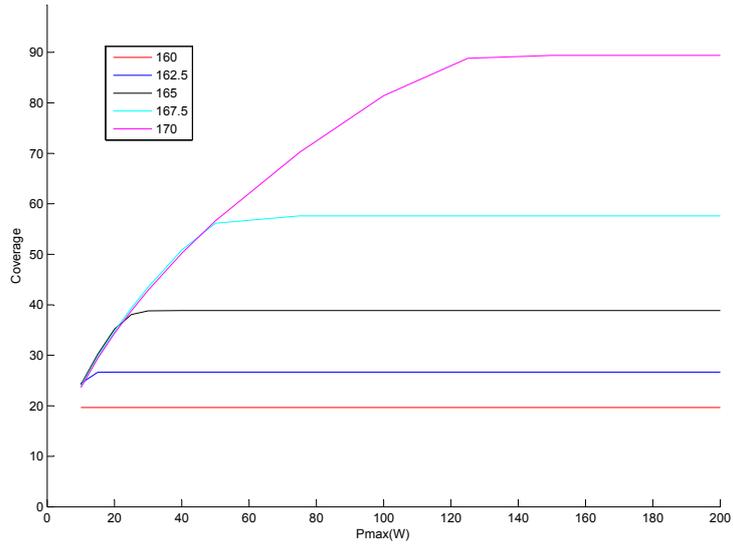
A scan over the range of beamwidths between 160° and 170° and P_{max} between 10W to 200W at an altitude of 10000 ft allows the sensitivity of the operating point to be evaluated.

Figure 3.7a shows consistent coverage of 38.8 mobiles at a beamwidth of 165° for values of P_{max} between 30W and 200W, indicating that the payload is footprint limited over this range of P_{max} .

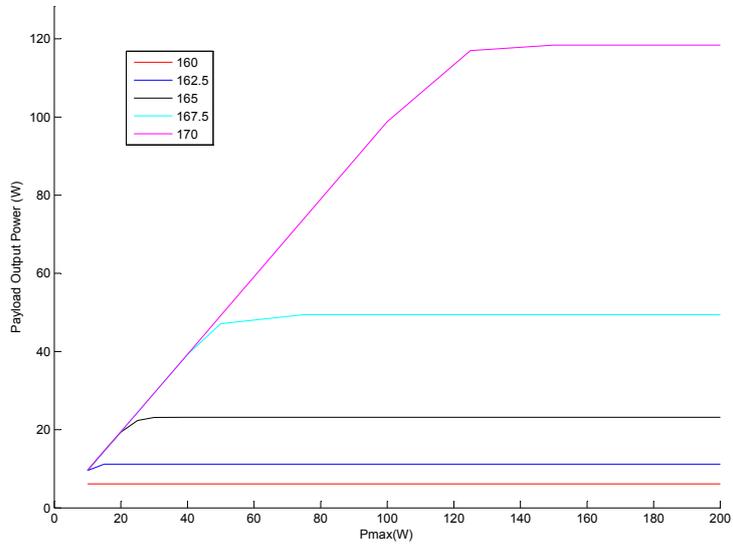
Increasing the beamwidth to 167.5° takes the payload to the edge of the footprint limited region, while increasing it to 170° takes the payload into a power limited region. In practice this increase to 170° also causes the footprints to overlap. We can thus use a beamwidth of 170° as the basis for a definition of a power limited operating point.

Reducing the beamwidth below 165° reduces coverage, but the payload continues to be footprint limited.

Inspection of figure 3.7b shows that the power output of the amplifier has a constant value of 23.2W over the range of values of P_{max} between 30W and 200W. It can be shown that a mobile at the edge of cover requires 1.12W, thus there is spare capacity for many additional mobiles. This spare capacity confirms that the payload is footprint limited.



(a) Coverage Characteristic



(b) Power Characteristic

Figure 3.7: Coverage and Power for Operating Point Selection

3.8.2 Sensitivity of the Operating Point to Altitude

It was expected that coverage would be sensitive to altitude. Lower altitude will reduce the footprint area and hence the number of mobiles inside the footprint. As altitude increases two factors become significant:

1. The area of the footprint will increase with altitude, thus more mobiles will come into view.
2. The slant range to the mobiles will increase, thus greater power will be required to support the same number of mobiles.

It is experimentally desirable to maintain a roughly constant mean coverage of 20% of the mobiles for each UAV. One way of achieving this is for the UAV payload to become power limited as it climbs above 10000 ft.

A sweep through the range of altitudes was conducted. Mean coverage was calculated at 1000ft steps across the range 1000ft to 22000ft. The results are shown in figure 3.8.

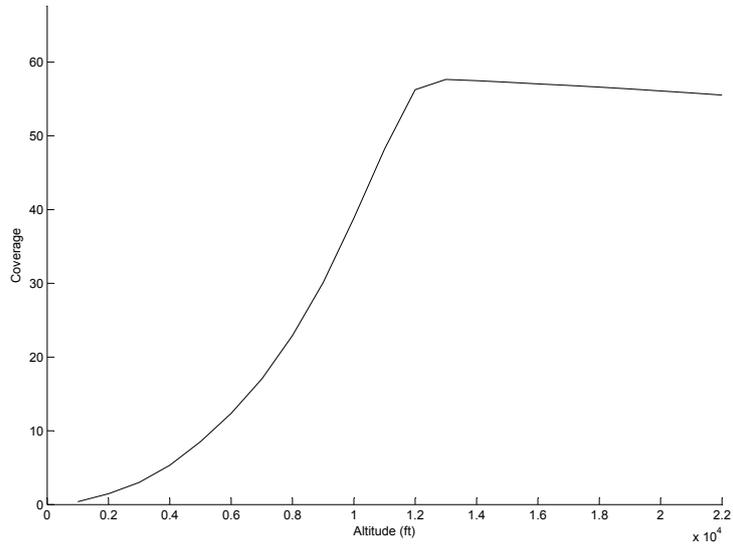
It can be seen from figure 3.8a that, as expected, the coverage increases with altitude from 1000ft to 13000ft, reflecting the increase in footprint area. Over this range of altitudes the payload is clearly footprint limited. Coverage peaks at about 60 mobiles at an altitude of 13000ft. This peak coincides with the transition from footprint limited to power limited, as can be seen in figure 3.8b.

As the altitude increases above 13000ft the power limited state of the payload causes coverage to slowly reduce with altitude.

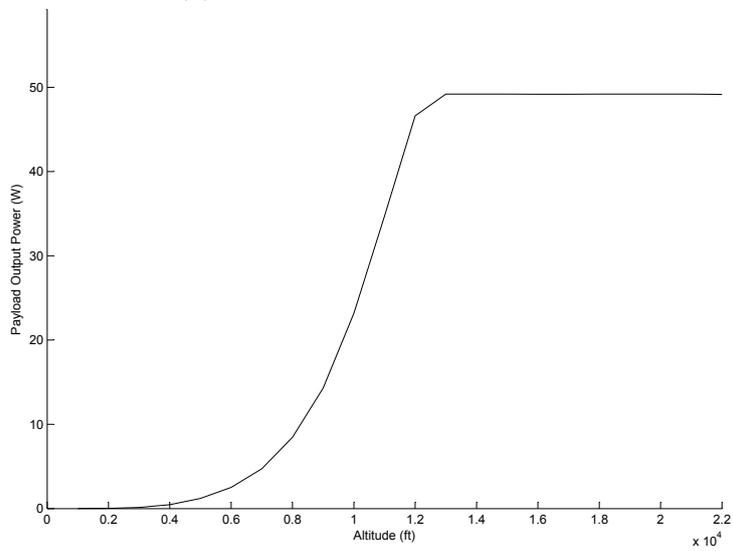
3.8.3 Statement of the Operating Point

Two operating points can be defined for the UAV payload; a footprint limited operating point and a power limited operating point. The parameter values are shown in table 3.3.

These two operating points are used throughout this research. Each chapter introduction contains a statement about which operating point pertains to that chapter.



(a) Mean Coverage v Altitude



(b) Mean Power v Altitude

Figure 3.8: Mean Coverage and Mean Power v Altitude at the Operating Point

Parameter	Footprint Limited	Power Limited	Units
OP Altitude	10000	10000	ft
OP Beamwidth	165	170	degrees
OP Power P_{max}	50	50	W

Table 3.3: Characteristics of the Operating Points

3.9 UAVs with Relocation

Before testing any algorithms it is necessary to establish whether any benefits result from relocating a communications UAV. This section considers the trivial case of a single UAV and compares the performance metrics if the UAV is circling at the centre of the scenario or is allowed to relocate according to some algorithm.

A basic relocation algorithm was proposed by the author in 2012 [20]. The algorithm, called “Select Locally Optimum Waypoint” (SLOW), allowed the UAV to select one waypoint from a defined set. It was clear from this simple algorithm that moderate improvements in coverage could be obtained by allowing a single UAV to relocate. Further improvements to this algorithm were devised, for example weighting the payoffs to favour continuity of coverage for some prioritised mobiles [18].

The SLOW algorithm did not lend itself to development into a multi-player game, but gave confidence that the principle of relocating the UAV was worth further exploration.

The OP has been selected for a single UAV circling at the centre of the scenario. Allowing the UAV to periodically relocate opens the possibility of developing algorithms that allow it to find a better location. There are four basic questions to be addressed in devising a relocation algorithm:

- Why should a UAV relocate?
- Where should the UAV be allowed to relocate?
- How should it decide when to relocate?
- How should the “better” location be identified?

The continuous motion of the mobiles results in local clusters forming and dispersing, providing opportunities for a UAV to increase its coverage by moving

towards these clusters. The UAV needs to relocate to track these clusters and improve its coverage.

A UAV with a defined location, heading and speed can relocate to a large number of points, limited only by the kinematics of the platform and the time permitted to execute the manoeuvre. Some optimisation techniques, for example evolutionary algorithms and non-linear programmes, will allow the UAV to converge on an optimal location by maximising a fitness function with a set of applied constraints. In this research it is required that any algorithm has a finite set of possible locations that define the action set for the UAV. The UAV can only relocate to one of these locations.

With a single UAV relocation can occur at any time. It could be triggered by the expiration of a timer, crossing of a threshold value, or randomly. If multiple UAVs are involved in a game that requires simultaneous decisions then all UAVs need to decide to relocate simultaneously. If a threshold crossing or random method was used then the UAV making that decision would need to communicate its need to relocate to the other UAVs. One of the drivers of this research is to maximise autonomy by minimising communications, so threshold crossing and random decisions are relatively weak solutions. This leads the research towards relocation decisions that are based on the expiration of a timer, mainly because this requires no communication between the UAVs and therefore satisfies the aspiration for high autonomy.

One disadvantage of using a timer to synchronise the decisions of all UAVs is that the maximum distance travelled by UAVs will be limited. This means that the method will tend to favour locally optimum solutions.

3.9.1 Scenario Layout

For algorithmic purposes, the airspace over the area is configured as a pattern of hexagonal cells. The UAV is constrained to circle the centre of cells at a fixed altitude of 10000 ft.

Figure 3.9 shows a section of the cellular pattern. At the start of each run of the game the UAV is circling in the central cell. This provided a benchmark set of data that could be used to quantify whether relocation of the UAV offered any advantages over a “static” UAV.

A relocating UAV has the choice of continuing to circle this location or moving

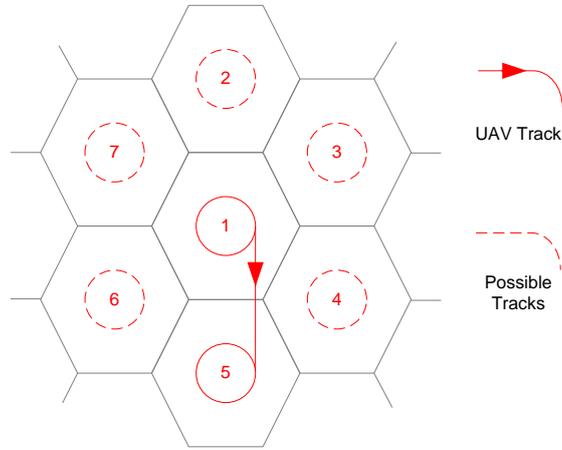


Figure 3.9: UAV Relocation Options

to circle the centre of any of the six adjacent cells.

The cellular system has a symmetry about the current location that allows all potential locations to be reached in the same time.

3.9.2 Relocation Periods

The value of the timer has to be sufficient for the UAV to complete its next manoeuvre before making its next choice. This time is calculated so that the UAV can start at any point in cell 1 and still complete one circle in cell 1, a straight flight to any adjacent cell, and a complete circle in the new cell. This ensures that the UAV can start at any point of the circle and be sure to complete at least one circle on arrival at the new location. The UAV's path can be seen in red in Figure 3.10.

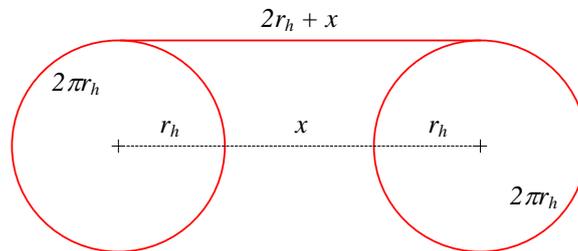


Figure 3.10: UAV Relocation Manoeuvre

If the UAV's turn radius is r_h , and the cell circumferences are x radii apart,

the total path length of two circuits and a straight flight is $(4\pi + 2 + x)r_h$.

If the UAV has a mean velocity v m/s the time taken to execute the complete manoeuvre is:

$$t_{choice} = \frac{(4\pi + 2 + x)r_h}{v} \quad (3.3)$$

We can use this to calculate the velocity at which the UAV moves across the scenario v_{sc} . The distance between cell centres divided by the time to execute a manoeuvre gives a measure of the speed at which the UAV can move across the scenario as a function of cell separation. Taking the defined UAV with velocity $v = 75$ kts and turn radius $r_h = 1500$ m, $t_{choice} = 312$ seconds. If the cells are separated by a distance r_h then the cell centres are 4500m apart so $v_{sc} = \frac{4500}{312} = 14.4$ m/s.

Cell Spacing x	Manoeuvre Time (s)	v_{sc} m/s	Mean Coverage
0	292	10.2	41.4
r_h	312	14.4	43.6
$2r_h$	332	18.2	44.7
$3r_h$	352	21.3	44.4
$4r_h$	372	24.2	43.2

Table 3.4: Cell Separation and Coverage

It can be seen from table 3.4 that coverage is relatively insensitive to cell spacing and the corresponding value of v_{sc} . The value of x used in this thesis is r_h (1500m). This has been chosen so that the velocity at which UAV transits cells is comparable with the velocity of the mobiles.

3.10 Relocation v Continuous Circling

Any potential improvement in coverage due to relocation can be quantified by comparing a UAV circling at the centre of the scenario with a UAV that is able to relocate. Using identical mobility data for the mobiles, and identical initial conditions, ensures a direct comparison.

This set of experiments took a single UAV with its payload in a footprint limited state. A default mobility data set was used for the 200 mobiles. Six runs

were made with it circling in a randomly selected cell at a fixed altitude of 10,000 feet. Ten runs were made in which it started at the centre of the scenario but was allowed to relocate to adjacent cells.

The choice of cell for relocation was based purely on seeking the location that allows support to the maximum number of mobiles. If two or more cells had the same number of mobiles, then one of them was chosen at random.

3.10.1 Coverage and Power

The mean coverage was calculated by taking the average of all the runs. Plots of coverage and power v time can be seen in figure 3.11. The data is based on the mean values of 1 second samples taken from a set of missions lasting 20000 seconds. The mean value of each plot is shown as a dotted line.

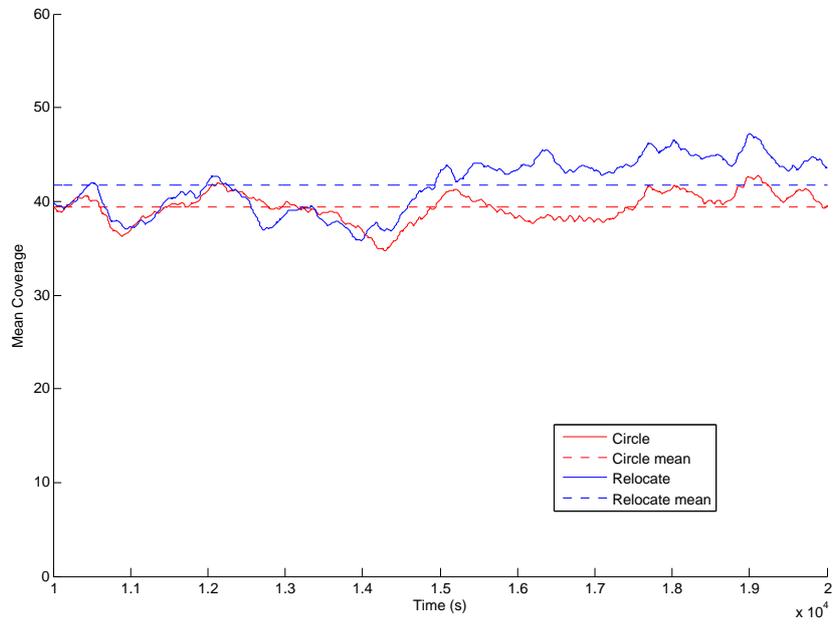
It is clear from figure 3.11a that the mean coverage for the relocating UAV (43.6 mobiles) is higher than that of the circling UAV (38.8 mobiles). The increase is mainly due to the UAV relocating itself towards clusters of mobiles and position itself so that the cluster is close to the centre of its footprint. This allows it to increase the number of mobiles that it can support. A circling UAV lacks the ability to move towards clusters of mobiles and can only support the mobiles that randomly move into its footprint.

Note that there is a period at the start of each of the runs when the coverage for circling and relocating UAVs are similar. This “settling period” is seen in many sets of results in this thesis, and is related to the initial conditions. At the start of any run it is unlikely that any UAV is located over a cluster of mobiles so the coverage is low. Eventually the UAV will find and track a cluster of mobiles, giving it a higher coverage. Figure 3.11a shows the coverages as being similar until about 1.4×10^4 seconds, after which the relocating UAV shows a coverage that is about 8 mobiles better than the circling UAV.

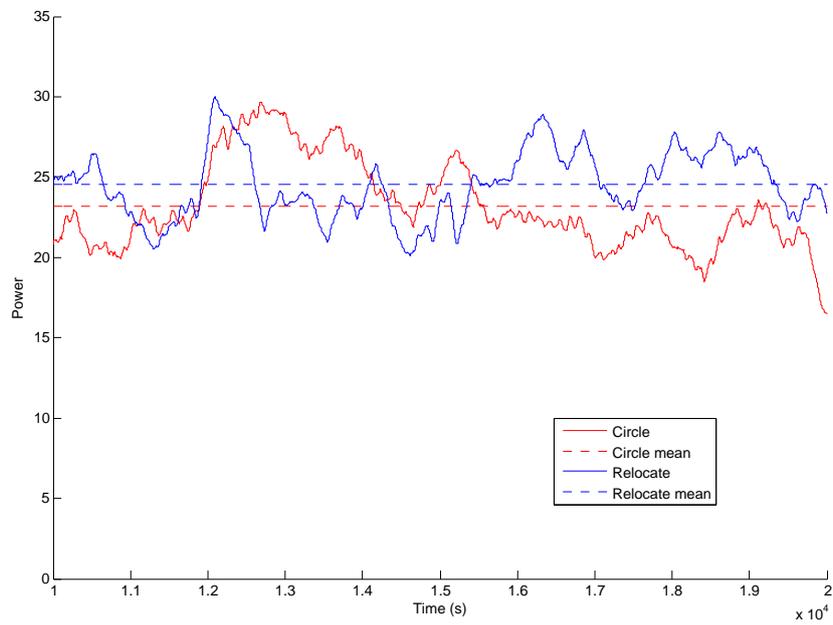
Over the duration of the mission the coverage for the relocating UAV is usually higher than that of the circling UAV. Relocating clearly improves coverage.

Encouraged by the improvement in coverage the next step was to see how efficiently the RF power was being used. Figure 3.11b shows plots of power v time.

Over the duration of this mission the mean power for the circling UAV is 23.2W, while the power for the relocating UAV is 24.5W. Recalling that the operating point



(a) Coverage v Time



(b) Power v Time

Figure 3.11: Coverage and Power

sets $P_{max} = 50\text{W}$, it can be seen that both UAVs operate below P_{max} for all of the time. This suggests that both are footprint limited, as expected from the choice of operating point.

The mean power for the relocating UAV is clearly higher than the circling UAV. The mean specific power shows how effectively the power is being used. In this example the mean specific power for the circling UAV is 0.60 W/mobile whereas the mean specific power for the relocating UAV is 0.56 W/mobile . Relocation thus makes better use of the available power.

3.10.2 Flightpath for Relocating UAVs

The statistics of the movement of the mobiles give some insight into the flight paths taken by the relocating UAVs. Figure 3.12 shows two examples of how the flight paths evolve for a single UAV and how they relate to a mobility data set.

Figures 3.12a and 3.12b show two examples of flightpaths overlaid on charts that show the probability of finding a mobile at any location over the time period. These were derived from two exemplar mobility data sets.

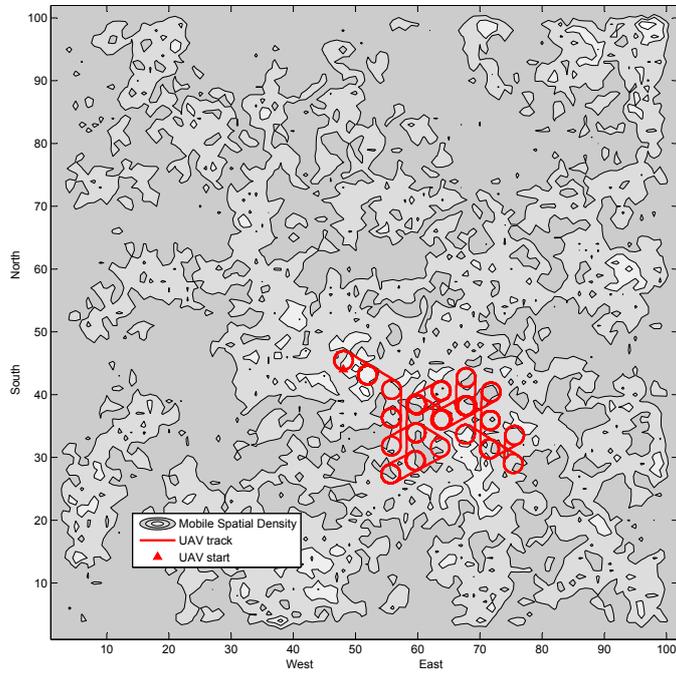
It can be seen in figure 3.12a that the probability of finding mobiles is generally higher in the south-east quadrant of the region than elsewhere. The UAV flight path in figure 3.12a shows the UAV operating exclusively in the south-east quadrant. Similar behaviours can be seen in figure 3.12b with the UAV operating mainly in the north-east quadrant as this is the region with the greatest concentration of mobiles.

It can be seen that, occasionally, the UAV moved to an adjacent cell then returned to the original cell. This “flip-flop” between adjacent cells was observed throughout this research. It indicated that marginal improvements were possible by moving between cells and was not seen as a significant issue. Indeed, it simply showed that the algorithm was working correctly.

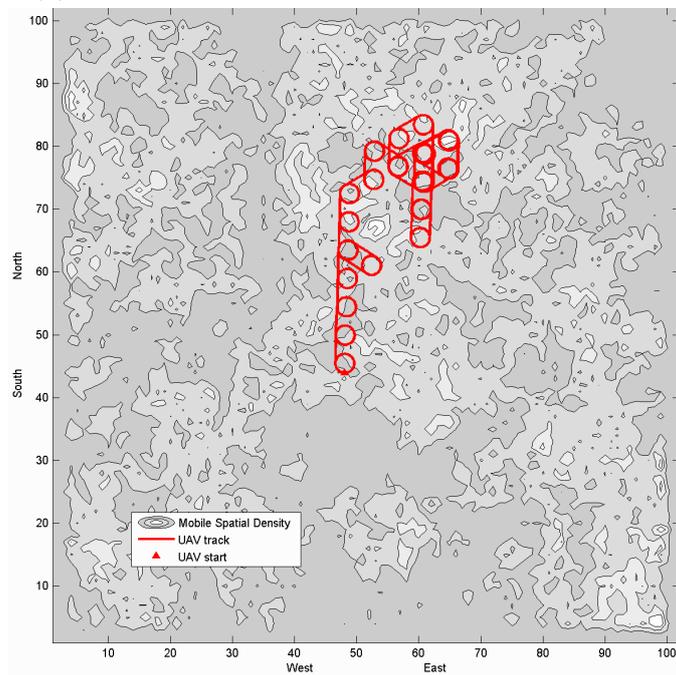
The behaviours and results shown in this section indicate that the simple algorithm to move to cells to maximise cover is effective. Subsequent chapters of this thesis will consider these behaviours with two or more UAVs in competition.

3.10.3 Comparison of Results

A comparison of the two sets of results for the default mobility set can be seen in table 3.5.



(a) Flight Path for Relocating UAV - Example 1



(b) Flight Path for Relocating UAV - Example 2

Figure 3.12: Flight Paths and Probability Distributions of Mobiles

Metric	Mean Coverage	Mean Power	Specific Power
Circle	38.8	23.2	0.60
Relocate	43.6	24.5	0.56

Table 3.5: Performance Metrics for the Default Data Set

It can be seen that the specific power improves from 0.60 W/mobile to 0.56 W/mobile if the UAV is permitted to relocate.

3.11 Chapter Discussion and Conclusions

This chapter has established some basic principles that are used throughout this thesis. It explores the relationship between beamwidth, power and coverage. It also sets an operating point for the payload that can be used in this research. It explains why UAVs will be inclined to relocate in order to track local clusters of mobiles. It establishes a movement pattern based on circling inside cells, and sets some criteria for how often mobiles should consider moving.

It is clear from the results that there are advantages to be gained by allowing the UAV to relocate periodically, especially if the UAV’s coverage is limited by the size of its RF footprint. Relocation enables a higher coverage and reduces specific power, indicating that the available RF power is used more efficiently. The advantages of relocation become less apparent if the UAV’s coverage is limited by the available power.

It can be seen that the mean coverage is not dependent on the spacing of the cell centres.

The movement of the UAVs is clearly dependent on the spatial distribution of the mobiles, however different mobility data sets will provide similar coverage and power. A single default mobility data set can be used in future experiments and results would be valid for other mobility data sets with similar statistics.

It was observed that the UAV occasionally “flip-flopped” between adjacent cells. This showed that the algorithm was working correctly by finding, and exploiting, minor improvements in coverage.

This lays the groundwork for considering what happens if two UAVs are used. This is considered in the next chapter.

This page is intentionally blank

Non-Cooperative Games with Two UAVs

4.1 Introduction

The previous chapter considered the case of a single, circling, UAV and identified the benefits of allowing a single UAV to relocate in response to movement of the mobiles.

This chapter considers systems with two UAVs and proposes an approach in which they compete with each other to maximize their individual, unique, coverage. This can be achieved when both UAVs attempt to find unique clusters of mobiles while minimizing the overlap of their footprints. Each UAV seeks to generate its own plan for maximising its coverage, mindful of what it believes the other UAV will do. This selfish and competitive approach was implemented by having both UAVs participate in a non-cooperative game.

In any run of the game the best achievable coverage for both UAVs can be obtained by finding the Nash Equilibrium (NE) and using it to define which action each UAV should choose from its action set.

The game runs as a memoryless “one shot” game as it has no knowledge of the outcome of any previous runs of the game. This constrains the UAVs to consider only the current state of the game when making decisions about the next move.

This chapter describes the structure of the two-player game that is played by the UAVs, showing how strategies are defined and how payoffs are calculated. It

evaluates five methods for finding Nash Equilibria. Four of these are normal-form games in which the game is expressed in matrix form. The fifth method presents the game in extensive form as a game tree. This approach, and some early results, were published in [21].

From the previous chapter it can be estimated that two UAVs should have a coverage of 80 to 90 mobiles if the payload operates at the footprint limited operating point. The scenario in this chapter was deliberately sized so that there are 200 mobiles available; more than could possibly be supported by two footprint limited UAVs. This ensures that the coverage of the UAVs will never be limited by the number of mobiles, only the performance of the algorithms. Throughout this chapter the payload was set at the *footprint limited* operating point.

It should be noted that there is no collision avoidance mechanism running for the two UAVs. The game should ensure separation between the two UAVs making collision avoidance unnecessary.

4.2 Game Description

Recall that the action set for the single UAV was defined using a cellular pattern. A similar cellular system was used to define the locations of the two UAVs. Figure 4.1 shows a section of the cellular pattern.

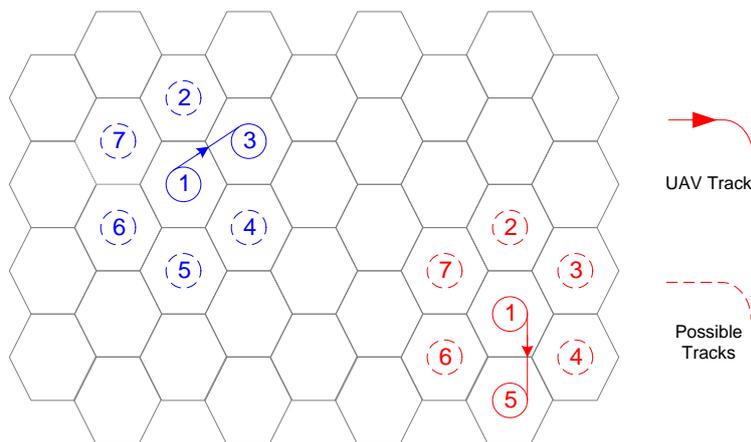


Figure 4.1: The two-UAV System

From figure 4.1 the pure strategies available to UAV j are to stay in the current cell or to move to one of the six adjacent cells. Remaining in cell 1 is strategy s_1 ,

moving to cell 2 is strategy s_2 and so on. We can thus define an action set S_j for UAV j as follows:

$$S_j = \{s_1, s_2 \dots s_7\} \quad (4.1)$$

Each pure strategy can be represented as a vector of the probability of moving to a specific cell, hence:

$$\begin{aligned} s_1 &= (1, 0, 0, 0, 0, 0, 0) \\ s_2 &= (0, 1, 0, 0, 0, 0, 0) \\ &\vdots \\ s_7 &= (0, 0, 0, 0, 0, 0, 1) \end{aligned}$$

The payoff matrix for each player is a 7×7 matrix. It has the same number of dimensions as players, the same number of rows and columns as strategies for players 1 and 2. It is convenient to represent the payoff matrices for both players as a bimatrix in which each cell of the matrix contains the payoffs for both players.

Consider the case where UAV1 selects strategy s_{row} and UAV2 selects strategy s_{col} . We can denote the payoff for UAV1 as $\pi_1(s_{row}, s_{col})$ and UAV2 as $\pi_2(s_{row}, s_{col})$. An example of the bimatrix for the game described at figure 4.1 is shown at figure 4.2.

		UAV2						
		1	2	3	4	5	6	7
UAV1	1	$\pi_1(s_1, s_1), \pi_2(s_1, s_1)$	$\pi_1(s_1, s_2), \pi_2(s_1, s_2)$	$\pi_1(s_1, s_3), \pi_2(s_1, s_3)$	$\pi_1(s_1, s_4), \pi_2(s_1, s_4)$	$\pi_1(s_1, s_5), \pi_2(s_1, s_5)$	$\pi_1(s_1, s_6), \pi_2(s_1, s_6)$	$\pi_1(s_1, s_7), \pi_2(s_1, s_7)$
	2	$\pi_1(s_2, s_1), \pi_2(s_2, s_1)$	$\pi_1(s_2, s_2), \pi_2(s_2, s_2)$	$\pi_1(s_2, s_3), \pi_2(s_2, s_3)$	$\pi_1(s_2, s_4), \pi_2(s_2, s_4)$	$\pi_1(s_2, s_5), \pi_2(s_2, s_5)$	$\pi_1(s_2, s_6), \pi_2(s_2, s_6)$	$\pi_1(s_2, s_7), \pi_2(s_2, s_7)$
	3	$\pi_1(s_3, s_1), \pi_2(s_3, s_1)$	$\pi_1(s_3, s_2), \pi_2(s_3, s_2)$	$\pi_1(s_3, s_3), \pi_2(s_3, s_3)$	$\pi_1(s_3, s_4), \pi_2(s_3, s_4)$	$\pi_1(s_3, s_5), \pi_2(s_3, s_5)$	$\pi_1(s_3, s_6), \pi_2(s_3, s_6)$	$\pi_1(s_3, s_7), \pi_2(s_3, s_7)$
	4	$\pi_1(s_4, s_1), \pi_2(s_4, s_1)$	$\pi_1(s_4, s_2), \pi_2(s_4, s_2)$	$\pi_1(s_4, s_3), \pi_2(s_4, s_3)$	$\pi_1(s_4, s_4), \pi_2(s_4, s_4)$	$\pi_1(s_4, s_5), \pi_2(s_4, s_5)$	$\pi_1(s_4, s_6), \pi_2(s_4, s_6)$	$\pi_1(s_4, s_7), \pi_2(s_4, s_7)$
	5	$\pi_1(s_5, s_1), \pi_2(s_5, s_1)$	$\pi_1(s_5, s_2), \pi_2(s_5, s_2)$	$\pi_1(s_5, s_3), \pi_2(s_5, s_3)$	$\pi_1(s_5, s_4), \pi_2(s_5, s_4)$	$\pi_1(s_5, s_5), \pi_2(s_5, s_5)$	$\pi_1(s_5, s_6), \pi_2(s_5, s_6)$	$\pi_1(s_5, s_7), \pi_2(s_5, s_7)$
	6	$\pi_1(s_6, s_1), \pi_2(s_6, s_1)$	$\pi_1(s_6, s_2), \pi_2(s_6, s_2)$	$\pi_1(s_6, s_3), \pi_2(s_6, s_3)$	$\pi_1(s_6, s_4), \pi_2(s_6, s_4)$	$\pi_1(s_6, s_5), \pi_2(s_6, s_5)$	$\pi_1(s_6, s_6), \pi_2(s_6, s_6)$	$\pi_1(s_6, s_7), \pi_2(s_6, s_7)$
	7	$\pi_1(s_7, s_1), \pi_2(s_7, s_1)$	$\pi_1(s_7, s_2), \pi_2(s_7, s_2)$	$\pi_1(s_7, s_3), \pi_2(s_7, s_3)$	$\pi_1(s_7, s_4), \pi_2(s_7, s_4)$	$\pi_1(s_7, s_5), \pi_2(s_7, s_5)$	$\pi_1(s_7, s_6), \pi_2(s_7, s_6)$	$\pi_1(s_7, s_7), \pi_2(s_7, s_7)$

Figure 4.2: Example of a Payoff Matrix with Two Players

Each cell of the matrix represents a selected location for each of the the two UAVs. From these locations, and with knowledge of the estimated locations of the mobiles, it is possible to create an allocation matrix $X(s_{row}, s_{col})$:

$$X(s_{row}, s_{col}) = \begin{pmatrix} x_{11} & x_{12} \\ x_{21} & x_{22} \\ \vdots & \vdots \\ x_{m1} & x_{m2} \end{pmatrix} \quad (4.2)$$

Recall that the payoff for each UAV is the sum of the corresponding column of X , hence the payoffs for the two UAVs for strategies (s_{row}, s_{col}) are:

$$\pi_1 = \sum_{i=1}^m x_{i1} \quad (4.3)$$

$$\pi_2 = \sum_{i=1}^m x_{i2} \quad (4.4)$$

The payoff matrix generated using this method can be solved to determine the existence of Nash Equilibria.

4.3 Benchmarks

4.3.1 Two Circling UAVs

A pair of circling UAVs was used as a benchmark for the two-UAV experiments. A UAV circling a fixed point at a constant altitude requires less airspace than one that is allowed to relocate. The predictability of circling makes it a useful benchmark from an airspace management viewpoint. Allowing the UAVs to relocate will increase the total airspace that will be occupied during a mission. The benefits of improved coverage that result from allowing the UAVs to relocate can be compared with the increased airspace required during a mission.

A sweep across the range of powers provides a useful check for the edges of the power limited and footprint limited regions of the joint system. It was clear from figure 4.3 that the operating point of 50W was in the *footprint limited* region.

A six-hour mission was simulated with two UAVs, circling at an altitude of 10000ft with non overlapping footprints. The mission was run six times, each time with a different mobility data set.

Figures 4.4a and 4.4b show mean coverage vs time and power vs time for two circling UAVs. The graphs are generated by averaging the coverage and power for

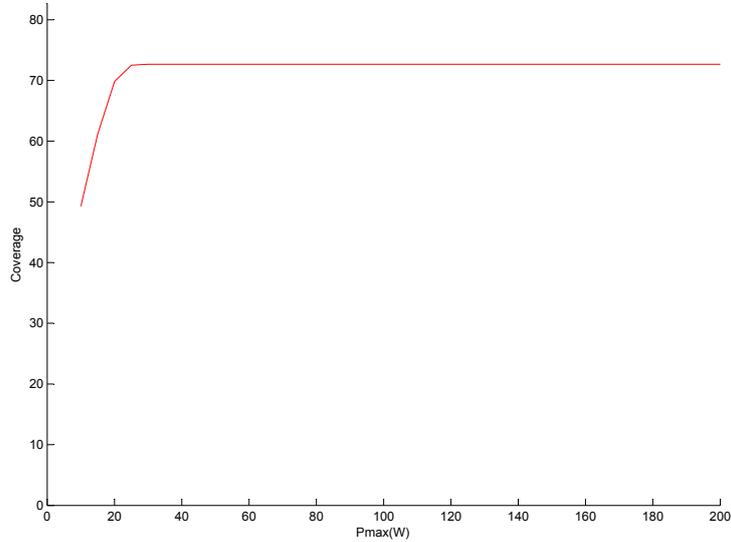


Figure 4.3: Coverage Characteristic for Two Circling UAVs

each of the six missions. Error bars set to 1 standard deviation have been added every 2000 seconds to show the variability of the data from the six missions.

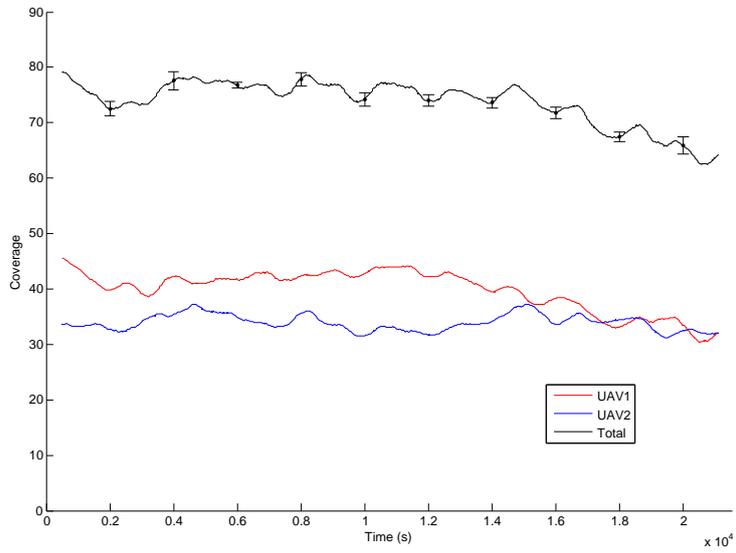
It can be seen from figure 4.4a that the mean coverage C_{mean} coverage was 73.7 mobiles, and that its it evenly divided between the two UAVs. The power vs time plot in figure 4.4b shows that both UAVs radiate similar power levels from their payloads, and that the total power for each payload was less than the full 50W.

The mean coverage C_{mean} of 73.7 has a standard deviation of 4.8. Recall that the single circling UAV had a mean coverage of 38.8 so the coverage of two UAVs was approximately double the coverage of a single UAV. The mean power P_{mean} of 41.4W has a standard deviation of 2.6W. The mean power required by the single UAV was 20.7W.

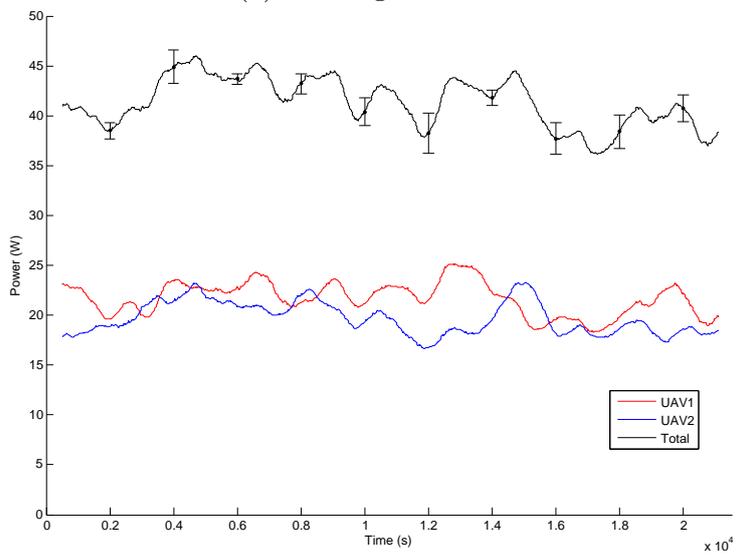
These results are consistent with the expectation that two circling UAVs with non-overlapping footprints should support twice as many mobiles, and that each should require approximately the same power to support its share. Figure 4.4 confirms that the power and coverage are similar for both UAVs.

4.3.2 Select Random Strategies

Selection of random strategies by both UAVs should result in unpredictable outcome. This provides a benchmark as any effective algorithm should always provide



(a) Coverage v Time



(b) Power v Time

Figure 4.4: Coverage and Power v Time for Two Circling UAVs

a better solution than random choice.

A six-hour mission was simulated with two UAVs making random strategy choices each time the game was run. The UAVs flew at an altitude of 10000ft, initially with with non overlapping footprints. The mission was run six times, each time with a different mobility data set.

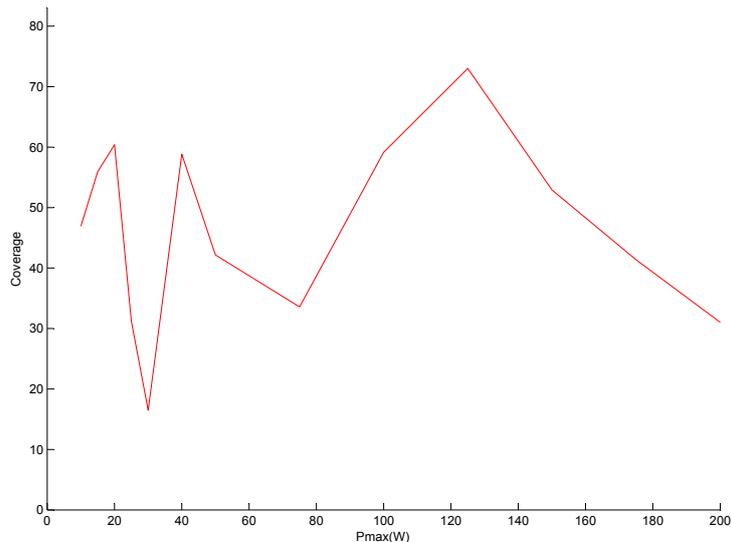


Figure 4.5: Coverage Characteristic for Random Strategies

Figure 4.5 shows the power characteristic when random strategies were selected. It was impossible to discern the power limited and footprint limited regions of this graph by inspection.

At low P_{max} the curve shows low coverage, as expected, but increasing power sometimes results in decreasing coverage. The payload limited section is poorly defined as adjacent points have significantly different coverage.

The random movement of the UAVs during each run creates difficulty in predicting trends in coverage. Each UAV, instead of moving towards clusters of mobiles, would roll a dice to decide its strategy. Coverage could increase or decrease randomly each time a strategy was selected. Furthermore, it was possible for the UAVs' footprints to overlap, reducing the coverage of both, or to leave the scenario area resulting in a coverage that trends towards zero.

Re-running the same simulation should generate different results. Figure 4.6 presents four plots of mean coverage vs time for the random strategies method. These runs were conducted with identical configurations and initial conditions.

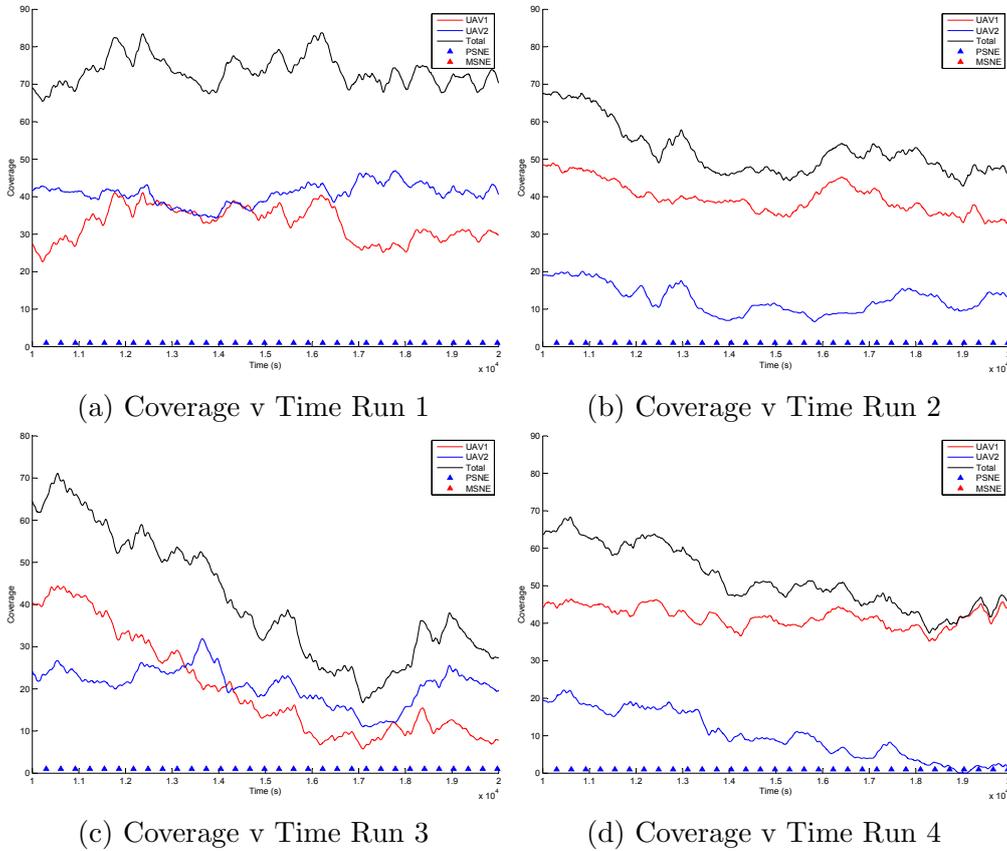


Figure 4.6: Examples of Coverage with Random Strategies

Visual inspection of the graphs indicates that the results are very different. ANOVA of the four sets of results confirms that the difference between the sets of results was significant.

It can be concluded that selecting a strategy at random was a poor method of finding solutions and forms an unsatisfactory benchmark.

4.4 Comparison of Methods for Solving Normal-Form Games

There are many published algorithms and methods for solving normal-form games. The research in this thesis uses exhaustive search, fast search, the Lemke-Howson algorithm and Chatterjee's method. These are described in more detail in appendix

A.

In this application, coordinating UAVs, an ideal method for solving the game would find a solution in a short and predictable time. The time should be very much less than the time required to execute the next manoeuvre, 312s, as this would allow relatively recent data to be used to calculate the payoffs.

The ideal method for solving the game should also give solutions that offer high mean coverage C_{mean} and low mean power P_{mean} when compared to other methods. This section compares the performance of these four methods for solving games.

The solution to the game must comprise one or more Nash Equilibria (NE). Nash's theorem predicts that every finite game has at least one equilibrium, giving rise to the possibility of multiple NE. This game was finite as it has a finite number of players and strategies, thus at least one NE will be found every time the game was run. Each method for finding a solution has its own way of addressing how to select one NE if multiple NE are found.

The NE could be Pure Strategy Nash Equilibria (PSNE) or Mixed Strategy Nash Equilibria (MSNE). Essentially a MSNE is a probability distribution across all the possible pure strategies for one or more players. A PSNE is simply a special case of the MSNE when one strategy has a probability of 1 and all others are zero. The difference between PSNE and MSNE is explained in greater depth in appendix A.

Figure 4.7 shows the power characteristic for the four methods for solving normal form games that are considered in this thesis.

It is clear from figure 4.7 that the operating point was in the *footprint limited* region. The curves for all methods have the same general shape as previous characteristics.

Exhaustive search shows consistently good results with a flat power characteristic for values of P_{max} from 30W to 200W. Exhaustive search finds all possible pure strategies and randomly selects one of them, thus there is a probability of finding one with a higher payoff. In practice it has been found that several of the pure strategies will have similarly high payoffs so randomising tends to favour finding solutions with higher payoffs. Fast search terminates when it finds the first NE so there is a possibility that it will sometimes find one with a lower coverage.

Most practical methods for solving games are capable of finding one PSNE quickly, or finding one MSNE by converging on a solution. Methods that converge

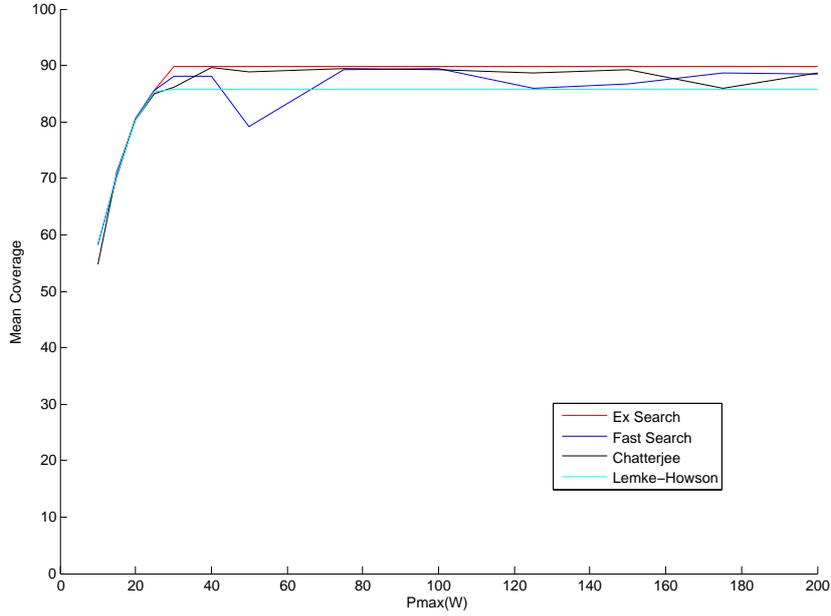


Figure 4.7: Coverage Characteristic for All Solution Methods

on a solution take an indeterminate time which may not be a good basis for a real-time system that is guiding UAVs. Several hundred simulations were run to compare the execution times of the methods with two UAVs. The results for these simulations can be seen in table 4.1.

Algorithm	Parameter		
	Mean Time (s)	Max Time (s)	Number of NE
Exhaustive Search	3.21×10^{-2}	3.61×10^{-2}	All PSNE
Fast Search	1.67×10^{-4}	6.35×10^{-4}	One PSNE
Lemke Howson	2.28×10^{-4}	1.38×10^{-2}	One MSNE
Chatterjee	0.11	0.18	One MSNE

Table 4.1: Mean Execution Times for Solution Methods (s)

Recall that the time required for a UAV to relocate from one cell to another is 312 seconds. Consequently the maximum time taken to reach a decision must be significantly less than 312 seconds. All methods generated solutions in sufficient

time to meet that requirement.

The mean times for fast search and Lemke-Howson indicate that these should be very effective methods for solving normal form games. Lemke-Howson was not consistently fast as the longest time was two orders of magnitude longer than the mean time.

Exhaustive search was relatively slow, but the variation in time was quite small. Chatterjees method converges on the solution in an indeterminate time and was the slowest method.

In terms of relative speed, fast search was clearly the best method if a PSNE existed. If, however, the solution was not a PSNE but an MSNE then the Lemke-Howson method would always find a solution.

4.4.1 Exhaustive Search for a Pure Strategy

This method assumes the existence of one or more PSNE in every payoff matrix. This is not always the case, so the method needs to consider three possible outcomes:

1. If no PSNE exist the UAVs will remain at the current location (the “do nothing” option).
2. If one PSNE exists then select those strategies
3. If multiple PSNE exist then select one by randomisation.

The first state, no PSNE exists, can occur in games where only a mixed strategy exists. If no PSNE existed then exhaustive search would find an indeterminate strategy. Assigning a “do nothing” decision to this state was a pragmatic solution.

A six-hour mission was simulated with two UAVs operating at an altitude of 10000ft. At the start of the mission the UAVs have non overlapping footprints. The mission was run ten times.

Figures 4.8a and 4.8b show coverage vs time and power vs time. The graphs are generated from the average coverage and power for each of the ten missions. Error bars set to 1 standard deviation have been added every 2000 seconds to show the variability of the data from the ten missions.

The coverage plot has been modified and shows a line of tick marks along the x -axis. The ticks correspond to the times when the one-shot game was run,

and the colour of the tick shows whether the solution was a pure strategy or a mixed strategy. In this game all the solutions were pure strategies; the “do nothing” option can be classed as a pure strategy as it was a clear decision to select strategies (s_1, s_1) .

Figure 4.8a, and similar graphs in subsequent sections, are based on a set of 200 mobiles. It is clear that coverage and power are not constant over time, and the the standard deviation for coverage also changes with time. The variation in coverage can be illustrated by the simple example shown in figure 4.9.

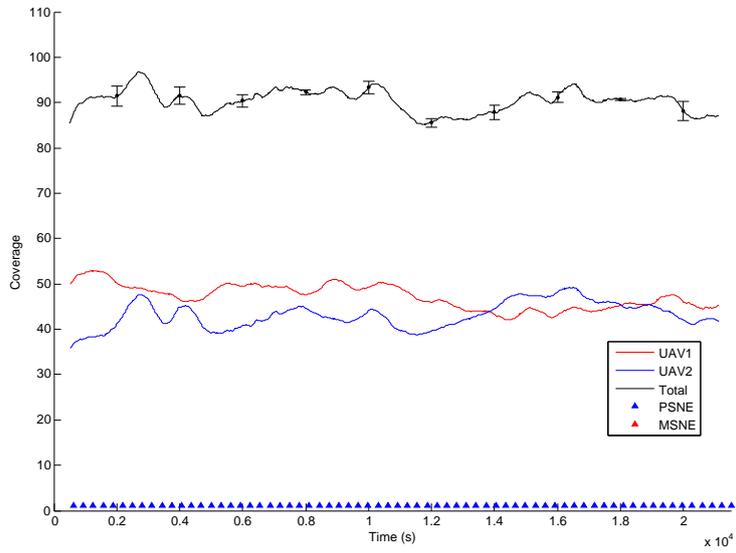
It can be seen in figure 4.9a that the UAV has acquired a small cluster of mobiles. While this cluster exists the UAV has 10 mobiles in its footprint. The cluster is near the centre of the footprint so the mobiles in the cluster require less power than the mobiles at the edge of the footprint.

Movement of these mobiles causes the cluster to disperse leaving the UAV with 7 mobiles in its footprint, as shown in figure 4.9b. Meanwhile the movement of mobiles has formed a new cluster at the edge of the footprint. This causes a reduction in coverage that is not always matched by a reduction in power. It can be seen that some of the the mobiles are located close to the edge of the footprint so they require higher power.

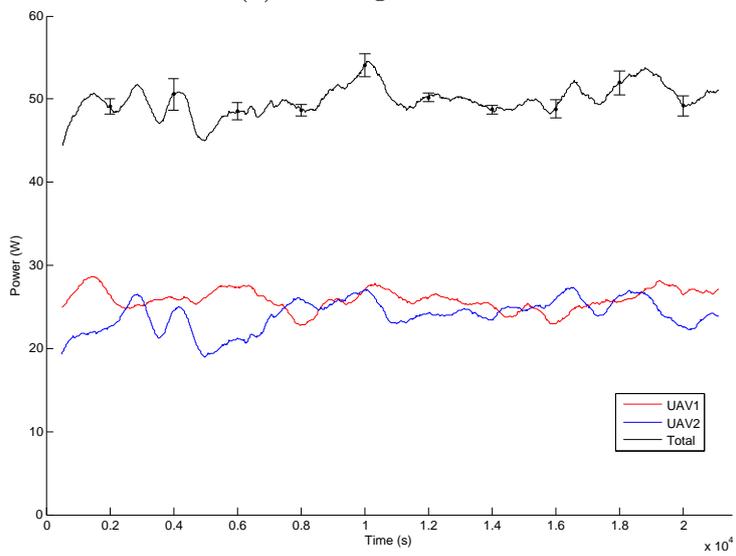
The location broadcast system allows the UAV to become aware of the existence of this cluster so the UAV starts to move towards it. In figure 4.9c we see that the mobile has acquired this new cluster, increasing its coverage to 11 mobiles. By relocating itself, the UAV can put the cluster close to the centre of its footprint. This reduces the power used to support the cluster, and allows the UAV to support a few more mobiles at the edge of its footprint.

Returning to the graphs at figure 4.8, this type of activity is happening continuously to all the UAVs. The game allows the two UAVs to identify clusters of mobiles and compete over them. A reduction in coverage does not necessarily correlate with a reduction in power; the distribution of the mobiles in the footprint determines the power that is necessary to support them.

We can see in figures 4.8a and 4.8b that coverage and power change over time. There is a tendency for power to track coverage, but there are several occasions, for example time 1.8×10^4 , when the total power increases for no real change in coverage. Close inspection shows that the power for UAV2 increases at this time while its coverage is slowly decreasing. This is indicative of a cluster dispersing and the mobiles moving towards the edge of UAV2’s footprint.



(a) Coverage v Time



(b) Power v Time

Figure 4.8: Coverage and Power v Time for Exhaustive Search

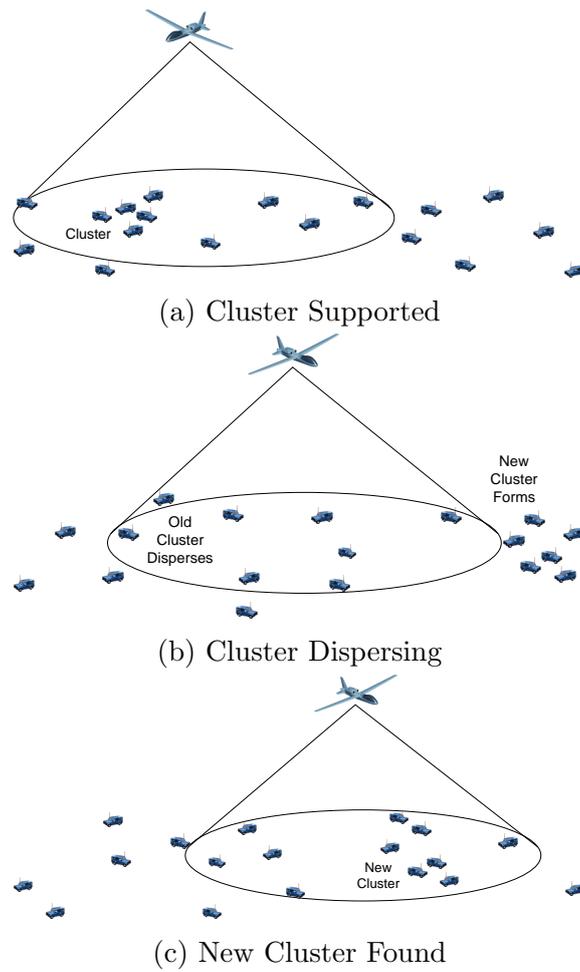
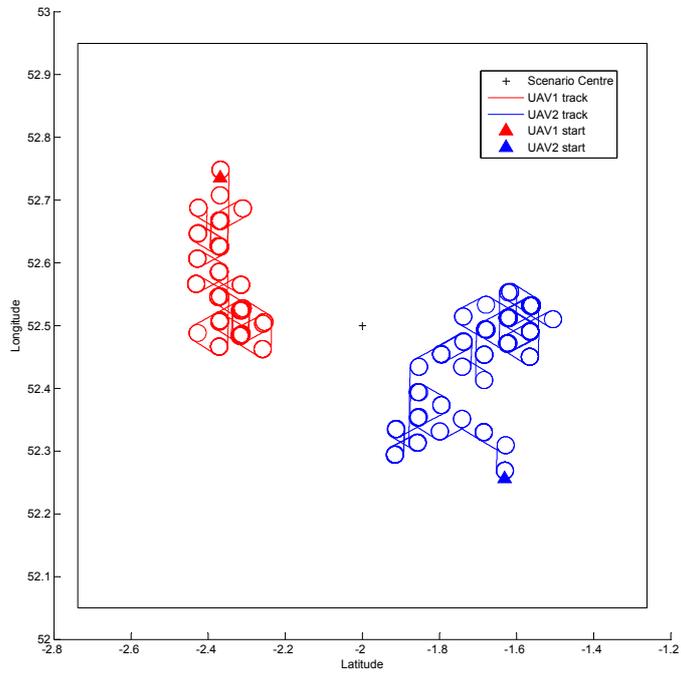
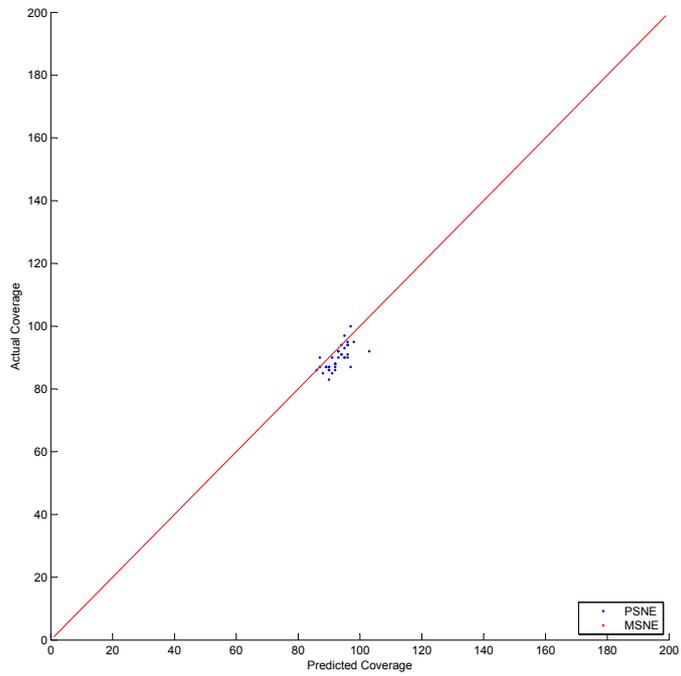


Figure 4.9: Behaviour of Mobile Clusters



(a) Flightpath



(b) Predicted v Actual Coverage

Figure 4.10: Flightpaths and Estimated Payoffs for Exhaustive Search

The mean coverage, C_{mean} , was 90.1 mobiles with a standard deviation of 3.8. This represented a 22.3% increase in mean coverage over two circling UAVs. A t-test for the exhaustive search and circling results confirms that this was a significant difference.

The mean power, P_{mean} , required to support this improvement in coverage was 49.8W with a standard deviation of 2.7W. This was a mean power of 24.9W for each UAV, 4.2W more than the power if the same UAVs were circling.

Figure 4.10a shows the flightpaths taken by the two UAVs on one exemplar run. UAV1 started in the north-west quadrant and UAV2 started in the south-east. As expected, the two UAVs moved towards each other and contested the clusters in the central region of the game. During this competition they maintained sufficient separation that their footprints did not significantly overlap.

Recall that the payoff matrix was calculated from estimated locations of the mobiles 312 seconds into the future. The effectiveness of the game depends on the accuracy of the estimated payoffs as much as the method of solving the payoff matrix. Figure 4.10b plots the estimated payoffs against the actual payoffs achieved.

The points represent the estimated coverage 312 seconds into the future and the actual coverage at that time. If estimation was perfect then the points should lie on the red line. A method that overestimates the coverage will generate more points below the line than above. It can be seen from figure 4.10b that exhaustive search was prone to overestimate coverage. The points are coloured to show whether the estimate was obtained from a pure strategy or a mixed strategy.

Overall, this method of seeking pure strategies by exhaustive search significantly improved coverage over two circling UAVs. The method was quite good at estimating coverage, and maintained a good separation between the UAVs.

4.4.2 Fast Search for Best Response

A pure strategy Nash Equilibrium is, by definition, the mutual best response by both players. It can thus be seen that finding any mutual best response thus finds one PSNE.

Fast Search was developed by the author as a possible improvement to exhaustive search. It attempts to find a mutual best response without exhaustively searching the payoff matrix. In the limiting case it will take the same time to

find a solution as exhaustive search. In general it should be faster than exhaustive search. If multiple pure strategy Nash Equilibria exist then it will only find one; which one it finds depends on the initial conditions of the search.

This method has two possible outcomes:

1. If no PSNE exist the UAVs will remain at the current location (the “do nothing” option).
2. If one PSNE exists then select that strategy

A six-hour mission was simulated with two UAVs operating at an altitude of 10000ft. At the start of the mission the UAVs have non overlapping footprints. The mission was run ten times.

Figure 4.11 shows that coverage and power are divided evenly between the two UAVs. The graphs are generated from the average coverage and power for each of the ten missions. Error bars set to 1 standard deviation have been added every 2000 seconds to show the variability of the data from the ten missions. The mean coverage, C_{mean} , was 89.7 mobiles with a standard deviation of 4.6. This was very similar to the exhaustive search method and shows a 21.7% improvement in coverage over the two circling UAVs. As before, the t-test shows that this improvement was significant.

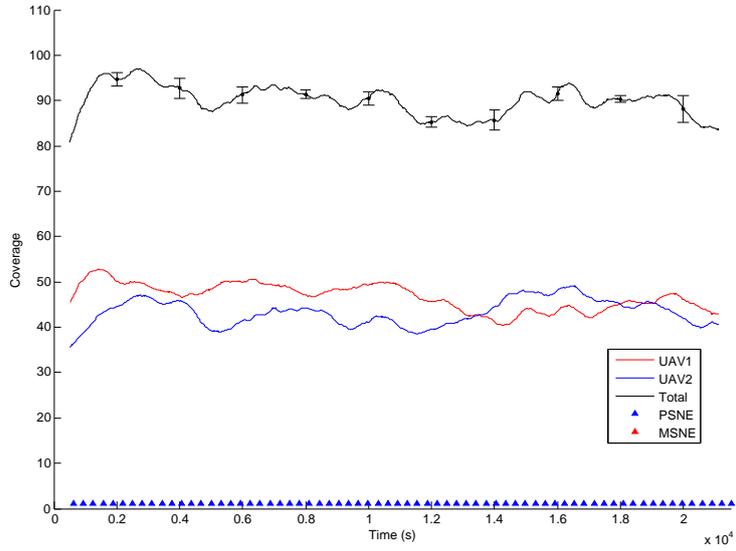
The mean power P_{mean} required to support this improved coverage was 49.0W with a standard deviation of 2.8W. The mean power for each UAV was thus 24.5W which was very similar to the exhaustive search method.

The flightpaths, shown in figure 4.12a, are also similar to those obtained by the exhaustive search method.

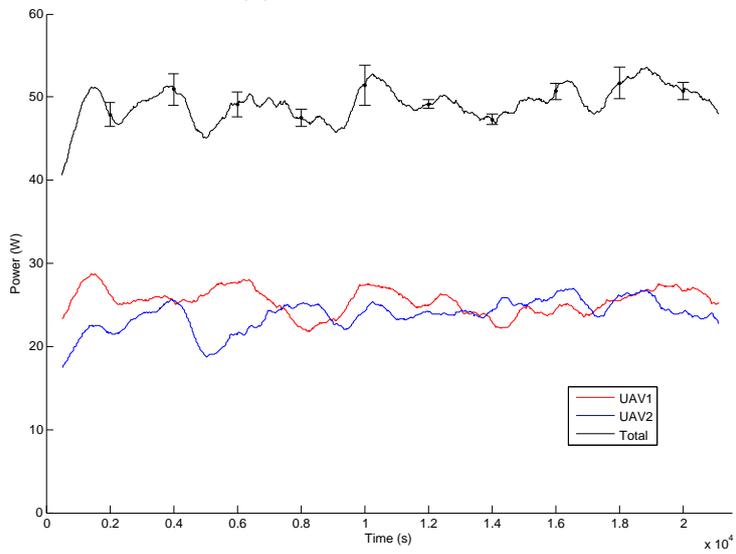
It is clear from figure 4.12b that fast search slightly overestimates coverage. These estimates are similar to exhaustive search. This is unsurprising as fast search finds one possible solution from the set of those found by exhaustive search.

These results should be unsurprising as fast search is, essentially, a subset of the exhaustive search method. The marginally poorer coverage statistics suggest that there could be occasions when there are two or more pure strategy Nash Equilibria. If the one found by fast search has a lower payoff than the one found through randomisation after exhaustive search, then a slightly lower payoff would be expected.

This method seems to perform very acceptably. It can be seen from table 4.1 that fast search is faster than exhaustive search. The improved speed of execution

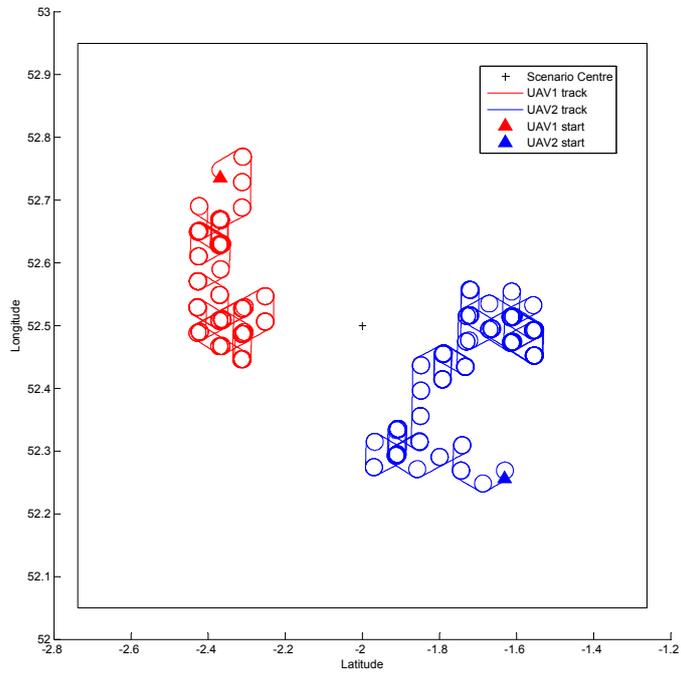


(a) Coverage v Time

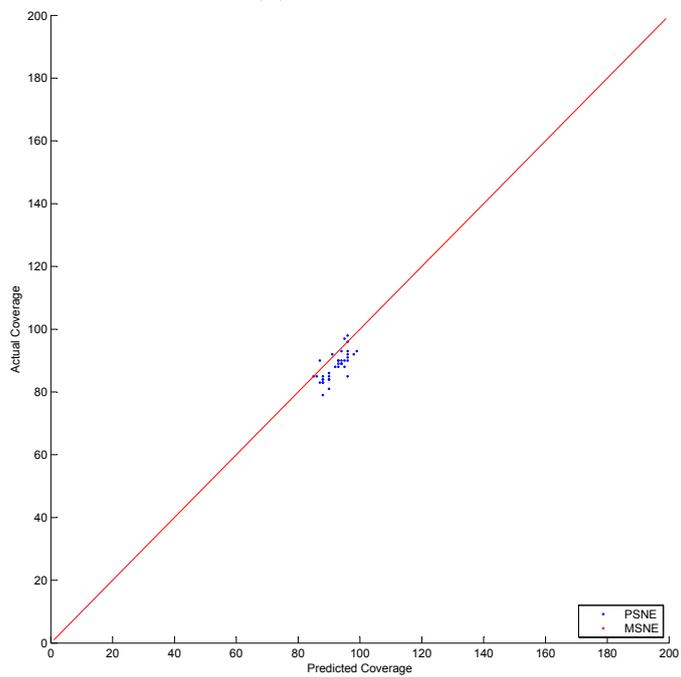


(b) Power v Time

Figure 4.11: Coverage and Power v Time for Fast Search



(a) Flightpath



(b) Predicted v Actual Coverage

Figure 4.12: Flightpaths and Estimated Payoffs for Fast Search

is obtained at the cost of the quality of solutions that are found.

4.4.3 Find One Mixed Strategy using Lemke-Howson

The previous two methods only sought PSNE. If no PSNE was found then the solution was to “do nothing”. This method offers a first attempt to find a solution to every instance of the “one-shot” game.

Nash’s theorem tells us that there is always at least one mixed strategy. This method looks for any mixed strategy Nash Equilibrium using the Lemke-Howson algorithm.

A six-hour mission was simulated with two UAVs operating at an altitude of 10000ft. At the start of the mission the UAVs have non overlapping footprints. The mission was run ten times.

Figure 4.13 shows plots of mean coverage vs time and power vs time. The graphs are generated from the average coverage and power for each of the ten missions. Error bars set to 1 standard deviation have been added every 2000 seconds to show the variability of the data from the ten missions.

The mean coverage, C_{mean} , was 87.4 mobiles with a standard deviation of 5.4. This represents a 18.6% improvement in mean coverage over two circling UAVs. This was lower than the mean coverage obtained from the exhaustive search and fast search methods.

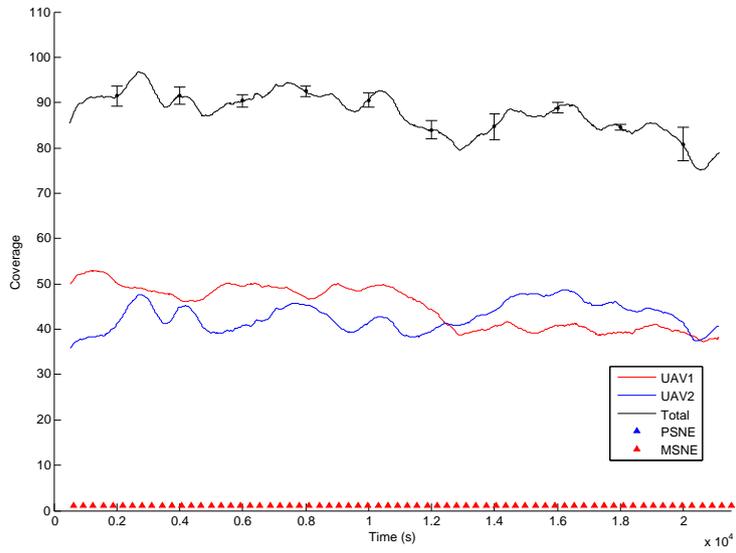
The mean power, P_{mean} , required to support this improvement in coverage was 48.0W with a standard deviation of 3.0W. The mean power for each UAV was thus 24.0W which was similar to other methods.

Figures 4.13a and 4.13b show that coverage and power are balanced between the two UAVs.

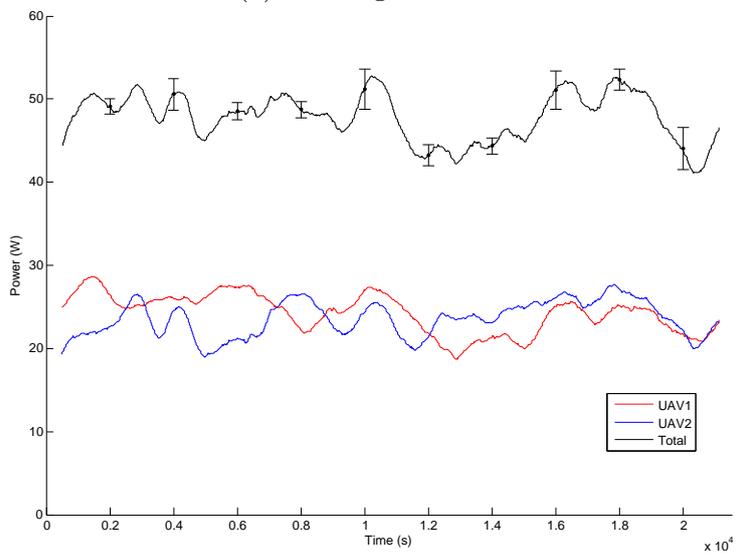
The flightpaths for the two UAVs are shown in figure 4.14a. The path for UAV2 was similar to those for the exhaustive search and fast search. UAV1 takes a dissimilar flightpath and spends the second half of its mission in the north of the scenario area.

Figure 4.14b shows that this method consistently overestimates coverage.

There was no “do nothing” option so a solution was always found. Each UAV had the option to relocate when, in previous methods, it would have stayed at its current location. This resulted in more frequent movement of the UAVs, sometimes to locations with a lower payoff. It was therefore unsurprising that a

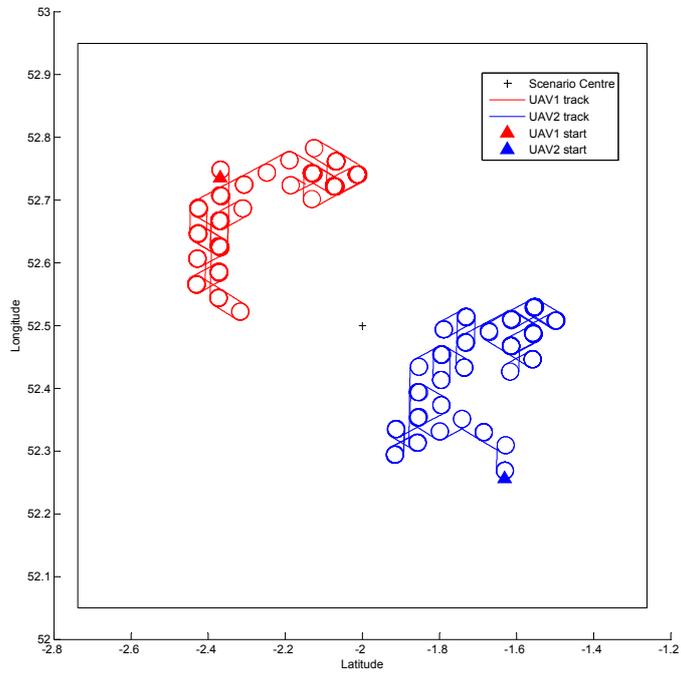


(a) Coverage v Time

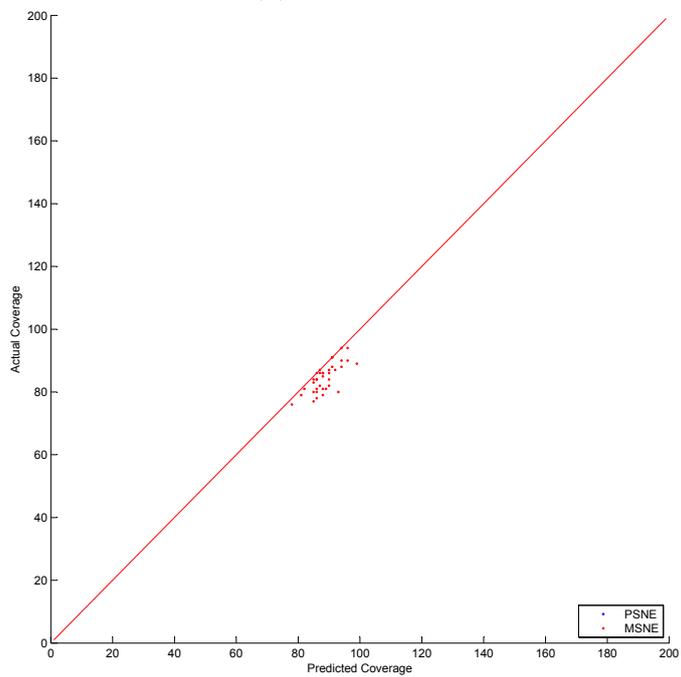


(b) Power v Time

Figure 4.13: Coverage and Power v Time for Lemke-Howson



(a) Flightpath



(b) Predicted v Actual Coverage

Figure 4.14: Flightpaths and Estimated Payoffs for Lemke-Howson

method designed to find a mixed strategy tended to produce slightly lower payoffs, move more frequently, and overestimate the payoffs.

4.4.4 Find One Mixed Strategy using Chatterjee’s Method

Chatterjee’s method finds one MSNE out of all those that could exist. It starts by assuming a random solution, then progressively refines that solution until the error between successive iterations is less than a given threshold. Which MSNE it finds, and the time taken to find it, depends on how close the assumed solution is to any MSNE.

A six-hour mission was simulated with two UAVs operating at an altitude of 10000ft. At the start of the mission the UAVs have non overlapping footprints. The mission was run ten times.

Figures 4.15a and 4.15b show a balanced coverage and power between the two UAVs. The graphs are generated from the average coverage and power for each of the ten missions. Error bars set to 1 standard deviation have been added every 2000 seconds to show the variability of the data from the ten missions. The mean coverage, C_{mean} , was 89.1 mobiles with a standard deviation of 4.6. The mean power, P_{mean} , required to support this improvement in coverage was 48.2W with a standard deviation of 3.0W. The mean power for each UAV was thus 24.1W which was similar to other methods.

This was better than the results for Lemke-Howson, but came at the cost of an execution time that was three orders of magnitude slower.

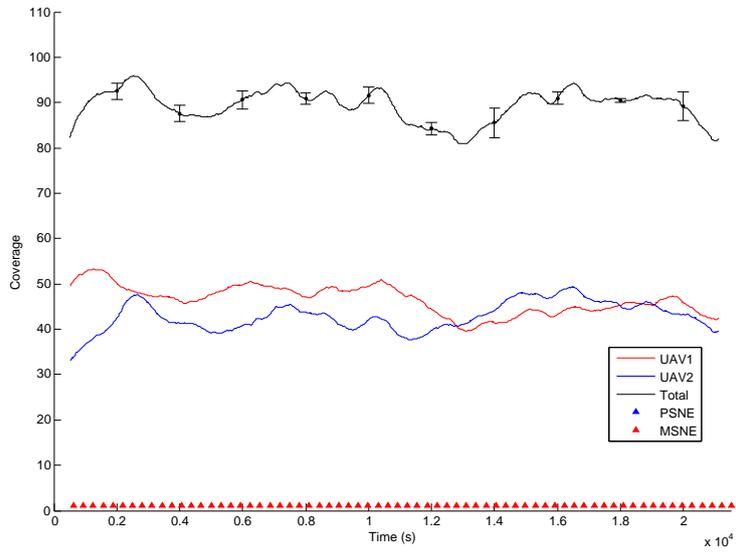
The flightpath for the two UAVs are shown in figure 4.16a and tends towards greater mobility than Lemke-Howson. This can be ascribed to the randomness of the start points.

Like Lemke-Howson, Chatterjees method generally overestimates the coverage as can be seen in figure 4.16b.

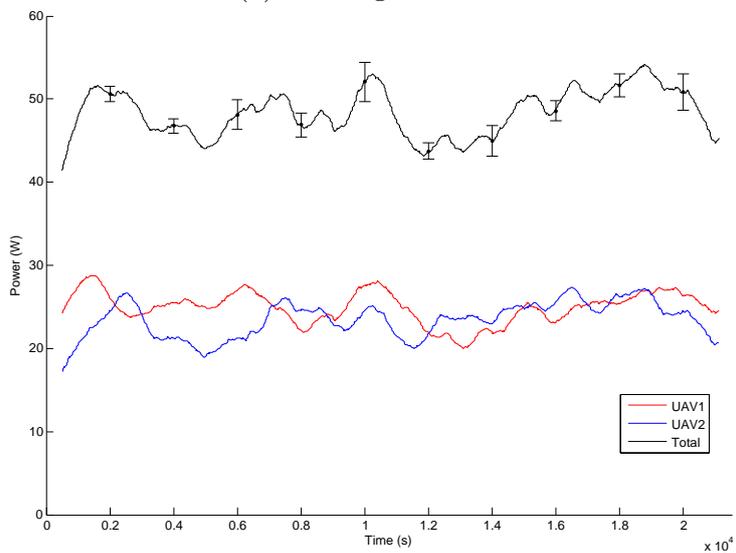
4.5 Extensive Form Game

The games considered in sections 4.4.1 to 4.4.4 have all be normal form games. In the normal form the players make their decisions simultaneously and independently, and the payoffs are expressed as a bimatrix.

In extensive form games the players make their decisions sequentially. The

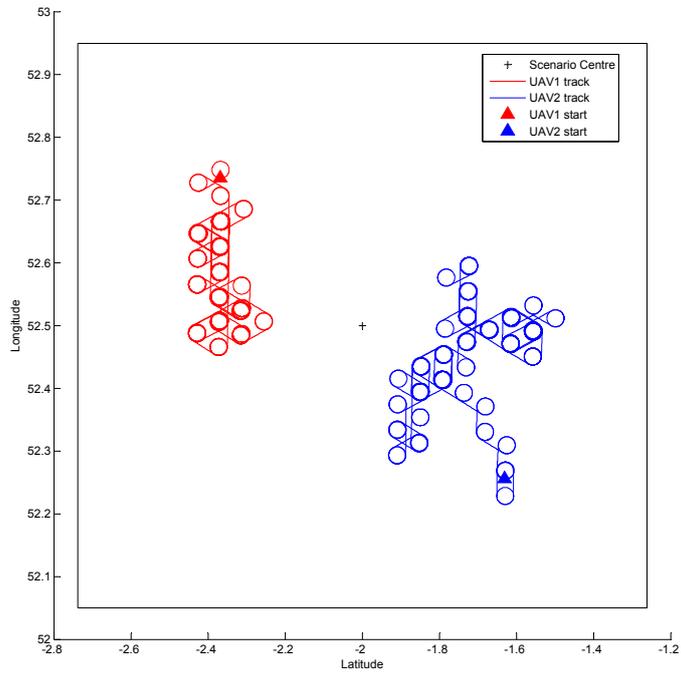


(a) Coverage v Time

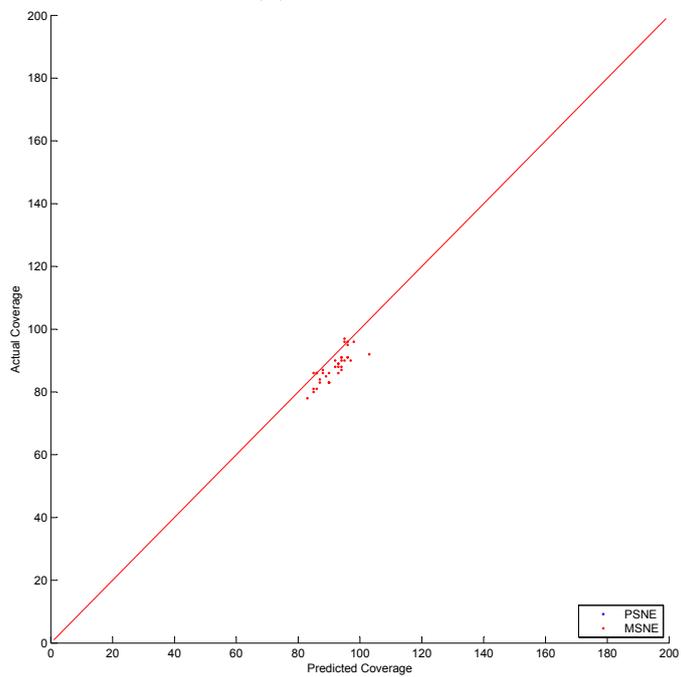


(b) Power v Time

Figure 4.15: Coverage and Power v Time for Chatterjee's Method



(a) Flightpath



(b) Predicted v Actual Coverage

Figure 4.16: Flightpaths and Estimated Payoffs for Chatterjee's Method

payoffs are expressed as a game tree with the leaves representing the payoffs of all players if a particular set of actions have been selected. Extensive form games are generally solved by backward induction, with the root of the tree representing the PSNE. The payoff for an extensive form game is calculated in a similar manner to a normal form game.

In this game UAV1's payoffs are calculated at all 7 of its possible locations. UAV2 then calculates its payoffs at all 7 of its possible locations, taking account of all the possible moves UAV1 could have made.

Figure 4.17 shows that this method of solving game generates a flat characteristic across a wide range of P_{max} confirming that there is no randomisation between PSNE. The payload was operating in the footprint limited region at $P_{max} \geq 30W$.

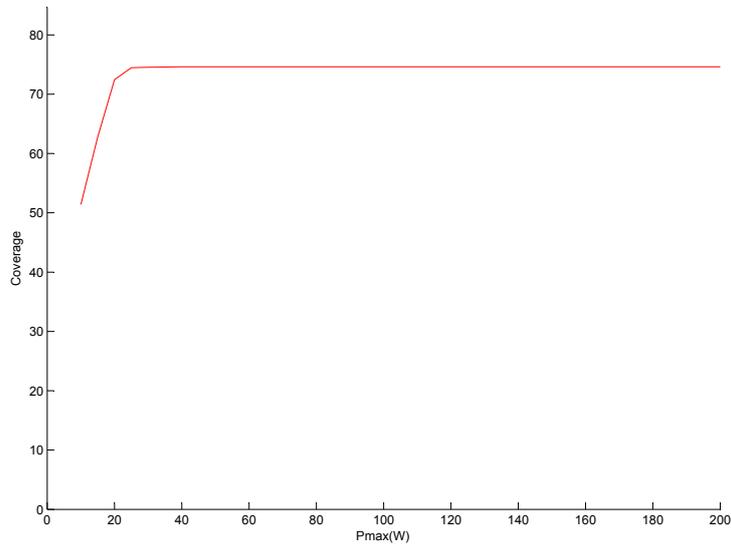


Figure 4.17: Coverage Characteristic for Extensive Form Game

A six-hour mission was simulated with two UAVs operating at an altitude of 10000ft. At the start of the mission the UAVs have non overlapping footprints. The mission was run ten times.

Figures 4.18a and 4.18b show mean coverage vs time and power vs time. The graphs are generated from the average coverage and power for each of the ten missions. Error bars set to 1 standard deviation have been added every 2000 seconds to show the variability of the data from the ten missions.

It is clear from figure 4.18a that there was an imbalance in coverage between the two UAVs. UAV1 has a significantly greater coverage than UAV2. Figure

4.18b shows that a power imbalance also exists.

The imbalance between coverage and power for the UAVs is unsurprising as UAV1 had the opportunity to pack its payload before UAV2, thus any overlapping coverage between the two UAVs was assigned to UAV1. This limited the potential coverage for UAV2, suggesting that its best strategy might be to move away from UAV1.

The mean coverage, C_{mean} , was 75.4 mobiles with a standard deviation of 4.2. The mean power, P_{mean} , required to support this coverage was 41.0W with a standard deviation of 4.5W. The mean power for each UAV was thus 20.5W.

It can be seen in figure 4.19a that UAV1 moves towards the centre of the scenario and tracks clusters of mobiles. UAV2, being unable to compete with UAV1 in the centre of the scenario, was limited to operating in a small region.

The coverage estimates for this game are generally good. Figure 4.19b shows a tight cluster of estimates.

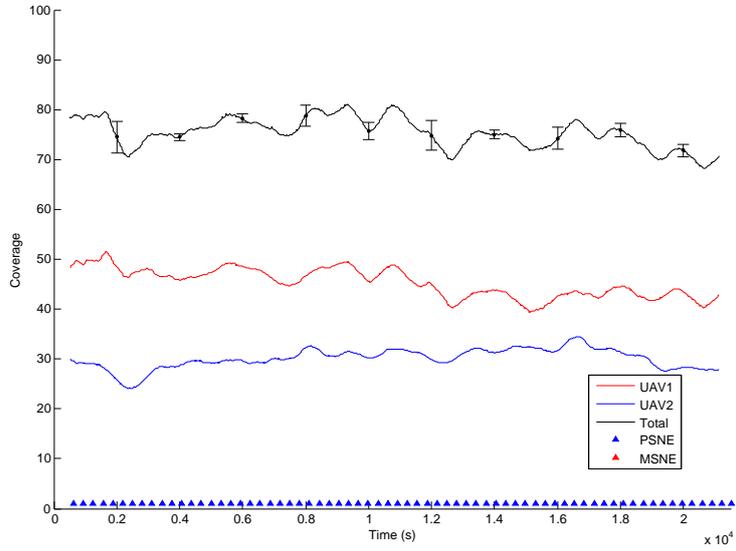
Table 4.2 shows the execution times for backward induction. The execution time for backward induction is relatively fast compared to normal form games. Furthermore, there is little variation in time between the mean and maximum times.

Algorithm	Parameter		
	Mean Time (s)	Max Time (s)	Number of NE
Backward Induction	4.36×10^{-4}	1.10×10^{-3}	One PSNE

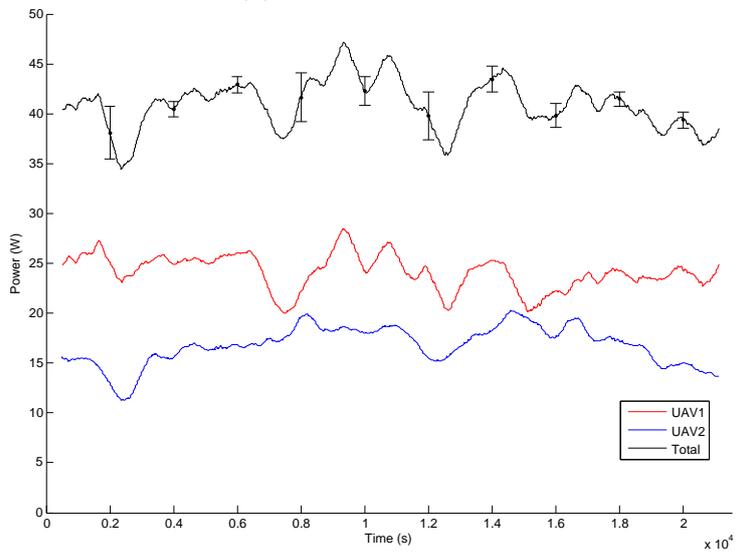
Table 4.2: Mean Execution Times for Backward Induction

4.6 Effects of Initial Conditions

In previous sections the two UAVs started in different quadrants of the scenario, thus ensuring that their footprints did not overlap. The observed behaviour was that both UAVs chose actions that would take them towards temporary clusters of mobiles. Attempts to share a cluster of mobiles resulted in a drop in coverage by one or both mobiles, thus the payoffs reduced. Any solution that generated an overlap in coverage was irrational as both players could find higher payoffs by choosing a different pair of actions.

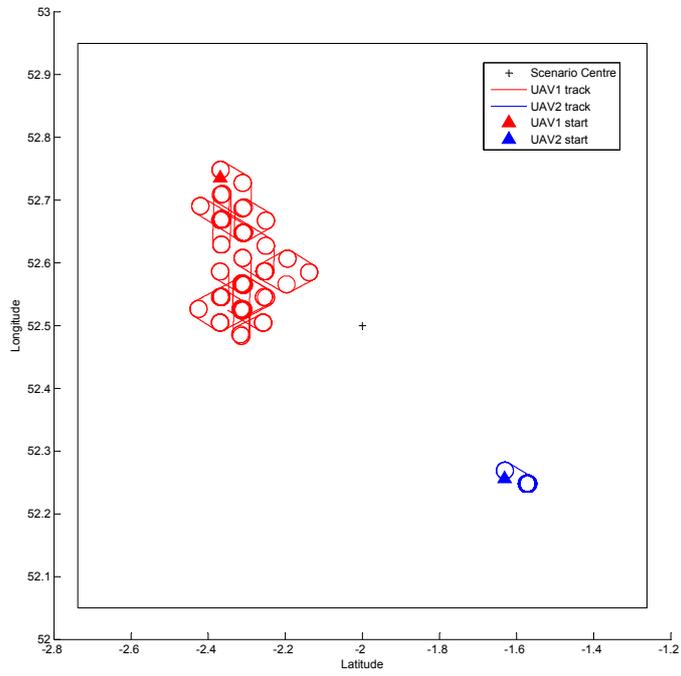


(a) Coverage v Time

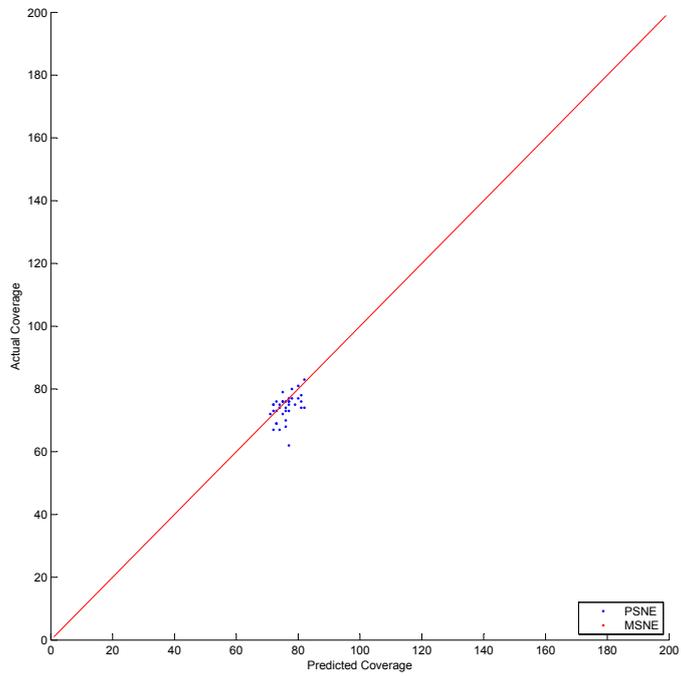


(b) Power v Time

Figure 4.18: Coverage and Power v Time for Extensive Form Game



(a) Flightpath



(b) Predicted v Actual Coverage

Figure 4.19: Flightpaths and Estimated Payoffs for Extensive Form Game

An example can be seen at figure 4.20. This shows the flight paths for two UAVs overlaid on the probability of finding a mobile at any location. There are many areas where there is a higher probability of finding a mobile, and these areas can be shown to be the areas where clusters are more likely to be found in this data set. The UAVs can be seen to favour such areas.

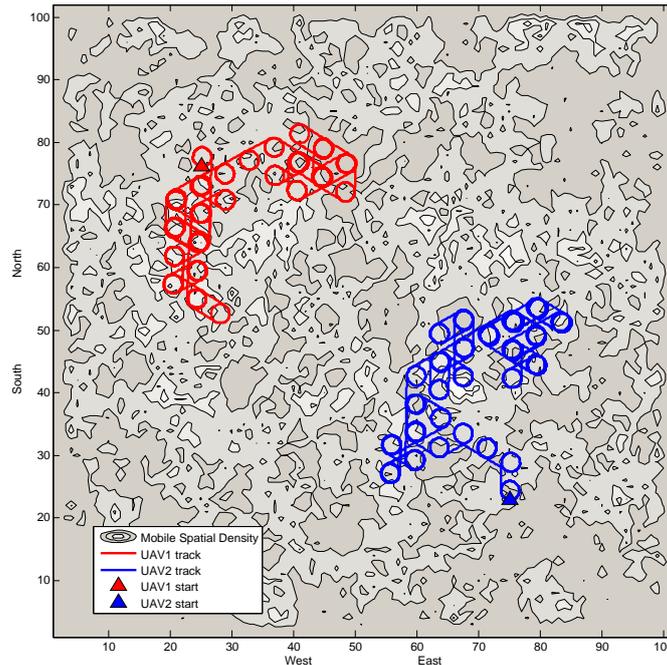


Figure 4.20: Flightpaths Overlaid on Mobile Probabilities: UAVS Start Separated

Because the footprints do not initially overlap each UAV can quickly find a unique set of mobiles to cover. Each will find a cluster in the first few manoeuvres and continue to maintain a consistent coverage.

This can be compared with examples where the two UAVs start in the same location. An example of this is shown in figure 4.21.

It can be seen if figure 4.21 that the UAVs start in the same location and immediately attempt to separate. Each of the UAVs will initially follow a sequence of manoeuvres that will take it to an area with a locally higher payoff. Over time they will separate sufficiently to find unique sets of mobiles and will track clusters in those areas.

Figure 4.22a shows the coverage when the UAVs start in different quadrants.

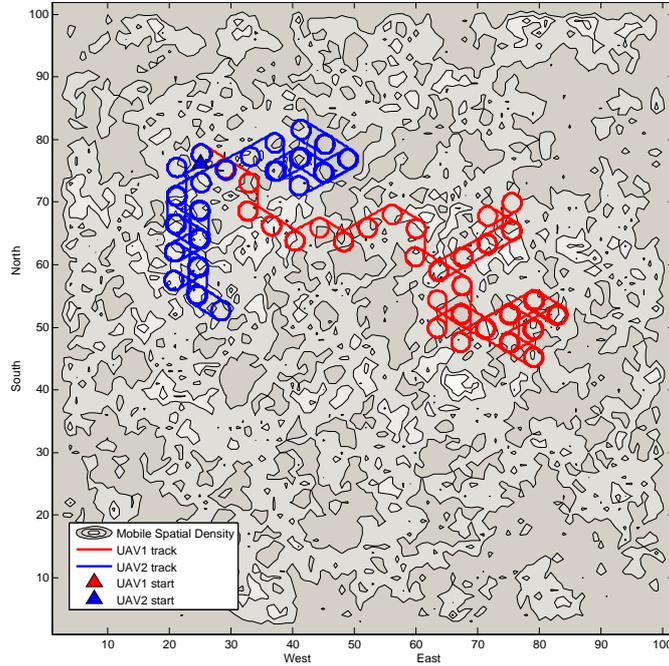


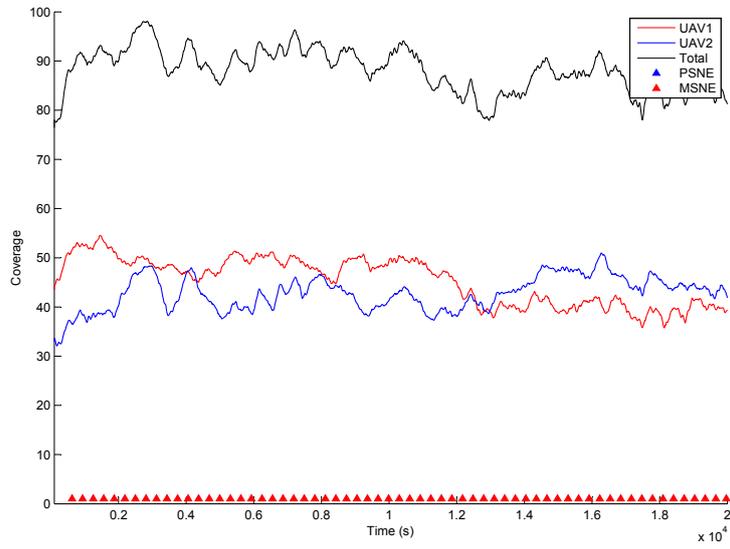
Figure 4.21: Flightpaths Overlaid on Mobile Probabilities: UAVS Start Collocated

The joint coverage settles to a steady value over about 1000 seconds as each UAV acquires a nearby cluster of mobiles. The settling time equates to 3 or 4 runs of the game.

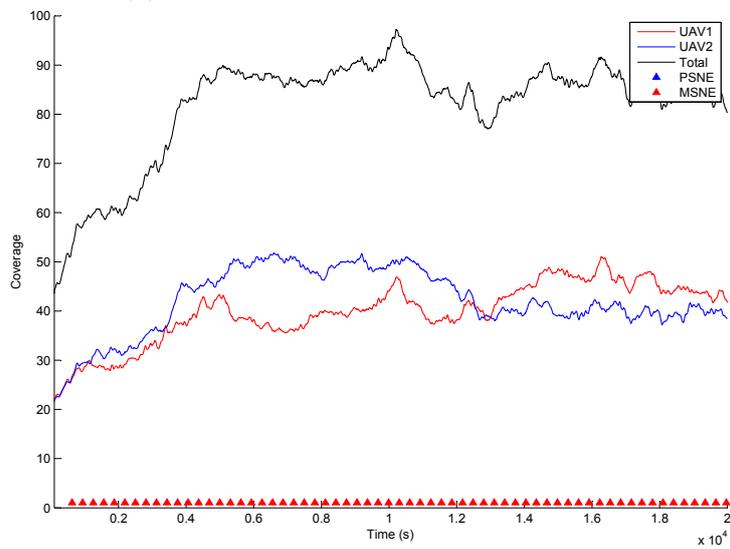
The effect on a collocated start on coverage can be seen in figure 4.22b. It can be seen that two UAVs are competing for the same mobiles so their individual coverages are low. Each UAV's coverage increases as they separate, and the joint coverage progressively improves until it reaches a steady state value comparable to that of the UAVs which started separately. We can see that the settling time for collocated UAVs was of the order of 5000 seconds, or about 15 runs of the game.

These exemplar graphs are typical of other results for collocated starts. Settling times from 2500 seconds to 6000 seconds were observed in runs with different data sets and different initial conditions.

The radius of a UAV footprint at 10000ft with a beamwidth of 165° is 22328m. Figures 4.22a and 4.22a show mean coverage over the whole scenario period. Figure 4.23 gives some additional insight into the behaviour of the UAVs over time by plotting the separation between the UAVs against time. The red plot shows the



(a) UAVs Start in Different Quadrants



(b) UAVs Start Collocated

Figure 4.22: Coverage v Time: UAVs start in NW and SE Quadrants

separation between the UAVs when they start in the NW and SE corners, while the blue plot shows what happens if both UAVs start in the NW corner. A black line shows the distance between the UAVs were separated by two footprint radii, in other words their footprints just touched each other.

In the case where the UAVs started in different quadrants we can see that they initially move towards each other, then maintain a separation of at least two footprint radii. This allows them to track clusters of mobiles while maintaining unique coverage.

When the UAVs start collocated in the same quadrant they move apart very rapidly until they are separated by at least two footprint radii. This separation was reached at about 5000 seconds into the scenario, the same time as the joint coverage starts to settle. After this time the separation and coverage for both sets of initial conditions are very similar.

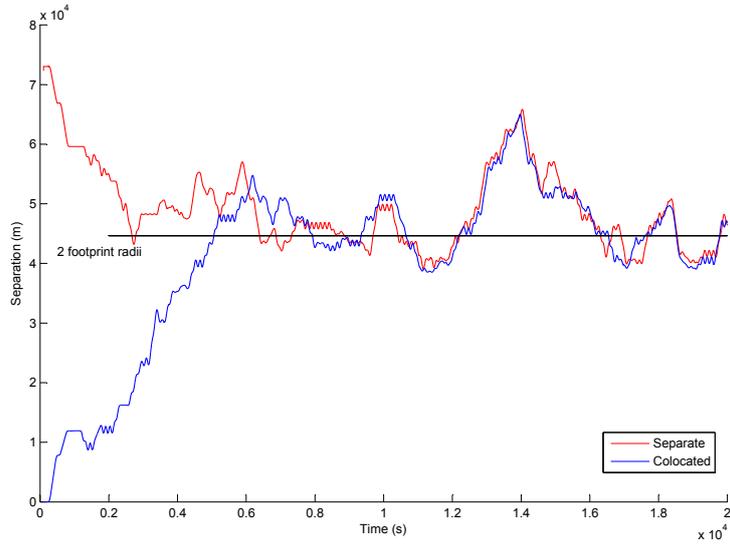
It is possible to estimate the minimum settling time. Assume that the coverage will settle when the UAVs are two footprint radii apart. It is known that, in this game, the UAVs move across the scenario at approximately 14.4m/s. If both UAVs move apart at this rate then they should separate in about 1550 seconds, whereas if one UAV remains fixed over a cluster this should take 3100 seconds. These estimates are compatible with the observed data that showed settling times over the range 2500 seconds to 6000 seconds.

The evidence indicates that the UAVs converged on the same solution from different sets of initial conditions.

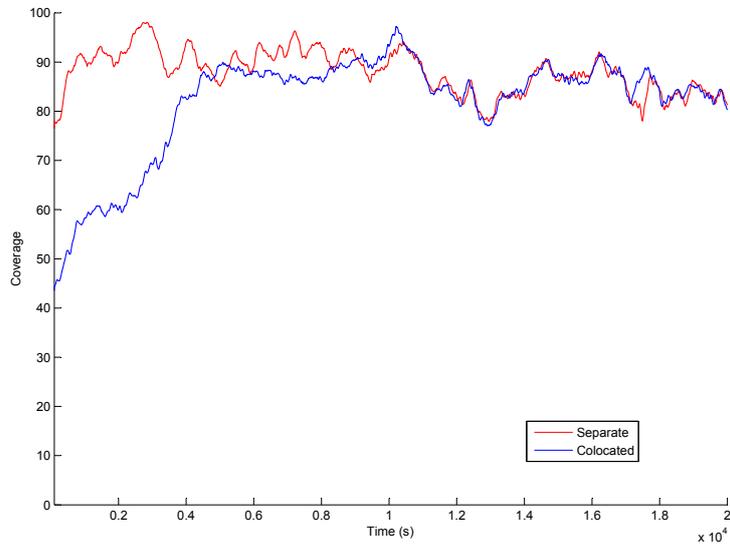
4.7 Initial Conditions in which one UAV is “Cornered”

The previous section suggests that the game will converge on a stable coverage regardless of the initial conditions. This is not always the case. An exception was found when there are very few mobiles and the UAVs start in the corner of a scenario. An example of this can be seen in figure 4.24. The black box denoted the edge of the 100km×100km region in which mobiles operate.

Figure 4.24a shows the flightpaths of the two UAVs. UAV1 moves towards the centre of the scenario, while UAV2 leaves the scenario and never re-enters. The corresponding plot of coverage vs time shows that the total coverage was identical



(a) UAV Separation v Time



(b) Joint Coverage v Time

Figure 4.23: UAV Separation and Coverage v time

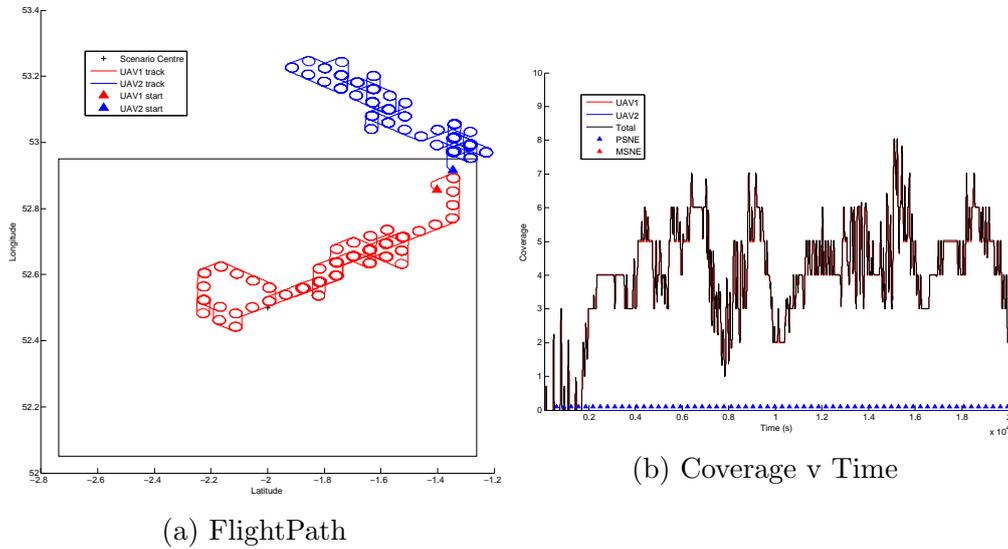


Figure 4.24: An Example of Cornering in which the UAV Fails to Recover

to that of UAV1, so the coverage of UAV2 was zero throughout this run.

This example was taken from a game with 20 mobiles. At the start of the run UAV1 was located between UAV2 and the centre of the scenario. UAV1 effectively blocks access to the mobiles by UAV2, so the payoff for UAV2 was very small. Whatever strategy was adopted by UAV2 it was unable to increase its payoff, whereas UAV1 simply needed to move towards the centre of the scenario. UAV2 has been “cornered”.

For cornering to occur one of the UAVs, in this case UAV2, must start with zero payoff for all combinations of strategies. In such circumstances every payoff for UAV2 was its best response and the PSNE was the strategy that gave the maximum payoff to UAV1. UAV2’s movement was entirely dictated by the strategy adopted by UAV1 so, if UAV1 selected a strategy that resulted in UAV2 remaining outside the game, UAV2 continued to have a payoff of zero.

Experiments showed that the situation shown in figure 4.24 was uncommon and that UAV2 recovered from cornering on over 90% of occasions. The trigger for recovery was when UAV1 selected a strategy that allowed UAV2 to have a non-zero payoff for some strategies. This creation of a non-zero best response by UAV2 would stimulate competition between the two UAVs.

An example of recovery can be seen at figure 4.25. UAV2 re-enters the game after a period and progressively increases its payoff until, by the middle of the

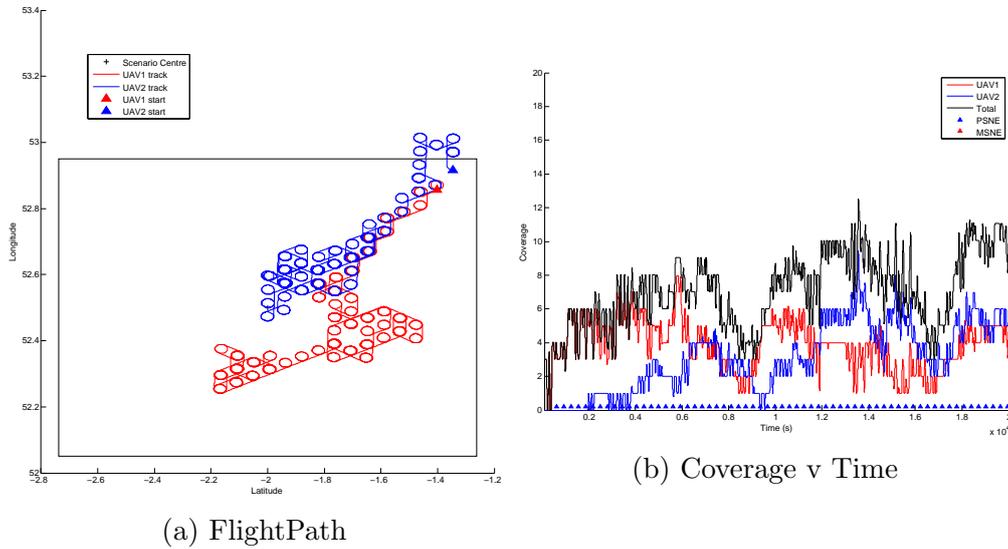


Figure 4.25: An Example of Cornering in which the UAV Recovers

period, it starts to compete with UAV1.

Cornering requires a specific set of circumstances to occur:

1. There must be a small number of mobiles so that it is possible for one of the UAVs to have a payoff of zero.
2. One of the UAVs must be separated from the mobiles by the other UAV.
3. The initial sequence of strategies must move the cornered UAV away from the scenario edge on a path that maintains a zero payoff.

This combination of circumstances is always detectable. A simple test in the code can be used to identify the existence of zero payoffs, demonstrating that there is a risk that cornering may occur. If this test shows a positive result on successive runs of the game then cornering can be assumed.

Countermeasures to cornering are straightforward. It can be prevented by ensuring that the two UAVs do not start close together. If it occurs, then the simplest countermeasure is for the cornered UAV to move towards the other UAV for the next few runs of the game as this will inevitably increase its payoff. An example of the successful applications of countermeasures can be seen in figure 4.26.

Figure 4.26a shows that UAV2 started to leave the scenario, however the existence of four successive zero payoffs was detected. By this time UAV1 was located

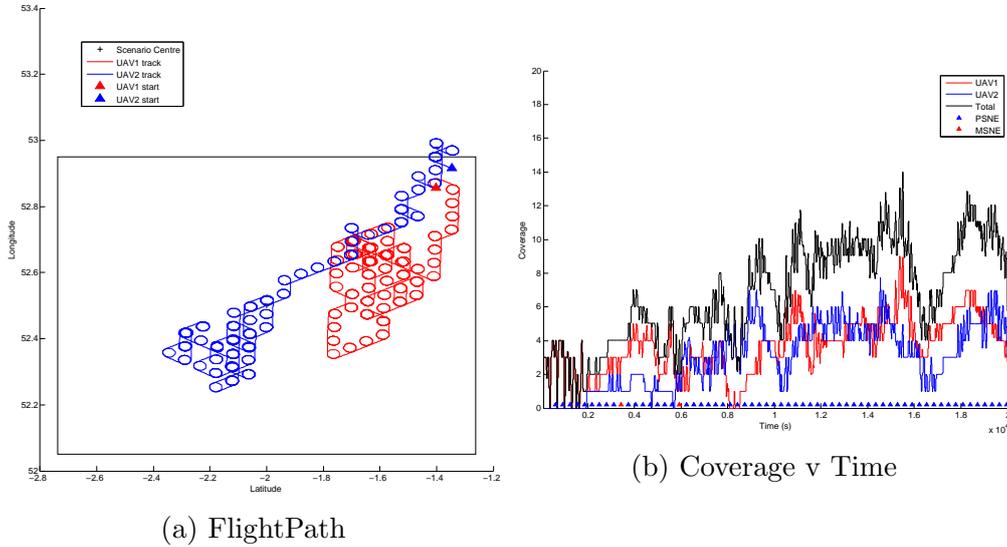


Figure 4.26: An Example of Countermeasures to Cornering

directly south of UAV2. Having recognised that it was cornered UAV2 selected strategies that took it south, towards UAV1, until it started to develop non-zero strategies. At this point the countermeasures ceased operating and UAV2 rejoined the competition with UAV1. The payoffs for both UAVs can be seen in figure 4.26b.

4.8 Chapter Conclusions

This chapter has considered five methods for solving non-cooperative games and has compared these to the benchmark of two circling UAVs. It is required that the coverage obtained for relocating the UAVs should exceed the two benchmarks. A table of results comparing these methods can be seen at table 4.3. The sampling time for the calculation of C_{mean} , P_{mean} and SP_{mean} was six hours.

It is clear that relocating the UAVs by solving a one-shot game will give significantly better coverage than allowing the UAVs to circle fixed points. Exhaustive search, fast search, the Lemke-Howson algorithm and Chatterjee’s method all gave significant improvements over the circling UAVs. The extensive form game solved by backward induction is only marginally better than the circling UAVs. The adoption of random strategies was also considered but the results were unpredictable.

The first method for solving the normal form game, exhaustive search for

Method	Mean Coverage C_{mean}	Mean Power P_{mean}	Mean Specific Power SP_{mean}	Quality of Predictions	Execution Time (s)
Circle	73.7	41.4	0.56	N/A	N/A
Exhaustive Search	90.1	49.8	0.55	Over estimates coverage	3.21×10^{-2}
Fast Search	89.7	49.0	0.55	Slightly over estimates coverage	1.67×10^{-4}
Lemke-Howson	89.7	48.0	0.53	Over estimates coverage	1.98×10^{-4}
Chatterjee	89.1	48.2	0.54	Over estimates coverage	0.11
Extensive Form Game	75.4	41.0	0.54	Slightly over estimates coverage	4.36×10^{-4}
Random Strategies	Indeterminate				

Table 4.3: Comparison of Solution Methods

PSNE, was consistent over a range of power levels and beamwidths and generated good results. A faster search method was considered, and this gave good results that appeared to be less consistent than exhaustive search. Both these methods ignored the possible existence of MSNE, adopting a “do nothing” approach if no PSNE was found. A third method used the Lemke-Howson algorithm to search for one MSNE, guaranteeing a solution and removing the need for the “do nothing” option. The results for the MSNE produced lower payoffs than the previous methods, mainly because of the randomisation process. Chatterjee’s algorithm, while relatively slow, was also able to find one MSNE and delivered better coverage than Lemke-Howson. The power efficiency of these methods was very similar as most methods had a specific power close to 0.55W/mobile. The power demand was always significantly below 50W, thus the system was in a footprint limited state. The extensive form game gave only a marginal improvement over the circling UAVs and, with imbalanced coverage, was considered unlikely to generate useful results.

Extensive form games based on game trees and backward induction generally

give poorer results than methods based on analysis of a payoff matrix. The equivalence of extensive and normal forms of a game is described in section A.5. The main difference between the two forms of game is that a normal form game assumes simultaneous decisions whereas an extensive form game assumes sequential decisions. For any normal form game there can be multiple, equivalent, game trees reflecting the different order of play. Solving each possible game tree by backward induction can give a different solution to the game, whereas solving the equivalent normal form game generally gives the same solution. Furthermore, extensive form games assume that a PSNE exists and look for the best responses to each player to all choices by previous players. This will always generate a result, but the overall payoff is usually lower than the payoff from the equivalent normal form game.

Table 4.3 shows that a normal form game, regardless of how it is solved, should always give better results than circling. The table shows that the normal form games give consistently good coverage irrespective of the method used to solve the game. The main differences between the four methods are the quality of the predictions and the execution times. All methods for solving normal form games tend to overestimate the coverage, but fast search appears to give the best estimates of coverage and have the best execution time.

The extensive form game slightly overestimates the coverage. The extensive form game was considered unusable and was dropped.

In a perfect scenario the UAVs would start with non-overlapping footprints in different regions of the scenario, In a real scenario it is possible that they can start with overlapping footprints, even to the point where the UAVs are collocated, and they will find solutions that maximise coverage. Separation of the UAVs from a collocated start can take a significant time. At the very least it will be 1550 seconds, and is more likely to be in the range 2500 to 6000 seconds.

It can be seen from the flight paths that, apart from the case when the UAVs were deliberately started together, both UAVs maintained good separation. This game maintains good separation between the UAVs.

It is uncommon, but possible, for one UAV to be cornered by another. When this occurs one of the UAVs will exhibit zero payoffs for all possible strategies for successive runs of the game. This property allows cornering to be detected and countermeasures to be applied.

This page is intentionally blank

Chapter 5

Some Implementation Issues

5.1 Chapter Introduction

The previous chapter demonstrated that repeated application of a memoryless two-player game can be used to relocate two UAVs. It was clear that similar solutions were obtained for normal form games, regardless of how the game was solved, and that different methods for solving the game required different execution times. The initial conditions of the game influenced how a scenario evolved over time. A steady state could be obtained if the initial conditions were controlled to avoid the “cornering” of one UAV by the other. The simplifying assumption of a footprint limited state allowed the different methods to be compared in a controlled manner.

This chapter seeks to develop the methods used in the previous chapter with the aim of obtaining solutions that are not just numerically satisfactory but which satisfy the needs of a real UAV network. A real UAV payload would try to maximise the use of the available power to support as many mobiles as possible. This chapter considers the *power limited* case to assess how well the game performs in a more realistic environment.

Civil communications systems often require a prioritised access scheme to ensure important users, or essential messages, have the highest probability of being delivered. Such a scheme could be implemented on this system by introducing a preferential packing scheme. A similar method for a single UAV was described by the author in reference [19].

Coverage has, so far, been treated as the total number of connected mobiles.

No account has been taken of which mobiles are connected, or how the connectivity is changing over time. Changes in link state, and hence network topology, can have a detrimental effect on network performance. It is thus desirable to maintain links for longer periods to maintain a stable network topology. This chapter starts to address how the game could be implemented to address the stability requirements of the network, not just instantaneous connectivity.

The frequency of operation has been constrained to 5GHz in earlier chapters. Radio path loss, and hence power and coverage, are frequency dependent. This chapter briefly considers how choice of frequency band affects coverage.

The previous chapter considered the UAVs participating in two-player games with the UAVs at fixed altitudes, resulting in a 7×7 payoff matrix. Larger payoff matrices will be generated if the number of UAVs increases, or if UAVs are allowed to move in three dimensions. This chapter expands the size and complexity of the payoff matrices and assesses how this affects the methods of solution, the total coverage, and timeliness of solutions.

5.2 Prioritised Access Schemes

A prioritised access scheme allows some mobiles to be given preference for gaining and retaining access to the UAV. The need for preferential access could be operational, for example the incident commander in a disaster must never be blocked and is therefore given priority access. Alternatively priority access could be required for technical reasons, for example to provide continuous access in order to maintain QoS for a particular information exchange.

Such a scheme was implemented by manipulating the payload packing method. A priority value was assigned to each mobile through a priority vector V :

$$V = (v_1, v_2, \dots, v_i, \dots, v_m)$$

This defined a simple two-layer scheme in which, for mobile i , the priority can be defined as:

$$v_i = \begin{cases} 1 & \text{if mobile } i \text{ has preferential access} \\ 0 & \text{if mobile } i \text{ does not have preferential access} \end{cases}$$

If, for example, mobiles 3, 8 and 11 are given priority access:

$$V = (0, 0, 1, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0)$$

We can prepend a column vector V^T to the power matrix P to indicate priority. This creates the matrix $P_{mod} = (V^T|P)$

$$P_{mod} = \left(\begin{array}{c|cccccc} v_1 & p_{11} & p_{12} & \cdots & p_{1j} & \cdots & p_{1n} \\ v_2 & p_{21} & p_{22} & \cdots & p_{2j} & \cdots & p_{2n} \\ \vdots & \vdots & & \vdots & & \vdots & \\ v_i & p_{i1} & \vdots & \cdots & p_{ij} & \cdots & p_{in} \\ \vdots & \vdots & & \vdots & & \vdots & \\ v_m & p_{m1} & p_{m2} & \cdots & p_{mj} & \cdots & p_{mn} \end{array} \right) \quad (5.1)$$

Sorting the rows of this matrix by column 1 will separate the mobiles with and without priority access. Matrix P_{mod} can then be partitioned into two submatrices, one with priority and one without. We can then apply a normal packing algorithm to each submatrix, starting with the priority mobiles and allocating any remaining power to the remaining mobiles.

Note that the priority vector V is not fixed and could be time variant. This would allow the relative priority of mobiles to change dynamically. It is also possible to add multiple priority levels by changing the piecewise function as follows:

$$v_i = \begin{cases} 2 & \text{if mobile } i \text{ has overriding access} \\ 1 & \text{if mobile } i \text{ has priority access} \\ 0 & \text{if mobile } i \text{ does not have priority access} \end{cases}$$

Multiple-level priority schemes, and dynamic schemes, are operationally interesting but offer similar performance to two-level schemes. They are therefore not explored further in this thesis.

This section gives some exemplar results for a system in which 40 of the 200 mobiles were randomly selected as having priority access. The mission was run twice, the first time with priority access and the second time without. Figure 5.1 compares the coverage of those mobiles in both missions.

It is clear that coverage improved access for those mobiles if they were given priority access. For a random sample of 40 of the 200 mobiles, the mean coverage,

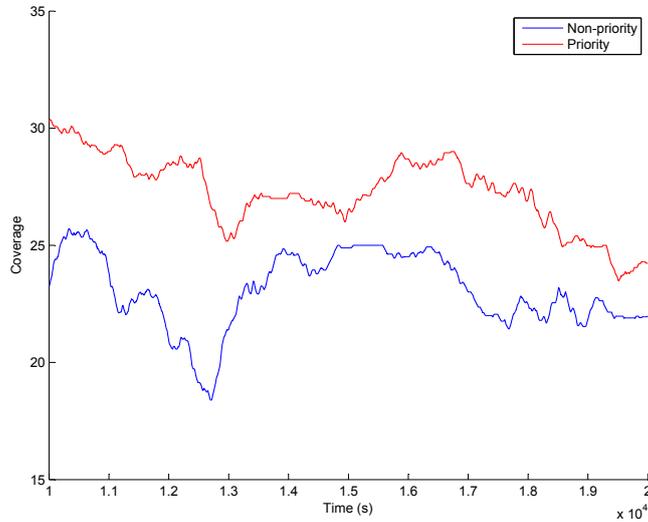


Figure 5.1: Coverage v Time for Mobiles with Prioritised Access

C_{mean} , over the sample period was 27.3 with priority access and 23.1 without priority.

The change in coverage for the priority mobiles was accomplished by changing the flight paths to maximise their access time at the expense of other mobiles. The mean coverage, C_{mean} , for all mobiles was 120.8 with priority access and 119.6 without, indicating that the overall coverage was not significantly affected.

5.3 Modifying Behaviour by Biasing Payoffs

The network described in this thesis has a tree topology with the inter-UAV links forming the branches and the mobiles forming the leaves. In a tree network there is some potential for interruption to ongoing information exchanges when a mobile is added or dropped by the UAV. Game solutions that are likely to minimise interruptions are desirable. The effect of this was to increase the duration of ongoing accesses, improving the stability of the network topology.

A method was devised which compared new allocations with those made during the previous run of the game. Changes were made to the method of calculating the payoff matrix. This allowed the behaviour of the UAVs to be modified without affecting the functionality of the game. A higher weighting was given to alloca-

tions that had not changed since the previous run, and a penalty was awarded to allocations that had changed. For this experiment the game was allowed to have some memory of previous results.

Recall from chapter 3 that the allocation of mobiles to UAVs is defined in an allocation matrix X , and that the column j of this matrix is the vector x_j that shows the allocation of mobiles to UAV j .

The allocation matrix for the time at which the game is run is X_{now} and the estimated assignments can be defined as X_{next} . We can derive the current and future assignments for UAV j from these matrices such that:

$$x_{jnow} = (x_{1now}, x_{2now} \dots x_{inow} \dots x_{mnow})$$

$$x_{jnext} = (x_{1next}, x_{2next} \dots x_{inext} \dots x_{mnext})$$

In chapter 4 the method used for calculating payoffs was simply to take the sum of the columns of matrix X_{now} . Calculating network stability requires that X_{now} and X_{next} are compared. This will establish four possible states for each mobile:

- **Null.** It is currently unconnected and is expected to remain unconnected after the current manoeuvre.
- **Retained.** It is currently connected and is expected to remain connected after the current manoeuvre.
- **New.** It is currently unconnected and is expected to become connected after the current manoeuvre.
- **Lost.** It is currently connected and is expected to become unconnected after the current manoeuvre.

The **NULL** state is of no interest to a retention policy, however the **Retained**, **New** and **Lost** states are of interest. We can use X_{now} and X_{next} to create scalar values C_r for the retained, C_n for the new and C_l for the lost mobiles.

$$C_r = \sum_{j=1}^m x_{jnow} \wedge x_{jnext}$$

$$C_n = \sum_{j=1}^m x_{jnow} \wedge \neg x_{jnext}$$

$$C_l = \sum_{j=1}^m \neg x_{jnow} \wedge x_{jnext}$$

We can generate a bias vector $B = (b_r, b_n, b_l)$ that weights the importance of each of these classes. The payoff π_j for UAV j can be calculated as:

$$\pi_j = b_r C_r + b_n C_n + b_l C_l$$

Within the vector B the value of b_r gives the weighting for a connection that is retained from the previous run of the game. Normally the value of $b_r = 1$. If, for example, $b_r = 2$ then any mobile that is connected for more than one run of the game would give a payoff of 2 instead of the default payoff of 1. A similar scheme was used for b_n , the weighting for new connections. Normally $b_n = 1$ but, if the intention is to encourage expansion, b_n could be set to a value of 2 or more. The value of b_l is intended to discourage UAVs from losing mobiles that they already support. the default value is $b_l = 0$, however the behaviour of the UAV can be modified by reducing its payoff if it loses a previously supported mobile. In this case the payoff needs to be reduced, so a penalty scheme could have negative values, for example $b_l = -1$.

The baseline against which this was compared was a game played using exhaustive search with $B = (1, 1, 0)$. The distribution of access durations can be seen at figure 5.2.

It is clear from figure 5.2 that the probability distribution function for access duration is long tailed. A method was needed to establish the statistics of access durations, but it was recognised that mean or median access duration would have little relevance with a long tailed distribution. It was decided to establish a cutoff point beyond which the longer accesses would be ignored.

It was noted from figure 5.2 that most of the accesses with durations greater than 250 seconds were single accesses, while there was a clearly defined distribution for accesses with durations of 250 seconds or less. Further analysis showed that 80% of the accesses had a duration of 274 seconds or less, as shown in figure Figure 5.3, indicating that the main lobes of the PDF encompassed 80% of the accesses. The cumulative distribution function indicated that the 20% of the accesses for

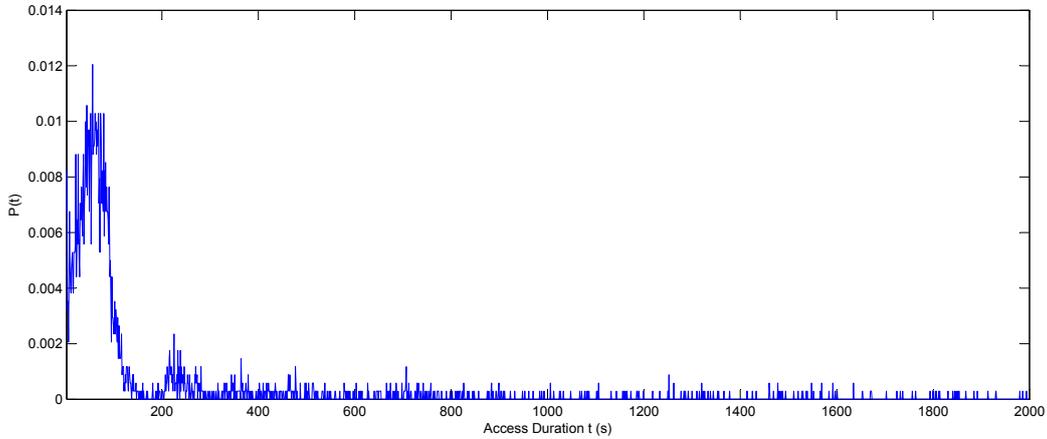


Figure 5.2: PDF for Access Duration

$B = (1, 1, 0)$ lie in the tail of the distribution. These are, by definition, the longer accesses so improving their duration may have little impact on the stability of the network topology.

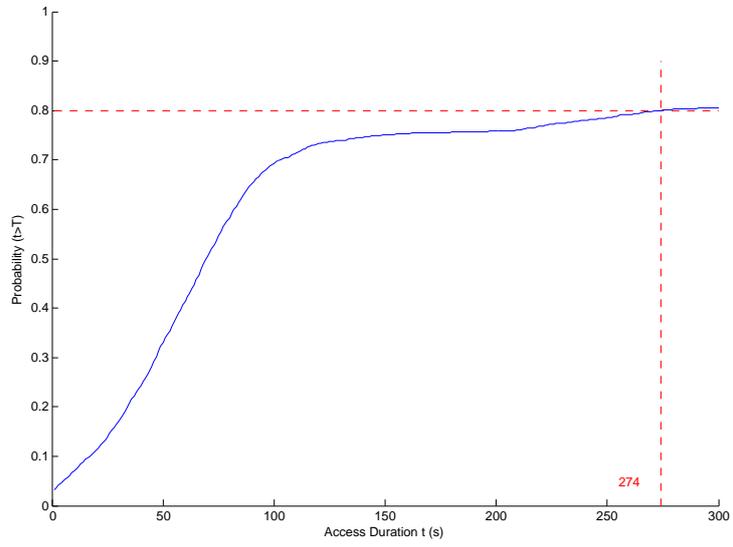
Recall from chapter 3 that each run of the game estimates the payoff 312 seconds into the future, and it can be seen that the largest impact on network topology will be achieved by focussing on increasing the shorter accesses rather than those in the tail.

Together these factors lead to the conclusion that the region of interest for evaluating the effects of biasing payoffs is likely to be 312 seconds. The cutoff point was thus established at 312 seconds.

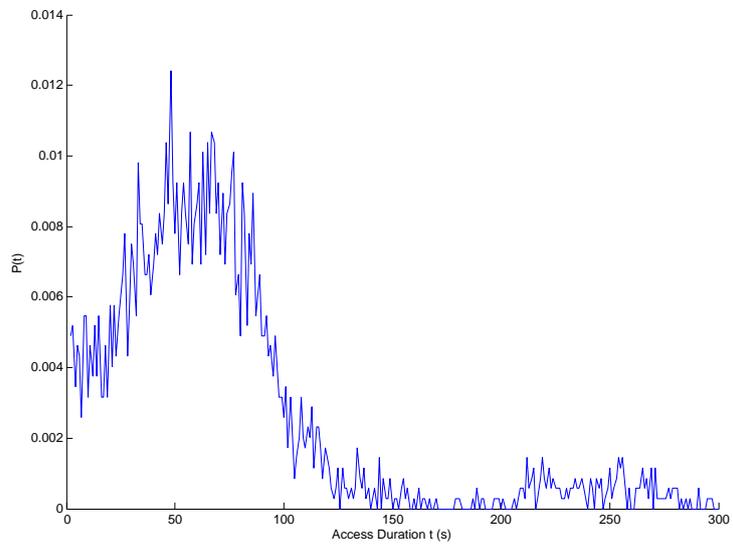
The default bias that corresponds to the games from chapter 4 is $B = (1, 1, 0)$. By changing B we can emphasise specific behaviours. Emphasising the importance of retention of mobiles can be implemented by increasing b_r so that, for example, $B = (2, 1, 0)$. We could also penalize the UAV for losing coverage of a mobile by setting $b_l = -1$ so that $B = (1, 1, -1)$.

The results are shown in table 5.1. Note that the mean duration of accesses only considers accesses of 312 seconds or less; it ignores the accesses in the tail.

It can be seen that biasing the payoff calculation to encourage retention increases the mean access duration by about 8%, whereas encouraging the UAV to find new mobiles causes a reduction of over 4%. Reducing the payoff if the UAV loses coverage of mobiles gives a modest 5% improvement in access duration, but is not as effective as encouraging retention. Combining the two effective techniques,



(a) CDF



(b) PDF

Figure 5.3: PDF and CDF for the first 300 seconds

Bias vector B	Mean Coverage	Access Duration (s)	Intention
(1, 1, 0)	89.8	42.4	Default settings for all games
(2, 1, 0)	84.6	46.1	Encourage retention
(1, 2, 0)	85.0	40.6	Encourage expansion
(1, 1, -1)	82.3	44.7	Penalize loss
(2, 1, -1)	86.8	48.8	Encourage retention, punish loss

Table 5.1: Effect of Biasing Payoffs

encouraging retention and penalizing loss, gives a combined 15% improvement, better than the sum of their individual improvements.

The improvement in access duration comes at the cost of reduced mean coverage. The mobiles are diverging over time, hence the cluster of mobiles that the UAV is attempting to track will dissipate. The UAV will maintain contact for longer, but with fewer mobiles.

The data shows that the default bias (1,1,0) has 110 accesses of 1 second duration, caused when a mobile grazes the edge of the footprint. These accesses are suppressed by encouraging retention (2,1,0), and eliminated by encouraging retention and penalizing loss (2,1,-1). Suppression of short accesses occurs in accesses in the range 1-10 seconds, and the data indicates that the reduction in coverage occurs at the expense of these short accesses.

At a first glance, the default bias vector $B = (1, 1, 0)$ should give similar results to $B = (2, 1, -1)$ as the difference between retention and punishment are the same for both vectors. The results in table 5.1 are significantly different, a difference that requires some explanation. The default bias $B = (1, 1, 0)$ is indifferent to which mobiles are supported, provided that the total payoff remains the same., so it can be completely opportunistic when selecting mobiles to be supported. The bias vector $B = (2, 1, -1)$ rewards the retention of specific mobiles and punishes the loss of mobiles. The UAV is thus rewarded for supporting specific mobiles over a longer period.

The difference in mean coverage and access duration reflect the difference between these two bias vectors. Indifference about specific connections tends to give higher coverage as the UAV is opportunistic in its search for mobiles. When the UAV is rewarded for supporting specific mobiles the coverage tends to be lower as the UAV tries to find locations that can maintain coverage of connected mobiles.

Opportunism results in shorter access times, while support for specific mobiles tends to give longer access times. Biasing the payoffs appears to be an effective way of improving the stability of the network topology.

5.4 Choice of Frequency Band

One component of the radio path loss, free space loss, is frequency dependent. The loss between transmitter and receiver is proportional to f^2 so increases in frequency will increase the path loss. A corresponding increase in transmitted power is needed to satisfy the required signal to noise ratio at the receiver. The choice of frequency band affects overall system performance as higher frequency requires higher power, resulting in fewer mobiles being covered by each UAV.

The useful range of frequencies for this application tend to occur between 1GHz and 10GHz, that is wavelengths between 30cm and 3cm. The longer wavelength at frequencies below 1GHz causes practical antennae to become too large to fit into the airframes of MALE UAVs. At frequencies of 10GHz and above the propagation path is increasingly affected by rain fade; The scattering of RF energy from raindrops introduces additional path losses resulting in the need for active power control to maintain the same signal to noise ratio at the receiver.

This thesis uses the 5GHz frequency band for its calculations. This band has been allocated by the International Telecommunications Union Radio (ITU-R) for use of links for UAV control. It is anticipated that adjacent frequency bands will be allocated for use by UAV payloads.

An experiment was conducted using two UAVs to assess the effect of using different frequency bands. Figure 5.4 plots mean coverage against P_{max} at several frequency bands. The 5GHz curve represents the default performance used in this thesis.

It can be seen that choice of frequency band affects coverage. Figure 5.4 shows a series of power characteristics, a familiar graph from previous chapters.

The position of the curve changes with frequency. At frequencies above 5GHz the payload remains in a power limited state but with reduced coverage whereas at frequencies below 3 GHz the payload is footprint limited.

The explanation is obvious. The power required to support an access falls with decreasing frequency, so the payload can support more accesses at a lower frequency. Consequently there must be a frequency at which all the accesses within

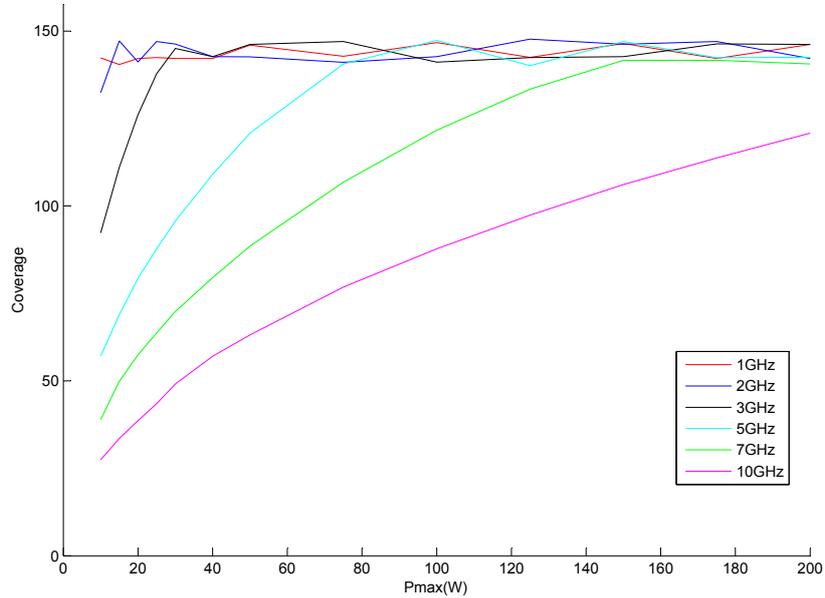


Figure 5.4: Coverage Characteristic at Several Frequency Bands

the footprint can be supported by the power available to the payload.

5.5 Games with Larger Payoff Matrices

Consider a game with n players, each of which has K possible pure strategies. The payoff matrix must contain K^n cells, so increasing either the number of players or pure strategies increases the size of the payoff matrix.

The results so far have considered games with $n=2$ and $K=7$. This section addresses two-dimensional games with multiple players in which $n \in \{2, 3, 4\}$ and $K = 7$. It then considers three-dimensional games in which $n \in \{2, 3, 4\}$ and $K = 21$.

The solution methods used in this chapter are exhaustive search and Chatterjee's method. This is because fast search and Lemke-Howson methods can only solve two player games.

5.5.1 Two-Dimensional Games with Multiple Players

This game is similar to the two-player game seen in earlier chapters. An example of an action set for three players in the two-dimensional game can be seen at figure 5.5.

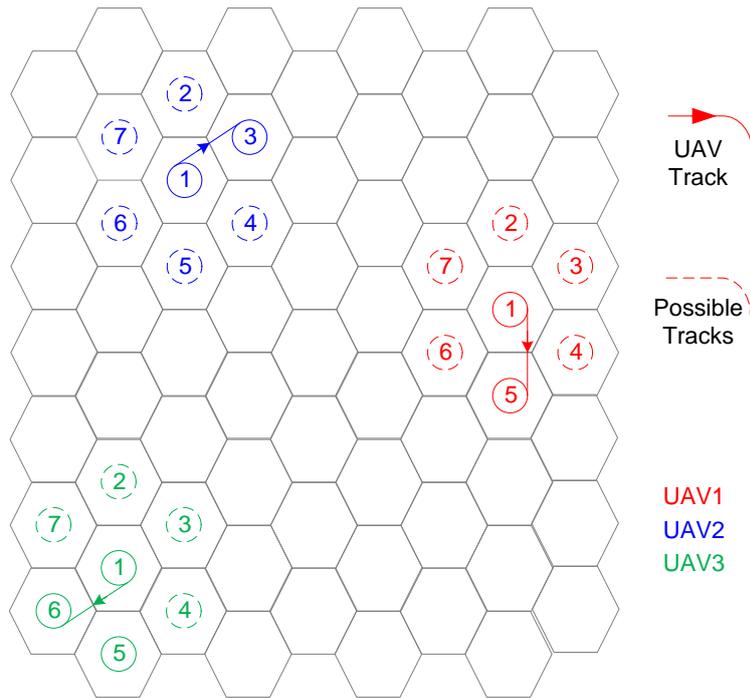


Figure 5.5: The Three UAV System in Two Dimensions

This is a normal form game so the players all move simultaneously to their next location. The path lengths for each move are identical to those in earlier two-dimensional games so the duration of the manoeuvre is unchanged at 312 seconds.

UAVs start at the footprint limited operating point. In 2D games the operating point remains constant, but in 3D games changes in altitude will affect the size of the footprint, allowing UAVs to move to power limited states.

5.5.2 Three-Dimensional Games

The three dimensional game extends the action set for each UAV. In addition to the 7 actions available at the current altitude, the UAV has another set that can be reached by climbing at the maximum rate of climb, and a further set that can

be reached by descending at the same rate. An example of an action set for one player in the three-dimensional game can be seen at figure 5.6.

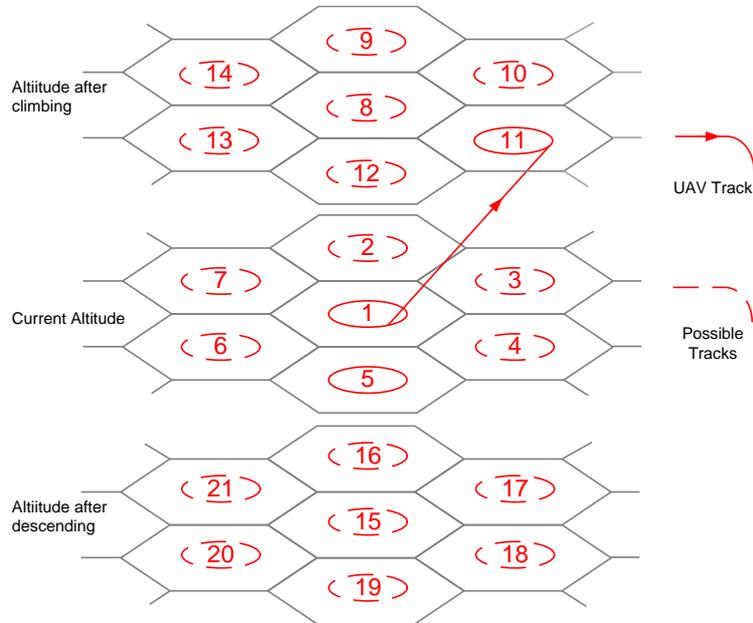


Figure 5.6: Action Set for a Single UAV in Three Dimensions

As before, the time allocated to complete the manoeuvre is sufficient for the UAV to complete one circle at its current altitude, move to the next cell, and complete one circle at its current altitude. The move to the next cell may involve a climb, level flight or descent. The flight path for a climb or descent will be longer than that required if the current altitude is maintained. In practice, the relatively benign rate of climb ensures that the slight increase in path length can be absorbed by a marginal but tolerable increase in airspeed. This allows the duration of the manoeuvre to remain constant at 312s.

5.6 Execution Times

The size of the payoff matrix, the method used to solve the game and the available computing resources all affect the time taken to find a solution to the game. The age of the data from which the payoff is calculated is a consequence of the time taken to solve the game. This section illustrates how these factors can affect the quality of the result.

Recall that these simulations were run in MATLAB on a dual core processor with 20 GB of RAM using Windows 7. The operating system was stripped of non-essential processes to improve speed. The results in this section are only valid for this configuration but they illustrate how system implementation affects the quality of the results.

In a real system the UAV cannot stop to await completion of the game; it would fall out of the sky. It is therefore required that the process of creating the payoff matrix and finding a solution is completed before the end of the current manoeuvre. The time taken to calculate the payoff matrix and solve it is defined as the execution time for the algorithm, t_{ex} .

Figure 5.7 illustrates the significance of execution time. The timeline is marked with red triangles that show the times at which the game is run. The current run is filled red while previous and future runs are unshaded. The runs occur at intervals of t_{sol} seconds.

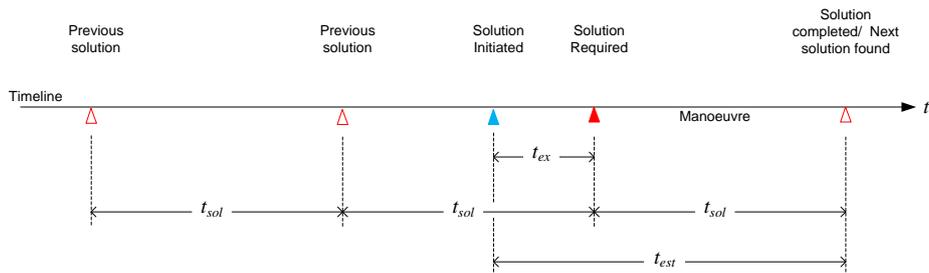


Figure 5.7: Timeline for Finding Solutions: $t_{ex} < t_{sol}$

A solution is required from the current run of the game. It is based on an estimate of the locations of the mobiles t_{sol} seconds in the future, the time at which the desired manoeuvre has been completed. If the method used to solve the game has an execution time t_{ex} seconds then the solution is based on data estimates that are $t_{est} = t_{sol} + t_{ex}$ seconds old.

Increasing the size of the payoff matrix increases the execution time t_{ex} of the solution method. Figure 5.8 plots the algorithm execution time t_{ex} to find a solution for multiple UAVs in 2D and 3D scenarios. The manoeuvre time of $t_{sol} = 312s$ is plotted on the graph for reference. Note the logarithmic scale on the y -axis.

As expected, the execution time increases with the size of the payoff matrix. The increase in execution time appears to be exponential for exhaustive search.

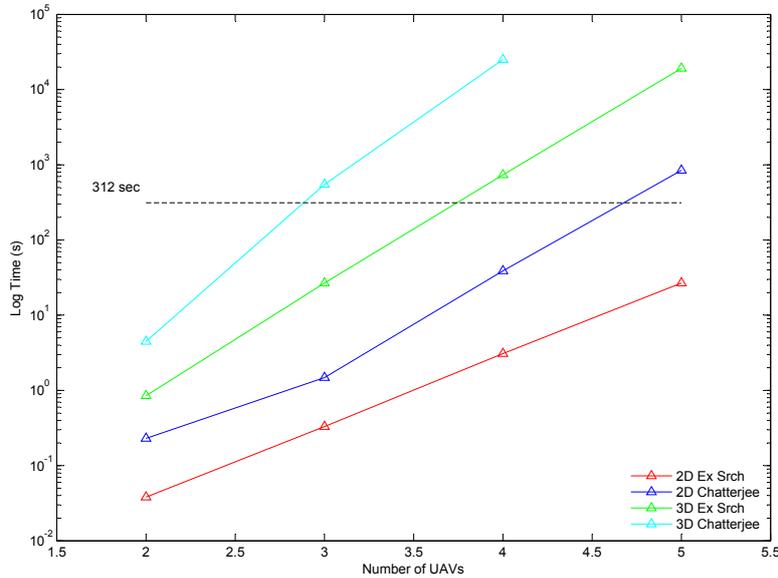


Figure 5.8: Algorithm Execution Time t_{ex} for Multiple UAVs

Chatterjee’s method appears to scale at a rate close to exponential, a result that is predicted in his paper [24].

It is clear that, for up to 4 UAVs, the 2D game is solvable using both exhaustive search and Chatterjee’s method, with exhaustive search being about an order of magnitude faster. With the 3D game both methods can produce solutions for two UAVs, however only exhaustive search can deliver a solution for 3 UAVs within the manoeuvre time.

5.7 Sensitivity to the Age of the Location Data

Clearly increasing t_{ex} will have some impact on the quality of the solutions. The game design fixes the value of t_{sol} , whereas the value of t_{ex} depends on the method used to solve the game.

The impact of execution time on solution quality was quantified by a simple two-UAV experiment that explored the quality of results as the age of the data increased. In this experiment the payoff calculations for the next step were estimated from progressively older data. The payoff matrix was populated with these estimates and the strategies chosen, after which the actual mean coverage was

calculated.

A fast solution method with a fast execution time (Lemke-Howson) was used, but the age of the data t_{est} was changed in multiples of t_{sol} . An example is shown on the timeline in figure 5.9 in which $t_{est} = 4t_{sol}$

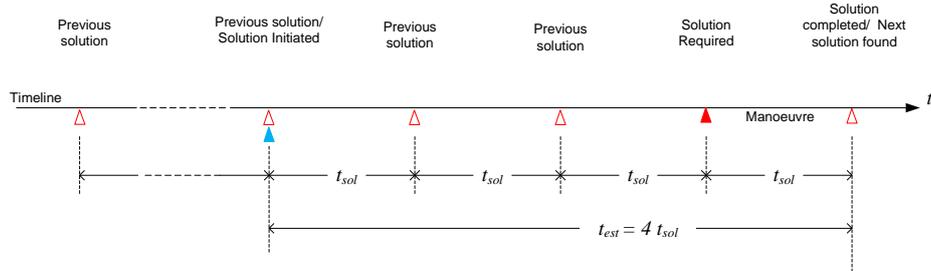


Figure 5.9: Timeline for Evaluating Solution Quality

The results can be seen in figure 5.10. The mean coverage remains consistent at about 85 mobiles until $t_{est} > 7t_{sol}$, at which time it starts to decrease. From time $t_{est} = 13t_{sol}$ coverage appears to stabilise at about 73 mobiles.

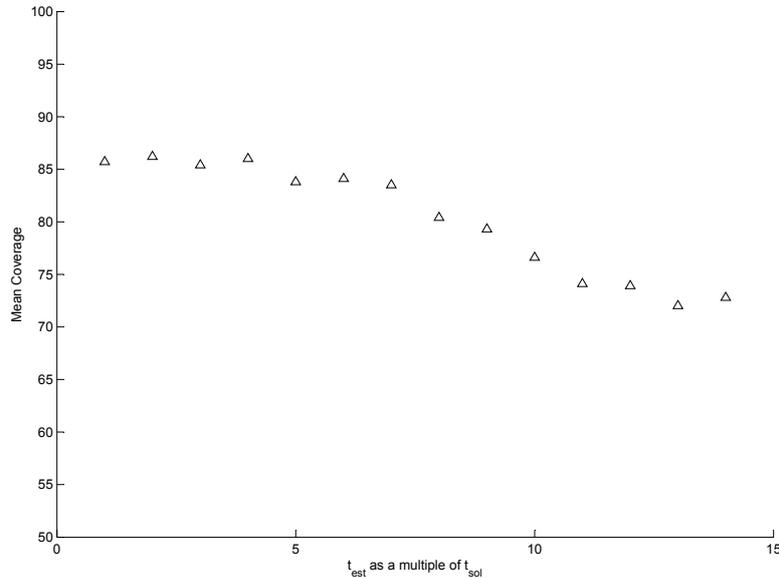


Figure 5.10: Coverage v Data Age, Expressed as a Multiple of t_{sol}

This result suggests that the method could produce useful results with data that is up to $7t_{sol}$ old. In the current game this is over 2000 seconds. The smooth shape of the points in figure 5.10 indicates that the game behaviour was stable, even with $t_{est} = 14t_{sol}$.

It is possible to gain some insight into figure 5.10 by considering a single cluster of mobiles in a single UAV footprint. A footprint-limited UAV at 10000ft has a footprint radius of 22328m, and the distance between cell centres is 4500m. A mobile at the cell centre could, if it travelled in a straight line, reach the edge of the cell in a time of $5t_{sol}$. Similarly, a mobile at the edge of the footprint could, theoretically, reach the centre in the same time. These figures indicate the timescale over which a single cluster could form and disperse.

Figure 5.10 shows that the coverage corresponds with that of a pair of UAVs tracking clusters until $7t_{sol}$. The coverage starts to diminish when the data used for estimating coverage has an greater than $7t_{sol}$. This decrease ends after $12t_{sol}$ when the coverage settles at around the value of associated with circling UAVs.

This shows that the non-cooperative game is good at estimating solutions based on information that is out-of-date, and that the quality of those estimates degrades gently. This is significant as it indicates that a UAV that receives no updates for a period of time will still generate a solution that is compatible with the solution found by another UAV. The algorithm has some inbuilt robustness against temporary loss of communications, and allows UAVs to operate autonomously.

5.8 Number of Solutions

Chatterjee’s method will always find a solution as there is always at least one MSNE. If exhaustive search is to be effective it must find a PSNE every time the game is run, failing which it adopts a “do nothing” strategy resulting in the UAV continuing to circle in the same cell.

The two-player games investigated in previous chapters always found at least one PSNE, therefore there was always at least one move that the UAVs could make to track clusters.

An experiment was run to evaluate the percentage of times that exhaustive search finds at least one PSNE. Table 5.2 shows the results.

As expected there is always a PSNE found if there are two UAVs. If this increases to 3 UAVs then a PSNE is found in 94% of the times that the 2D game is run, but only 7.4% of the occasions when the 3D game is run. This is not a concern for the 2D game as, on most occasions, there is a solution that give a positive decision. The 3D game, however, has the problem that positive decisions are rare and 92.6% of the time the “do nothing” option is taken.

Dimensions	Number of UAVs	Percentage PSNE
2D	2	100.0%
2D	3	94.1%
2D	4	48.5%
3D	2	100.0%
3D	3	7.4%
3D	4	0.0%

Table 5.2: Exhaustive Search: Percentage of Solutions that are PSNE

If there are 4 UAVs the “do nothing” option is adopted in over 50% of cases in the 2D game and 100% of cases in the 3D game. In the latter case, the game is producing no meaningful results and the UAV is circling.

Recall Nash’s statement that every finite game has at least one NE, and that a PSNE is a special case of an MSNE in which one action has a probability of 1 and all others are 0. It is clear that PSNE represents a subset of all possible solutions. As the number of players increases it becomes more likely that one player will have a mixed strategy as his solution, hence the solution to the game cannot be a PSNE. We see this clearly in table 5.2; as the number of players increases the percentage of solutions that are PSNE decreases.

Changing from a 2D game to a 3D game increases the size of the action set for each player from 7 to 21. A larger action set reduces the probability that any pair of players will find best responses that coincide. Table 5.2 shows that, for games with more than two players, the games with the larger action sets have fewer PSNE than games with smaller action sets. Larger action sets give more possible pure strategies for each player; in this case each player has a choice of 21 actions instead of 7.

Combining the previous two paragraphs we can start to understand table 5.2.

- The existence of a PSNE requires that the solutions for all players are pure strategies. If one player has a solution that is a mixed strategy then the solution to the game cannot be a PSNE.
- As the number of players increases it becomes more likely that one of them will have mixed strategy.
- As the size of each players action set increase it becomes more likely that the

best responses of one player will not coincide with those of another player.

This can be illustrated by the simple example in figure 5.11.

Consider a two-player game in which the players all have 2 actions. The best responses to the other players' actions can be stored in an array that has the same dimensions as the payoff matrix. Figure 5.11 gives an example of this in which each player stores the location of his best responses (in red) in a 2×2 array.

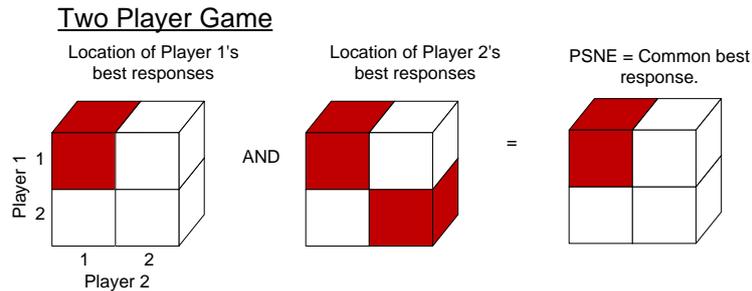


Figure 5.11: Logic for Locating PSNE - 2 Players

Player one's best response to player two's actions is (1,1). Player two has two best responses to player one's actions: (1,1) and (2,2). By definition, a PSNE is the set of actions that are best responses to both players, so (1,1) must be a PSNE.

Consider the three player game in figure 5.12 . Player one's best responses are (1,1,1) and (2,2,2), Player two's best responses are (1,1,1) and (2,2,1), and player three's best responses are (1,2,1) and (2,2,2). It is clear that there are no best responses common to all players, therefore there is no PSNE.

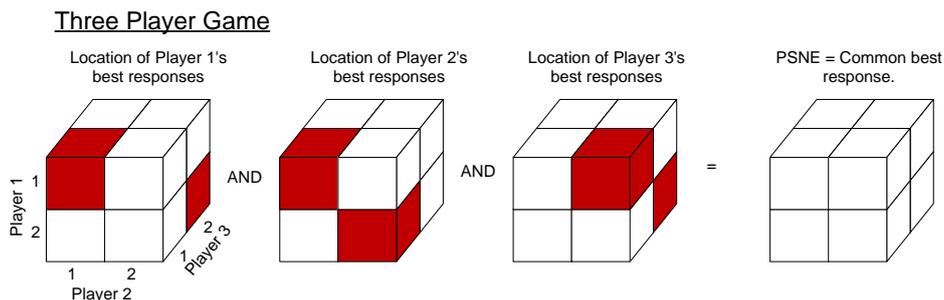


Figure 5.12: Logic for Locating PSNE - 3 Players

The location of PSNEs in an array is generally sparse. As an example, the two-UAV, 2D, game with 7 strategies has 49 possible locations for pure strategies

to exist, but an average of only 2 PSNE every time the game is played. This is mainly due to the low probability of two player's best responses coinciding. As more players are added, and as the array grows, the probability of finding PSNE rapidly diminishes. It is thus unsurprising that the 3 UAV, 3D, game has PSNE on only 7.4% of runs.

Algorithms that only find PSNE have obvious limitations. As the size of the payoff matrix increases the probability of finding PSNE reduces with consequent impact on performance. If, for example, there is a long sequence of games that fail to find a PSNE then the consequent "do nothing" choices would be equivalent to the UAV circling a fixed point. The expected coverage, as shown in chapter 3, should be about 38.8 mobiles per UAV.

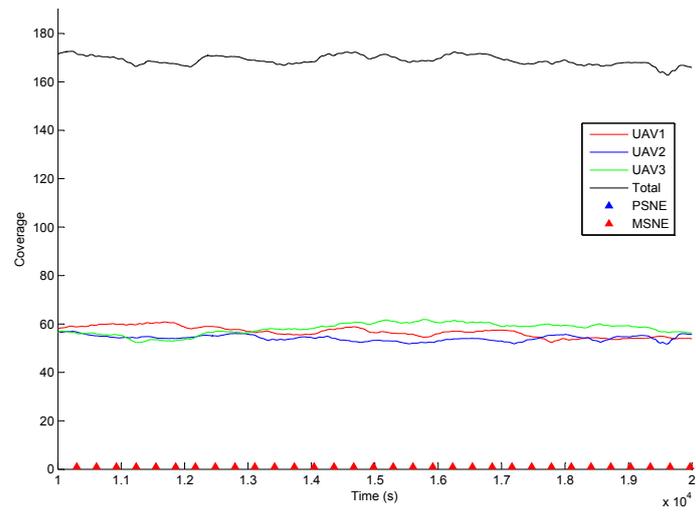
We can see an example of this behaviour in figure 5.13, a game run with 3 UAVs in which the UAVs are allowed to change altitude. Figure 5.13a shows a sample of the data for Chatterjee's method, while figure 5.13b shows the corresponding sample for exhaustive search.

The difference between the samples is immediately apparent. Figure 5.13a shows a constant joint coverage of about 170 mobiles, while figure 5.13b shows a step change in joint coverage from about 110 mobiles to 160 mobiles. Inspection of the tick marks on the x -axis of figure 5.13b shows that the the game runs a long sequence of "do nothing" options, so we would expect the mean coverage over this period to be 38.8 mobiles/UAV. Unsurprisingly, this is about the coverage that we see for the three UAVs.

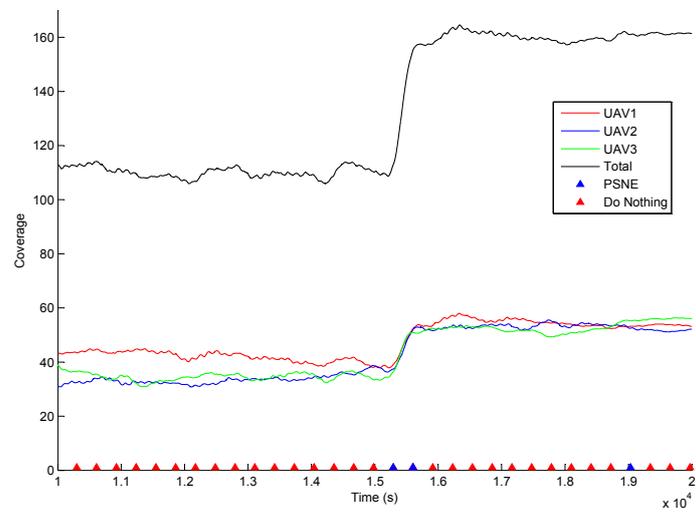
The UAVs continue to circle a fixed point until about 15300 seconds, at which time there are two consecutive PSNE. These trigger the step change in coverage.

The flight paths for the three UAVs at the time of the change in coverage can be seen at figure 5.14b. It can be seen that all three UAVs can obtain better payoffs by climbing and hence increasing their footprint area. The first PSNE requires UAVs 1 and 2 to climb and relocate, while UAV3 climbs in the same cell. The second PSNE requires UAVs 1 and 2 to relocate at the same altitude, while UAV3 continues to climb in the same cell.

The increased coverage of about 160 mobiles continues until the end of the scenario. From figure 5.14b we can see that two of the UAVs have climbed to about 13000 ft, while the third has climbed to 16000ft. Figure 3.8 indicates that a single circling UAV is power limited at these altitudes and has a mean coverage of about 55 mobiles. Three UAVs should thus have an expected coverage of about

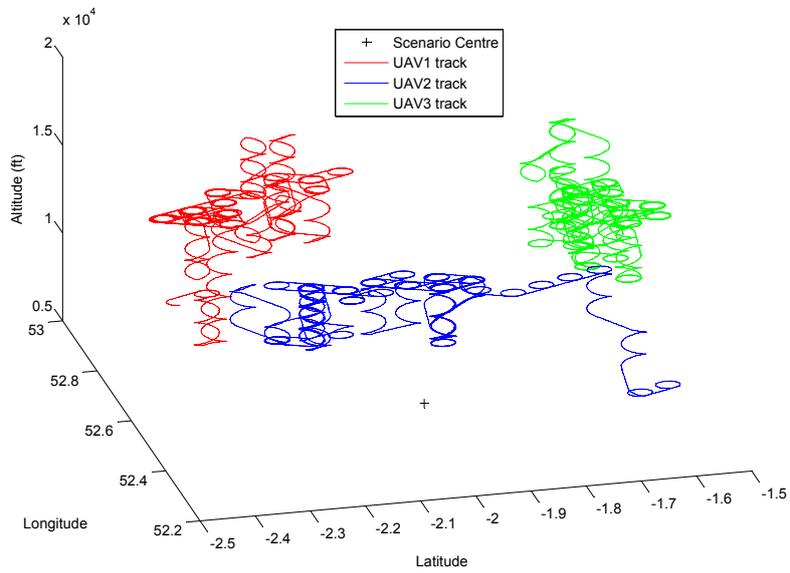


(a) Chatterjee

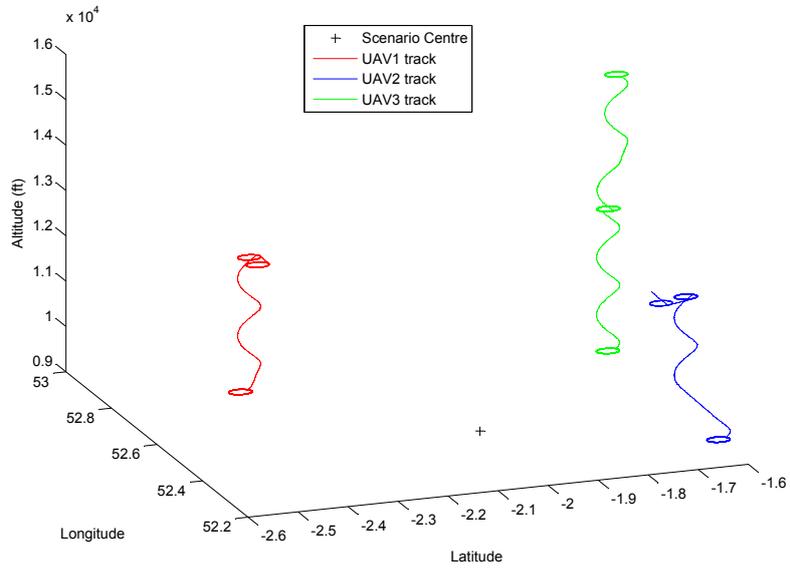


(b) Exhaustive Search

Figure 5.13: Coverage v Time for 3 UAVs



(a) Chatterjee



(b) Exhaustive Search

Figure 5.14: 3D Flight Paths for 3 UAVs

165 mobiles, which is close to the observed coverage of about 160.

We can see that, with larger payoff matrices, exhaustive search is no better than circling. Most of the time the UAV will circle at a fixed altitude unless a particular conjunction of best responses triggers a change. Searches for MSNE, however, generate continuous change in location and altitude, allowing the UAVs to track clusters of mobiles, obtain the best possible coverage, and use power efficiently.

Figure 5.14a shows the flight path for the same mission, but with the three UAVs using Chatterjee’s method to solve the game. It is obvious that each UAV has many changes of location and altitude.

5.9 Coverage with Multiple UAVs

This section considers how coverage is affected by the number of UAVs and the size of the action set for each UAV.

Figure 5.15 plots mean coverage against the number of UAVs for each of the methods. It also, as a benchmark, plots the equivalent coverage for the same number of UAVs if they were circling at a fixed altitude of 10000ft.

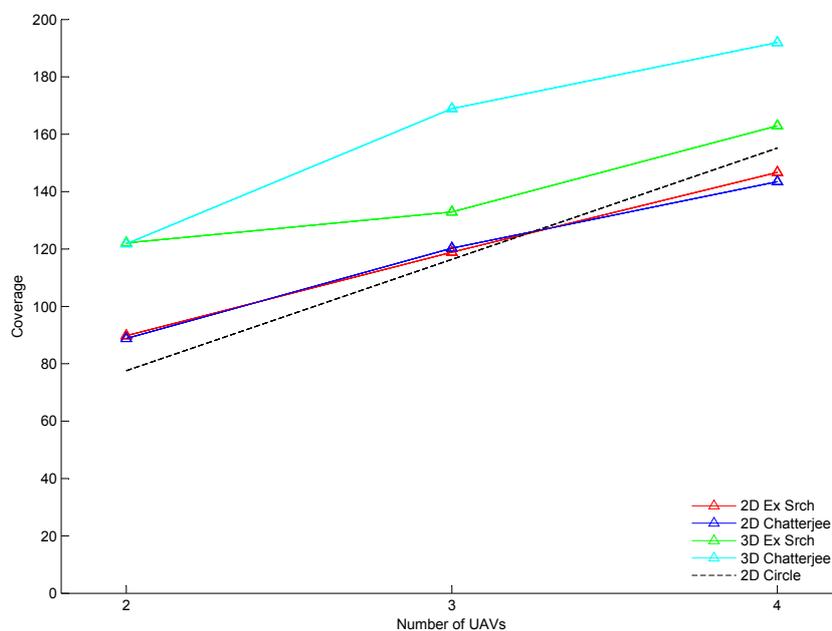


Figure 5.15: Mean Coverage for Multiple UAVs

It can be seen that the UAVs that are operating in two dimensions, in other words relocating at 10000 ft, provide comparable coverage but this diminishes as the number of UAVs increases. Evidence indicates that this is due to the sum of the areas of the footprints becoming a significant fraction of the area of the whole scenario. Whenever the number of footprints exceeds the number of clusters then some UAVs will have smaller coverage.

Movement in 3D is much better than operating at a fixed altitude. UAVs can both relocate in pursuit of clusters, and climb until they enter a power limited state. These two effects should combine to provide very high coverage, as we can see that Chatterjee's method shows this effect in practice. The coverage approaches the maximum achievable coverage of 200 mobiles with 4 UAVs.

Exhaustive search shows good performance with 2 UAVs and improves slowly as the number of UAVs increases. With 4 UAVs the coverage is comparable with 4 UAVs circling at a fixed altitude. Recall that, with 4 UAVs, exhaustive search fails to find any PSNE so the UAVs never change from their initial location. If the UAVs start circling at a fixed altitude then, as no solution causes them to move, their coverage will reflect that of circling UAVs.

5.10 Chapter Conclusions

Manipulating the payoff calculations offers a useful technique for modifying the behaviour of the UAVs. A prioritised access scheme increased the coverage of the priority mobiles by favouring coverage of some UAVs.

It can be seen that churn in the network topology can be reduced by comparing the previous and current connectivity of each UAV. Biasing the payoff matrix to favour continuity of each connection can increase the mean duration of connections, at the cost of reducing coverage. Schemes that reward continuity of connection are more effective than schemes that penalise loss of connections.

One consequence of manipulating payoffs is that mean coverage is reduced. The data infers that the accesses that contribute most to this reduction in coverage occur at the edge of the footprint, and tend to have durations in the range 1-10 seconds.

The choice of frequency can have a significant effect on the total coverage. Path losses increase with frequency so more power is needed to support each mobile. As frequency increases the coverage reduces, and the payload will move from a

footprint limited state to a power limited state.

Increasing the number of UAVs, or actions available to each UAV, results in larger payoff matrices that take much longer to solve. The execution time for exhaustive search increases exponentially with the number of UAVs. The execution time for Chatterjees method increases at about the same rate, but is less predictable as it converges on a solution rather than executing a search. In all cases, Chatterjees method takes longer than exhaustive search to find a solution.

Increasing the dimensions of the payoff matrix exposes a weakness in exhaustive search. The probability of finding any PSNE decreases rapidly with the number of UAVs. Eventually this probability becomes so small that the UAV rarely relocates, and it behaves as if it is circling at a fixed location. The coverage becomes comparable with that of an equivalent number of circling UAVs.

A similar effect is observed with Chatterjee's method in 2D games. As the number of UAVs increases the coverage, initially high, starts to fall. Observation of the data indicates that this occurs when the number of UAVs is greater than the number of clusters. This suggests that the mobility model may be a significant factor in the power limited case.

The 3D game allows the UAVs to transition from payload limited to power limited state. It is clear that the coverage increases with the number of UAVs and converges on full coverage of all UAVs. This reinforces the view that the results from the 2D game are related to the limitations of operating in a footprint limited state.

Underlying all of this data is the message that methods that can find MSNE in multi-player games are the only satisfactory approach to finding good solutions in a scalable game. The limitation to finding such solutions is the timeliness in which they can be found. Research in this area by Govindan and Wilson has started to produce new algorithms, but they are currently unwieldy and uncoded [45]. Chatterjee's algorithm remains the best current option.

This page is intentionally blank

Chapter 6

Conclusions

The research reported in this thesis addresses a new problem: the provision of area communications coverage by UAVs. It also considers a novel approach to determining the best locations for these UAVs by having them compete for coverage of a community of mobile communications nodes.

6.1 Research Challenges

The research challenges identified in section 1.4 were:

1. To define a structure for a suitable non-cooperative game. This needed to reflect realistic behaviour of objects, for example the kinematics of UAVs and mobiles, power limitations of payloads, and the selection of suitable antennae.
2. To investigate possible methods for solving the game, looking particularly at the quality of the solution and the time taken to reach that solution.
3. To consider how the game would be affected by real-world issues. These included choice of frequency band, imperfect knowledge of system state, requirements for some mobiles to have priority access to UAVs and the need for a stable network topology.
4. Understanding how the performance of the game scaled if the number of UAVs, or the size of each UAV's action set, increased.

6.1.1 Defining a Suitable Non-Cooperative Game

The approach was implemented as a non-cooperative game with the UAVs as the players. A grid of hexagonal cells was superimposed on the region, and UAVs were required to circle within each cell. The action set for each UAV was to remain in the current cell or relocate to an adjacent cell. The payoff for each UAV was the number of mobiles that it could support from the current or adjacent cells. This game satisfied the definition of a game as comprising players, strategies and payoffs.

UAVs estimated the number of mobiles that they could support, and that other UAVs could support, taking account of the possibility that two UAVs could be competing to cover the same mobile. Each UAV's payload was constrained by the area of the footprint it could project at the earth's surface, and the RF power that was available to support any mobiles within that footprint. The Nash Equilibrium for the game was the set of strategies that adopted by the UAVs that allowed each UAV to support the maximum number of mobiles within these constraints.

This approach allowed UAVs to determine their best locations by finding a solution to the game. To arrive at a solution they needed a common understanding of the locations of the mobiles and other UAVs but, significantly, they did not need to share or distribute a plan. This gave each UAV the ability to operate with a high level of autonomy.

In non-cooperative games the players need to make their decisions simultaneously. The synchronisation of decisions resulted in the need for all UAV flight paths between cells to have the same duration, a measure that was implemented by ensuring that all flight paths had the same length and that UAVs cruised at the same speed.

A benchmark case was considered of a two-UAV system in which both UAVs circled in different areas of the scenario. The use of UAVs circling a fixed point made a useful benchmark as it required the minimum amount of airspace. Furthermore, the coverage of the UAVs was dictated by the number of mobiles that happened to enter their footprints. In running the game, the UAVS would be allowed to relocate and find clusters of mobiles that would allow them to improve their coverage. The game should always give better coverage than the benchmark; the issue was to quantify the improvement.

6.1.2 Solving the Game

Several methods for solving the games were considered. An exhaustive search of all combinations of actions by all UAVs provided a simple method for searching the payoff matrix for PSNE. If no PSNE existed then the UAVs would not relocate, but would continue to circle in their current cells. This “do nothing” option was not ideal but was a pragmatic way to respond to a game with no outcome. A faster version of the exhaustive search, known as “fast search”, was developed; This searched a path through the payoff matrix until it found a mutual best response by all UAVs. Fast search included the ability to terminate the search before expending more time than an exhaustive search, allowing it to reach a “do nothing” decision if the game had no PSNE.

Mindful that PSNE are a special case of MSNE, methods were considered that allowed MSNE to be found. The first of these was the Lemke-Howson algorithm, a well established method for finding a solution to two-player games. The second of these was a method by Chatterjee that found solutions to games with n players, where $n \geq 2$. If multiple MSNE existed then both methods would find one of those MSNE; initial conditions for the search determined which one was found.

There is an equivalence between normal and extensive forms of a game. The game was thus redefined from a matrix to an equivalent game tree and solved by backward induction.

Experiments showed that all forms of game showed improvements in coverage and power over the benchmark. The improvement in performance of the two-UAV normal-form game was largely independent of the method used to solve the game. Faster methods such as Lemke-Howson were ideal for finding good mixed strategies in two-UAV normal-form games. Extensive form games had fast execution times but did not perform as well as normal form games. The main problem was that the solutions were depended on the order of play, thus coverage was biased towards the first player. Every method used to solve the games results in a moderate improvement in power efficiency, typically of the order of 7%. This improvement in coverage and power efficiency would aid spectrum planners in a real system by reducing the interference footprint.

6.1.3 Practical Implementation of the Game

Investigations based on a two-player game demonstrated that relocation of the UAV facilitated improvements in coverage and specific power when compared to the benchmark. Transferring this method to the real world introduced some practical issues, for example:

1. Consideration of how well the game would perform at other frequency bands
2. Resilience of the method to changes in initial conditions
3. Prioritisation of some users for operational or technical reasons
4. Requirements to provide a stable network topology when all agents are in relative motion

It is assumed in section 3.7 that the system would operate at 5GHz. This choice was seen as a likely frequency band for future systems because of its proximity to the current ITU-R allocation for UAV control. It is possible that other frequency bands could be used if, for example, a military deployment chose to use a dedicated military frequency band. An experiment was conducted to establish how the method performed at frequencies between 1GHz and 10 GHz, these being the likely limits of practical systems. As predicted, the system was in a power limited state at 10GHz and a footprint limited state at 1 GHz. The system transitioned between these states at about 3GHz.

The initial conditions of many experiments placed the UAVs in difference regions of the scenario. In a real system the UAVs would probably operate from the same airfield, and are likely to enter and leave the scenario at the same point. Starting two UAVs at the same location was therefore an interesting, and important, test for the method. Experiments usually showed that the UAVs naturally separated and occupied different regions of the scenario. If both UAVs started in a corner of the scenario then, occasionally, one UAV would leave the scenario rather than compete, an effect referred to as “cornering”. Investigation revealed that this occurred with a predictable set of circumstances. A modification to the method allowed these circumstances to be detected and a countermeasure was implemented in the code.

In real military and security operations some users are given a higher priority than others. In the military communications systems, for example, commanders

are allowed to pre-empt ordinary systems users. In police and emergency services there is a hierarchy of users with Gold, Silver and Bronze Commands having the ability to override other communications users. The ability to support priority access was incorporated into the game by pre-pending a priority vector to the power matrix and modifying the payload packing algorithm. This approach was shown to be effective.

The provision of a stable network topology is essential for IP networks. Short duration accesses and repeated hand-offs between UAVs will generate routing table changes that could destabilise the IP network and adversely affect any applications using the network. Modifications to the method for calculating the payoff matrix were explored. By introducing memory of the previous network topology, and “rewarding” UAVs that supported the same connections for successive runs of the game, it was possible to increase the duration of accesses. These changes also reduced the number of short duration accesses; these caused the greatest problems to network stability.

6.1.4 Scalability

The scalability of games exposes weaknesses in some solution techniques. Exhaustive search for PSNE becomes less reliable as the number of UAVS, or size of their action sets, increases. The only algorithm that successfully generates solutions is Chatterjee’s method, however the execution time of the code makes it impractical for use in this game with more than 3 UAVs, unless the size of the action set is reduce by limiting the UAV’s movement to two dimensions.

Current research in game theory is starting to develop new algorithms for solving n-player games, but currently only Chatterjee’s method has been successfully coded.

6.2 Application of this Work

This research was sponsored by Airbus Group, a leading manufacturer of UAVs, as part of a wider programme of research into UAV applications. The corporate investment in time and resources indicates a significant commitment to this research area.

This thesis is limited to communications missions by two or more UAVs. The

basic algorithm can be applied to other types of mission such as area surveillance by simply changing the method of calculating the payoffs.

Using this algorithm on a real aircraft is not a simple matter. The algorithm runs on the communications controller, a device that does not yet exist, and interfaces directly with the navigation systems. Development and certification of the communications controller will require significant investment. The connection of the communications controller to the navigation system offer a potential security vulnerability, allowing malicious users or third parties to attack this system through the communications payload.

Collision avoidance is not directly addressed in this thesis, but one aspect of this algorithm is that it enforces separation between UAVs. It does not assure separation from other air users, therefore a separate collision avoidance system is still required on the UAV.

6.3 Contributions to Knowledge

This thesis addresses two new ideas:

1. The use of more than one MALE UAV on communications missions is a subject that has seen little research.
2. The application of game theory as a means of coordinating several communications UAVs is a new approach to the problem.

This thesis investigates these two ideas, addresses some of the practical issues of implementing them on real UAVs, and considers the scalability of this method for swarms of more than two UAVs. The research described in this thesis shows that non-cooperative games are an effective way of coordinate a small group of UAVs on area coverage missions.

Future Research

The last decade has seen a steady increase in research into communications UAVs. The main research focus has been on small, short range, systems. These offer an affordable platform for academic research but are largely unrepresentative of MALE UAVs; their mission duration, communications range and low altitude prevent them from flying missions that are representative of MALE UAVs. There is almost no current research into MALE communications UAVs, but the UAV industry has started to recognise and address this shortfall. The high cost of manufacturing and flying MALE UAVs means that research is likely to be based on simulations for some years. Eventually, as the flying platforms become available and algorithms start to mature, industry will start test flying real systems.

This thesis represents the start of this development process and presents a candidate approach for coordinating several communications UAVs. It is unclear whether the game described in this thesis is the best method for solving this problem. Classical optimisation techniques such as evolutionary algorithms (EA) could provide comparable results. Discussions during this research have resulted in collaboration with a post doctoral researcher at Aberystwyth University to compare this with an EA approach. An initial joint paper outlining the EA method has been presented [40]. A further paper [41] presented initial results from this research, showing that the EA gave marginally better results than the game but at the cost of increase processing time and dependence on a central planning agent. Work is ongoing to write a more comprehensive journal paper comparing EA and game approaches to the problem.

The radio spectrum issues have not been addressed in this research. The as-

sumption that all communications occurs in the 5GHz frequency band ignores the problem of frequency planning and interference. As part of the wider research around this topic, the author has been considering game theory as a technique for dynamic frequency planning. Two MSc student projects have been supervised. The first, in 2012/13, looked at possible techniques for real-time frequency management in mobile networks. The second, in 2014, has started to address the scalability of the problem by decomposing the larger game into many small, and therefore quickly solved, subgames. Ongoing research by the author is looking at extensive form games as a method for solving subgames.

This topic, the use of UAVs to provide area communications cover, and this approach, non-cooperative games, are new areas of research. Many problems remain to be addressed, and some new problems will undoubtedly be identified as this research continues. One example is the design of the security interface between the Communications Controller and Navigation System. Another is the integration of this algorithm with collision avoidance systems to allow UAVs running this, or similar, algorithms to share airspace with other air users.

The growing commercial and government interest in communications UAVs is already adding urgency to the resolution of these problems. Solutions are starting to emerge for some problems, for example spectrum allocation and airspace regulation. There is still considerable scope for innovative research into the applications of UAVs to real problems. This area of research should remain lucrative for many years.

Appendix **A**

Introduction to Game Theory

Gibbons defines game theory as “the study of multiperson decision problems” [42]. Game theory has its origins in economics where it is used to model the behaviour of conflicting or collaborating parties [91]. It offers a toolset for analysing interactions between competing players that has been successfully applied to problems in evolutionary biology, politics, control systems, and numerous other disciplines. In recent years it has successfully been used to analyse communications systems and is becoming a mainstream technique [47].

This appendix addresses a class of game in which *rational* players make decisions that will maximize their share of a resource, cognizant of how the other players are expected to behave while maximizing their share of the same resource. Such games are *non-cooperative* as the players do not cooperate or form coalitions to find a solution, and the players have *perfect information* as they are all able to determine each others’ strategies and payoffs. Such behaviour should lead to a division of the resource such that no player can increase his share. This state is known as a *Nash Equilibrium* (NE).

This appendix introduces the basic concepts of game theory which are applied in this thesis. The notation is derived from that used by Webb [94] as his notation allows discrimination between the strategy and the choice of strategy, an important distinction which is overlooked in some texts.

A.1 Key Concepts

A game can be comprises a set of players \mathbb{N} , a set of available strategies \mathbb{S} and a payoff functions Π . A game G can thus be defined as:

$$G = \{\mathbb{N}, \mathbb{S}, \Pi\} \quad (\text{A.1})$$

The set of players can be defined as $\mathbb{N} = \{1, 2, \dots, j, \dots, N\}$. The strategy set \mathbb{S} can be defined as the Cartesian product of the strategy sets for all the players $\mathbb{S} = S_1 \times S_2 \times \dots \times S_j \dots \times S_N$. Player j has a strategy set S_j where $j \in \mathbb{N}$. The payoff function for player j , denoted π_j , maps the strategy to some real numerical value that represents the consequences of that strategy. Thus $\pi_j : \mathbb{S} \rightarrow \mathbb{R}$.

It is possible to define an optimal strategy set S^* for all players such that:

$$\pi(S^*) \geq \pi(S) \quad \forall s \in \mathbb{S} \quad (\text{A.2})$$

Note that S^* does not have to be unique, thus there is a weak inequality.

The set of k strategies for player i can be defined as:

$$S_j = \{s_1, s_2, s_3 \dots s_j, \dots s_k\} \quad (\text{A.3})$$

where each strategy is a vector:

$$\begin{aligned} s_1 &= (1, 0, 0 \dots 0) \\ s_2 &= (0, 1, 0 \dots 0) \\ s_3 &= (0, 0, 1 \dots 0) \\ s_k &= (0, 0, 0 \dots 1) \end{aligned}$$

These can be interpreted as s_1 implies that the player chooses action 1 out of a possible k actions. This is referred to as a *pure strategy*.

There can be occasions when there is a probability distribution across the pure strategies. known as a *mixed strategy*. This is normally represented as a vector σ_i of probabilities across the pure strategies. For example:

$$\sigma_j = (p(s_1), p(s_2), \dots p(s_k)) \quad (\text{A.4})$$

It can be seen that the pure strategy is a special case of a mixed strategy where the probability of using one strategy is 1, and the probability of all others is 0.

To illustrate this, consider the game of rock-paper-scissors. Each player has three pure strategies: rock, paper and scissors. Over multiple games a player can never win outright, but can never lose by playing a mixed strategy with a distribution $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$

A.2 Best Response

Each rational player will select a strategy that maximizes his own payoff, given the strategies selected by other players. Considering a simple two-player game where player 1 has a set of mixed strategies Σ_1 , we can define a strategy $\hat{\sigma}_1$ that is player 1's best response to player 2's strategy σ_2 as:

$$\hat{\sigma}_1 \in \arg \max_{\sigma_1 \in \Sigma_1} \pi_1(\sigma_1, \sigma_2)$$

Similarly, strategy $\hat{\sigma}_2$ that is player 2's best response to player 1's strategy σ_1 if:

$$\hat{\sigma}_2 \in \arg \max_{\sigma_2 \in \Sigma_2} \pi_2(\sigma_1, \sigma_2)$$

This concept of best responses can be extended to any number of players.

A.3 Nash Equilibrium

The *best response* by each player to every other player leads to an equilibrium state as there is no incentive for any player to deviate from his selected strategy; such a move would reduce that player's payoff. This set of strategies is a Nash Equilibrium.

The Nash Equilibrium is the most popular solution concept for games. The aim is to identify the combination of strategies that each player would choose that lead to a NE.

In a two-player game the Nash Equilibrium can be defined as a pair of strategies (σ_1^*, σ_2^*) that satisfy the conditions:

$$\begin{aligned} \pi(\sigma_1^*, \sigma_2^*) &\geq \pi(\sigma_1, \sigma_2^*) & \forall \sigma_1 \in \Sigma_1 \\ \pi(\sigma_1^*, \sigma_2^*) &\geq \pi(\sigma_1^*, \sigma_2) & \forall \sigma_2 \in \Sigma_2 \end{aligned}$$

An equivalent version of these conditions states that the Nash Equilibrium exists when all players best responses coincide:

$$\sigma_1^* \in \arg \max_{\sigma_1 \in \Sigma_1} \pi_1(\sigma_1, \sigma_2^*)$$

$$\sigma_2^* \in \arg \max_{\sigma_2 \in \Sigma_2} \pi_2(\sigma_1^*, \sigma_2)$$

Nash's theorem predicts that every finite game has at least one equilibrium, giving rise to the possibility of multiple NE. Games in which there are multiple NE, and in which all players are working towards a common objective, are called *coordination games*. Coordination games are non-cooperative; the players have a common objective but do not form teams or coalitions to achieve that objectives.

In practice the sets of mixed strategies, Σ_1 and Σ_2 can be very large. This can lead to practical difficulties in as finding mixed strategies involves a very large search space. A useful survey of techniques for finding NE can be found at [72].

A.4 Games with Multiple Players

When there are multiple players a truncated form of notation is often used. The symbol σ_j is used for the mixed strategy of player j , and the symbol σ_{-1} is used to represent the strategies of all other players. As a trivial example, in a three player game:

$$\pi(\sigma_1, \sigma_{-1}) \equiv \pi(\sigma_1, \sigma_2, \sigma_3)$$

We can thus define the Nash Equilibrium of a multi player game as:

$$\sigma_j^* \in \arg \max_{\sigma_j \in \Sigma_j} \pi_j(\sigma_j^*, \sigma_{-j})$$

A.5 Extensive and Normal Forms

Non-cooperative games can be represented and analysed in two ways: as a graph or as a matrix. Games that are represented as graphs are known as *dynamic games* as each player moves in turn. They are represented in *extensive-form* as directed graphs. Games that are represented as matrices are said to be in *normal-form*.

A.6 Extensive-Form Games

The extensive-form of a three-player game can be seen in Figure A.1a. In this game player one calculates his payoffs for each possible action. Player two then calculates his payoffs for each of his actions, given the outcomes of player one's actions. This process continues until all players have calculated their payoffs given all the actions and payoffs of all previous players.

The extensive-form states the order in which players choose their actions. Figure A.1a shows player one moving first. Player two, knowing which action player one has chosen, can choose the action that gives his best response. Finally player three chooses his best response to the actions chosen by players one and two.

A.7 Normal-form Games

The normal-form of a game comprises a matrix with the same number of dimensions as the number of players. The actions of each player form the edges of the matrix. Each cell of the matrix corresponds to a particular set of actions by all players and contains the payoffs for each player in order. An example of a normal-form game matrix with three players, each having two actions, can be seen in Figure A.1b.

In the normal-form it is assumed that all players make their decisions simultaneously, thus they have no knowledge of any events that may already have occurred. The normal-form of the game in Figure A.1b does not indicate the order in which players make their decisions; it only recognizes the actions that each player can choose and the associated payoffs.

Solutions to a normal-form game require the analysis of the matrix. At its simplest, this can take the form of an exhaustive search of the matrix to identify the existence of any PSNE. If, for example, player i selects strategy s_i , then the remaining strategies for all other players, usually denoted s_{-i} lie in the hyperplane of the payoff matrix associated with player s strategy i . This hyperplane can be searched for each player's best responses.

Exhaustive searches need to evaluate all combinations of action by all players. This approach scales exponentially with the number of players. While cumbersome with large n-matrices, this can be acceptable with small numbers of players and actions. The weakness of exhaustive search is that it will only find PSNE. If where

there are only MSNE then exhaustive search will yield no solution. Exhaustive search can also reveal multiple NE, bringing the problem of how to selected the “best” solution of the many possibilities.

An MSNE implies that the player’s strategy is a probability distribution for all the action set. The player does not pick the action with the highest probability but generates a random number that is used to select one action from those with non-zero probabilities.

Techniques for finding MSNE are inevitably more complex than for finding PSNE. Futhermore, finding a single MSNE offers no assurance that the solution is efficient.

A.8 Methods for Solving Games

A solution to the game comprises a Nash Equilibrium and its associated payoff. In this thesis the payoff comprises the maximum individual coverage for each UAV and, by implication, the joint coverage for two or more UAVs.

Five methods for solving games were considered in this research:

1. Exhaustive search
2. Fast search
3. The Lemke-Howson algorithm
4. Chatterjee’s method
5. Backward induction

These methods, and their key properties, are briefly described in the following sections.

A.8.1 Exhaustive Search

Exhaustive search is based on the “Colonel Blotto” game, widely described in introductory game theory textbooks. It conducts an exhaustive search of all combinations of pure strategies to find the best response(s) by each player to the strategies played by the other player(s). The results of each search are compared and, if a set of strategies of is a best response for all players, then it must be a

pure strategy Nash Equilibrium (PSNE) by definition. If, on the other hand, there are no common best responses by all players there cannot be a pure strategy Nash Equilibrium and the result is discarded.

If multiple pure strategies are found then, following best practice on coordination games, one was chosen at random. Exhaustive search has one pass through all the strategies for each player. A practical difficulty occurs if no PSNE exists as the game has no apparent solution. Nash assures us that there must be at least one solution, so the absence of a PSNE means that the solution must be an MSNE. Exhaustive search is used, and no PSNE is found, then convention is to discard this run of the game and adopt a “do nothing” strategy. In the game described in section 4.2 the “do nothing” strategy instructs the UAV to stay at its current location by selecting strategy 1.

Exhaustive search can be used in general n player games. A flowchart of the exhaustive search can be seen in figure A.2a.

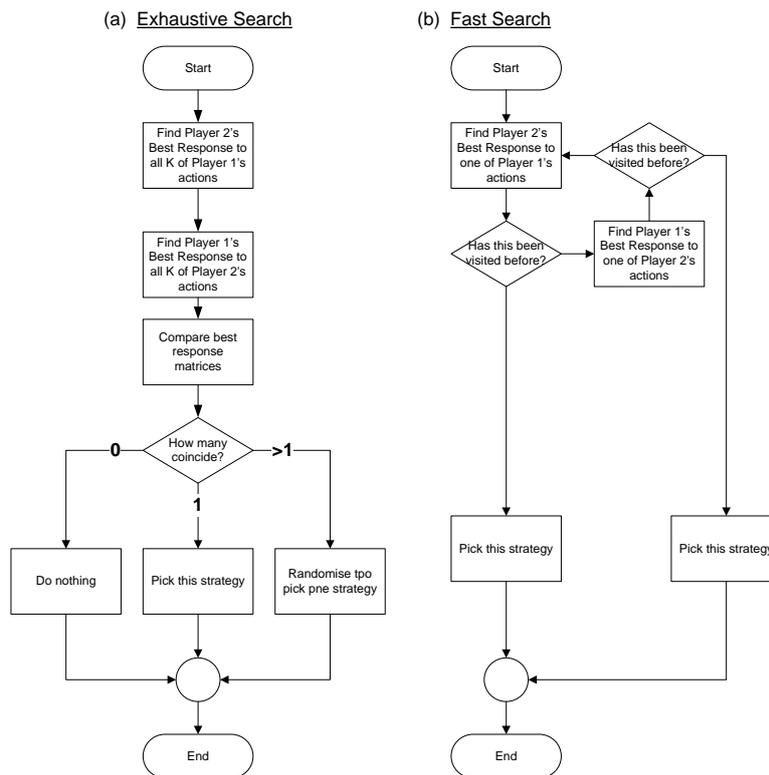


Figure A.2: Exhaustive and Fast Search Methods

A.8.2 Fast Search

A faster method for finding a single PSNE was devised by the author. This starts by selecting a player and strategy at random. It then searches for the other player's best response to that strategy and uses that as its start point. It continues to use alternate players' best responses as the start point until no better response is found. It is necessary to set termination criteria against the possibility that no PSNE is found as, in these circumstances, the algorithm enters a loop. The loop will be terminated if any of the following criteria are satisfied:

1. A PSNE is found.
2. Any strategy is revisited, as this implies a loop.
3. All strategies have been visited once; an exhaustive search.

The fast search algorithm is invariably faster than exhaustive search and, if a pure strategy exists, then it will be found. Fast search, like exhaustive search, adopts a "do nothing" strategy if no PSNE is found.

A flowchart of the fast search can be seen in figure A.2b.

A.8.3 The Lemke-Howson Algorithm

One method of avoiding the "do nothing" option is to search for MSNE. Mindful that PSNE are a special case of MSNE, any method that searches for MSNE will also find PSNE.

The Lemke-Howson algorithm [54], is the most commonly used method for solving two player games. It is essentially a pivoting algorithm that is similar to the simplex algorithm. The algorithm will find a single MSNE out of all possible MSNE.

Lemke-Howson can be run multiple times, each with different initial conditions, in an attempt to find more than one MSNE. Each pivot defines a start point for the algorithm, and each start point uniquely maps to a solution. Because the number of possible pivots is the sum of the number of strategies for each player number of possible MSNE that the algorithm can find is bounded by the sum of the number of pure strategies for the two players.

A.8.4 Chatterjee's Method

Another commonly used method for finding MSNE was devised by Chatterjee [24]. This method expresses the problem as a series of polynomial objective functions, the optimal solution to which is equivalent to the NE.

Chatterjee's method has become popular as it is easily coded and, unlike Lemke-Howson, is able to solve games with more than two players. It is used in this research for solving n -player games.

The initial conditions for this method comprises an assumed MSNE. In practice this is a randomly generated probability distribution for each player. Chatterjee's method, this probability distribution is found by iteratively solving a set of non-linear equations.

The number of iterations before a solution is found, and hence the execution time, is indeterminate. Chatterjee's paper, based on a specific PC configuration, show that the number of players and strategies per-player have a strong influence on execution time. Increasing either factor will significantly increase execution time.

The potentially infinite number of initial conditions suggests that, given enough time, Chatterjee's method could find all existing MSNE.

A.8.5 Backward Induction

The sequential nature of extensive-form games allows solutions to be found using backward induction. The general method is to work backwards from the last player in the sequence, and decide which strategy he would have chosen based on the available payoffs. By progressively eliminating strategies that each player would not choose, the remaining paths from root to leaves are the only choices that rational players would take. Kuhn's theorem [38] leads to the conclusion that the solution found by backward induction is always a PSNE.

Backward induction scales linearly with the number of players and the number of actions available to each player.

Appendix **B**

Technical Principles and System Models

B.1 Section Summary

This appendix describes the principles, models and assumptions used in this thesis. The sections are:

1. An overview of the communications system
2. Link budgets and associated models
3. Traffic models
4. Kinematic and Navigation Models for UAVs

B.2 Communications Overview

The communications payload comprises a radio receiver, transmitter, antenna and communications control system. The communications control system uses the downlink for communicating with the mobiles through short messages to manage power and access control. The structure of these messages is outside the scope of this thesis. It is assumed, for simplicity, that all control messages are correctly received.

The system uses a closed loop power control system between the payload and the mobiles. The comms controller calculates the power required to support each mobile and decides, based on a set of packing rules, which mobiles it will support and what power they require. The power calculation is based on the known characteristics of the signal, and the estimated E_b/N_0 of the received signal.

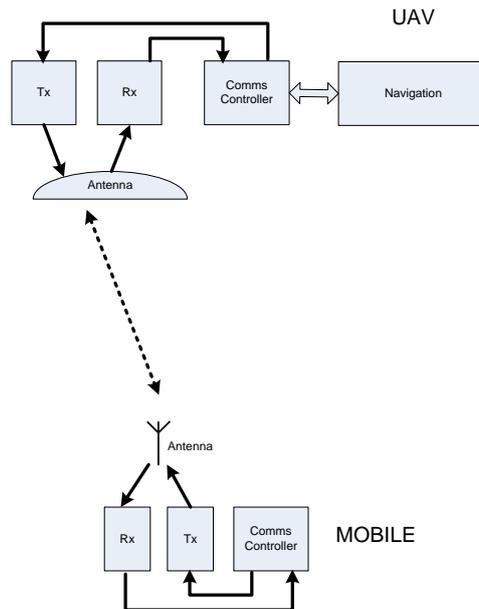


Figure B.1: Power and Access Management

The payload determines whether the link can be supported from knowledge of the location of mobile, its current E_b/N_0 , and an estimate of the required power P_t . If P_t cannot be supported by the payload then the power control message to the mobile can be modified to set P_t to 0 Watts, denying service to the mobile.

The antenna models are deliberately simple. Experiments showed that the most practical antenna system for area coverage is a uniformly illuminated conical horn. This type of antenna is used on the UAV, with the beamwidth θ_h . The antennae on the mobiles are assumed to have hemispherical cover. This removes any tracking requirements from the mobiles.

If the horn antenna is rigidly aligned to the z axis the antenna footprint will move with the airframe. This is desirable with horizontal flight as movement of the aircraft will move the antenna footprint and hence the payload's coverage. Climbing, descent or banking turns will offset the footprint by some angle, as

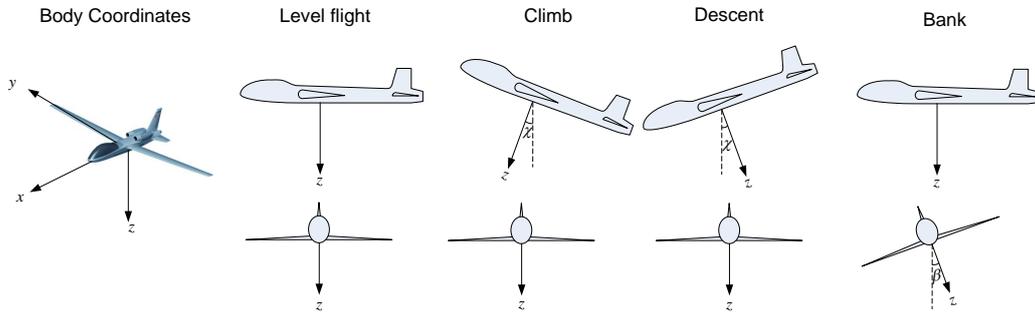


Figure B.2: UAV Attitude

shown in Figure B.2. Controlling this requires that a limit is placed on the bank, climb and descent angles. This is explored further in section B.5.

B.3 Link Budgets

Each communications link between UAV and a mobile is treated as independent from all others, thus a link budget can be created for each path. For any link the transmitter and receiver characteristics, including frequency data rate and modulation, are considered to be fixed. A simple link budget model is used to calculate the RF power required to support any single link through the UAV at this quality [61]. The total RF power presented to the antenna is simply the sum of the powers of the individual communications links.

It is possible to determine the slant range from link geometry, and hence the free space path loss can be calculated. All the elements are in place to calculate a dynamic link budget in clear sky conditions.

The gains, losses and noise temperatures used in this model are shown in Figure B.3.

The required link signal to noise ratio is expressed in a normalized form to provide a measure of independence from the choice of modulation scheme. It is normalized to 1 Hz of bandwidth and 1 bit/s, and is expressed as the ratio of the energy in one bit E_b to the noise in 1 Hz N_0 .

Starting from the Friis transmission equation, it can be shown that the ratio $\frac{E_b}{N_0}$ can be calculated from the equation:

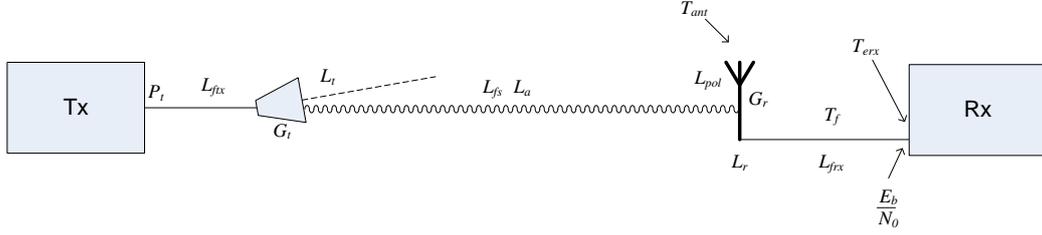


Figure B.3: Link Budget Model

$$\frac{E_b}{N_0} = \frac{\left(\frac{P_t G_t}{L_{FTX} L_t} \frac{1}{L_{FS} L_a} \frac{G_r}{L_{FRX} L_r L_{Pol}} \right)}{\left(\frac{T_{Ant}}{L_{FRx}} + T_f \left(1 - \frac{1}{L_{FRx}} \right) + T_{ERx} \right)} \frac{1}{KRb} \quad (\text{B.1})$$

The terms in equation B.1 are:

- G_t is the boresight antenna gain for this transmitter
- G_r is the boresight antenna gain for this receiver
- L_{FTx} is the transmitter feeder loss.
- L_t is the transmitter pointing loss
- L_{FS} is the free space loss
- L_a is the atmospheric loss
- L_{FRx} is the receiver feeder loss
- L_r is the receiver pointing loss
- L_{Pol} is the polarisation loss, normally assigned to the receiver.
- T_{Ant} is the receiver antenna temperature (Kelvin)
- T_f is the receiver feeder temperature (Kelvin)
- T_{ERx} is the effective receiver temperature (Kelvin)
- P_t is the transmitter power for this transmitter which is intended for the target receiver (W)

- K is Boltzmanns constant 1.38×10^{-23} (J/K)
- R_b is the information rate on the link (bit/s)

Note that all gains and losses are real numbers which are ≥ 1 . For the purpose of this research we can assume that some of these losses can be ignored. It is assumed that no energy is lost in the transmitter and receiver feeders as the antennae are perfectly coupled, hence $L_{FTx} = L_{FRx} = 1$. The communications model assumes an omnidirectional antenna on the mobiles, thus $L_r = 1$. The model ignores polarization losses as these will depend on the implementation of the system, thus $L_{Pol} = 1$. Finally, it is assumed for simplicity that the atmospheric losses can be ignored, thus $L_a = 1$.

The system noise temperature T_{sys} can be simplified because some the losses have been set to unity:

$$T_{sys} = T_{Ant} + T_{ERx}$$

The receiver antenna on the mobile is assumed to have a hemispherical polar pattern and a gain $G_r=1$.

The model for the UAV's transmitter antenna profile L_t will depend on the antenna selected for the UAV. Earlier experiments showed that the optimum antenna for the payload is a conical horn. The antenna radiates its signal in a conical region with half-angle θ_h and that no signal is radiated outside this cone. If the off-boresight angle θ_k is less than θ_h a signal is received, otherwise no signal is received, as shown in Figure B.4.

Such an antenna can be defined by the piecewise function:

$$L_t = \begin{cases} 1, & \text{if } \theta_k < \theta_h \\ \infty, & \text{if } \theta_k \geq \theta_h \end{cases}$$

The gain of the transmitter antenna is a function of the beamwidth. It can be shown that the gain of a conical horn antenna with efficiency η and a half angle χ_h is:

$$G_t = \frac{2\eta}{1 - \cos \theta_h}$$

The free space loss L_{FS} for a path of length dm at a frequency f Hz can be calculated from the well-known equation:

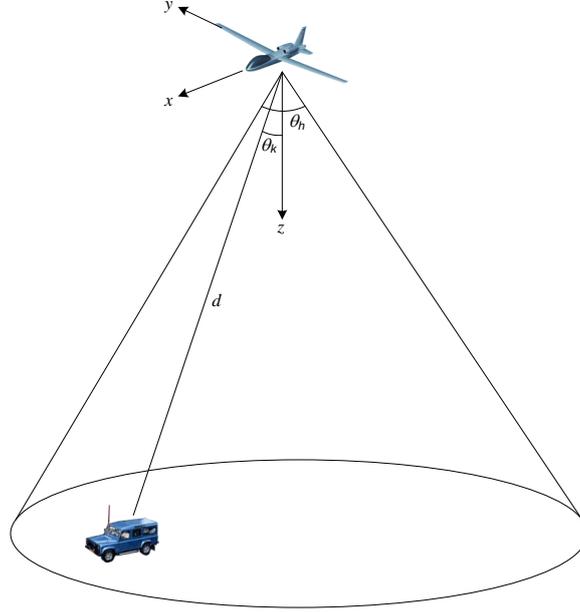


Figure B.4: Antenna footprints

$$L_{FS} = \left(\frac{4\pi f d}{c} \right)^2$$

In this equation c is the speed of light, 3×10^8 m/s.

The required signal to noise ratio can be normalized as $\frac{E_b}{N_0}$, where E_b is the signal energy in one bit. If we have a signal with a fixed data rate and frequency we can rearrange equation B.1 to calculate the power required to provide a specified $\frac{E_b}{N_0}$. This gives a simplified equation:

$$P_t = L_t d^2 \left(\frac{E_b}{N_0} \frac{1}{G_r G_t} \left(\frac{4\pi f}{c} \right)^2 T_{sys} R_b K \right) \quad (\text{B.2})$$

The slant range d can be calculated from the relative positions of the UAV and mobile. For accuracy, this is calculated in spherical coordinates and not over a “flat” earth.

This research considers the UAV resources which, for communications links, comprises the RF power required to support two-way communications. In a practical system most of the power will be used by the power amplifier PA in the UAV, and a small amount of power will be used by the receivers. The power consumed

by the PA will be changed by the automatic power control system to overcome the changing path losses due to the relative movement of the UAV and mobile nodes. The power consumed by the receivers is expected to be relatively small and constant. It is thus assumed that the power consumed by the receivers will be negligible and can be ignored, thus the research need only consider the transmitter power.

B.4 Traffic Models

The calculation of link budgets requires that the data rates for each link must be known. The data rates used in this series of experiments are:

- UAV-Mobiles: $R_b=2$ Mbit/s symmetric links
- Inter-UAV Links (IUAVL): $R_b=20$ Mbit/s symmetric links

The default data rates can be seen in Figure B.5.

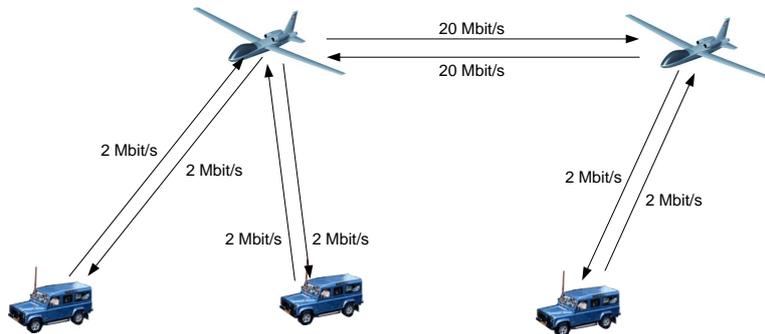


Figure B.5: The Data Rates R_b Used in these Experiments

For the majority of experiments the data rates were held constant over time. The use of continuous transmissions minimized any uncertainty of the locations of the mobiles as all location estimates were current.

B.5 Kinematic and Navigation Models for UAVs

The modeling of aircraft motion and attitude can be complex as it implies simulation in 6 degrees of freedom. For the purpose of these simulations, whose emphasis is on communications, the kinematic models of the UAV have been simplified to

two independent kinematic models, one for motion in the horizontal plane and one for motion in the vertical plane.

These models assume that manoeuvres can be decomposed into simple turns of fixed radius r_h for horizontal manoeuvres and r_v for vertical manoeuvres.

A UAV in horizontal flight has equal and opposing lift and weight forces, denoted L and W . To execute a horizontal or vertical manoeuvre requires the addition of a centripetal force F . We can define the wing's load factor n as the ratio of the lift to the weight $n = \frac{L}{W}$. In horizontal flight $n = 1$. In any type of manoeuvre $n \neq 1$, for example in a horizontal turn $n > 1$ and the aircraft is pulling positive "gee".

Horizontal motion

Consider the UAV in a horizontal turn of radius r_h at a bank angle β , as shown in FigureB.6

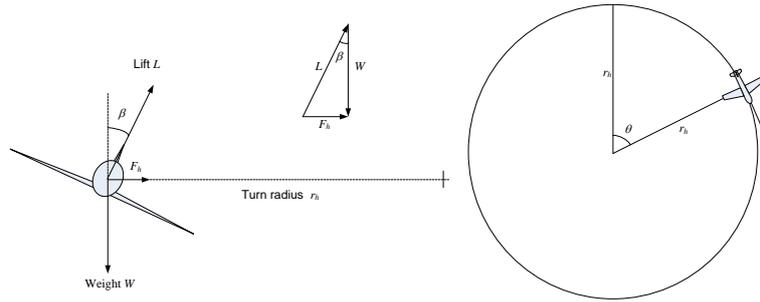


Figure B.6: Banking Turn

For a constant rate of turn the lift and weight components must be balanced so $L \cos \beta = W$. We already know that $n = \frac{L}{W}$ so we can relate the angle of bank to the wing load factor:

$$n = \frac{1}{\cos \beta} \quad (\text{B.3})$$

The resulting centripetal force can be calculated from:

$$F_h = \sqrt{L^2 - W^2} = W \left(\sqrt{n^2 - 1} \right)$$

The weight of the UAV is its mass m multiplied by the acceleration due to gravity g . We can thus write:

$$F_h = mg \left(\sqrt{n^2 - 1} \right)$$

This gives some insight into the meaning of the equation. The term $\sqrt{n^2 - 1}$ can be considered as a multiplier for the acceleration due to gravity. The multiplier increases as the centripetal force increases, thus a tight turn requires high centripetal force and generates high wing load factor.

By definition we can relate the rate of turn $\dot{\theta}$ to the turn radius and velocity. Manipulation shows that:

$$\dot{\theta} = \frac{v}{r_h} = \frac{g\sqrt{n^2 - 1}}{v}$$

The term $\dot{\theta}$ is the rate of change of the angle θ , thus:

$$\dot{\theta} = \frac{\Delta\theta}{\Delta t} = \frac{g\sqrt{n^2 - 1}}{v} \quad (\text{B.4})$$

The angular rate of turn is thus a function of the velocity, wing load factor and acceleration due to gravity g .

By setting a maximum load factor we can calculate the maximum bank angle from equation B.3.

Equation B.4 allows us to calculate the angle $\Delta\theta$ through which the UAV turns in time Δt .

B.5.1 Vertical motion

A similar approach can be used to calculate the vertical rate of turn.

The system in Figure B.7 shows that the centripetal force required to initiate a pull-up is:

$$F_h = L - W$$

We can show that the radius of the pull-up turn r_v is:

$$r_v = \frac{g(n - 1)}{v}$$

Thus the angular rate of pull-up is:

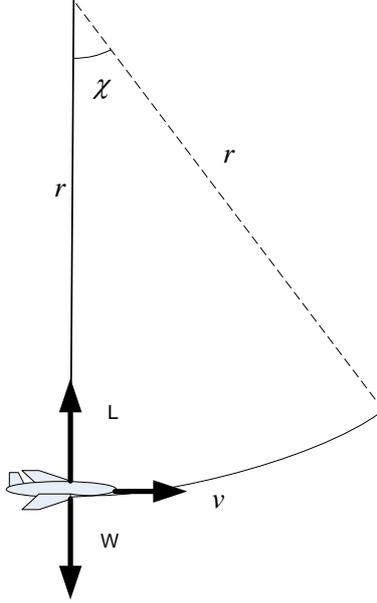


Figure B.7: Pull-up

$$\dot{\chi} = \frac{\Delta\chi}{\Delta t} = \frac{g(n-1)}{v} \quad (\text{B.5})$$

The UAV has a maximum rate of climb ρ . This is usually defined in feet/min but for the purpose of this research it is in SI units. We can define the climb angle κ as:

$$0 \leq \kappa \leq \arcsin\left(\frac{\rho}{v}\right)$$

B.5.2 Change of Latitude and Longitude

At the start of a manoeuvre a UAV's location can be defined as a latitude, longitude and altitude $(\varphi_1 \lambda_1 h_1)$. The UAV also has a velocity vector that comprises a speed v and a heading θ_1 . On completion of the manoeuvre the UAV has a new latitude, longitude, altitude and heading $(\varphi_2 \lambda_2 h_2 \theta_2)$.

The manoeuvre is constrained by the kinematics of the UAV, for example its turn radius for any given angle of bank and speed. The shortest path between the two locations and headings is the Dubins path [34]. A Dubins path comprises two sections of constant curvature and a straight section. Further information on

Dubins paths can be found in [88]. An example of a Dubins path is at Figure B.8.

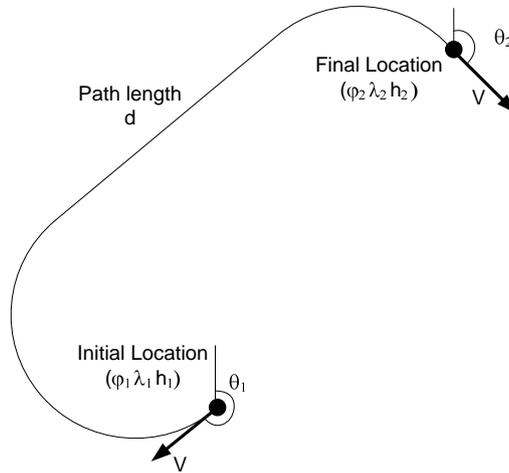


Figure B.8: Dubins Path Example

The total path has a length d that can be approximated as a sequence many short straight line path segments of length Δd . The start of each path segment is a waypoint with known latitude, longitude, altitude and heading. Simple geometry allows the next waypoint to be estimated as each of these straight line sections has a duration Δt and we know the average velocity v of the UAV and its rate of turn ρ ° s.

If the UAV starts at waypoint a and travels to waypoint b in time Δt , then the distance travelled is:

$$\Delta d = v\Delta t \quad (\text{B.6})$$

The new heading after time Δt is:

$$\theta_b = \theta_a + \rho\Delta t \quad (\text{B.7})$$

We can assume that the change of heading is not continuous but occurs at waypoint b . This gives the model shown in Figure B.9.

During time Δt the UAV follows a great circle path. We can therefore estimate its latitude, longitude and heading at waypoint b from our knowledge of waypoint a using spherical trigonometry.

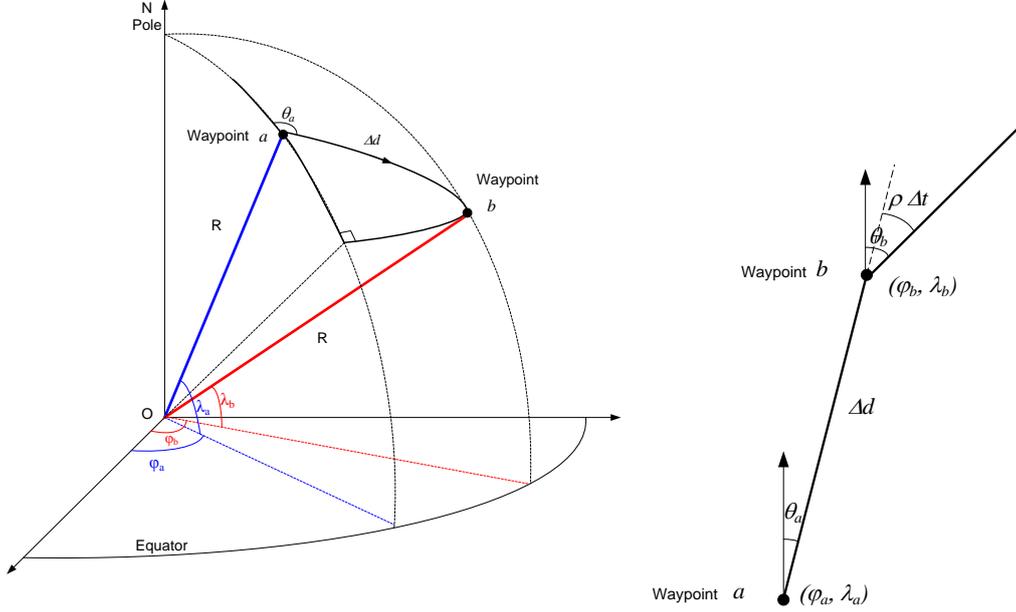


Figure B.9: Path Segment

$$\varphi_b = \arcsin\left(\sin(\varphi_a) \cos\left(\frac{\Delta d}{R}\right) + \cos(\varphi_a) \sin\left(\frac{\Delta d}{R}\right) \cos(\theta_a)\right) \quad (\text{B.8})$$

$$\lambda_b = \lambda_a + \text{atan2}\left(\sin(\theta_a) \sin\left(\frac{\Delta d}{R}\right) \cos(\varphi_a), \cos\left(\frac{\Delta d}{R}\right) - \sin(\varphi_a) \sin(\varphi_b)\right) \quad (\text{B.9})$$

B.5.3 Altitude

Changes in altitude are more complex than horizontal manoeuvres because there can be three initial states: the UAV can be climbing, in level flight or descending. A UAV with a limited wing *load factor* cannot instantly transition between these states. Furthermore, these manoeuvres introduce a pointing offset into the antenna footprint as shown in figure B.2.

In addition to the three initial states there are upper and lower altitude limits that must be considered. The upper limit could result from the flight characteristics of the UAV or be imposed as part of an air traffic management system. The lower limit could result from requirements for ground clearance or air traffic management

limitations.

The three initial states, and the upper/lower altitude limits, create a table of logical choices. If, for example, the UAV is close to its altitude ceiling then a command to climb might not be feasible and must be refused. Alternatively, if a UAV is descending then it must pull up before reaching its lower limit. Figure B.10 illustrates the logic of the manoeuvres.

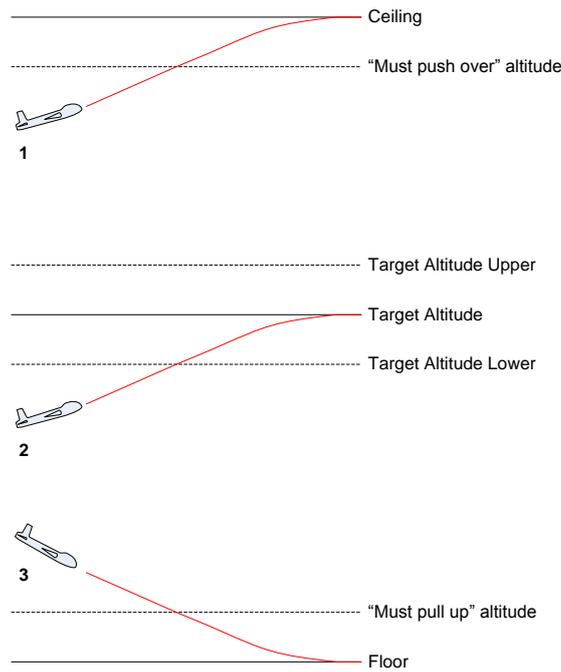


Figure B.10: Altitude Management

In example 1 the UAV is climbing towards its ceiling at its maximum rate of climb. Sometime before its ceiling is reached it must “push over” so that it is in level flight on reaching its ceiling. This critical altitude, referred to in the simulation code as the “must pushover altitude” (MPOA), is a function of the UAVs velocity, load factor and rate of climb. If the UAV is not climbing at its maximum rate of climb ρ it will still push over at this altitude, causing it to level out just below its ceiling.

A similar but subtly different case is shown in example 2. The UAV is climbing towards its target altitude (TA) and needs to know when to start to push over. The simulation has been written to produce upper and lower limits to the target altitude that will determine when the UAV starts to push over. Unlike MPOA

these limits are not fixed, but vary with rate of climb and velocity. The UAV will thus finish the manoeuvre somewhere between the target altitude upper (TAU) and target altitude lower (TAL). TA is set as a real number but no altitude manoeuvres are undertaken is the UAV enters level flight between TAU and TAL.

Example 3 shows a UAV descending towards its floor. As with the ceiling, there is a fixed altitude called the “must pull up altitude” (MPUA) that forces the UAV to pull up.

These principles are included in the simulation code. The logic starts to become complex if the TA is above MPOA or below MPUA. If, for example, the UAV is in a shallow climb towards its ceiling and the TA is above MPOA it is possible that TAL is also above MPOA. MPOA will always override TAL, thus the UAV will push over into level flight at MPOA and could level out below TAL. Similarly, MPUA always overrides TAU in descent manoeuvres near the floor.

B.5.4 UAV Attitude and Footprint

Recall that the downlink antenna on the UAV is a conical horn with a boresight direction aligned to the UAVs z axis. As a UAV climbs, descends or banks its pitch and roll angles change. This affects the boresight pointing of the antenna causing changes in cover to mobiles at the edge of the footprint. A model is required to estimate the changes in cover due to changes in aircraft attitude.

By setting a maximum load factor we can calculate the maximum bank angle $\hat{\beta}$ from equation B.3. We can also calculate the maximum climb angle $\hat{\chi}$ from the rate of climb ρ . For simplicity, let's assume that the maximum rates of climb and descent are identical.

Pitching the aircraft through $\hat{\chi}$ has the effect of reducing the antenna beamwidth by an angle $2\hat{\chi}$ around of the y axis. Similarly, rolling the aircraft by $\hat{\beta}$ has the effect of reducing the antenna beamwidth by an angle $2\hat{\beta}$ around of the y axis.

The effective area of the footprint is thus less than the theoretical area, as shown in Figure B.11. The corrected footprint dimensions are used when calculating antenna coverage.

In practice, this correction is only significant at high payload powers. If the payload is packed in increasing order of power the mobiles nearest the boresight are favoured. Mobiles at the edge of the footprint are usually the last ones to be considered by packing algorithms. Mobiles at the edge of the footprint are

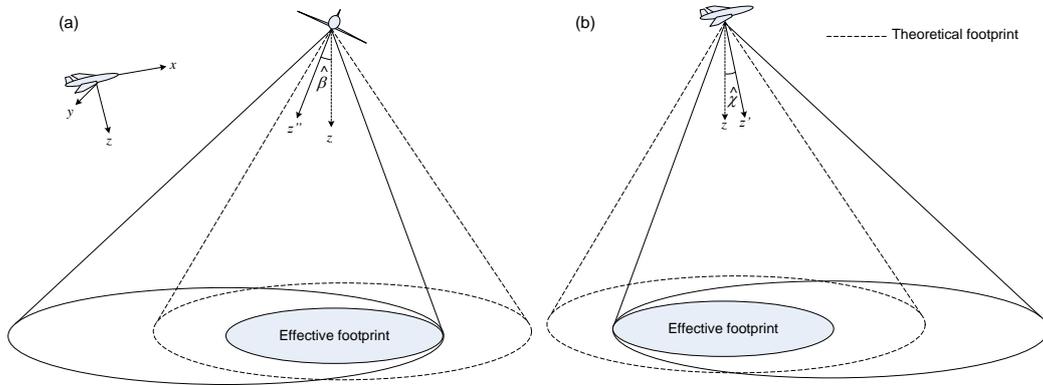


Figure B.11: Effective Footprint Area

only affected by this reduced footprint size if they are likely to be considered for packing, in other words when the UAV has relatively high values of P_{max} .

B.5.5 Navigation Accuracy

Experience in running the simulations showed that simulations based on timesteps had a slow accumulation of positional errors over long UAV missions. The main sources of location errors were:

- Choice of earth model and datum
- Model assumptions
- Computational accuracy

Early in the research there was a lot of effort expended to minimize these errors.

UAV navigation was based on the WGS 84 spheroid. This model is acceptably accurate for radio systems mounted on terrestrial and airborne objects. It also allows space objects to be included in future simulations, but should be used with caution if space/earth optical communications paths are modelled. The properties and accuracy of this earth model are well known and described in [89].

Common simplifying assumptions, for example considering the earth to be locally flat, were found to be a significant source of errors. These errors were reduced by about 2 orders of magnitude by using spherical trigonometry for all location and pointing calculations. For the purpose calculation we treat the earth

as a sphere of radius $R = 6371009$ meters but project location onto a WGS 84 spheroid. This approach introduces negligible errors but greatly simplifies the maths.

Another source of error came from calculating the changes of a UAV's location on the surface of the earth and projecting the new latitude and longitude to the UAV's operating altitude. These errors were significantly reduced by calculating all changes in location as differences in latitude, longitude and altitude. The resulting slant distance gave a more accurate estimate of the distance travelled during a timestep.

A smaller source of errors originated from the conversions between angles. Early implementations of the simulation calculated latitude and longitude as degree:minute:second, headings in decimal degrees, and radial distances in radians. Continuous conversions between these systems led to the accumulation of small errors. The code was re-written so the all angles were represented in decimal degrees. When conversion was necessary, the "master" value of the angle was retained in decimal degrees.

Figure B.12 shows that the distance travelled in time Δt can vary with the altitude of the UAV and its climb angle. The early versions of the simulation showed calculated horizontal distances at the surface of the earth and ignored differences due to climbing and descending. The code was modified in early 2013 to reduce this source of inaccuracy.

In executing a climb or descent manoeuvre the horizontal distance travelled is reduced by a factor corresponding to the cosine of the climb angle κ . Furthermore the radius of the earth needs to be modified to take account of the additional altitude. We thus need to modify equations B.10 and B.11 as follows:

$$\varphi_b = \arcsin(\sin(\varphi_a) \cos\left(\frac{\Delta d \cos(\kappa)}{R + h_a}\right) + \cos(\varphi_a) \sin\left(\frac{\Delta d \cos(\kappa)}{R + h_a}\right) \cos(\theta_a)) \quad (\text{B.10})$$

$$\lambda_b = \lambda_a + \text{atan2}\left(\sin(\theta_a) \sin\left(\frac{\Delta d \cos(\kappa)}{R + h_a}\right) \cos(\varphi_a), \cos\left(\frac{\Delta d \cos(\kappa)}{R + h_a}\right) - \sin(\varphi_a) \sin(\varphi_b)\right) \quad (\text{B.11})$$

We can estimate the final altitude as:

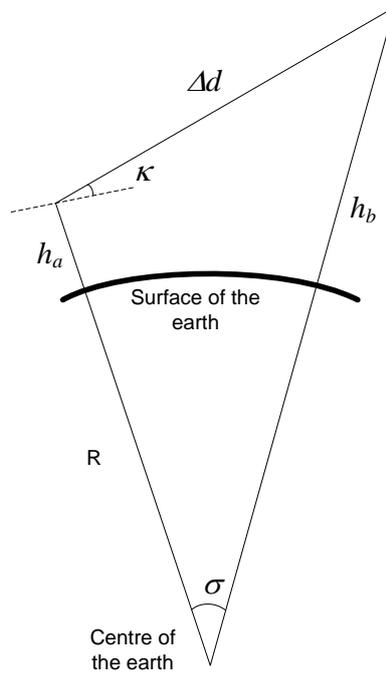


Figure B.12: Change of Altitude

$$h_a = h_a + \Delta d \cos(\kappa) \quad (\text{B.12})$$

The combined effects of these improvements was calculated. The benchmark was a 6 hour mission by a UAV operating between 10,000 ft and 35,000 ft. Without these assumptions the cumulative location errors over the mission were of the order of 10-50km. With these assumptions the cumulative location errors were $\leq 1\text{km}$, sometimes as low as 100m.

This page is intentionally blank

Appendix **C**

Bibliography

Bibliography

- [1] A. Agogino, C. HolmesParker, and K. Tumer. Evolving large scale UAV communication system. In *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference*, pages 1023–1030. ACM, 2012.
- [2] K. Akkarajitsakul, E. Hossain, and D. Niyato. Distributed resource allocation in wireless networks under uncertainty and application of Bayesian game. *Communications Magazine, IEEE*, 49(8):120–127, 2011.
- [3] E. Altman, T. Boulogne, R. El-Azouzi, T. Jimeenez, and L. Wynter. A survey on networking games in telecommunications. *Computers and Operational Research*, 33:286–311, Feb 2006.
- [4] V. G. Ambrosia, S. S. Wegener, J. A. Brass, and S. M. Schoenung. The UAV western states fire mission: Concepts, plans and developmental advancements. In *AIAA 3rd "Unmanned Unlimited" Technical Conference, Workshop and Exhibit*, 2004.
- [5] M. Arjomandi, S. Agostino, M. Mammone, M. Nelson, and T. Zhou. Classification of unmanned aerial vehicles. Technical report, University of Adelaide, Australia, 2007.
- [6] N. Avlonitis and . P. Charlesworth. Performance of retro-reflected modulated links under weak turbulence. *IET Optoelectronics*, 6:290–297, Dec 2012.
- [7] P. Basu. Coordinated flocking of UAVs for improved connectivity of mobile ground nodes. In *Nodes, Proceedings of IEEE MILCOM 2004*, pages 1628–1634, 2004.
- [8] I. Bekmezci. Connected multi UAV task planning for flying ad hoc networks. In *IEEE International Black Sea Conference on Communications and Networking (BlackSeaCom)*, 2013.

- [9] I. Bekmezci, O. K. Sahingoz, and S. Temel. Flying ad-hoc networks (FANETs): A survey. *Ad Hoc Networks*, 11(3):1254 – 1270, 2013.
- [10] S. Bhattacharya and T. Basar. Game-theoretic analysis of an aerial jamming attack on a UAV communication network. In *American Control Conference (ACC), 2010*, pages 818–823. IEEE, 2010.
- [11] W. Bolton. Operational experience with UAV payloads for climate research applications. In *2nd AIAA "Unmanned Unlimited" Systems, Technologies, and Operations*, 2003.
- [12] S. Bortoff. Path planning for UAVs. In *American Control Conference, 2000. Proceedings of the 2000*, volume 1, pages 364–368. IEEE, 2000.
- [13] T. X. Brown, B. Argrow, C. Dixon, S. Doshi, and P. Nies. Ad hoc UAV-ground network (AUGNet) test bed. In *4th Scandinavian Workshop on Wireless Ad-hoc Networks*, 2004.
- [14] T. X. Brown, B. Argrow, C. Dixon, S. Doshi, R.-G. Thekkekkunnel, and D. Henkel. Ad hoc UAV ground network (AUGNet). In *AIAA 3rd Unmanned Unlimited Technical Conference*, 2004.
- [15] O. Burdakov, P. Doherty, K. Holmberg, and P.-M. Olsson. Optimal placement of UV-based communications relay nodes. *Journal of Global Optimization*, 48:511–531, Dec 2010.
- [16] CAA. Unmanned aircraft system operations in UK airspace - guidance, April 2010.
- [17] D. E. Charilas and A. D. Panagopoulos. A survey on game theory applications in wireless networks. *Elsevier Computer Networks*, Computer Networks 54:3421–3430, 2010.
- [18] P. Charlesworth. Performance of the SLOW algorithm for a communications UAV supporting mobile subscribers with a prioritized access scheme. Technical report, EADS Innovation Works, 2012.
- [19] P. Charlesworth and S. Allen. Use of dynamic flight paths to enhance support to priority subscribers on a communications UAV. In *Military Communications Conference, 2012 - MILCOM 2012*, pages 1–6, 2012.
- [20] P. B. Charlesworth. Path planning for communications UAVs based on payload requirements. In *Bristol International UAV Systems Conference*, 2012.

- [21] P. B. Charlesworth. A non-cooperative game to coordinate the coverage of two communications UAVs. In *2013 - MILCOM 2013 Track 4) System Perspectives (MILCOM 2013 Track 4)*, San Diego, USA, Nov. 2013.
- [22] P. B. Charlesworth. Simulating missions of a UAV with a communications payload. In *UKSim 15th International Conference on Computer Modelling and Simulation, UKSim2013 (UKSim2013)*, Cambridge, United Kingdom, Apr. 2013.
- [23] P. B. Charlesworth. Using non-cooperative games to coordinate communications UAVs. In *IEEE Wireless Networking and Control for Unmanned Autonomous Vehicles*, 2014.
- [24] B. Chatterjee. An optimization formulation to compute Nash equilibrium in finite games. In *Methods and Models in Computer Science, 2009. ICM2CS 2009. Proceeding of International Conference on*, pages 1–5, 2009.
- [25] C.-M. Cheng, P.-H. Hsiao, H. Kung, and D. Vlah. Performance measurement of 802.11 a wireless links from UAV to ground nodes with various antenna orientations. In *Computer Communications and Networks, 2006. ICCCN 2006. Proceedings. 15th International Conference on*, pages 303–308. IEEE, 2006.
- [26] J. Child. Embedded processing brings more functionality to small UAVs. *COTS Journal*, December 2011.
- [27] L. Cong, L. Zhao, K. Yang, H. Zhang, and G. Zhang. A Stackelberg game for resource allocation in multiuser cooperative transmission networks. *Wireless Communications and Mobile Computing*, 11(1):129–141, 2011.
- [28] T. H. Cox, C. J. Nagy, M. A. Skoog, and I. A. Somers. Civil UAV capability assessment. Technical report, NASA, 2004.
- [29] L. G. M. Curran. UAVs: A critical multiplier for current and future forces. *RUSI DEFENCE SYSTEMS*, Summer:64–66, 2005.
- [30] K. Dalamagkidis, K. P. Valavanis, and L. A. Pieg. *On Integrating Unmanned Aircraft Systems into the National Airspace System*. Springer, 2012.
- [31] C. Danilov, T. R. Henderson, T. Goff, J. H. Kim, J. Macker, J. Weston, N. Neogi, A. Ortiz, and D. Uhlig. Experiment and field demonstration of a 802.11-based ground-UAV mobile ad-hoc network. In *Milcom*, 2009.
- [32] D. A. Day, J. M. Logsdon, and B. Latell. *Eye in the sky: the story of the Corona spy satellites*. Smithsonian, 1998.

- [33] J. R. Dixon. UAV employment in Kosovo: Lessons for the operational commander. Technical report, Naval War College, Feb 2000.
- [34] L. E. Dubins. On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, pages 497–516, 1957.
- [35] EC. Towards a european strategy for the development of civil applications of remotely piloted aircraft systems (RPAS). Technical Report SWD(2012) 259 final, European Commission, September 2012.
- [36] EEA. Mapping the impacts of natural hazards and technological accidents in Europe. Technical Report No 13/2010, European Environment Agency, 2010.
- [37] Frost and Sullivan. Study analysing the current activities in the field of UAV. Technical report, European Commission, 2007.
- [38] D. Gale, H. Kuhn, and A. Tucker. *Contributions to the theory of games*, volume 1, chapter On symmetric games, pages 81–87. Princeton University Press, 1950.
- [39] Y. Gao, Y. Chen, and K. Liu. Cooperation stimulation for multiuser cooperative communications using indirect reciprocity game. *Communications, IEEE Transactions on*, 60(12):3650–3661, december 2012.
- [40] A. Giagkos, E. Tuci, M. S. Wilson, and P. B. Charlesworth. Evolutionary coordination system for fixed-wing communications unmanned aerial vehicles. In *Towards Autonomous Robotic Systems 2014*, 2014.
- [41] A. Giagkos, E. Tuci, M. S. Wilson, and P. B. Charlesworth. Autonomous coordinate flying for groups of communications UAVs. In *Austrian Robotics Workshop*, May 2015.
- [42] R. Gibbons. *A Primer in Game Theory*. Prentice, 1992.
- [43] N. Goddemeier, S. Rohde, J. Pojda, and C. Wietfeld. Evaluation of potential fields mobility strategies for aerial network provisioning. In *GLOBECOM Workshops (GC Wkshps), 2011 IEEE*, pages 1291–1296, dec. 2011.
- [44] M. C. Gonzales, C. A. Hidalgo, and A. Barabasi. Understanding individual human mobility patterns. *Nature*, 453:779–782, 2008.
- [45] S. Govindan and R. Wilson. A decomposition algorithm for n-player games. Technical Report Research Paper No 1967, Stanford Graduate School of Business, August 2007.

- [46] R. Gunasekaran, E. Niranjani, S. Suganya, D. Vivekananthan, and K. Raja. Cross-layer optimization using game theory to alleviate unfairness in wireless networks. In *Knowledge, Information and Creativity Support Systems (KICSS), 2012 Seventh International Conference on*, pages 7–13, Nov.
- [47] Z. Han, D. Niyato, W. Saad, and A. Hjørungnes. *Game Theory in Wireless and Communication Networks: Theory, Models, and Applications*. Game Theory in Wireless and Communication Networks: Theory, Models, and Applications. Cambridge University Press, 2011.
- [48] D.-T. Ho, E. Grtli, P. Sujit, T. Johansen, and J. Borges Sousa. Cluster-based communication topology selection and uav path planning in wireless sensor networks. In *Unmanned Aircraft Systems (ICUAS), 2013 International Conference on*, pages 59–68, May 2013.
- [49] A. Holmberg and P.-M. Olsson. Route planning for relay UAV. In *26th International Congress of the Aeronautical Sciences*, 2008.
- [50] J. P. How, C. Fraser, K. C. Kulling, L. F. Bertuccelli, O. Toupet, L. Brunet, A. Bachrach, and N. Roy. Increasing autonomy of UAVs. *IEEE Robotics and Automation Magazine*, pages 43–51, June 2009.
- [51] D. Jia. Parallel evolutionary algorithms for UAV path planning. In *in Proceedings of AIAA 1st Intelligent Systems Technical Conference*, 2004.
- [52] M. C. A. Jones. Unmanned aerial vehicles (UAVs) an assessment of historical operations and future possibilities. Master’s thesis, Air Command and Staff College, March 1997.
- [53] S. Kim, P. Silson, A. Tsourdos, and M. Shanmugavel. Dubins path planning of multiple unmanned airborne vehicles for communication relay. *Journal of Aerospace Engineering*, 255:12–25, 2011.
- [54] C. Lemke and J. Howson. Equilibrium points of bimatrix games. In *Journal of the Society of Industrial Applied Mathematics*, volume 12, pages 413–422, June 1964.
- [55] J. Leonard, A. Savvaris, and A. Tsourdos. Towards a fully autonomous swarm of unmanned aerial vehicles. In *UKACC International Conference on Control*, 2012.
- [56] A. O. Lim and Y. Kado. Using game theory for power and rate control in wireless ad hoc networks. In *SICE, 2007 Annual Conference*, pages 1166–1170. IEEE, 2007.

- [57] C. Luo, S. McClean, G. Parr, Q. Wang, X. Wang, and C. Grecos. A communication model to decouple the path planning and connectivity optimization and support cooperative sensing. *Vehicular Technology, IEEE Transactions on*, 63(8):3985–3997, Oct 2014.
- [58] A. MacKenzie and L. DaSilva. *Game Theory for Wireless Engineers*. Morgan and Claypool, 2006.
- [59] A. MacKenzie and S. B. Wicker. Game theory in communications: Motivation, explanation and application to power control. In *Globecom 2001*, 2001.
- [60] A. B. MacKenzie and S. B. Wicker. Game theory and the design of self-configuring, adaptive wireless networks. *Communications Magazine, IEEE*, 39(11):126–131, 2001.
- [61] G. Maral and M. Bousquet. *Satellite Communications Systems*. Wiley, 2009.
- [62] R. Menon, A. MacKenzie, R. M. Buehrer, and J. H. Reed. Game theory and interference avoidance in decentralised networks. Technical report, MPRG, Virginia Tech, 2004.
- [63] R. Menon, A. Mackenzie, J. Hicks, R. Buehrer, and J. Reed. A game-theoretic framework for interference avoidance. *IEEE Transactions on Communications*, 57:1087–1098, 2009.
- [64] J. Neel, A. Mackenzie, R. Menon, L. Dasilva, J. Hicks, J. Reed, and R. Gilles. Using game theory to analyze wireless ad hoc networks. *IEEE Communications Surveys & Tutorials*, 7:46 – 56, 2005.
- [65] J. O. Neel, J. H. Reed, and R. P. Gilles. Game models for cognitive radio algorithm analysis. In *SDR Forum Technical Conference*, pages 15–18, 2004.
- [66] L. R. Newcome. *Unmanned Aviation*. American Institute of Aeronautics and Astronautics, 2004.
- [67] N. Nie, C. Comaniciu, and P. Agrawal. A game theoretic approach to interference management in cognitive networks. *Proc IMA*, Volume 143:199–219, 2006.
- [68] Office of the Secretary of Defense. Unmanned aircraft systems roadmap 2005-2030. Online, August 2005.
- [69] G. Pavlidou, F-N abd Koltsidas. Game theory for routing modelling in communications networks - a survey. *Journal of Communications and Networks*, 10:268–286, 2008.

- [70] F. Pinkney, D. Hampel, S. DiPierro, B. Abbe, and M. Sheha. UAV communications payload development. In *Milcom1997*. IEEE, 1997.
- [71] F. J. Pinkney, D. Hampel, and S. DiPierro. Unmanned aerial vehicle (UAV) communications relay. In *Milcom*, 1996.
- [72] R. Porter, E. Nudelman, and Y. Shoham. Simple search methods for finding Nash equilibrium. Elsevier Science, March 2005.
- [73] D. Rathbun, S. Kragelund, A. Pongpunwattana, and B. Capozzi. An evolution based path planning algorithm for autonomous motion of a UAV through uncertain environments. In *Digital Avionics Systems Conference, 2002. Proceedings. The 21st*, volume 2, pages 8D2–1. IEEE, 2002.
- [74] B. Rich and L. Janos. *Skunk Works*. Back Bay, 1994.
- [75] K. Ro, J.-S. Oh, and L. Dong. Lessons learned: Application of small UAV for urban highway traffic monitoring. In *45th AIAA Aerospace Sciences Meeting and Exhibit*, 2007.
- [76] V. Roberge, M. Tarbouchi, and G. Lebonte. Comparison of parallel genetic algorithm and particle swarm optimization for real-time UAV path planning. *IEEE Transactions on Industrial Informatics*, 9:132–141, 2013.
- [77] M. Ryan and M. Frater. The utility of a tactical airborne communications subsystems in support of future land warfare. Australian Land Warfare Studies Centre Working Paper No. 112, Jan 2001.
- [78] D. S. M. Schoenung, S. S. Wegener, J. Frank, C. Frost, M. Freed, and J. Totah. Intelligent UAV airborne science missions. In *Infotech@Aerospace*, 2005.
- [79] D. Shen, G. Chen, J. Cruz, and E. Blasch. A game theoretic data fusion aided path planning approach for cooperative UAV ISR. In *Aerospace Conference, 2008 IEEE*, pages 1–9, march 2008.
- [80] H.-S. Shin, C. Leboucher, and A. Tsourdos. Resource allocation with cooperative path planning for multiple UAVs. In *UKACC International Conference on Control*, 2012.
- [81] D. Shiyou, Z. Xiaoping, and L. Guoking. Cooperative planning method for swarm UAVs based on hierarchical strategy. In *International Conference on System Science, Engineering Design and Manufacturing Informatization*, 2012.
- [82] R. M. Taylor, S. Abdi, R. Dru-Drury, and M. C. Bonner. *Engineering Psychology and Cognitive Ergonomics: Aerospace and transportation systems*, volume Vol 5. Ashgate, 2001.

- [83] M. Tortonesi, C. Stefanelli, E. Benvegna, K. Ford, N. Suri, and M. Linderman. Multiple-UAV coordination and communications in tactical edge networks. *Communications Magazine, IEEE*, 50(10):48–55, october 2012.
- [84] T. Tozer. High altitude platforms for broadband services. In *The IEEE Wireless Broadband Conference*, 2006.
- [85] T. C. Tozer and A. C. Smith. High altitude platforms (HAPs) to meet future capacity requirements. In *IET Milsatcom*. IET, June 2010.
- [86] W. Truszkowski, L. Hallock, C. Rouff, J. Karlin, J. Rash, M. G.Hinchey, and R. Sterritt. *Autonomous and Autonomic Systems*. Springer, 2009.
- [87] J. Tsitsiklis. *Problems in Decentralised Decision Making and Computation*. PhD thesis, MIT, 1984.
- [88] A. Tsourdos, B. White, and M. Shanmugavel. *Cooperative Path Planning of Unmanned Aerial Vehicles*. Wiley, 2011.
- [89] USDoD. World geodetic system 1984 (third edition). Technical Report 7643-01-402-0347, National Imagery and Mapping Agency, January 2000.
- [90] USDoD. Department of defense global information grid architectural vision. Online, June 2007. Version 1.0.
- [91] J. Von Neuman and O. Morganstern. *Theory of Games and Economic Behaviour*. Princeton University Press, 1947.
- [92] W. Wagner. *Lightning Bugs and other Reconnaissance Drones*. Aero Publishers, 1982.
- [93] G. Walden. China deploys communications UAV in naval exercise. Flight Global, July 2011.
- [94] J. Webb. *Game Theory Decisions, Interaction and Evolution*. Springer, 2007.
- [95] B. Weiß, M. Naderhirn, and L. del Re. Global real-time path planning for UAVs in uncertain environment. In *Computer Aided Control System Design, 2006 IEEE International Conference on Control Applications, 2006 IEEE International Symposium on Intelligent Control, 2006 IEEE*, pages 2725–2730. IEEE, 2006.
- [96] B. White, A. Tsourdos, I. Ashokraj, S. Subchanx, and R. Zbikowski. Contaminant cloud boundary monitoring using UAV sensor swarms. In *AIAA Guidance, Navigation and Control Conference and Exhibit*, 2007.

- [97] S. H. Wong and I. J. Wassell. Application of game theory for distributed dynamic channel allocation. In *IEEE Vehicular Technology Conference*, 2002.
- [98] E. Yanmaz. Connectivity versus area coverage in unmanned aerial vehicle networks. In *Communications (ICC), 2012 IEEE International Conference on*, pages 719 –723, june 2012.
- [99] E. Yanmaz, R. Kuschnig, and C. Bettstetter. Channel measurements over 802.11a-based UAV-to-ground links. In *GLOBECOM Workshops (GC Wkshps), 2011 IEEE*, pages 1280–1284. IEEE, 2011.
- [100] E. Yanmaz, R. Kuschnig, M. Quaritsch, C. Bettstetter, and B. Rinner. On path planning strategies for networked unmanned aerial vehicles. In *Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on*, pages 212 –216, april 2011.
- [101] S. Zaloga, D. D. Rockwell, and P. Finnegan. World unmanned aerial vehicle systems, market profile and forecast 2013. Technical report, Teal Group, 2013.