A Nature Inspired Guidance System for Unmanned Autonomous Vehicles Employed in a Search Role

Alec Banks

Software Systems Research Centre

School of Design, Engineering and Computing

Bournemouth University

In Partial Fulfilment of the Requirements
for the Degree of Doctor of Philosophy
awarded by Bournemouth University

June 2009

BEST COPY

AVAILABLE

Poor text in the original thesis.
Some text bound close to the spine.

PAGE

NUMBERING

AS ORIGINAL

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with its author and due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Acknowledgements

I would like to thank my colleagues at the Tornado In-Service Software Maintenance Team (TISMT), Royal Air Force, for their support; particularly Squadron Leader Darren Goldstraw, former Officer Commanding, for his enthusiastic response at the beginning of the work and Squadron Leader Paul McKernan and Flight Lieutenant Michael Tomlinson, Navigators, for their aircraft handling advice.

The guidance and encouragement of my supervisory team has been constantly welcome and appreciated; Dr Richard Stillman for ethological advice, Dr Keith Phalp for ensuring progress could continue during periods of project uncertainty and, most of all, Dr Jonathan Vincent, who has inspired and believed in me from the outset.

Finally, but without whose unconditional love and support the work would not have been possible, my love and thanks go to my children, Levi and Mia, and to my wife Clare.

Abstract

Since the very earliest days of the human race, people have been studying animal behaviours. In those early times, being able to predict animal behaviour gave hunters the advantages required for success. Then, as societies began to develop this gave way, to an extent, to agriculture and early studies, much of it trial and error, enabled farmers to successfully breed and raise livestock to feed an ever growing population. Following the advent of scientific endeavour, more rigorous academic research has taken human understanding of the natural world to much greater depth. In recent years, some of this understanding has been applied to the field of computing, creating the more specialised field of natural computing. In this arena, a considerable amount of research has been undertaken to exploit the analogy between, say, searching a given problem space for an optimal solution and the natural process of foraging for food. Such analogies have led to useful solutions in areas such as numerical optimisation and communication network management, prominent examples being ant colony systems and particle swarm optimisation; however, these solutions often rely on well-defined fitness landscapes that may not always be available. One practical application of natural computing may be to create behaviours for the control of autonomous vehicles that would utilise the findings of ethological research, identifying the natural world behaviours that have evolved over millennia to surmount many of the problems that autonomous vehicles find difficult; for example, long range underwater navigation or obstacle avoidance in fast moving environments.

This thesis provides an exploratory investigation into the use of natural search strategies for improving the performance of autonomous vehicles operating in a search role. It begins with a survey of related work, including recent developments in autonomous vehicles and a ground breaking study of behaviours observed within the natural world that highlights general cooperative group behaviours, search strategies and communication methods that might be useful within a wider computing context beyond optimisation, where the information may be sparse but new paradigms could be developed that capitalise on research into biological systems that have developed over millennia within the natural world. Following this, using a 2-dimensional model, novel research is reported that explores whether autonomous vehicle search can be enhanced by applying natural search behaviours for a variety of search targets. Having identified useful search behaviours for detecting targets, it then considers scenarios where detection is lost and whether natural strategies for re-detection can improve overall systemic performance in search applications.

Analysis of empirical results indicate that search strategies exploiting behaviours found in nature can improve performance over random search and commonly applied systematic searches, such as grids and spirals, across a variety of relative target speeds, from static targets to twice the speed of the searching vehicles, and against various target movement types such as deterministic movement, random walks and other nature inspired movement. It was found that strategies were most successful under similar target–vehicle relationships as were identified in nature. Experiments with target occlusion also reveal that natural reacquisition strategies could improve the probability of target redetection.

Preface

Several journal publications have been produced based on the material in this thesis. Initially, a two-part survey of Particle Swarm Optimisation exploring the development of the field and indicative applications was published in Natural Computing (Banks *et al.* 2007 and 2008a). Following this, a study into the application of findings from ethological research into search behaviours to artificial search problems, also published in Natural Computing (Banks *et al.*, 2008c), which forms a large component of background material. During this time, an article exploring the application of particle swarms in the guidance of autonomous unmanned aerial vehicles was published by the Royal Institute of Navigation (Banks *et al.*, 2008b), based on the author's previous work that led to this thesis; this article was awarded the Royal Institute of Navigation's prestigious Michael Richey Medal for the best article published in The Journal of Navigation throughout 2008. More recently, this was further developed to include the nature inspired search at the centre of this thesis, once again published by the Royal Institute of Navigation (Banks and Vincent, 2009), and covering much of the method and findings reported here.

Banks A, Vincent J and Anyakoha C (2007) A review of particle swarm optimization. Part I: background and development. Natural Computing, Vol. 6, No.4, pp 467–484. Springer.

Banks A, Vincent J and Anyakoha C (2008a). A review of particle swarm optimisation. Part II: hybridisation, combinatorial, multicriteria and constrained optimization, and indicative applications. Natural Computing, Vol. 7, No. 1, pp. 109–124. Springer.

Banks A, Vincent J and Phalp K (2008b). Particle swarm guidance system for autonomous, unmanned aerial vehicles in an air defence role. The Journal of Navigation, Vol. 61, No.1, pp. 9–29. The Royal Institute of Navigation.

Banks A, Vincent J and Phalp K (2008c). Natural strategies for search. Natural Computing. Springer.

Banks A and Vincent J (2009). Exploring the performance of natural search strategies for the control of unmanned autonomous vehicles. The Journal of Navigation, Vol. 62, No.2, pp. 283–301. The Royal Institute of Navigation.

Contents

1 Introduction **1**
1.1 Aim 2
1.2 Motivation 3
1.3 Application 4
1.4 Contribution to Knowledge 4
1.5 Thesis Overview 5
2 Recent Developments in Autonomous Vehicles **6**
2.1 Aerial Vehicles 7
2.2 Ground Vehicles 9
2.3 Maritime Vehicles 10
2.4 Space Vehicles 11
2.5 Micro and Nano Air Vehicles 13
2.6 Sensing, Controlling and Communicating 14
2.7 Autonomous Vehicle Search 18
2.8 PSO in Autonomous Vehicle Control 20
3 Natural Search **23**
3.1 Natural Cooperation and Group Behaviour 23
3.2 Search Strategies 26
3.2.1 Search Environments 26
3.2.2 Sit-and-Wait 27
3.2.3 Continuous Motion Search 27
3.2.4 Systematic Searches 30
3.2.5 Path Integration 30
3.2.6 Area Concentrated Search 32
3.2.7 Saltatory Search 33
3.3 Knowing When to Give Up 34
3.4 Capitalising on Environmental Cues 35
3.5 Variable Strategies and Patterns 35
3.6 Repeat Searches 36
3.7 Competitor Avoidance 37
3.8 Communication Strategies 37
3.8.1 Indirect Communication 38
3.8.2 Direct Communication 38
3.9 Example Behaviours 40
3.9.1 Ants 41
3.9.2 Bees 42
3.9.3 Bacteria 43
3.10 Prey Reacquisition 44
3.11 Discussion 48
3.11.1 Ponerine Ant Inspired Combinatorial Optimisation 50
3.11.2 Trophallaxis and Swarming Robots 51

4 Method **53**
4.1 Context and Purpose 53
4.2 Approach 56
4.3 Simulation Environment 59
4.3.1 Search Environment 59
4.3.2 Agent Characteristics 61
4.3.3 Target Characteristics 63
4.3.4 LSZ Neighbourhood 64
4.3.5 Predictive Memory System Strategies 66
4.3.6 Occlusion 72
4.3.7 Communications 73
4.4 Search Algorithms 74
4.4.1 Initial Search 74
4.4.2 Reacquisition 81
4.4.3 Detected Target Interception 86
4.4.4 Known Target Interception 86
4.4.5 Behaviour Hierarchy 87
4.4.6 Giving Up Time 88
4.5 Metrics 89
4.5.1 Losses 90
4.5.2 Time 90
4.5.3 Coverage 90
4.5.4 Efficiency 92
4.6 Statistical Analysis 92
4.7 Schedule of Experiments 98
5 Research Findings **101**
5.1 General Considerations 101
5.2 Tuning Results 104
5.2.1 LSZ Neighbourhood 104
5.2.2 PMS Neighbourhood 105
5.2.3 ACS Giving Up Time 108
5.2.4 PMS Giving Up Time 110
5.3 Static Target Search 116
5.3.1 Locate-only 116
5.3.2 Locate-and-destroy 118
5.4 Dynamic Target Search 122
5.4.1 Deterministic Targets 122
5.4.2 Non-deterministic Targets 127
5.5 Effect of Swarming 137
5.6 Coverage Quality 140
5.7 Energy Efficiency 144
5.8 Predictive Memory System Where No Occlusion or Communication Failures Occur 146
5.9 Predictive Memory System Where Targets Become Occluded 149
5.9.1 Permanently Occluded Targets 150
5.9.2 Temporarily Occluded Targets 153
5.10 Predictive Memory System When Communications Failure Occurs 155

5.10.1 Permanent Failure 156
5.10.2 Temporary Failure 157
5.11 Statistical Significance 159
6 Discussion **161**
7 Summary and Conclusions **168**
7.1 Summary 168
7.2 Conclusion 171
7.3 Contribution To Knowledge 172
8 Further Work **173**
References **176**
Appendices **188**
Appendix 1 - Analysis of Bacterial Search Behaviour 188
Appendix 2 - Results of Initial Search Experiments 192
Appendix 3 - Results of Reacquisition Search by PMS Enabled Agents Where No Occlusion or Communication Failures Occur. 202
Appendix 4 - Results of Reacquisition Search by PMS Enabled Agents Where Permanent Occlusion Occurs 206
Appendix 5 - Results of Reacquisition Search by PMS Enabled Agents Where Temporary Occlusion Occurs. 213
Appendix 6 - Results of Reacquisition Search by PMS Enabled Agents Where Permanent Communication Failures Occur. 220
Appendix 7 - Results of Reacquisition Search by PMS Enabled Agents Where Temporary Communication Failures Occur. 227
Appendix 8 - Mann-Whitney U Significance Results for Static Target Experiments. 234
Appendix 9 - Mann-Whitney U Significance Results for Dynamic Target Experiments. 236
Appendix 10 - Mann-Whitney U Significance Results for Efficiency Experiments. 238
Appendix 11 - Mann-Whitney U Significance Results for Predictive Memory System Experiments. 239

1 Introduction

In a recent article to mark 100 years of British aviation (Aerospace International, 2008), a selection of leading aviation stakeholders were invited to predict its development over the next century. Perhaps unsurprisingly, autonomous behaviour featured among several of the opinions, including: "In the long term self-organising 'swarms' of military 'aircraft' will become feasible. Even more of the battle will become electronic rather than physical." (Alex Dorien, chief executive, Thales UK). In a recent submission to the House of Commons Defence Select Committee, the Royal Aeronautical Society (Royal Aeronautical Society, 2008) called for a commitment to the development of unmanned aerial systems in general. With respect to autonomy, the Society argued that, whilst it could reduce costs, it could also reduce communication bandwidth requirements, which in operational theatres can be problematic for other users and applications.

As acceptance of unmanned vehicle use grows, however, the need to provide autonomous behaviours has increased. For example, it was recently reported (Reilly, 2008) that Irish peacekeepers working in Chad had lost a non-autonomous UAV due to a communications failure. The default behaviour for the vehicle under such circumstances was to return to its operating base; however, it is believed that operators may have failed to update the home coordinates following training sorties in the Republic of Ireland. Following the failure, it is believed that the vehicle attempted to return to a base several thousand miles beyond its maximum operating range. Had the vehicle been capable of operating autonomously, communication loss would have been less problematic; the vehicle could have completed its mission, which would have included a final destination, usually its operating base.

The field of Natural Computing is increasingly being utilised by engineers and scientists wishing to find acceptable solutions to problems that are intractable using conventional processes that would otherwise require prohibitive computing resources. Such problems are often, though not exclusively, formulated as search problems. Generally, an acceptable solution is discovered in the following manner: an initial position is identified in the search space (typically at random), and its acceptability assessed through the application of a fitness function; a position change strategy, consistent with the paradigm in use, is then iteratively applied with the hope of improving the solution quality; the final solution is identified either through achieving an acceptable level of fitness or on the completion of a preset amount of

computation (either chronologically or in terms of the number of cycles executed). Examples of such approaches include *inter alia* evolutionary algorithms (Friedberg, 1958), which further inspired techniques such as genetic algorithms (Holland, 1962 and 1975); evolutionary strategies (Rechenberg, 1965); evolutionary programming (Fogel *et al.*, 1964) and, more latterly, genetic programming (Koza, 1992); ant colony systems (Colorni *et al.*, 1991), and particle swarm optimisation (Kennedy and Eberhart, 1995).

Sometimes, however, the search landscape is very flat. That is, the information returned by the fitness function is of little use until the selected position is within a given distance of a local or global optimum. In such circumstances the techniques typically associated with optimisation may not be viable, since they often rely on informative landscapes to guide search progress. A synthetic example might be the well-known multimodal test function referred to as Shekel's Foxholes, as incorporated into the five function test suite used by De Jong (1975) in an early exploration of the properties of genetic algorithms, whilst an application example might be autonomous robots searching for avalanche victims (Sandia National Laboratories, 2001). In other cases, the target may not be static, or may be actively trying to avoid detection, such as in a wide area security patrolling system (Sandia National Laboratories 2001). In such cases, a random walk or a deterministic search pattern may prove to be ineffective or inefficient, hence, the need to consider alternative search strategies. Therefore, this thesis considers search in a broader sense than just optimisation and, although many of the techniques may be applicable to optimisation problems, some are only of interest in the more general search problems where a paucity of information may exist, such as might be encountered by autonomous agents in a real-world scenario. Further, autonomous agents must also be able to communicate their findings to the group. In a perfect world this is a fairly simple requirement, but in harsh environments where, for example, the physical landscape or electromagnetic interference are problematic, the group of agents may require alternative techniques.

1.1 Aim

The aim of this work is look to nature and the search strategies that have evolved over millennia to produce the effective predatory behaviour observable today, apply them in a simulated environment and assess whether such strategies could be used to guide autonomous vehicle search. This does not suppose that there is one strategy for all occasions, but, where

there is *a priori* knowledge regarding the type of targets expected to be encountered, nature may well have already developed an effective solution. Where *a priori* information is unavailable, then it may be useful to have an array of strategies that can be selected according to environmental information, either sensed by the vehicle (or, more generally, 'agent') or passed to it by a cooperating ally.

1.2 Motivation

As enabling technologies advance, the field of autonomous vehicles is developing rapidly, targeting a variety of applications and operating in a wide range of environments that are difficult, dangerous or mundane for humans. Further, autonomous vehicles are increasingly being utilised where existing control paradigms have been recognised as providing sub-optimal solutions to modern problems. For example, where vehicles become too difficult for direct human control (Robinson, 2006) or provide inappropriate strategies (Miasnikov, 2005). The benefits are such that a great deal of effort is being applied to their development; however, as the drive to push the boundaries of autonomous vehicle behaviour continues, real-world considerations continue to restrict levels of autonomy. At the forefront of such considerations are the difficulties encountered in building the real-time models of the environment, that are needed to plan forthcoming behaviour, caused by imperfect sensors and / or communications with cooperating agents.

The developing field of autonomous defensive system control is vast with many facets; the motivation for this work is to develop robust strategies that may lead to systems that require minimal human intervention and also reduce inter-agent communication levels and, where sensing difficulties are such that target detection is lost, the system can produce appropriate behaviour to recover. Instinctively a developer of such a system may attempt to coordinate search agents to provide efficient coverage of the search area; this may have several weaknesses that could be exploited by an intruder, such as predictability (Gage, 1990) and a dependence on reliable communication, for example (Jin *et al.*, 2006). This latter point was also highlighted in a NATO technical report (NATO, 2003) that looked at urban operations in the future; in the report, potential counter-measures against all envisaged platforms are presented and all autonomous vehicles included susceptibility to communication jamming.

Autonomous vehicles are often considered for military roles, and such a system should be able to operate under a variety of defender / attacker relationships from stationary targets to

those capable of operating several times faster. Defenders should be able to effectively locate targets, cooperate only when necessary (minimising communication requirements) and react intelligently where initial detection is lost, since real world sensors are not perfect. Nature has developed effective strategies to secure food supplies from a given environment and this work considers some of these as underlying metaphors to provide simple yet effective (defensive) search strategies. This does not suggest that the strategies developed are optimal in a perfect world but that they could offer alternatives for difficult environments or potential reversionary modes to support more elaborate systems.

1.3 Application

Perhaps the most obvious application of this research is in the defence sector; returning to the British aviation industry's vision of the future (Aerospace International, 2008), the uptake of autonomous behaviour such as that described in this work may not only be a technological challenge. Sir Brian Burridge's view of the future envisages autonomous behaviour for civil applications but warns that military applications '...will be limited by the appetite (or otherwise) in international law for robotic warfare...'. Such ethical arguments are beyond the scope of this work, since its findings are not restricted to combative military applications. Indeed, the findings of this research would, it is envisaged, be applicable to a wide range of autonomous search applications. Examples that could benefit include: ship hull inspection (Reed *et al.*, 2006), which is an application that is increasing its use of Unmanned Underwater Vehicles (UUVs) to avoid dangerous manual inspection by human divers; and odour location applications, where research has generally concentrated on tracking plumes rather than on initial detection, a recent example being the odour plume tracking algorithms described by Lochmatter and Martinoli (2008).

1.4 Contribution to Knowledge

This research represents the first comprehensive study of the application of natural search strategies to the navigation of autonomous agents in an area that does not contain information concerning the existence or whereabouts of targets. Whilst there is a wide range of biological literature, the consideration within the computing field has been largely restricted to a few well known natural phenomena, such as evolution, swarming, and ant colony behaviour. This work goes beyond this by exploring a wide range of natural search behaviours and

considering how they may be exploited in a natural computing context. Hence, the investigation of natural search behaviours is particularly novel.

This work also presents the first study of its kind into the properties of natural search when applied to autonomous agents working under a variety of target / agent relationships; these may be utilised as primary search navigation mechanisms or as reversionary systems where coordinated search breaks down due to environmental issues such as high levels of electro-magnetic noise that may make communication problematic. They could also be used in conjunction with existing bio-inspired autonomous vehicles to improve initial location of sensory information; for example, a robot with chemical odour source location behaviours must first locate a trace of the odour. The findings of this research, applying the natural search behaviours to a simulation of agent based search, are a novel contribution, and cannot be found within the existing body of knowledge. Additionally, the development of a systemic working memory provides new behaviours that could assist groups of autonomous agents in the redetection of targets that have escaped or been lost from previous detection.

1.5 Thesis Overview

Sections 2 and 3 review related research. Section 2 considers recent developments within the field of autonomous vehicles, establishing the context for this work. Section 3 then explores a wide range of natural search behaviours, drawing on the biological literature and considering this in a computational context. Finally, natural solutions to prey reacquisition are examined.

Section 4 outlines the research method, first discussing the approach to the aims of this research, to assess whether natured inspired search behaviours may be useful in an artificial agent based system. A simulation environment is presented as a means to assess the search behaviours, and search algorithms and performance metrics are described. An outline of the empirical study is presented along with findings from configuration / tuning procedures.

Section 5 presents the research findings for a range of agent and target configurations, considering static targets, dynamic targets, search space coverage, energy consumption and finally the predictive memory system that addresses target reacquisition when targets are lost, for example, through limited sensor capabilities or occlusion.

Section 6 provides a discussion of the research findings, whilst Section 7 concludes and Section 8 suggests areas of further work.

2 Recent Developments in Autonomous Vehicles

Remotely controlled vehicles have been in operation for many years; however, more recently there has been increasing interest in vehicles that operate autonomously. Despite their relative youth, investment in the vehicles and related enabling technologies has been substantial, with levels of permitted autonomy ranging from minor automatic assistance for remote human operators to completely autonomous decision making. As part of ONERA's (France's Office National d'Etudes et Recherches Aérospatiales) long term autonomous vehicle search and rescue project (ReSSAC - Search And Rescue (SAR) with a Co-operative Autonomous System) Fabiani *et al.* (2007) describe a hierarchy of autonomy in which all air vehicles may be located (see Figure 1), although similar hierarchies could be developed for all autonomous vehicles. At the lowest level, only vehicle stabilisation is provided autonomously with the human controller dictating the vehicle's coarse grained movement remotely. This level is satisfactory where time delays are not problematic and loss of communication with the vehicle even for a short period is unlikely. However, time delays have been cited as the rationale for increased levels of autonomy in British Aerospace's jet powered Unmanned Aerial Vehicles (UAVs) (Robinson, 2006). As the level of autonomy increases, the level of human involvement is reduced. In terms of the control of faster, more complex, vehicles this is desirable since the risk of human error is reduced. However, this can introduce safety, as well as ethical, problems, although these issues are beyond the scope of this review.

It is not the intention of this review to provide complete coverage of autonomous vehicle development. Rather, the following sections will provide insight into recent state-of-the-art developments and applications in the field, beginning with some developments in aerial, ground, maritime and space vehicles before moving on to micro and nano vehicles. Finally, the some of the cornerstone technologies associated with vehicle sensing, control and communication are investigated to provide an insight in indicative technologies that could enable the implementation of the autonomous behaviours described in this work.

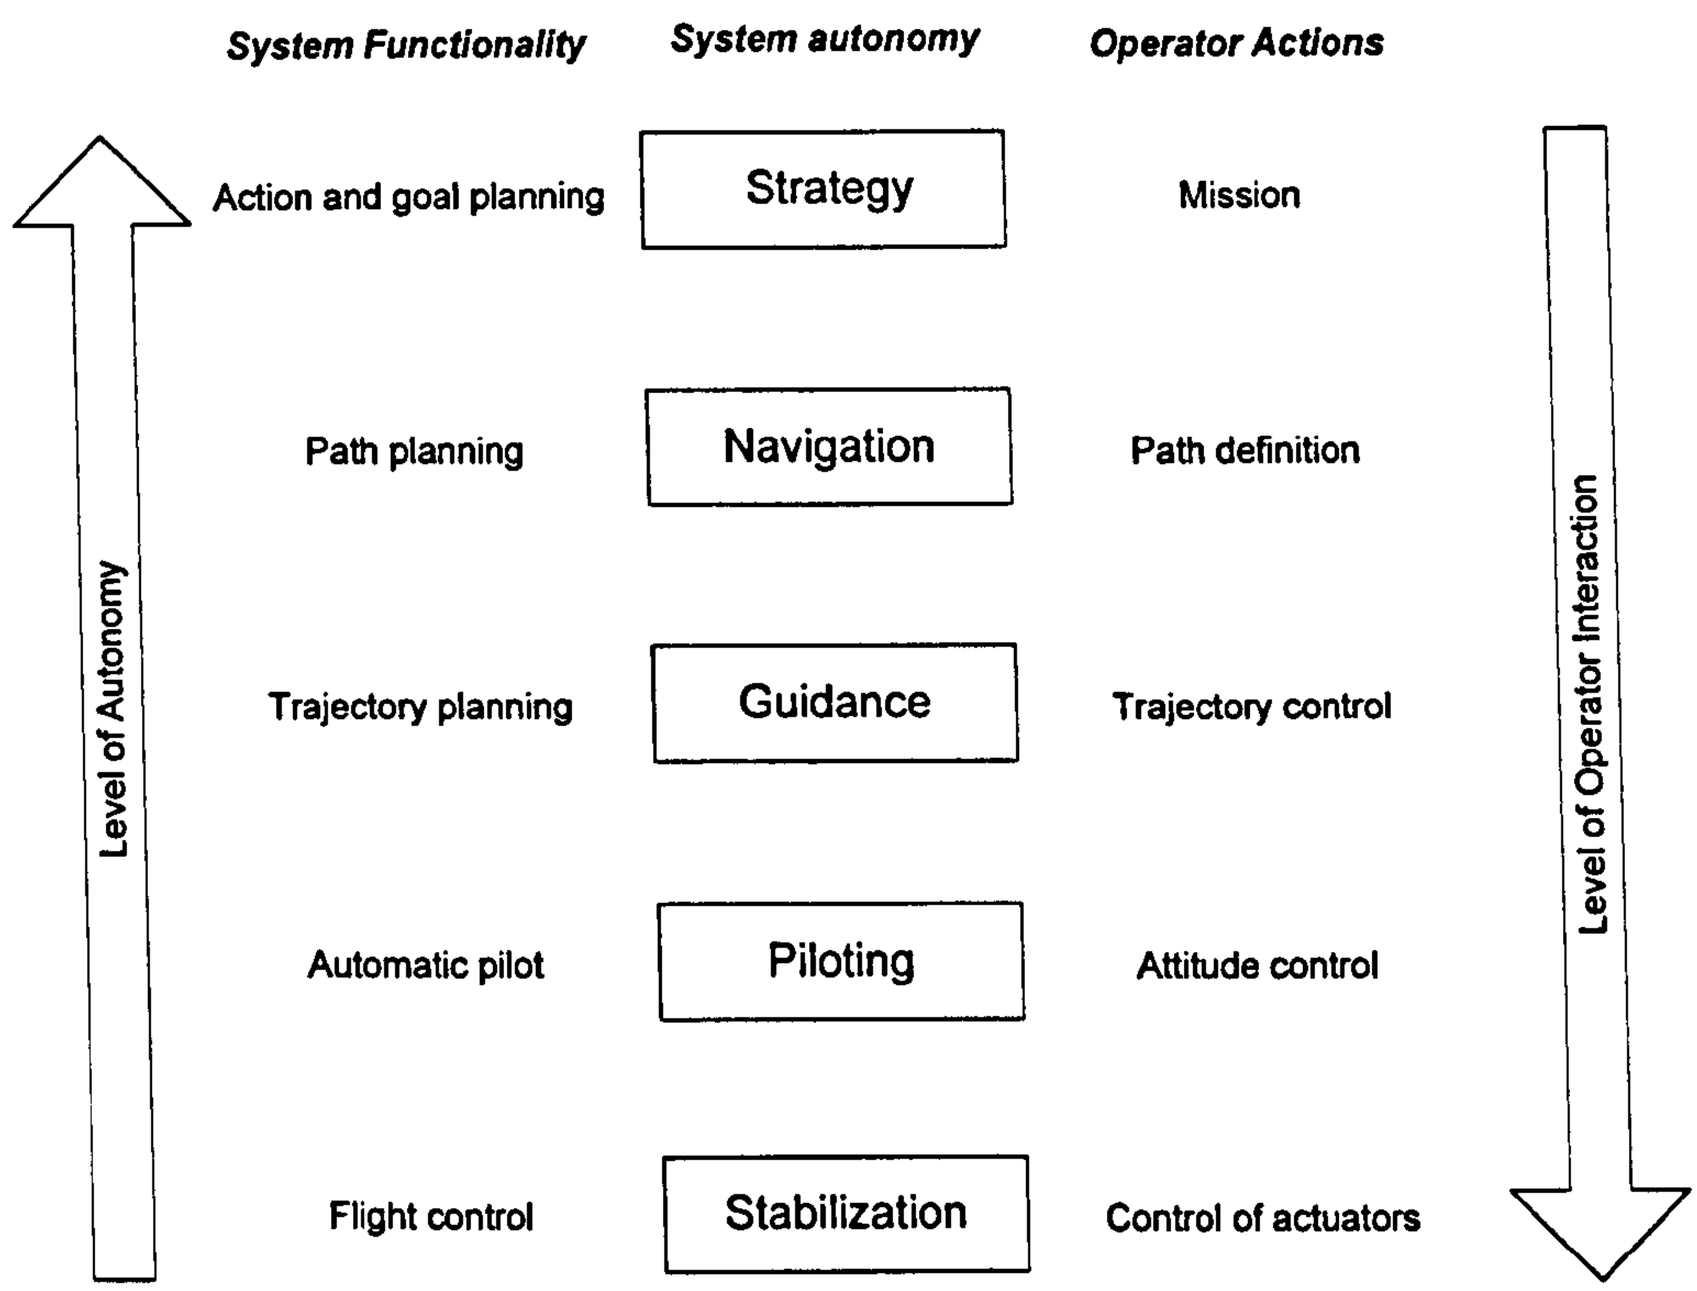

Figure 1. Levels of UAV autonomy and user interaction (adapted from Fabiani et al., 2007)

2.1 Aerial Vehicles

The development of a complex system, such as an UAV, could be structured in a myriad of ways. Hsiao *et al.* (2006) described the development of a small UAV that was capable of beyond visual range autonomous flight. To achieve this, the system was divided into three elements: the air vehicle; the onboard avionics; and the ground control station. The avionics suite consisted of: an onboard computer (providing the autonomous control); GPS for localisation; an air data sensor that was used to provide more accurate altitude and airspeed data than could be provided by the GPS; an attitude-heading reference system; and an actuator and servo control system. The autonomous controllers were further subdivided into guidance, navigation and flight control. Due to safety considerations, tight flight envelope restrictions were placed on the vehicle, which were acceptable for short range tests but proved problematic for the long range over sea trial where, despite careful selection of prevailing weather conditions, a reduction in engine power led to the test being safely aborted due to the system not being able to maintain altitude. This highlights the difficulty in the control of autonomous vehicles; in the presence of such uncertainty, they generally have to

abort, whereas a human operator may be able to overcome the difficulty using a novel approach.

Take-off and landing are amongst the riskiest activities for UAVs. Lange *et al.*, (2008) developed an autonomous landing system for multi-rotary wing UAVs that used optical detection and identification of the landing site. Whilst other autonomous landing systems had already been proposed, *e.g.* Saripelli *et al.* (2002), Lange *et al.*'s system had the advantage of being able to detect the target from higher altitude and not requiring the landing target to be fully visible. This was made possible by the use of a novel target consisting of unequally sized concentric rings (*i.e.* differing ring thickness) which meant that targets large enough to be seen from high altitude could also be detected at close range since the target identifier is simply searching for a circle. Once a set of target rings have been detected, the vehicle controller can then ascertain its height above the target by the ratio of the inner radius of the outermost detected circle to its outer radius and control the vehicle landing accordingly.

In a further application of a vision-based system to facilitate UAV autonomy, Conte *et al.* (2008) utilised geo-referenced images to provide reversionary vehicle localisation in the absence of GPS. This was significant since most UAVs rely on GPS for navigation and where the system becomes unavailable (due to failure or jamming) the UAV should be able to continue with its mission or be able to navigate home. The reversionary system used data from a standard IN system, air pressure sensor and vision system (replacing the GPS) combined using a Kalman filter to provide vehicle state information. The positional data provided by the vision system was produced by matching the video camera imagery with a geo-referenced image database and visual odometer.

Temporary loss of GPS data and GPS latency (slow speed compared to IN data) was also provisioned for in the lightweight inertial attitude and reference system developed by Jung and Tsiotras (2007) for use in small scale UAVs (wingspan of approximately 50cm). In such applications the processing power required for complex navigation systems is not available, so the developed system reduced Kalman filtering dimensionality to reduce computational overheads and used a simple magnetic compass to compensate for short term GPS loss.

For a search role, Quigley *et al.* (2005) developed a fixed wing mini UAV that applied a spiral search pattern, having a focal point defined by the operator, for open area searching and linear search for enclosed areas such as narrow valleys, although this latter element

contained unresolved issues such as dealing with GPS loss generated by local landscape features. In later work, Goodrich *et al.* (2008) applied the approach to wilderness search and rescue where the operator defined waypoints were selected based on likely locations for the target, these were then used by the UAV as the start of the spiral search. This general technique was found to be unsatisfactory for some terrain and so other techniques were investigated such as contour following.

2.2 Ground Vehicles

Although arguably operating in the least difficult environment, because failure of the vehicle would be unlikely to result in its loss, safe operation of ground based vehicles can still be problematic. Keviczky *et al.* (2006) investigated a Non-Linear Model Predictive Controller (NLMPC) that had previously been presented by Borrelli *et al.* (2005) to assess the approach's performance in side wind rejection and lane changing on slippery surfaces. The controller combined several models (such as vehicle properties, traction and cornering) to successfully reject side winds of up to 10.1 m/s and perform double lane change manoeuvres at vehicle speeds up to 17 m/s.

Obstacle avoidance is not the only solution to dealing with complex and unstructured environments, Vincent and Trentini (2007) defined a method for a robot (the Shape-shifting Tracked Robot Vehicle, STRV) to learn to negotiate obstacles rather than avoid them. The STRV utilised data from GPS, Inertial Navigation (IN) and laser range finder sensors for localisation, vehicle movement sensing and to build a 3D picture of the terrain to be negotiated respectively; further internal sensors provided track velocity and axle position data.

The behavioural learning mechanism of the STRV consists of a two-stage reinforcement learning architecture. In the first layer, information from the terrain mapping and IN were input into the learning algorithm for computing axle angular position; the second layer then selects additional behaviours from previously controlled runs based on obstacle recognition (for example steps, ramp, *etc*). In the work reported the vehicle remained semi-autonomous in that an operator controlled the direction of the vehicle; however, in the additional work identified the authors advised that it should be equipped such that full autonomy should be available should it be required (*e.g.* for reversionary control).

Autonomous navigation of urban vehicles could be critical in addressing current issues surrounding inner city congestion and pollution. Courbon *et al.* (2008) developed an efficient visual memory management system that permits real-time navigation in a large-scale outdoor environment, a university campus, using a single vehicle mounted camera and natural landmarks.

Further recent developments in Autonomous Ground Vehicle (AGV) technology include Ghumman *et al.*'s autonomous highway overtaking planning (2008), Kurt and Özgüner's (2008) hybridisation of a finite state machine controller with a continuous state controller to improve vehicular control across a range of urban scenarios and Lee *et al.*'s (2009) utilisation of an *e. coli* metaphor to develop a novel autonomous vehicle traffic data harvesting system.

2.3 Maritime Vehicles

The combination of unacceptably high diver casualty rates, increased need for integrated force protection and waterside security, combined with advances in UUV technology have, according to Reed *et al.* (2006), been the driving forces behind an increased effort to provide autonomous underwater surveys of artefacts such as pipelines and ship hulls. In their automatic ship hull inspection system, a 3D model was produced of the hull using a Doppler velocity log (DVL), which was then overlaid with a mosaic of imagery produced using a Didson sonar camera. These techniques were facilitated by autonomous control and navigation modules, the former utilising a modified Zeigler-Nichols method to maintain constant position despite underwater currents and the latter uses a SLAM-RTS (Simultaneous Localisation And Mapping - Rauch-Tung-Striebel) architecture that combined an extended Kalman filter and a Rauch-Tung-Striebel smoother to conduct manually defined landmark to landmark navigation.

The results of a practical application of an Autonomous Underwater Vehicle (AUV), named Autosub, were recently reported by Nicholls *et al.* (2008), who utilised it to further the understanding of the Fimbul Ice Shelf in Antarctica. Previously, scientific studies of ice melt activity below the shelf had been very limited, restricted to ship based measurements along the ice front, making ice bores on the surface and inverting glacialogical data. To give an idea of scale, the AUV was small compared to manned vehicles, having a length of 6.8m and 0.9m diameter, with a range of 400km and maximum depth of 1600m. Navigation was provided by a combination of GPS whilst on the surface and a combination of a downward

looking Doppler system (when close to the seabed) and a Photonic IN system when beneath it. To collate data, the AUV followed a pre-programmed route taking measurements using an array of onboard sensors. Results of the study indicated that previous studies, made using sub-optimal surface techniques, had been misleading and highlighted the need for a re-examination of existing theories on the influence of the ice melt on the climatically important Filchner-Ronne ice melt, which is further downstream of the study's location.

2.4 Space Vehicles

Space exploration is an obvious application for autonomous vehicular technology, although it does present its own difficulties such as, *inter alia*, relatively low processing speeds, gravity and atmospheric pressure (in the case of planetary exploration).

The first autonomous spaceflight was performed in 1994 during the United States Air Force Phillips Laboratory's Technology for Autonomous Operational Survivability (TAOS) program (Guinn *et al.*, 1995). To provide orbit determination a combination of GPS and positional information provided by a space ground link were fed to an autonomous navigation system. This was improved on, in terms of level of autonomy and flight efficiency by a fully autonomous system (Gurevich *et al.*, 2000) that did not require updates from ground stations (due to improved position prediction). The improvement in efficiency was gained due to the satellite being able to fire its thrusters more often but with less power. Later, Wertz (2003) described a fully autonomous navigation system that did not require an existing GPS network, affording the possibility of using the system for autonomous inter-planetary navigation and remote planet orbiting. This system used improvements in optical sensor technology to identify planets and stars and then measured distance and inertial direction to a nearby planet. This system was first used successfully in NASA's Deep Space 1 (DS1) mission, a high-risk technology test mission (Rayman *et al.*, 2000), which used navigation data that included all solar system planets (except Pluto) and a catalogue of 250,000 stars. DS1 also conducted other autonomous behaviour tests: a beacon monitor that would check for other spacecraft in the vicinity, where other spacecraft were detected the system would apply a less demanding protocol for mission reporting thus reducing communication network load; a high-level remote agent was also tested that operated a planning and scheduling engine, which took vehicle measurements and re-planned to achieve high-level goals and fault rectification.

Closer to home, NASA's Autonomous Sciencecraft Experiment, ASE, (Chien *et al.*, 2005) has been orbiting the Earth since 2003, utilising several autonomous components to analyse world environmental data, such as earthquakes and volcanoes detected by onboard sensors, and respond by replanning its activities to maximise observation and thus scientific returns of the event. This is achieved through the application of a traditional three-layer architecture: at the highest level of abstraction a mission planner schedules activity based on information from scientific sensors and spacecraft monitoring (timescales being in the tens of minutes due to limited spacecraft processor speeds); these plans are then fed to a spacecraft scripting layer that responds within seconds to produce low-level commands to the conventional flight software layer that responds reflexively. The advantage of this use of autonomous behaviour is that the spacecraft can decide how to respond to target events and downlink data only when necessary, thereby allowing the limited bandwidth to earth stations to be used more efficiently.

Autonomous near-surface navigation in outer space presents additional difficulties to those found in Earth based or low orbit satellite applications. Thakoor *et al.* (2003) applied a bio-inspired approach, based on Dragonfly Ocelli eyes (that are dedicated to horizon sensing), to develop a horizon sensing system for use in an attitude reference system to assist control of UAVs that will explore the surface of Mars. The development is important to such missions due to environmental differences on Mars, such as reduced pressure and gravity fields, that render conventional attitude reference systems unusable. Where existing systems are viable such as inertial navigation systems the Ocelli system was reported to be advantageous due to its very small size and weight (quoted as being 40 times lighter than a comparative IN system). The system is very simple and operates by measuring the difference in light between two sensors each detecting UV and green light; the two light channels being spectrally processed in opposition to filter potential bias inducing features such as the Sun or clouds. Output from the system can then be fed directly to the control system to provide attitude stabilisation in pitch and roll. An indicative application for such a referencing system might be the Mars rotorcraft proposed by Young *et al.*, (2005) although the DelFly flapping MAV (De Wagter *et al.*, 2007) could offer an attractive alternative. The DelFly uses a conventional lightweight camera, applying vision-based technology that stitches images of the horizon together to form a 360° panorama, which when used in conjunction with a Sun detection algorithm can then provide a reference system for the vehicle to follow.

2.5 Micro and Nano Air Vehicles

Whilst other micro and nano vehicles are being developed, e.g. Bergbreiter and Pister (2007), aerial vehicles (MAVs and NAVs), driven by an acceptance of the technology by military forces around the world, have become the most commercially viable platforms.

Networks of autonomous vehicles are often limited by the communication networks in which they operate; in a reversal of this Hauert *et al.* (2009) proposed a self-organising swarm of MAVs to establish, and maintain, wireless communication networks in areas that would otherwise be problematic due to factors such as remoteness, environmental features or congestion of existing networks. Such factors can also make position sensing difficult (*e.g.* GPS signals may not be available) and therefore a requirement for the system was that the agents should be able to operate without externally provisioned positional information. To overcome these challenging conditions, genetic algorithms were applied to evolve neural controllers that could maintain vehicle positions such that the wireless network was maintained. Following experimentation using the most evolved controller, several design principles were identified: circular trajectories were the most efficient movement patterns to maintain the network; disconnected MAVs should backtrack to attempt reconnection; alternating disconnected and connected (with the base station) phases provide the basis for synchronisation; and MAVs that are connected to both the base and user stations should follow the smallest possible circular trajectories to maintain the link.

The advantages of MAVs over conventional vehicles are their ability to operate in locations not viable for their larger counterparts. For example, Honeywell (Honeywell, 2009) have developed a MAV system called the T-Hawk (Figure 2) that consists of a pair of duct fan operated vehicles, each with an outside diameter of less than 36cm, meaning that an individual can carry one of the pair in a backpack. However, the small size does

Figure 2. Honeywell T-Hawk in operation.

not mean a lack of functionality: the system consists of a pair of vehicles, one carrying an electro-optical sensor and the other an infra-red sensor. They can operate under manual control or fully autonomously, able to follow a route of up to 100 waypoints for 40 minutes, due to a fully functioning flight management system containing an inertial navigation unit, GPS, barometric altimeter and air data sensors.

Additionally, the size of MAVs can also allow operating features not possible for larger vehicles. For example, Green and Oh (2006) developed a fixed wing micro UAV that can transition from conventional horizontal flight to an autonomous vertical hover due to its thrust to weight ratio of greater than one. This allows for the more efficient and speedier flight associated with fixed wing vehicles, coupled with the flexibility offered by rotary wing aircraft.

Powered by a miniature two stage rocket in its wing tip that provides a range of approximately 1000m and a top speed of up to 10m/s, Lockheed Martin's Maple Seed, as shown in Figure 3, is arguably one of the most novel NAVs under development to date (reported by Lockheed Martin, 2007; prototyping discussed by Jameson *et al.*, 2007). Its planned reconnaissance role is to be fulfilled by a rotating camera module that maintains its focus in flight and transmits video images back to the operator in real-time. This module is expected to be interchangeable for other modules such as infrared imaging, biological and chemical sensors. Notably, whilst present development is concentrating on individual vehicles, Lockheed Martin believes that the power of the technology lies in the application of massively parallel autonomous swarms of NAVs in the battlefield. Other recent advances in miniature technologies include the development of flocking in a swarm of mini-helicopters (De Nardi and Holland, 2007) and autonomous glider perching (Cory and Tedrake, 2008).

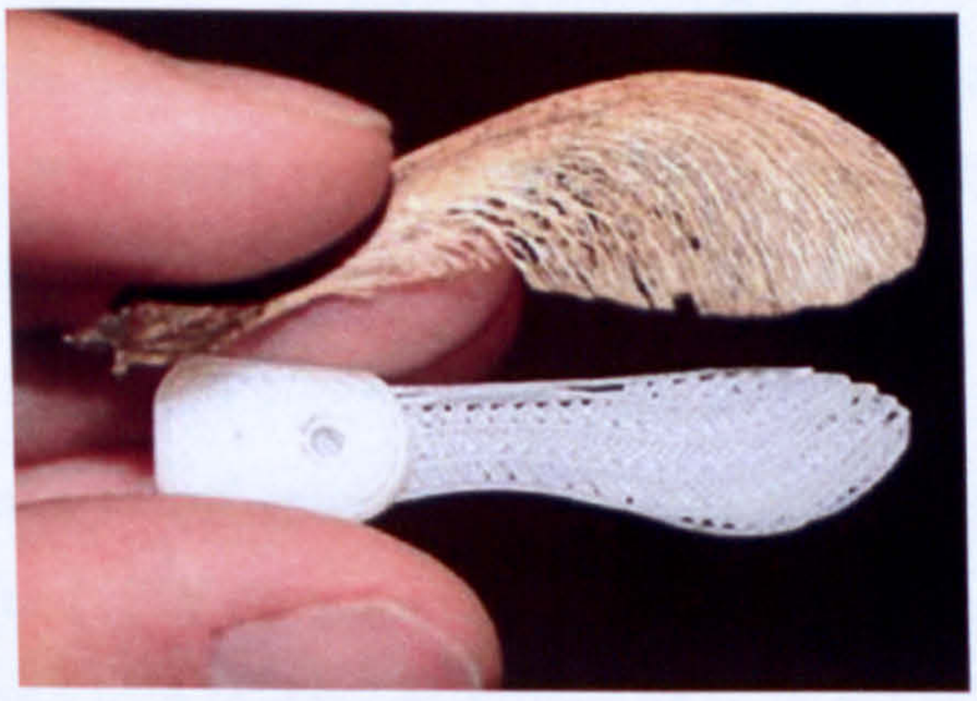

Figure 3. Lockheed Martin's Maple Seed NAV.

2.6 Sensing, Controlling and Communicating

Where autonomous vehicles operate in well defined environments it may be possible for them to simplify navigation through identifying, and reacting to, known landmarks; in

complex environments this remains a non-trivial task. Vacek *et al.* (2007) investigated a system for analysing markings on rural German roads to assist behaviours such as lane changing, obstacle avoidance and turning off. Lane identification was the first stage in the process and was achieved by defining four parameters associated with detection: longitudinal offset of the vehicle with respect to the middle of the lane; rotational angle between vehicle and lane; lane width; and tilt angle of the camera to the road surface. Each sample taken of the road surface was then recorded in terms of these parameters and pattern matched against a series of known cues such as lane widths and road edges. The weighted sums of all the matches to cues were then tested against a threshold to decide whether a lane was being followed. Once lane could be successfully followed using this technique road marking analysis was accomplished based on two types of marking: lines and arrows. Lines were analysed by overlaying orthogonal lines across detected lines and matching the number of intersections (0, 1 or 2) with neighbouring orthogonal lines. Arrows were analysed by bounding the detected marking, which occur in the centre of the road and once again pattern matching. The system proved promising during trials where road markings were well defined (as laid down by strict German law) but further work was still identified to increase the robustness of the system by further systemic fusion that would allow more holistic decision making.

In less well defined application environments, such as off road navigation, the features to detect and possibly avoid may be less well defined. For an autonomous off road military vehicle a feature as simple as a muddy puddle could lead to mission failure or even result in increased burden for the operators since it could require rescuing. Rankin and Matthies (2008) reported that whilst work on remotely sensing water moisture was mature, existing methods did not provide the resolution required for local navigation. To more accurately detect muddy areas a range of passive sensors were utilised including: colour stereo, in which isolated mud was detected by segmenting regions that were significantly darker than their surroundings whilst being of the same colour hue; multi-spectral sensors which used the differences in wet, dry and shadowed soil reflectance in the Near Infra Red (NIR) and red bands; Short Wave Infra Red (SWIR) camera, using a similar phenomenon as the previous camera but comparing SWIR with NIR reflectance; thermal imagery using a Long Wave Infra Red (LWIR) camera, wet soil being consistently cooler throughout the day; and a polarisation camera, which uses differences in the degree of linear polarisation of light

reflected wet areas with increased levels being an indicator of the presence of water. A combination of polarisation and stereo colour cameras were mounted on a vehicle, the polarisation camera providing primary mud detection whilst the stereo cameras were used to improve localisation and reject false positives such as vegetation that were above the ground. Output from the sensors was then combined with a terrain map to produce a world map that the vehicle could use for path planning.

Biomimetic robots model nature in an attempt to gain the benefits of form and sensory capabilities evolved over millennia. Erwin (2004), building on earlier work modelling bat sensory–motor behaviour, implemented algorithms for sonar tracking of dynamic targets. These algorithms included non-predictive tracking, which was found to be only suitable for stationary targets; predictive tracking based on a Benedict-Bordner α/β filter, which does not include an acceleration element; and several predictive filters that included acceleration (a second Benedict-Bordner α/β filter, spherical and generalised unscented Kalman filters, and a simplified Segmenting Track Identifier model). Only the last of these models performed reasonably well, the others working poorly due to the having insufficient reaction time, the target moving with unsteady acceleration and, in the case of the Kalman filters, there being insufficient time for the filters to develop enough information to predict acceleration. Conclusions from the work suggested that using existing predictive technologies was insufficient for the application and that modelling natural behaviours more closely may be more effective.

High level decision making is arguably one of the most intractable areas of autonomous vehicle development and several existing artificial intelligence paradigms have been applied to the problem. Vacek *et al.* (2007) proposed the use of case based reasoning for autonomous road vehicle guidance. Particularly important to the application was the bounding of individual cases since in a real dynamic environment it would be difficult to identify the start and end of each case. To alleviate this, cases were seen as snapshots of situations (containing core components such as vehicle behaviour and that of other participants) rather than time intervals. A given situation could then be compared with previously experienced situations and an appropriate reactionary behaviour selected. A major drawback acknowledged by the authors, was the prediction of vehicle behaviour where novel situations arose that could not be associated with any existing cases.

As future systems develop, particularly in the defence sector, it can be expected that communications between heterogeneous elements would be required. Where systems were developed by a single agency, agent interoperability can be designed from the outset using standard protocols. However, where systems evolve, sometimes in an *ad hoc* manner, even small networks of networks can be difficult to administer, especially where systems involve multi-national elements that have been developed separately. Liang *et al.* (2008) have developed an Ada based architectural model called the Quantitative Interoperating Model (QIM), for C4ISR (Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance) that provides an enabling layer for new and legacy systems to interoperate on a CORBA basis. Such a model would ensure that systems know their respective communication roles, and a latency feature would allow agents to know whether cooperating agents are unavailable or unreliable.

At a lower level, Chitre (2008) also tackled the requirement for autonomous vehicle interoperability through the development (initially for AUVs, but also applicable to other autonomous vehicles) of a 4 layer distributed software architecture. At the lowest level, an internal communications layer provides an interface to whatever internal communications medium is in use, *e.g.* Ethernet, UDP/IP for WiFi, *etc.* Above this sits a Remote Procedure Call (RPC) layer that implements the remote procedure call semantics with the third layer, which contains services such as core vehicle services, sensor / actuator drivers and external communications links. Finally at the top layer operates vehicle command and control components.

Some environments pose particular communication difficulties; for example, in networks of AUVs. AUV network links are usually maintained using acoustic modems which, due to the properties of the water in which they operate, have a relatively narrow optimal operating range. To overcome this, Shahabudeen *et al.* (2007) developed a medium access control (MAC) protocol that would enable the simultaneous operation of short, medium and long range modems using a DLL (data link layer) architecture. This development produced increased data rates and efficiency in a simulated AUV network.

Similar problems present themselves in networks of surface based acoustic beacons that can be used to facilitate low cost AUV navigation; the quality of the signals available to the AUV being dependent on the network geometry (Matos *et al.*, 1999), a problem that is compounded by signal attenuation increasing markedly with distance. To overcome these

problems, Santos *et al.* (2008) implemented a mobile network that placed autonomous surface vehicles that could be programmed with mission data that would position the acoustic beacons they carried in the correct place to coordinate AUV movement. The system improved AUV mobility but with additional complexity introduced by the movement of the surface vessels.

2.7 Autonomous Vehicle Search

In nature, animals apply a variety of decision making mechanisms to produce behaviours that maximise the chances of survival either for the individual or their species (broadly reviewed by Stephens *et al.*, 2007). Andrews *et al.* (2007) exploited such mechanisms found in foraging theory to produce high level decision control for autonomous vehicles. To achieve this, targets were given a point value and a handling time and agents were given to the goal of attaining maximum scores available. The work, using agents that followed 'lawnmower searches', indicated that the agents could decide, given *a priori* knowledge of the environment and available targets, whether, upon detection, a target was worth pursuing or whether, in order to gain maximum return, it should continue to search for more profitable targets known to exist in the search area. The adoption of foraging theories such as these could be criticised for using perfect world models, however, the work does point to adaptations to the model that could be made to align it more to real implementations; for example, a penalty could be included to accommodate non-zero recognition times or imperfect resemblance such that targets that were more difficult to recognise would incur the penalty during the agent's determination of whether a target warranted pursuit.

With the recent rise in the popularity of nature inspired computing, it is unsurprising that work into applying ethological research findings to autonomous vehicles has been undertaken. What is noticeable in the work reviewed is that it assumes the presence of sensory information from the outset. This brief indicative review highlights the range of sensory that has already been successfully exploited to improve the capabilities of autonomous vehicles across a range of applications.

Many naturally occurring sensory modalities have been exploited in the drive to improve the autonomous search in robotic applications. Morse *et al.* (1998) harnessed the behaviour of a nematode worm's chemotaxis (chemical following) based foraging to equip a robot with a

phototaxis (light following) spatial orientation ability that proved robust to large scale changes in robot speed and turning bias.

Anderson *et al.* (2004) implemented a phonotaxis (sound following) robot by modelling the structure of the human head as a near sphere, using two microphones mounted, diametrically opposed, as ears to localise the source of a single broadband sound. The 'head' was fixed to a mobile robot and the output of the localisation algorithm fed into the robotic navigation system. The algorithm calculated sound direction using pressure measurements (phase and amplitude) made at each microphone, with the robot manoeuvring left or right to achieve minimization in the inter-aural difference – the sound being directly ahead or behind when the difference between the two were minimal.

Edwards *et al.* (2005) successfully exploited the pheromone plume tracking behaviour of male moths to construct an airborne chemical sensing 'robo-moth' (a robotic gantry located in a wind tunnel) that has the ability to locate the source of an airborne chemical plume. When introduced into a plume of pheromone bearing odour, a male moth will actively sample the air making upwind surges, each consisting of rhythmical left / right tacking movements. When the odour is lost, due to the patchy distribution found in airborne chemicals, the moth begins to fly a crosswind pattern, known as casts, in an attempt to reacquire the plume.

Extending the principle of modelling single sense search behaviour, Pyk *et al.* (2006) constructed a robot capable of harnessing multi-modal sensory inputs to locate the source of a chemical odour whilst avoiding obstacles. These modalities were the anemotaxis (wind sensing) behaviour of male moths, *i.e.* airborne search behaviour in response to pheromone odours, and the optomotor collision avoidance behaviours of locusts (the overall behaviour being generally invariant across flying insect species). The search strategy employed by the male moth is described later, whilst the object avoidance mechanism of the locust, known as the lobula giant movement detector, is outside the scope of this work. To combine the two modalities, a process not yet fully understood biologically, each operated as parallel competing schemes with integration occurring at the motor stage, with the optomotory element having precedence.

2.8 PSO in Autonomous Vehicle Control

Brief consideration of the use of Particle Swarm Optimisation (PSO) in autonomous vehicle control is included here for two reasons: first, it is an approach that bears most similarities to the work reported in this thesis; second, the author's previous work in this area, which provided the initial starting point for this thesis, resulted in the first study in this area (Banks, 2005; Banks *et al.*, 2008b).

Initially developed by Kennedy and Eberhart (1995), Particle Swarm Optimisation (PSO) is a cooperative population based search algorithm. An initial population of potential solutions (particles) are randomly located throughout the problem space. The fitness of each particle is computed according to the given problem, then a combination of the particle;s current velocity vector, previous best solution and the best solution found throughout the swarm determines the particle's next position. The basic algorithm for calculating the next particle position, x, in the ith dimension is:

$$v_{t+1} = v_i + \varphi_1\beta_1(p_i - x_i) + \varphi_2\beta_2(p_g - x_i) \tag{1}$$

$$x_{t+1} = x_t + v_{t+1} \tag{2}$$

where constants φ_1 and φ_2 determine the balance between the influence of the individual's knowledge (φ_1) and that of the group (φ_2); β_1 and β_2 are uniformly distributed random numbers defined by some upper limit, β_{max}, that is a parameter of the algorithm; p_i and p_g are the individual's previous best position and the group's previous best position; and x_i is the current position in the dimension considered. This algorithm was later modified to prevent the particles from accelerating out of the problem space (Shi and Eberhart 1998) through the inclusion of an inertia weight, ω, as follows:

$$v_{t+1} = \omega v_t + \varphi_1\beta_1(p_i - x_i) + \varphi_2\beta_2(p_g - x_i) \tag{3}$$

Although a significant amount of further research into understanding and improving the technique has taken place, the most significant with regard to the cooperative behaviour exhibited by the agents in this work was Kennedy's (1997) investigation of the PSO learning metaphor. Kennedy produced several alternative learning models: the cognition-only model in which the individual only uses its own prior experience (Equation 4); the social-only

model, in which the individual uses the group's experience (Equation 5) and the self-less model, where the individual's experience is not included in the social-only model, *i.e* $i \neq g$.

$$v_{t+1} = v_i + \varphi_1 \beta_1 (p_i - x_i) \tag{4}$$

$$v_{t+1} = v_i + \varphi_2 \beta_2 (p_g - x_i) \tag{5}$$

Pugh and Martinoli (2007) applied PSO to multi-robot search by utilising the mechanics of PSO to move simulated robots (each robot being a particle in the swarm) in search of a target within a bounded area. The fitness function for the simulation was provided by the intensity of the target, *I(d)*, using the function given in Equation 6 where P_0 is the source power, d is the distance between the robot and source and η is a sampling of Gaussian noise.

$$I(d) = \frac{P_0}{d^2} + \eta \tag{6}$$

This approach was tested with robots that had perfect knowledge of their position via GPS and also with robots that could only know their position relative to neighbours. In the latter scenario, swarm members could not communicate their previous positions because of not having a coordinate system or knowing where the targets were positioned. To overcome this, they communicated the current and immediate past target intensity to near neighbours who then calculated the relative target location. Both GPS and relative position systems showed success, however, no indication is given as to performance without target detection or in the presence of multiple targets.

The issue of multiple targets was tackled by Banks (2005), later reported in Banks *et al.* (2008b), by applying Kennedy's cognitive models described by Equations 4 and 5 to PSO controlled agents. These models were important to resolve conflicts where several agents concurrently detected separate and remote targets (if targets were close the effect would have less impact). In a full PSO model such conflicts may cause an agent to be drawn away from a potential target engagement by externally detected targets. To resolve this, the agents were endowed with three basic modes of operation: search mode, in which no targets are detected within a range such that pursuit is viable; swarm mode, where targets are detected but not by the individual, and the social-only model is applied; and engage mode, where the individual detects targets and forsakes all other group knowledge through the application of the cognition-only model.

Whilst the application of these models provided resolution to individual dilemmas, all agents responded to target detections even when they had no chance of intercepting the target. Systemically, this was problematic, since a target detected early could divert all defenders to a particular region of the defended space and leave other regions uncovered and susceptible to exploitation, *i.e.* the system would over-commit to early detections. To alleviate the problem a dynamic neighbourhood scheme known as the Launch Success Zone (LSZ) was developed to enhance the effectiveness of cooperation. Essentially the LSZ controlled the switch between search-mode and swarm-mode; the agent must be inside the LSZ to enter swarm-mode.

Although empirical results suggest that PSO based guidance was beneficial, when compared with systematic alternatives, two weaknesses were noted: the success of the system was hampered by its limited ability to initially detect targets, and, where detection was lost, the system 'forgot' the target existed and made no attempt to reacquire it. These findings formed the foundation for the research presented in this thesis.

3 Natural Search

This section reviews relevant literature concerning natural strategies for search. Various forms of cooperation and group behaviour are introduced, followed by in-depth consideration of various search behaviours found in the natural world. Common methods of communication to facilitate cooperation within a group are then examined, followed by a brief consideration of indicative foraging behaviour, to illustrate how some of the behaviours highlighted in the preceding sections combine to produce effective cooperative predation. Finally, there is a discussion of the exploitation of these behaviours, providing examples of how some have been utilised to develop natural computing paradigms; some less well known examples are deliberately selected, to highlight the breadth of research already taking place beyond autonomous vehicle exploitation.

3.1 Natural Cooperation and Group Behaviour

The manner in which (non-human) animals cooperate has fascinated biologists for many years. Cooperation among individuals of the same or different species may at first seem contradictory to the evolutionary principle of the survival of the fittest. Stevens *et al.* (2005) argue that all forms of cooperation must benefit the cooperators, whether directly, for example in mutual grooming, or indirectly, where cooperation helps to ensure genetic survival (an example of this being the division of labour in eusocial insect colonies, where individuals forego breeding in order to provide some other function that benefits their siblings).

Dugatkin (1997) reviewed the four main, accepted, types of animal cooperation: reciprocity, where an individual performs an action with the expectation that the beneficiary will reciprocate (*e.g.* impala grooming); by-product mutualism, in which cooperation only occurs when it provides immediate benefits that outweigh possible benefits that could be accrued uncooperatively (*e.g.* lions hunting in packs); kinship based cooperation, where individuals cooperate for the benefit of their genetic relatives (*e.g.* ground squirrels alarm calling), and group selected cooperation where an individual works at personal cost, but for the benefit of the group (*e.g.* foraging ants). In addition to these main types, animals may also cooperate to avoid some form of punitive measure. For example, Hauser (1992) presented evidence that

rhesus monkeys announce food capture to avoid attacks that occur when they withhold the information.

Exactly why animals congregate and operate as a group is debateable, due to several interrelated factors, such as protection from predation, environment and availability of food sources (Morrison *et al.*, 1987). However, this work concentrates on the strategies and tactics that animals use to search (usually for food). First, it is useful to establish how individuals interact, and, especially whether the group consists of an aggregation of individuals, a group of individuals working together for mutual benefit, or a team of individuals each with their own specialised role.

Aggregations can occur where there is restricted access to a resource, and individuals are forced together; in such cases, neither the individual nor the group benefit from the grouping. Freeman and Grossman (1992) studied group foraging by stream minnows, and noted that groups at feeding sites were actually just an aggregation of individuals who arrived and departed individually or in small groups because that was where food was located. This type of behaviour does not accrue any added value and is therefore not considered further here.

The defining attribute of group behaviour is that the individuals perform the same task concurrently. This is relatively easy to achieve and can be very effective. Anderson and Franks (2001) categorised two broad groups as either those that form physical structures (known as self-assemblages) such as bridges, bivouacs and rafts (Anderson *et al.* 2002), or those that combine to form a more effective force, for example, ponerine army ants (Maschwitz *et al.*, 1989), that, upon encountering a suitable foraging location, overpower prey much larger than themselves through swarming.

Where animals work in teams, there is a requirement for a more structured approach than needed when working alone, since individuals or sub-groups have separate sub-tasks to perform, which must be coordinated to achieve the overall team goal. In this highly complex behaviour, the individuals may be working sequentially or concurrently. Anderson and Franks (2001) note that team members are not restricted to the performance of a single simple task; they can change their role to fulfil group needs, although performance can be degraded where an individual performs an unfamiliar task. Gordon (2003) reviewed how work was organised in colonies of harvester ants, and indicated that, whilst there was no central coordination, each individual based its behaviour on its encounters with the environment and

other colony ants, forming a form of *ad hoc* dynamical network. For example, to coordinate foraging, the following activities take place: initially, a group referred to as patrollers (pro-active foragers) leaves the nest in search of food sources; on encountering a food source, the patrollers return to the nest; reactive foraging ants leave the nest and forage along the trails taken by the patrollers; as foraging continues, the patrolling ants maintain the trails by encounters with the foragers; once the foragers cease to meet patrollers the foraging stops and all ants return to the nest. It has also been shown that group foraging may not always be the optimum strategy due to the potential cost of information sharing. Dechaume-Moncharmont *et al.* (2006) developed mathematical models of insect foraging that indicated that there may be times when, instead of waiting at the nest for foraging information, reactive foragers would fare better if they left the nest and searched for food proactively. This argument is supported on the grounds that, where food availability is transient, the time taken to return to the nest can be more costly than the returns made through cooperation.

Individuals that work both alone and in groups can also be seen to change their behaviour to optimise their current mode of operation. In studies of social monkeys, Boyer *et al.* (2004) observed that the trajectories of individuals altered to a more optimal Lévy search pattern when foraging in groups (from $\mu \approx 1.5$ in individual foraging to $\mu \approx 2.1$ in groups, see Section 4.3.1). This switch in pattern is believed to be due to the increase in resource information available to the group and is considered the most efficient strategy under those circumstances. It is important to note that, although foraging became more efficient with larger groups, it was resource availability that affected group size; fewer resources resulted in smaller groups, although the group members would change throughout the day. It is useful to note that competition for the available resources does not always dictate the size of the group, members of a group can accrue other benefits such as mate availability. Delestrade (1999) conducted experiments using alpine choughs to ascertain the effect of competition and food availability on flocking behaviour. She found that where food sources were small in size and plentiful (competition was low) the alpine choughs did not require social interaction and the efficiency of belonging to the flock was high. When the food sources were large and patchily distributed, this efficiency was reduced because the number of birds visiting the site remained the same, but not all birds gained access to the food. In these experiments, where competition was high, the female members stayed with the group, which was dominated by the males, despite restricted access to food. In this case, Delestrade concluded that it was the mating

benefit that was more important. Thus, some caution must be applied when observing natural phenomena, such as group feeding, in search of new natural computing paradigms, since misleading associations could be made easily; in this case, are the group members acting cooperatively to gain mutual benefit or do individuals gather in conspecific groups simply because of the availability of food?

3.2 Search Strategies

The success of a species depends on many factors, including its ability to search effectively for resources, such as food and water, in a given environment. Pianka (1966) broadly classified search strategies into sit and wait (for ambush) or foraging widely (for active searchers). Butler (2005) reports that, in general, active foragers range widely, eating slower prey, and expend high levels of energy in the process; in contrast, sit-and-wait foragers capture fewer prey (which are usually the active foragers) but expend less energy in doing so.

Differing species utilise a variety of search strategies depending on their physiological characteristics, the environment, and the resources available to them. These can range from a continuous motion search (also called cruise search), such the long range foraging patrols of the albatross (Stahl and Sagar, 2000), through saltatory search, such as tyrant flycatchers (Fitzpatrick, 1980), to the ambush type predation of the scorpion (Skutelsky, 1995). It is important, therefore, to understand the types of environment, in order to establish which type of search is appropriate. The following subsections briefly examine pertinent characteristics of search environments before considering the strategies employed in nature to address search tasks within them. O'Brien *et al.* (1990) provide a useful summary of search strategies and the appropriate circumstances under which they should be adopted.

3.2.1 Search Environments

Resources can be distributed uniformly or regularly, and may vary temporally and spatially; where they are distinctly grouped they are termed patchy. Where distribution is regular and stable it is advantageous for the predator to be easily conditioned to the most effective strategy and then to maintain it. Papaj and Rausher (1983) describe two types of conditioning that may take place: conditioning of response, where behaviour is learned through repeated encounters with stimuli, and conditioning of perception, where the probability of perception is increased through repeated encounters. The effect of this was shown in experiments with

butterflies (Weiss and Papaj, 2003), where butterflies alighted more frequently on model flowers with colours that they had been conditioned to believe would be more profitable in terms of nectar foraging and ovipositioning. The work indicated that, whilst efficiency improved in same context activity, it led to mistakes occurring following both behaviour switches (between foraging and ovipositioning) and environmental context switching.

For patchy or irregular resources, a predator will avoid becoming conditioned because to do so would leave it too dependent on a particular, possibly unreliable, resource. Burns (2005) illustrates how this can occur in conspecific group foraging (in this case, with bees) where he argues that it can be beneficial for a group when individuals within it display different foraging behaviours to exploit the variable resources available, *i.e.* the group itself does not become conditioned.

3.2.2 Sit-and-Wait

There are times when active searching does not provide sufficient energy gains to warrant the effort required; for example, in a harsh environment, or where prey speed or agility make it highly unlikely that the predator would capture the prey in an open pursuit. Naturally, in a predator / prey situation, the predator must be sufficiently camouflaged to avoid being detected until the prey is close enough to attempt capture.

As with most natural search strategies the tactic is often used in conjunction with other tactics to provide maximal benefit, such as in the case of the sidewinder snake (Secor, 1994). In this study, sidewinders often performed either a straight line or random walk search, both in combination with chemoreception, to locate a suitable ambush site at the entrance to a prey burrow. Once located the Sidewinder would partially bury itself in the sand to provide camouflage, environmental protection and an anchorage from which an attack could be launched.

3.2.3 Continuous Motion Search

Bell (1991) offers four theoretical continuous motion search strategies: random walks, where the direction of each step is changing at random; straight line walks in a random direction for some distance; systematic strategies, such as spiral patterns; and kinesthetic input mappings. The latter case, in which the animal uses knowledge of its own previous movements to search for a return location such as a burrow, is part of a navigation system known as path

integration or dead reckoning, which may include idiothetic (internal cues) or allothetic (external cues) mechanisms (Merkle *et al.*, 2006). Animals may use a combination of these in a variable strategy determined by individual needs, group needs or environmental factors.

After examining several experiments designed to assess the affect of spatial distribution on search patterns, Bell concludes that the dispersal of resources affects the locomotory search pattern in that the search patterns of those insects under examination was sufficiently plastic to be constrained by the resource dispersion rather than inherent patterns.

The following subsections elaborate further on each of the continuous motion search tactics.

3.2.3.1 Random Walk

A random walk may consist of either Brownian motion, where the direction of one step is independent of its predecessor, or Lévy paths, where there is a lower probability of a change in direction. For hidden fixed targets, Lévy paths have been utilised by several species when foraging for food. Examples of such behaviour have been observed in spider monkeys (Ramos-Fernández *et al.* 2004) and jackals (Atkinson *et al.* 2002), amongst others. Lévy path lengths (l_j) can be described as being stochastic with continuous probability (P) having stationary and independent movement increments with distribution:

$$P(l_j) = l_j^{-\mu} \qquad \text{(where } \mu \approx 2\text{)} \tag{7}$$

Viswanathan *et al.* (1999) showed that Lévy search is a statistical random search that outperforms normal random search (*i.e.* Brownian motion, where $\mu \geq 3$) under certain searcher / resource conditions, particularly where prey density is low. This may appear to be counter-intuitive since Brownian motion paths produce more homogenous coverage and would be more likely to find widely dispersed targets. However, when the general overriding principle for foragers is that of maximising efficiency, where sufficient food is available, then more intense, systematic, local search would require less effort (see Figure 4). Viswanathan *et al.* (2001) further investigated the efficiency of the approach under such circumstances, assessing whether the argument held with short leg lengths and different numbers of dimensions. Findings were that the optimum of $\mu \approx 2$ still held under the scenarios tested. Further extending those findings, Marthaler *et al.* (2004) reported that, in scenarios where targets were removed, searches were more efficient when μ was approaching 1; *i.e.* less tortuous paths were more efficient.

Further comparisons between Lévy and Brownian motion searches in a biological context were made by Viswanathan *et al.* (2002) and Bartumeus *et al.* (2002), to assess predator success against varying prey densities and motion types. The relationships among these factors are complex, but for a given target density, they indicated that, for Brownian motion targets, predators that are larger and faster are better served by Lévy search, whilst for smaller and slower predators, Brownian search was a better option. Where targets moved using Lévy motion, size ratio was unimportant and velocity ratios dominated, with Lévy search being preferred for slower targets. Target density was regarded as the prime factor, with Lévy search only being suited to lower density targets, although a target with Lévy motion reduced the effectiveness. Further, where targets have a proportionally higher velocity (*i.e.* super-diffusive movement), Lévy search is less profitable; indeed, Brownian search, or even no search at all (sit-and-wait), is usually adopted, since the prey is just as likely to move into sensory range without the higher effort required for Lévy search.

3.2.3.2 Straight-line Search

A strategy that includes straight line search can be effective when there is no other information regarding resource location and provided the searcher does not pursue the same line for an inappropriate length of time. If the time period between direction changes were too short, the search area would be too constrained; if it were too long, the searcher may unnecessarily follow an unproductive path. It can also be shown that the strategy is inefficient when the resource density and sensory acuity are small, since the probability of target detection (T_d) is:

$$T_d = \frac{\arcsin(\alpha / r_0)}{180} \tag{8}$$

where α is the detection radius and r_0 is the distance between the start point and target. Using several case studies, Bell (1991) indicates that when the searcher does decide to move, it usually uses a strategy, such as alternating left and right turns (with a mean angle of 35°), that reduces the possibility of revisiting previously searched territory. Further, this strategy may be useful in group search scenarios since the decreased turning rate may make the search more efficient than strategies with higher turn rates, since the likelihood of each predator covering the same ground as another predator is also decreased.

3.2.4 *Systematic Searches*

Bell (1991) discusses three hypothetical systematic search tactics. The expanding square, the spiral, and parallel sweeps (Figure 4). To be most effective, the search path should be spaced at twice the sensory acuity of the searcher. These strategies are most viable where the target is moving more slowly than the searcher and there is no spatial or temporal information.

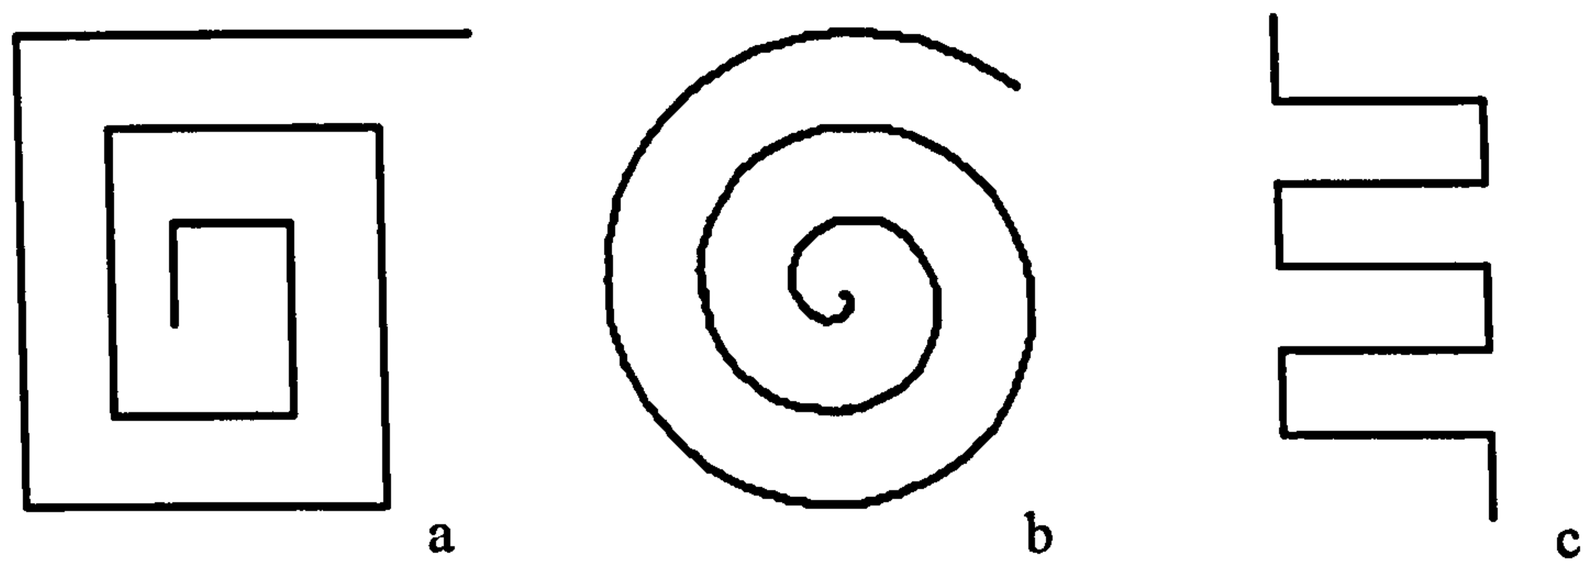

Figure 4. Systematic search strategies: (a) expanding square, (b) spiral and (c) parallel sweep (Bell (1991), Figure 6.1).

Whilst there is a paucity of evidence for such tactics in the literature, Müller and Wehner (1994) indicate that desert ants use spiral searching as part of a homing strategy that includes Path Integration (PI). In this case the ants use the PI strategy to calculate a global return vector, which is used to return to the nest after a foraging sortie. The return vector is not always sufficiently accurate, and when they do miss the nest they enter an expanding spiral search.

3.2.5 *Path Integration*

An animal will often need to search for a location that it has moved away from, for example, a previously lucrative location or its burrow. Possible tactics could be to employ one of the search patterns discussed, or to navigate using remembered landmarks, but these can be inefficient or error prone (a landmark might move, for example). Alternatively, the animal can map its outbound movements using idiothetic and allothetic mechanisms, and calculate a global return vector by an innate path integration (PI) ability (see Etienne and Jeffery, 2004, and Merkle *et al.*, 2006, for reviews). Two PI models have been proposed in both Cartesian and polar coordinate systems (only Cartesian models are shown here for comparison). Both models only require the animal to sense the input variables of forward speed (v) and angular turn rate (ω). The basic models are as follows.

Geocentric models are based upon the animal's perception of its movement in relation to its environment (allothetic inputs) where it continuously updates its global return vector through the integration of movements calculated using trigonometry. Research by Freake (2001) indicates that the parietal eye (a photo-sensitive organ present in some animal species) plays a significant role in the homing behaviour of the sleepy lizard, suggesting that a form of celestial compass could be in use; this could be used to facilitate a geocentric PI system. In the geocentric PI model (Mittelstaedt, 2000) the global return vector **G** is:

$$\mathbf{G} = \mathbf{P}_0 - \mathbf{P} \tag{9}$$

where $\mathbf{P}_0$ is the origin [0,0] and

$$\mathbf{P} = \int (\cos\phi, \sin\phi)\ ds \tag{10}$$

if integration is over walked distance (s is the arc length), or

$$\mathbf{P} = \int (v\phi)(t)\ dt \tag{11}$$

if integration is over time, and

$$\phi = \int \omega\ dt \tag{12}$$

In contrast to geocentric models, egocentric models require less complex calculations, since the animal uses idiothetic input signals to calculate its forward speed and angular turn rate. Merkle *et al.* (2006) argue that, although recent improvements to the geocentric model described above have included results consistent with mammalian neuro-activity, it has yet to be shown that animals have the ability to physiologically perform the necessary trigonometric calculations, the egocentric model is physiologically possible and simpler to calculate, and therefore more plausible. For the egocentric model the update to the global return vector (X, Y) is:

$$\frac{dX}{dt} = -v + \omega Y \tag{13}$$

$$\frac{dY}{dt} = -\omega X \tag{14}$$

These geo- and ego-centric models need not be mutually exclusive; Menzel *et al.* (2000) argue that some honeybee navigation can utilise both systems. Generally, the argument is

presented that honeybees only apply egocentric models, since most evidence does not show that novel shortcut behaviour takes place when bees have been trained to fly a route to and from feeding sites. Findings from the experiments of Menzel *et al.*, combined with supporting evidence from other work, suggests that they store relational information on landmarks and if they do not receive sufficient orientation to employ ego-centric based navigation they can use an alternative geo-centric system. This geo-centric system is believed to further use combinations of reference systems depending on the accuracy of available sources, such as scent deposits or visual cues for their 'odometer' (Vladusich *et al.* 2006).

Although these models present efficient mechanisms, they are not perfect, and a returning animal may not return to its start position. In this case another local search mechanism can be employed, such as a local random walk or a systematic search. As previously indicated, Müller and Wehner (1994) show this type of behaviour in desert ants, which perform a spiral search where PI fails to locate the return location of a foraging sortie.

3.2.6 Area Concentrated Search

Area concentrated search (ACS), also known as area restricted search, is a variable strategy in which a predator will attempt to exploit the natural tendency of prey to appear in patches, by concentrating their search effort in a particular area upon encountering an individual prey item. Searching often begins in a ranging mode, with long leg lengths and small turning angles, followed by intensive searching initiated by some kind of cue. This strategy can be termed *win-slow:fail-fast* since, when a successful encounter occurs the predator slows into intensive search, but speeds into extensive search when an encounter fails (Krakauer and Rodríguez-Gironés, 1995).

An example of this behaviour is that of the juvenile plaice (Hill *et al.*, 2000), which use localised cues, although it is not known what sensory apparatus is used, to alter their search pattern from extensive to intensive search. This switch, on average, shortens pause time by 15%, increases move duration by 14%, reduces speed by 22% and increases median absolute turn angle by 33%. The effect of this change is that the predator can effectively search a small area rapidly, since it knows that once it has discovered one item of prey, it is likely to discover others nearby.

Krakauer and Rodríguez-Gironés (1995) presented an ACS model to examine several foraging rules:

Tracking errors. These occur when the forager searches extensively where resources are good or intensively where resources are poor. With ACS it is possible to reduce tracking error by capitalising on patchy distributions through initiating intensive searching only on encountering prey items. The amount and types of error will therefore depend on the patchiness of the resource.

Learning. The forager is required to learn about the environment and distribution of the resources to maximise exploitation. This activity must be balanced with the time available for, and value to be gained from, exploitation.

Evaluation of information. The forager must decide whether there is any value in the information it gathers or to apply a general strategy instead. As the patchiness increases, so does the likelihood of prior assumptions being incorrect, and the value of time spent in assessing the prey distribution and selecting an appropriate exploitation strategy increases.

Functional response. The prey intake in response to its density. The selection of a resource dependent strategy can alter the response curve from one in which, at high prey density, the limiting factor is the amount of time spent handling the prey, to one in which the predator spends more time searching.

Further to these rules, Krakauer and Rodríguez-Gironés (1995) also assessed whether the standard *win-slow:fail-fast* ACS strategy was always optimal. They argued that, where the resource distribution regularly alternates between good and bad, the reverse could be true (*win-fast:fail-slow*), since the switch to intensive mode is only worthwhile where the chances of successive encounters are high.

3.2.7 Saltatory Search

Animals do not always need to have located prey to enter an intensive search phase; the intermittent scan and relocate foraging method, also known as saltatory search, is a widely adopted strategy when the targets are sparse and difficult to detect (Bénichou *et al.*, 2005). The strategy consists of two distinct phases, a search phase, in which the searcher pauses to explore its immediate vicinity (within the scope of its sensory acuity), and a motion phase, in which the searcher moves more quickly, such that it is not normally possible to detect hidden prey.

During the search phase, the predator may have to scan a given area many times in order to detect a well hidden target. Bénichou *et al.* modelled this phase using a diffusion process (random walk) that was notable in that it was a sensory search rather than the animal's physical movement. Results of the modelling closely matched that of several animal species, indicating that the diffusive search was an appropriate selection for the search phase.

The motion phase is notable in its similarity to the klinokinetic strategy used by the juvenile plaice (Hill *et al.*, 2000) when changing from intensive to extensive searching during an ACS, although Bénichou *et al.* (2005) adopted an efficient relocation strategy that had been observed in nature – a straight ballistic motion to the next search area. The results of the modelling indicated that the optimal relationship between the average duration of the search and motion phases was bimodal (motion phase duration equals search phase duration to the power 3/5 or 2/3).

The similarity of this strategy with ACS is striking, but the difference is vital to the potential exploitation of either in a natural computing context. To summarise, ACS is suitable where predator acuity and prey visibility are such that a prey item can be detected during extensive search, whilst saltatory search is suitable where sensory acuity is limited or prey visibility is low, and thus unlikely to be detected during extensive search.

3.3 Knowing When to Give Up

Once a predator has assumed a search tactic, it must then decide when to switch to another or move location, which is known as the giving up time (GUT). The decision could be based on factors such as estimated prey numbers, capture rate, or time (Bell, 1991). For example, juvenile plaice (Hill *et al.* 2000) use a short term klinokinetic strategy; that is, they will only exploit a patch for a short time before swimming away (by reducing their turn rate). This, it was argued, is because other external pressures are operating; if they loiter to fully exploit the patch, there is an increased possibility of predation. Charnov (1976) presented a more generalised departure strategy known as the Marginal Value Theorem (MVT). This states that, where an omniscient predator depresses prey availability, departure from a patch should occur when the capture rate drops to the average capture rate for the habitat.

A GUT also needs to be optimised for sit-and-wait predators since, in common with continuous motion predators, sit-and-wait predators often do not know prey density within a

patch, and, furthermore, prey size may not be known until capture, and to give up before any capture is a waste of the expended energy (Iwasa *et al.* 1981). Nishimura (1999) explored, using modelling, the behavioural variability of GUT in a stochastic environment for sit-and-wait predators and indicated that, when entering an unknown patch, the predator should employ a variety of strategies to learn the optimal GUT. Nishimura asserts that the number of decisions made in search of the optimal GUT strategy would relate to the uncertainty of the environment and from this it is possible to further assert that to be optimal in highly stochastic environments, a predator should have an array of potential strategies to explore, and the efficacy of decision making in the pursuit of the optimum is vital.

3.4 Capitalising on Environmental Cues

A priori information regarding certain search target attributes can provide additional cues that can be used to improve the likelihood of success. These cues might be topographical or other environmental factors. For example, harbour seals have been shown to exploit their knowledge that salmon migration runs occur under certain tidal conditions and that during these runs the salmon densities are increased in constricted channel areas (Zamon, 2001).

Environmental cues can also be used to avoid intra-specific competition. This is illustrated by the foraging behaviour of *Frugivorous lepidopteran* larvae (Marchand and McNeil, 2004). These larvae feed on cranberry fruits and exhibit plastic foraging behaviours that are cued on fruit size, spacing and colour. The first two are self-explanatory: the larvae will exploit a patch where the resource can provide sufficient return on foraging effort, *i.e.* select the largest fruit in the densest patches. The fruit colour is significant since the larvae feed from the inside of the fruit to avoid predation and the activity of feeding larvae affect the fruit by turning them prematurely red. Other foraging larvae use this cue to avoid wasting energy and possible intra-specific confrontation. The colour cue is so important that during laboratory controlled experiments, far higher numbers of larvae entered small green fruits than large red ones (even though the larger fruits offered potentially higher gains).

3.5 Variable Strategies and Patterns

The reasons for variations in search behaviour are diverse, but usually related to a changing environment or physiological state. There are a variety of reasons why a variable foraging strategy could be optimal. For example, Skutelsky (1995) studied the scorpion *Buthus*

occitanus isrealis and discovered that the subject had added active foraging to its usual sit-and-wait strategy, to maximise its predation of termites. The reason for this adaptation is that when a group of termites detects a predator they all retreat to the safety of the nest. Within seconds, by changing to an intensive search mode, the scorpion was observed to kill up to eight termites compared to the single kill that would be possible with a pure sit-and-wait strategy. This discontinuous variation is an example of an Evolutionarily Stable Strategy (ESS) as proposed by Maynard Smith (1982). Whilst the scorpions used prey defined variation, female red-faced warblers vary their strategy due to breeding restrictions (Dobbs and Martin, 1998). The warblers were observed capturing stationary prey from a perch (sit-and-wait) whilst egg laying or incubating, and hover-sallying for stationary prey during the nesting period. This variation was believed to be a result of changes in time and energy requirements, since the former strategy occurred whilst the female nested alone and the latter was during a period when the male assisted with the hatchlings.

It is not just the foraging strategy that can be adapted; the search pattern itself can change in response to dynamic changes to the environment, although this requires that an individual acquires information about the environment and the optimality of its current strategy, and then makes a decision about the proposed pattern's efficacy. In stable, homogenous environments, this is a fairly straightforward decision, but real situations are rarely simple and the learning / decision making process itself requires energy, and the individual must also decide, therefore, whether a net energy gain would result.

3.6 Repeat Searches

In dynamic environments it can be important that an area is revisited to search for new resource occurrences. Bell (1991) describes several strategies:

Cropping. Cropping is the process of depleting a resource, moving on, and then delaying a return until the resource has had time to renew. When implementing cropping, it is argued that a systematic approach is most efficient since resource renewal increases exponentially over time (notwithstanding unrelated environmental pressures).

Traplining. This is similar to cropping in that it involves returning to the locations of previously successful encounters. In traplining, however, the searcher learns a successful foraging route that allows resources to recoup between foraging visits.

Partial depletion. The predator can choose to only partially deplete the resource during a foraging visit, leaving resources available for return visits. This has the added advantage that the resource will replenish more rapidly due the exponential nature of resource rejuvenation.

All three share a common principle: that the resource is allowed to renew between visits. Whether they are employed, very much depends on the goal of the target application.

3.7 Competitor Avoidance

It is not beneficial for individuals to compete for the same resources. This can be seen, for example, in insect oviposition site selection; when the eggs hatch there should be sufficient local food for early larval survival. When searching for a site, the insect will examine the amount of resources available and maximise them by avoiding superparatism where other, better, resources are available. Cope and Fox (2003) conducted experiments that indicated that the seed beetle uses a combination of resource size, competition (*i.e.* existing eggs) and their own need to egg lay (older beetles were more likely to lay since it is better to lay sub-optimally than not to lay at all).

This principle can be applied to improve group searching, since the competitors for resources can be considered to be cooperating in the goal of ensuring the survival of the species. Thus, if a group of agents are searching for hidden resources, it is not efficient for them to cover the same ground (assuming a static environment) or to converge on small incidences of the resource.

3.8 Communication Strategies

The manner in which inter-group communication occurs not only affects the efficiency of information passing but can also produce emergent behaviours in its own right. Communication strategies are often multi-modal and can be either direct, where two or more individuals pass information and receipt occurs within the same time frame as transmission, or indirect, which is usually more subtle, and occurs where a message is left, and any group member can receive it at a later time.

3.8.1 Indirect Communication

Grassé (1959) first documented the method by which social insects communicate indirectly to achieve the complex emergent behaviour required to perform feats such as termite nest construction. He termed the process stigmergy (from the Greek *stigma*, goad, and *ergon*, work). To achieve stigmergy, multiple individuals interact when one modifies the environment and the others react in some way on encountering the modification. Essentially this description highlights the voluntary nature of response, for example, in the case of ants a receiver may choose to ignore a pheromone signal because the intended action is outside the responsibility of its caste (*e.g.* a soldier ant may well ignore food trail signals). Holland and Melhuish (1999) expanded on Grassé's concept of stigmergy occurring where activity is stimulated through changing the environment. They argued that it is the outcome of the behaviour that is important and can be manifest in three ways:

- It may affect an agent's choice of action.
- It may not affect the agent's choice of action but how it performs that action.
- It may affect neither of the previous points but influence the outcome. For example, by placing a physical barrier to alter the path of fellow group members.

In social insects the stigmergic reaction usually involves the application of low level rules; in the termite nest building scenario, a termite drops a piece of chewed earth when encountering a raised piece of ground (generated by other termites dropping pieces of chewed earth) until large, complex structures emerge. Stigmergy is most commonly associated with ants and termites but is also used by other insects such as bees (Dornhaus and Chittka, 2004) and wasps (Reed and Landolt, 2000; Steinmetz *et al.*, 2002). Stale data is often avoided via mechanism such as evaporation, where the importance of response time can be tied to evaporation rate.

3.8.2 Direct Communication

The obvious forms of direct communication in social groups are through aural and visual cues, although direct communication can occur using all sensory cues, including the use of pheromones. The methods discussed in this section are not exhaustive; however, they do cover some of the direct communication methods that have already been used in natural computing. Also notable is that the latter two strategies (tandem calling and information

centres) can be regarded as higher level strategies that make use of lower level forms of communication.

3.8.2.1 Trophallaxis

Some species of social insects communicate successful foraging information via the transfer of bodily fluids; this may be proctodeal (faeces sampling), or stomodeal (regurgitation). Dejean *et al.* (2005) report on research that indicates that trophallaxis is employed by tramp ants (and even simulated trophallaxis in less phylogenetically advanced sub-families) to pass information regarding good foraging sites located in the territory of interspecific competitors. In this instance, the odour of the competing group is contained within the transferred fluid, alerting the receiver, and therefore expediting exploitation of the site in order to avoid confrontation, and thus gain competitive advantage. Bees also use trophallaxis, as a learning medium for odour signalling (Gil and De Marco, 2005). This epigenetic learning occurs as the transferred liquid contains the odour used to mark the source, and the receiver associates the odour (via olfactory stimulation) with the reward. Experiments indicate that the process is as successful as experiential learning in effectuating long term memories with the properties of the liquid (odour and reward concentrations – the latter being sucrose solution during experimentation) affecting association intensity.

3.8.2.2 Honeybee Dancing

When a honeybee locates a good source of nectar, it returns to the nest, noting the resource's distance from the hive and its direction with respect to the sun. On entering the hive, it performs its dance: the distance is communicated by the speed of the dance and the direction by the dance orientation (Von Frisch, 1967). Several types of dance have been identified, including the waggle dance described here, and experiments have indicated that the precision of the dance correlates to the availability of resources (Dornhaus and Chittka, 2004). Further, Farina and Wainselboim (2005) examined the dancing behaviour of honeybees and indicated that some honeybees, upon witnessing the dancing performed by a successful forager, engaged in stomodeal trophallaxis in order to sample the nectar quality, possibly influencing the sampling bee's decision to visit the resource.

3.8.2.3 Tandem Calling

Tandem calling, also known as tandem recruitment or tandem running, is a communication method that can be used where an insect wants a specific individual to assist with a particular task. The actual contact method can vary, and multiple strategies can be employed, including chemical signals, touch or visual cues, and often comprise a combination of stimulation signals to enable the recruitment and directional signals (De Meyer *et al.*, 2003). The strategy can be expanded to include many recruits, which may be led using a visual cue as stimulation followed by pheromone trail laying for guidance, and has been identified in a variety of species such as ponerine ants (Fresneau, 1985) and stingless bees (Lindauer and Kerr, 1960). This may still be considered as direct communication, since the action of following has already been stimulated, and individuals would not follow the trail without such stimulation.

3.8.2.4 Information Centres

When animals forage unsuccessfully, it may well be in the best interest of the species that other, more successful, animals communicate likely foraging locations to them. One such communication method is the communal 'information centre' (Ward and Zahavi, 1973). In this hypothesis, animals are believed to communicate successful foraging whilst visiting central communal areas, allowing less successful members to follow them to the resource at a later time. Although this hypothesis has seen some controversy, evidence for its existence is accumulating. Wright *et al.* (2003) conducted field experiments on communities of ravens in which food carcasses were laced with coloured pellets. They noted that dominant birds feeding at the carcass would perform acrobatic displays at the communal information centre at pre-roost aggregations. Following this, increased numbers of birds would follow the lead bird to the feeding area.

3.9 Example Behaviours

The natural world is awash with different foraging behaviours and to explore all of them here is beyond the scope of this work. It is worthwhile, however, to sample a few species to highlight how the some of the behaviours discussed above are combined, and to examine some of the material available to natural computing for inspiration in developing potential search algorithms. For example, ponerine ant foraging has already been used, and the algorithm developed from it is discussed briefly in Section 3.11.1.

3.9.1 Ants

The many species of ants display an array of foraging behaviour; the strategies of two types of ponerine ant are described here, the army ant style foraging of the genus *Leptogenys* and the partitioned space search of *Pachycondyla apicalis*.

Maschwitz *et al.* (1989) studied the nocturnal foraging of the South East Asian *Leptogenys* sp.1. Their large colonies were found to forage in 3 stages termed exodus, main raid and retreat (see Figure 5 for a depiction of the overall search pattern formed). During the exodus phase, new trails are established in narrow, tentacle-like patterns extruding from the colony's current bivouac (they are not nest builders). The outgoing traffic follows (it is assumed) pheromone trails deposited by occasional ants travelling contra-flow. Once the ants on one of the tentacles finds a suitable raiding area, that tentacle becomes very attractive, and the traffic rate along it increases to a peak rate of over one thousand workers per minute in a column of up to 13 ants abreast. At the column forefront, the ants spread out in an organised swarming unit to begin the raiding phase, and, as this moves forward a supporting fan shaped network of trails forms behind it. Whilst the main raid is in progress, there is a balanced flow of traffic to and from the bivouac, with the majority of homebound ants carrying food. Once the raid is complete, the swarm front turns, almost synchronously, to return to the bivouac; those ants still heading for the swarm also turn for home once the homebound flow becomes dominant. This is referred to as the retreat phase.

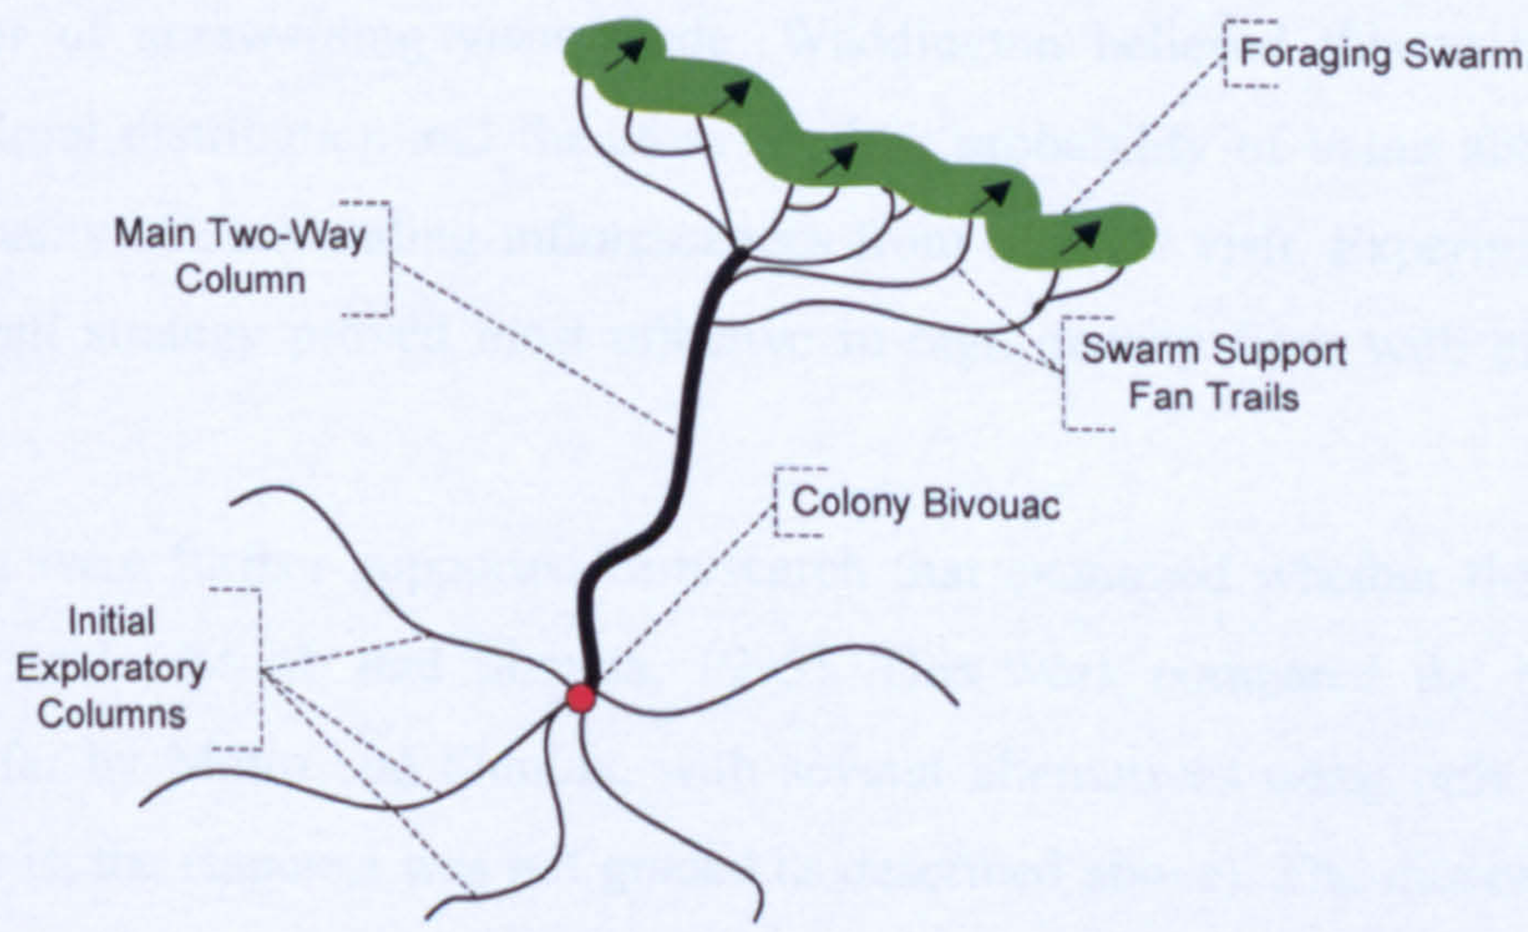

Figure 5. Foraging behaviour of *Leptogenys* sp.1, modelled from photographs and diagrams produced by Maschowitz *et al.* (1989).

The foraging strategy of the ponerine ant *Pachycondyla apicalis*, from Mexico (Fresneau 1985), involves partitioning the global search area into smaller local areas. The colony members (colonies are relatively small, numbering 20 to 100 individuals) begin by setting up hunting sites roughly uniformly distributed about 10m from the nest, each having an approximate radius of 2.5m. Foraging ants then go out to individual hunting sites to look for prey. If they are successful, they memorise the location from landmarks and later return using these landmarks rather than the pheromone trails used by many other ant species. When ants are unsuccessful at one site, they tend to return to other previously successful sites, indicating that multiple site memory is possible. Since the system relies on landmarks, and local foraging at the hunting sites is an individual based activity, interaction among ants during foraging is limited. Occasionally, the nest is moved to another location (providing a global search aspect), using tandem calling to recruit ants from the existing nest to the new one.

3.9.2 Bees

Bees forage for targets (flowers containing varying amounts of nectar) that are static and generally patchy. Initial bee foraging follows an ACS strategy slightly different from that described earlier (Pyke, 1978; Waddington, 1980). Following a visit to a nectar rich inflorescence, a bee will fly a short distance to the next inflorescence, but the direction is maintained; this, it is believed, is to avoid revisiting a site that it has depleted. When an inflorescence provides poor rewards the bee will extend its flight and increase its turn angles to move away from the area. This response is not immediate though, and is graded according to the number of unrewarding visits made. Waddington believed this to be due to the diversity of floral distribution and therefore the low probability of being able to correctly predict the quality of surrounding inflorescences from a single visit. Experiments indicated that this overall strategy proved most effective in high density flora with clumped nectar distribution.

These finding were further supported in research that examined whether the strategy was evolutionary stable (Motro and Shmida, 1995). This work compared the bees' strategy, termed near–far by Motro and Shmida, with several alternatives using only a single level memory (that is, the response was not graded as described above). The movement strategies compared were:

Random search. The next location could be any of the surrounding flowers, excluding the last visited, with uniform probability.

Near search. The next visited location is always the nearest flower to the last visited.

Far search. The next visited is never the nearest to the last visited and is selected randomly with uniform probability.

Near–random search. A combination of the first two strategies where the bee selects the nearest flower subject to sufficient reward from the last, otherwise a random flower is selected.

Near–far search. This is similar to near–random except the strategy switches to a far search following an unrewarding visit. This represents the non-graded version of the earlier work described above.

In a large area comparison with random search, the near–far strategy was shown to be most effective in an evolutionary stable environment where resources were patchy. It was also shown that the advantage increased with increased search area size. A second, smaller, area comparison that included all strategies showed that the near–far search discovered most flowers in fixed length runs in all but the shortest search, in which random search was most effective, since the resources were not depleted sufficiently for repeated visitation to be problematic. In an exhaustive search the near–far search again consistently outperformed the other strategies and was shown to be an evolutionary stable strategy. The worst performer in all cases was the near search because there was a high probability of repeat visits.

As previously reported, once profitable sites have been located the bees make use of path integration to return to hive and communicate the find to fellow bees via waggle dances.

3.9.3 Bacteria

Cooperative foraging behaviour is not limited to insect and high order species; bacteria have also been found to display emergent group foraging behaviours far beyond the scope of the individual. Passino (2002) drew together a wide variety of literature, in particular relating to *E. coli*, to explore the possibility of using bacterial foraging as inspiration for new techniques in distributed optimisation and control systems. Whilst the types of bacteria are numerous, it is the foraging behaviour of *E. coli* that is described here.

Physically, an *E. coli* bacterium consists of the main cell body, the pili (used for the transfer of DNA to other bacteria) and flagella (long, left-handed helix, whip-like projections that enable motor activity). When the flagella rotate counter-clockwise they are pushed forward in a swimming motion. Upon reversing the rotation, the bacterium is pulled, causing it to tumble, allowing direction change to occur. The activities required for such movements are very efficient and, despite being in almost continuous motion, cost less than 1% of the overall energy requirements. When placed into a neutral environment, that is, one without nutrients or harmful substances, the bacteria work independently, tumbling and swimming for equal time periods. Notably, the swimming movement produces a Brownian motion that induces approximately a 30° wander, and the independent tumbling allows the group to search, albeit in an uncoordinated fashion, for nutrients. Upon discovering the presence of a nutrient, the bacteria engage in a specialist movement behaviour known as chemotaxes. In an environment with a constant level of nutrient the chemotaxes is similar to the neutral case, except that the mean swim length and speed increase at the cost of tumbling time. When a positive nutrient gradient is encountered, this effect is further increased but only until the gradient is lost, either to a negative gradient or a flat area whereupon, after a short period, they will resume the equal swim and tumble search behaviour. In summary, the bacteria will always seek positive gradients, even in nutrient rich environments. This behaviour is reversed in the presence of a harmful substance where negative gradients are sought.

Whilst bacteria do not use obvious communication strategies, emergent behaviours are produced when a collection of cells is subjected to certain conditions. For example, where cells are placed in a substance containing high levels of succinate, it causes them to produce aspartate, which acts as an attractor, bringing the cells together to form a kind of swarm that moves together in a protective aggregate. This suggests that even such a basic form of indirect communication, when accompanied by adherence to the simple set of rules described above, can produce beneficial behaviour for the whole group.

3.10 Prey Reacquisition

Whilst in the pursuit phase of hunting, an animal may find that the sensory information about a target becomes obscured, for example, due to limitations in sensory capability (such as eyes failing to adjust to changes in light levels) or because the prey has hidden or escaped. At that point the predator may decide to give up or to apply a reacquisition strategy. There are many

components to target tracking and to fully review them is beyond the scope of this work. However, some elements have been identified as having potential as metaphors for improving the systemic response to loss of target detection, these are: prediction, as used to reacquire targets that can move faster than the visual system can track using smooth movement; anticipation, where an animal can use its knowledge of target movement, likely occlusion and own readiness; visual occlusion, where less certainty exists; and crosswind casting, used when tracking an airborne odour to reacquire the lost scent from an upwind source.

Usefully, and in support of techniques for tracking visually occluded targets, it can be argued that existing ethological research into how predators visually track targets can be applied to the problem of reacquiring targets after detection has been lost. The rationale being that when animals are visually tracking a target they must apply predictive techniques to overcome intrinsic sensorimotor delays and therefore tracking occurs in a cycle of smooth pursuit and predictive saccades. Orban de Xivry and Lefèvre (2007) provide insight into how the two elements combine in a single sensorimotor process in humans. In this case the smooth pursuit operates with a fixed delay and maximum tracking speed, and can operate successfully in scenarios that change within these boundaries. However, where abrupt changes occur or targets are moving quickly, corrective saccades are applied, albeit at the cost of visual acuity during the movement. It is important to note that saccadic movements do not utilise visual feedback from the target, they are totally predictive and are intended to place the target back into the fovea (the region of the retina with the highest acuity). The relationship between the two modes is dependent on the ability of a smooth pursuit system; for example, Orban de Xivry and Lefèvre (2007) compare the human process with that of cats and note that the latter rely more heavily on saccades due to the lower and more variable gain in their smooth pursuit feedback system. Klam *et al.* (2001) reported two saccadic techniques in the sensorimotor process of cats, the tracking saccade and the interception saccade. The former operates in much the same way as for humans and is utilised where target speeds are lower or positional error is deemed small. Interception saccades are more extreme in their movement and intentionally overshoot the predicted target position before attempting to operate smooth pursuit. Figure 6 shows indicative movements for all three strategies. In the first third of the graph, the sensorimotor process can keep track with the target; in the middle third, the target speeds up but the process can keep track by making small catch up tracking saccades; in the final third, where target speed increases still further, interception saccades are made. The

defining features of the overall process are that the least time possible is spent in saccadic motion, where visual acuity is low, and that the target is kept within the fovea for the maximum time during smooth pursuit.

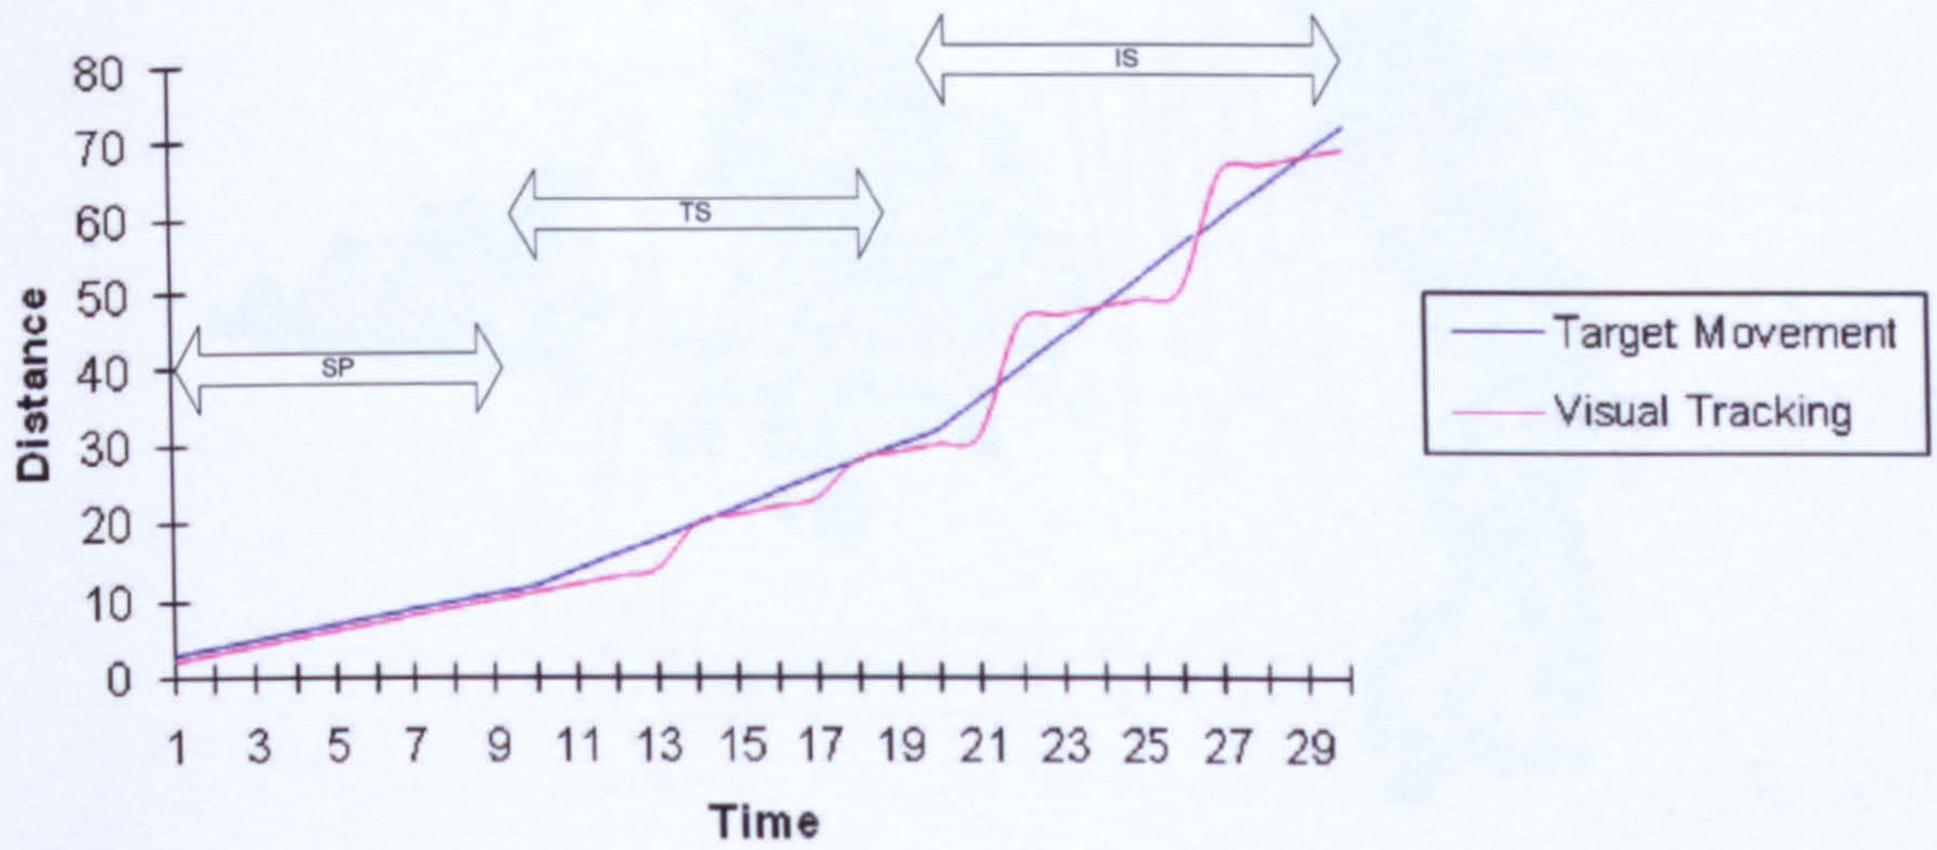

Figure 6. Indicative smooth pursuit (SP), tracking saccades (TS) and interception saccades (IS).

Crosswind casting is a physical movement pattern employed by many insect species that use chemical odours to locate prey, for example, fruit flies (Budick and Dickinson, 2006) and moths (Willis and Arbas, 1998). During anemotactic search in odour plumes insects, employ a tacking motion whilst surging upwind toward the source. The nature of odour plumes is such that the scent is often inconsistent, and when a gap is encountered, the insect turns perpendicular to the wind direction and crosswind casting is applied to reestablish a lost odour trail (Figure 7).

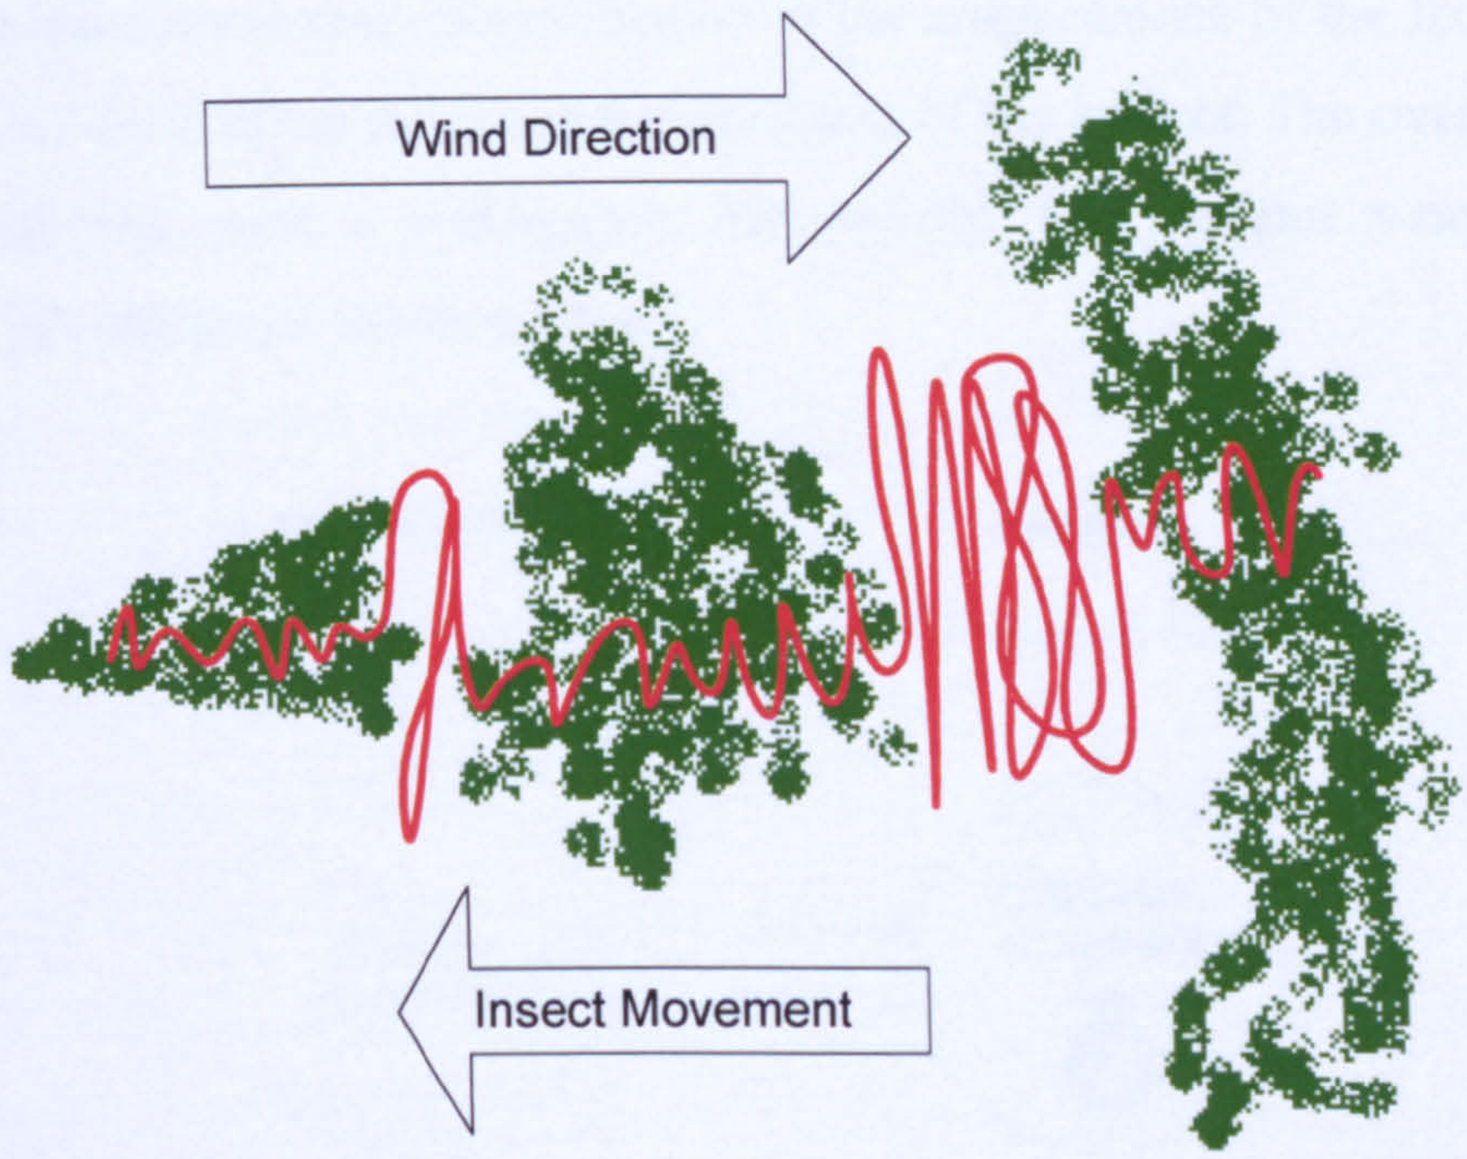

Figure 7. Indicative movement path of an insect (red trail) through an odour plume (green).

Direct search is a strategy whereby a predator will return to the last known location of a lost prey in cases of occlusion because the prey may have stopped to hide. Ogden (1974) observed this behaviour in the short tailed hawks of Florida. Generally, these hawks would adopt soaring, either 'balancing' in a fixed position on deflected air currents or spiralling upward in small circles. Following, an unsuccessful attempt to capture a prey item, they execute an alternating flap and sail direct search across the treetops for a short distance, endeavouring to relocate the prey before resuming their initial search strategy. Interestingly, this direct search technique was amongst a selection of initial search techniques, which included sit-and-wait (perched in trees), soar-and-stoop (high level soar followed by a stoop) and soar-and-sneak (high level soar followed by an inconspicuous direct line movement), adopted by a variety of Florida based hawks and described by Johnson and Peeters (1963), but was not observed as an initial technique by Ogden.

Badler and Heinen (2006) conducted laboratory experiments to investigate the anticipatory eye movements of monkeys tracking a target that was occluded for a given time. A conceptual model, Figure 8, was produced that had two main components, fixation and anticipation. The fixation component was tonically active and its influence on movement increased with visibility of the fixation point. This also acted as a filter to remove unwanted movements. When visibility was lost, the anticipation component was activated; this was a

timer that reactivated movement in anticipation of the reappearance of the fixation point and had knowledge of the fixation point and the readiness of the subject. The overall effect being that when employing such a mechanism successfully, the predator would increase its advantage through improved reaction time.

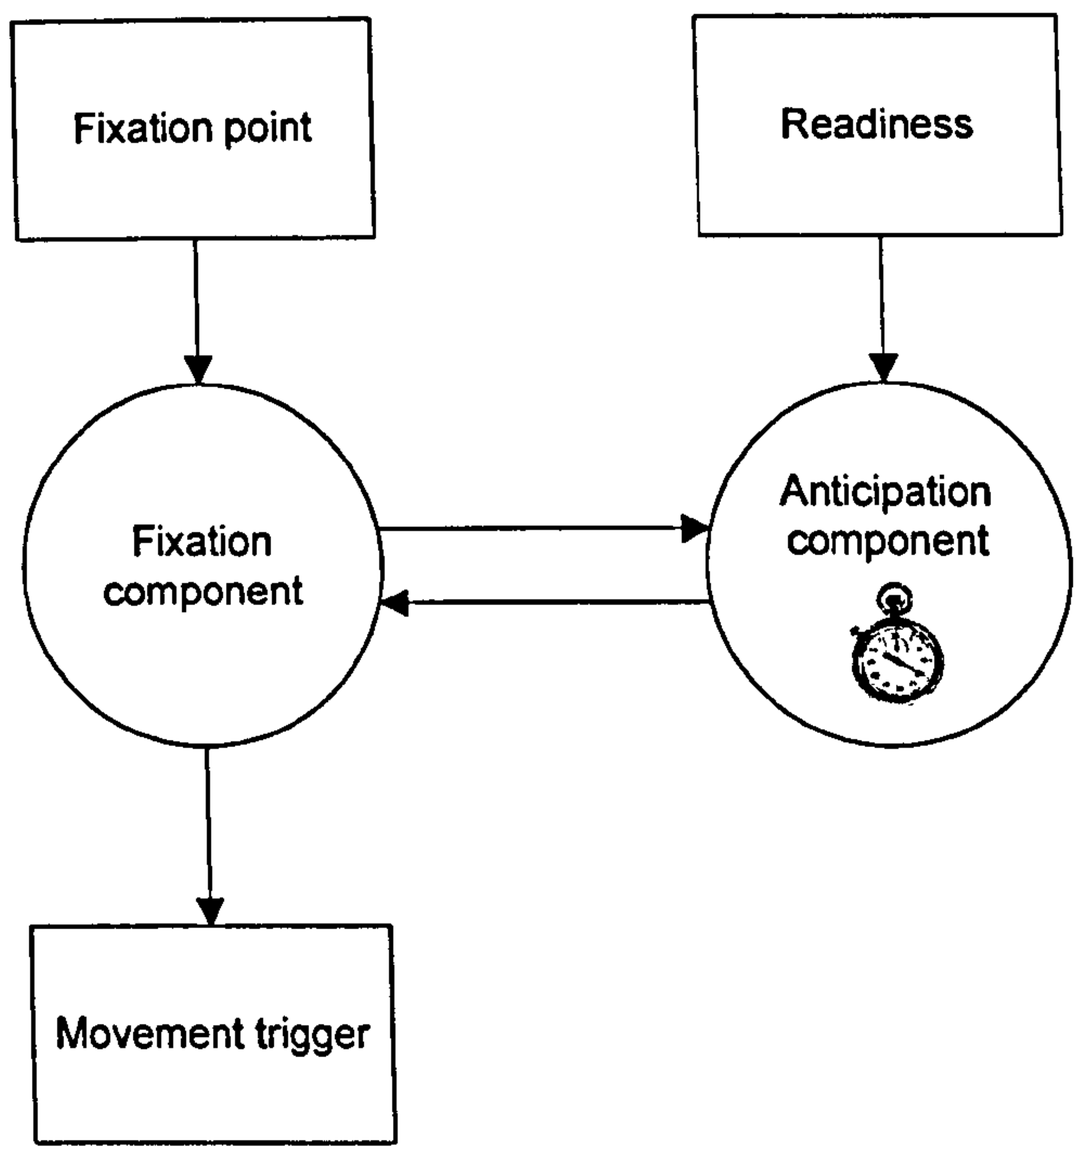

Figure 8. Conceptual model for pursuit initiation proposed by Badler and Heinen (2006).

3.11 Discussion

Within the field of natural computing there already exist several well established paradigms that rely on individual elements working cooperatively toward a common goal; such systems are generally termed 'agent based' systems (a term inspired by Minsky's book, The Society of Mind, 1986). Agent based systems are characterised by a set of computing entities that exist independently, but have a common goal, which they attempt to achieve either alone or as part of a team.

From the taxonomy of cooperation strategies indicated in Section 3.1, it can be seen that the type of cooperative mechanism that is employed in nature is dependent on the intended result. Agent based systems could utilise any of the forms of cooperation, but it should be noted that the choice directly affects the overall system behaviour. It could be argued, for example, that particle swarm optimisation employs by-product mutualism and genetic algorithms use

group-selected cooperation, and it is these aspects that define the very nature of the paradigms. Within the context of the cooperation strategy, the manner in which the agents' work can affect the efficiency with which the goals are achieved. Selection of specific individual roles for agents can enable more complex group goals to be achieved, but this could be at the cost of emergent behaviours, which are less likely, since systems that include agents with specific roles are generally more deterministic. Individual behaviours need not be so restrictive, however – it has been seen that some animals adapt their behaviours to individual or group needs. Ants are a classic example of this, where colony members will usually perform one of a number of tasks, such as foraging, but switch to another task following environmental or internal perturbations (Waibel *et al.*, 2006). The ability to vary behaviour could allow a system to be more adaptive and robust in hostile or unpredictable environments.

A wide variety of search behaviours have been examined, and it has been seen that nature has adapted its behaviour to suit the prevailing conditions. Similarly, it is important for the developer of a natural computing paradigm to consider the target application and the strengths (and weaknesses) of naturally occurring behaviours before selecting an appropriate strategy. Consider, for example, a reversionary navigation system for swarming search robots in harsh environments where limited bandwidth or electromagnetic noise could hinder communication. While communication is available the robots could search in a coordinated fashion, sharing information to increase search efficiency. If the system degraded, it could be more profitable for the robots to enter individual search patterns using the strategies described earlier; an example that has already been researched as a target application of autonomous robotic swarms is mine detection; see, for example, Fruergaard-Pedersen (2006). Such systems rely heavily on good communications, and loss of this facility may render the system ineffective; however, in such situations, the robots could enter a saltatory search mode to look for well hidden, patchily distributed mines. Such use of natural metaphors could also lend itself to division of labour, which is already suggested by Fruergaard-Pedersen, where some robots could be specialist detectors, whilst others could be disposable clearance drones, and, once the expensively equipped detectors were clear, destroy the mine.

Two examples of how natural behaviour has already inspired natural computing paradigms are also considered in this section, the examples have deliberately been selected outside of the more established techniques, such as particle swarm optimisation, to illustrate the breadth

of potential for naturally inspired paradigms. The first considers how ponerine ant (*Pachycondyla apicalis*) foraging behaviour has inspired the development of a search algorithm for use in combinatorial optimisation (Monmarche *et al.*, 2000), whilst the second examines how trophallaxis is being investigated as a potential communication strategy for miniature robots (Schmickl and Crailsheim, 2006).

3.11.1 Ponerine Ant Inspired Combinatorial Optimisation

The natural foraging behaviour of ponerine ants consists of a global and local strategy, as previously described. The search algorithm derived from it, known as API (from *Pachycondyla apicalis*), is as follows (with pseudo-code shown in Figure 9). An initial nest location is selected at random within the search space. Each of the ants then does one of the following: if they do not have a minimum number of hunting sites, they create one and explore it; otherwise, if the previously explored site was successful it is explored again, and if it is unsuccessful, one of the other sites is selected at random and explored. If a site is unsuccessful for a given period of time, it is removed from the ant's memory. Tandem recruitment is then performed, in which two ants are selected at random and the performance of their best hunting sites compared. The ant with the weakest site replaces it with the strongest site of the other, in a form of elitism that assists exploitation, because it increases the number of trials in the better sites. After a set number of iterations, the nest position is moved to the group best position and the last nest move and all ant memories are erased.

```
choose randomly the initial nest location N
for each ant a_i, i in 1..n:
    if a_i has less than p hunting sites in memory
    then
        create and explore a new site in the neighbourhood of N
    else
        if the previous site exploration was successful
        then
            explore same site again
        else
            explore randomly selected site (among the p hunting sites in memory)
        end if
    end if
    remove from the ants memories all sites which have been explored unsuccessfully more than
    P_local(a_i) consecutive times
    perform recruitment (best site copy between two randomly selected ants)
    if more than T iterations have been performed
    then
        change nest location N to the most successful site explored and reset memories of all ants
    end if
end loop or stop if a stopping criterion is satisfied
```

Figure 9. The API algorithm pseudo-code (adapted from Monmarche *et al.*, 2000, Figure 3).

Communication among the ants is limited in both the natural and artificial systems, but cooperation occurs both explicitly and implicitly. Explicit cooperation occurs via tandem calling and implied cooperation occurs due the nest being moved to a location that is based on information gathered by individual ant search behaviour. Of particular interest is the note made by the API authors that the search behaviour of the ants need not be random, but can be matched to the specific problem, where landscape information is known *a priori*. A potential problem with the algorithm is its dependency on initialisation – in a multimodal landscape, the group could easily be drawn away from the global optimum if the nest site is initialised to a remote local optimum.

3.11.2 Trophallaxis and Swarming Robots

Whilst the example above did not utilise extensively naturally based information sharing, many naturally inspired paradigms do require complex communication solutions. Ant Colony Optimisation (Colorni *et al.* 1991) is one of the most widely known and successful uses of a naturally inspired communication to facilitate group behaviour, but others have been exploited.

Schmickl and Crailsheim (2006) utilised a strategy inspired by trophallaxis to maintain communication in a simulated swarm of 200 miniature robots; the aim of the research being to develop a swarm of 1000 2mm x 2mm robots. The swarm communication works on the

ollowing premise: when a honeybee collects nectar from a source and flies back to the nest it vill enhance its communication of its find to other nest mates using trophallaxis. As this ccurs, a small amount of the crop is passed to the enquiring bee; thus, bees closer to the ource would be expected to have higher crop levels. The robots in the simulation were vorking in a cleaning scenario; when a robot located dirt it collected an amount and moved way to find a dump area, this amount was stored in the robot memory. On encountering mpty robots they would electronically pass an amount of dirt thus lightening their own irtual load. The empty robots would in turn pass on any information they had regarding a lumping area. Using this technique, empty robots could follow the gradient towards the ource, and a traffic system emerged. The approach has several interesting facets: sensor and ommunication system size can be very small, enabling small scale robots to be used ffectively, and the emergent path from the dirt to the dump area was directed in that it ontained no random walk behaviour even where there where obstacles along the most direct ath. Further, where multiple equidistant paths existed, the technique was seen to be more ffective than pheromone based techniques, because all paths were used equally, distributing he traffic flow more efficiently.

4 Method

This section details the method used to achieve the aims of the research. The context and purpose of the work are outlined, followed by a discussion of the approach taken to the study of the utility of natural search strategies for autonomous search control from a high level perspective. A simulation environment is then described in more detail and search algorithms and associated behaviours specified. The metrics used to evaluate performance are presented followed by a discussion of the approach to statistical analysis. Finally, a schedule of the experiments that follow in Section 5 is included.

4.1 Context and Purpose

In Section 2, the current state-of-the-art in autonomous vehicles was discussed. Whilst remotely controlled vehicles have been in use for some time, and have obvious advantages in hazardous or remote environments, autonomy has only recently become the subject of major research and development efforts. Autonomy, although raising safety and ethical questions, also has distinct advantages, particularly in environments where control lag is critical and renders remote control at least partially ineffective, or where communication is unreliable, as is often the case in military scenarios and harsh environments. One could also envisage large swarms of vehicles being deployed that would simply be too complex for effective human control. Hence, the need to find robust and efficient controls strategies for autonomous vehicles is becoming more pressing.

In searching for solutions to complex engineering problems, researchers are increasingly turning to the natural world for inspiration. Whilst there are some examples of well established natural computing paradigms, such as evolutionary computation, ant colony systems, and particle swarm optimisation, these barely scratch the surface of potentially valuable solution strategies. In Section 3, a more broad review of natural search was reported, indeed, the first of its kind, which revealed a wide range of behaviours that could be applied in autonomous vehicle search. Generally, in cases where natural phenomena have been applied, they have been shown to be robust and effective; where studied, by observation and/or simulation, the natural behaviours have often been found to be optimal in some sense. Given that a critical aspect of the survival of many animal species is successful search, there

is a clear motivation for exploring the extent to which these evolved search strategies may be usefully applied in the context of autonomous vehicle search.

In previous work, the author developed an approach to the control of UAVs based on particle swarm optimisation (PSO) (Banks, 2005; Banks *et al.*, 2008b). The approach was successful, in the sense that the UAVs that followed random search and applied PSO on target detection generally performed better than those following systematic search paths but did not apply PSO, and interesting swarming behaviour emerged that enabled the UAVs to locate and intercept targets travelling faster than the UAVs were capable of. Whilst a swarm inspired approach is a reasonable starting point for such an investigation, and, indeed, Pugh and Martinoli (2007) made a similar attempt, it was realised that the natural world has many examples of search behaviour, and that a more thorough study, returning first to the available biological research, was needed to form a sound basis for future development.

This background work formed the general motivation for this study, which therefore aims to ascertain whether there is any value in developing search strategies for autonomous vehicles based on metaphors drawn from nature. Few studies have been conducted in this area; even the review of natural search behaviours is significant and novel. It is, therefore, impossible to be exhaustive in this first foray into this new field. It is necessary to construct a scenario in which the search behaviours may be explored (of necessity, by simulation), and evaluated against non-nature inspired approaches, that is, the systematic approaches that a human operator might reasonably apply. Clearly, the range of behaviours – both natural and systematic – and the range of possible scenarios for autonomous vehicles are prohibitively large. However, the choice of scenario is, to a large extent, not especially relevant to answering the main research question; that is, if natural search strategies are understood to be of value in any scenario, then a significant step forward has been made.

The selection of scenario is not entirely arbitrary, however. Much of the autonomous vehicle research is driven by future military needs. Hence, the selection of air vehicles as the basis for the author's previous research, and the continuation of that theme here. It is worth considering, briefly, the purpose of a simulation environment. In general terms, we have a wide range of application scenarios, as discussed in Section 2, which share some common properties:

- Autonomous vehicles (agents) with the ability to move throughout the search space, whether air, land, sea or space.
- A search task, which may be anything from the detection and interception of military targets to the location of avalanche victims. Hence, targets may be moving or stationary, indeed, may even be faster than the agents, depending on the search task. Even within a given application domain, target properties are likely to vary, therefore, there is a need to explore a range of agent–target relationships.
- Limited sensory capability, such that a movement strategy is needed in order to bring the agent into proximity with the target.
- Sparse feedback, such that, unless within the proximity of a target there is no indication of where the most profitable search area might be. This separates the problem space from that of the more common search (optimisation) problems, which typically have landscapes that give some clue as to potentially useful search trajectories, which are then exploited by the various search algorithms (such as genetic algorithms and particle swarm optimisation).

Whilst the simulation environment developed for this study is described in some detail in later sections, it is important to bear in mind that this is only a vehicle with which to explore the search behaviours, and that many other, equally valid, simulation scenarios could be constructed. Hence, the choice of particular scenario is somewhat arbitrary, provided that it recreates the above conditions.

The general approach chosen for this research, therefore, is to construct a (military inspired) simulation of agents and targets, endow the agents with both systematic and natural search behaviours, and identify and evaluate the conditions under which the natural search behaviours may be advantageous. Whilst an attempt may be made to draw general conclusions from observation of the simulations, the critical result is ascertaining whether natural search behaviours are viable contenders for further research and development within the field of autonomous vehicles. At this stage of the research, it is essential to strip down the simulation to a fairly minimal form, such that one can be confident that variations in performance are due to the inherent properties of the search strategies, and not complicated by fine details in the world model. Hence, to achieve the aim of this research a simplified simulation environment is not only satisfactory, but to be preferred. For example, there is no

sense in providing complex environments and physics models, if the search behaviours are not first found to be useful. The design and implementation of even a simplified simulation environment is non-trivial, but not the subject of this research. Hence, only the properties of the environment are described in detail (sufficient for reproduction of the results), along with the rationale for various decisions in approach.

4.2 Approach

Of necessity, this exploratory investigation is conducted by simulation rather than implementing autonomous behaviours into physical vehicles. In simulation, all agents are identical, allowing differences in search performance to be attributed to the search behaviours rather than environmental factors, such as individual vehicle performance or local conditions impacting on movement. However, due to the coarse grained control of the agents in the simulation, it is not envisaged that individual differences in vehicles would affect significantly the findings regarding the efficacy of each strategy. The ability of any form of implementation vehicle (*e.g.* aerial / ground / marine based) to perform the movements required for a given search strategy could affect results; for example, a ground based vehicle should be able to apply a high turn rate strategy more effectively than a fixed wing aerial vehicle. For the purposes of fair comparison, therefore, the simulation environment permits all strategies to operate within a reasonable movement envelope (for example, one would not generally consider it feasible to allow the simulated agents to apply the turn rates observed in fruit flies). Moreover, it was decided to maintain agent speed, unless a strategy specifically includes changes, since the strategy itself should determine the agents' movement across the search area. This also applies to targets, due to the need to explore the performance of each search strategy against targets applying different forms of movement. In the literature reviewed, the influence of relative speed was also seen to be influential on the choice of search strategy, with predators using different strategies against prey that moved faster to those that moved more slowly relative to their own ability. In particular, it is of interest to explore whether a number of agents are able to cooperate in order to effectively detect and intercept targets travelling at significantly higher speeds.

Inter-agent communication and sensory imperfections could also affect results; initially, search strategies are compared without considering occlusion, and assuming that all communication is perfect. At lower target speeds, communication is less critical, because an

individual agent is capable of tracking and intercepting a dynamic target, even if cooperation could speed up the process. Where agent speeds cannot match that of the targets, communication becomes more critical as cooperation is then essential for interception. In recognition of potential communication difficulties, sensory imperfections and other such problems that may result in target detection being lost, a mechanism is needed to evaluate the ability to reacquire targets and assess the robustness of the system to communications failures.

Clearly, there are many conceivable scenarios for occlusion, and a comprehensive study of different scenarios is not within the scope of this work. It was noted that the impact of target loss could potentially be very different depending on whether the effect was permanent (*e.g.* through total sensor failure) or transient. Hence, two such mechanisms were implemented: the first considering an agent permanently losing the ability to sense a target shortly after initial detection; and the second, a more transient occlusion, that could occur, for example, where a target moves behind a physical structure. The former permits the robustness of the system to be evaluated, whilst the latter offers an opportunity to explore the effectiveness of various approaches to target reacquisition.

Similarly, communication failure is examined with both permanent and transient effects; with increased severity a graceful degradation is desirable. One might consider loss of communication to have a similar impact to occlusion, since, in both cases, the impaired agents are unable to share target information with the rest of the swarm. However, in cases where the speed of the target speed is less than that of the agent, communications failure, unlike occlusion, will not result in loss of the target.

The literature highlights a breadth of natural search behaviours; clearly, it is not possible to investigate the entirety of known behaviours and a sampling of representative behaviours is needed. The rationale for that sampling applied the general findings of ethological research, *i.e.* that search can be placed along the two behavioural axes of path tortuousness and pause lengths. This is then supplemented with notable species specific searches drawn from the review of relevant literature. Ambush behaviour (sit-and-wait initial search or the anticipation mechanism for prey acquisition) is not investigated here, having rather narrow applicability (indeed, being impossible for some vehicle types) and, in particular, incapable of searching for static targets. The two species specific searches, fruit fly and bacterial, were chosen because of their two phase nature and marked difference from the generalised strategies

common to many species, with the former having a mixture of long and short leg lengths containing distinct turn angles, and the latter slowing down and increasing turn rates when targets are not being detected. The key properties of interest in the search strategies are: how quickly agents can locate targets within the search area, whether static or dynamic; how effectively the agents can cover the search area; and whether the more successful searches are less efficient in terms of energy requirements.

With respect to the simulation environment, the search area is featureless, allowing an uncomplicated comparison of search behaviours. This is a similar approach to that taken by Walmsley and Sims (2004), who used a featureless landscape as a baseline for evaluating environmental impacts on ISTAR (Information, Surveillance, Target Acquisition and Reconnaissance) assessment. Oftentimes, a predator may use environmental clues to enhance their hunting prospects, and it is possible that artificial search behaviours could also benefit from this. However, at this early stage of investigation, it is the underlying properties of the search patterns that are of interest. It should be noted that this does deny the saltatory search strategy its most profitable scenario, where targets are potentially hidden. Detection of hidden targets, however, is rather application specific; that is, it is difficult to construct some general simulation of mechanisms of hiding and discovery that might have some practical relevance, hence, this is considered to be beyond the scope of this particular study, although an interesting area for further work.

Search applications and scenarios are diverse. For the purposes of this investigation, several simplified (military inspired) scenarios are considered. These are, however, deemed sufficiently general that findings are more broadly applicable. Scenarios involve swarms of agents of varying size pitted against varying numbers and distributions of targets in both non-destructive and destructive situations; that is, location only, or location and destruction (*i.e.* targets removed when detected by, and within a given range of, an agent). Basic search ability is evaluated against static targets. For dynamic targets, two main approaches are considered under various conditions. First, a scenario inspired by a military air defence role, as adopted in the author's previous work on PSO based control of UAVs (Banks, 2005; Banks *et al.*, 2008b), in which two locations within the search space are defined as 'defended locations', and the objective of the agents is to prevent targets from reaching those locations and exiting the search area. Up to four targets are explored in this scenario, each following one of four different planned routes to the defended locations. Hence, targets behave in a

deterministic manner. Second, to evaluate performance against non-deterministic targets, dynamic targets are allowed to move throughout the search space using various behaviours drawn from the same pool as the agents. Where dynamic targets are explored, cooperation among agents allows potential improvement in performance, in particular where targets are travelling faster than agents, and interception by cooperating agents is necessary. The impact of cooperation on performance is examined for selected scenarios.

4.3 Simulation Environment

A simulation environment was developed to allow the search strategies to be compared as fairly as possible, whilst minimising factors that could confound the results. It was constructed in Ada 95 using the John English Windows Library (JEWL) (English, 2000) and AdaCore's Gnat Programming System (AdaCore, 2005), along similar lines to that used in Banks (2005) and Banks *et al.* (2008b). This section describes the environment in sufficient detail for replication of this study, although the findings should not be particularly sensitive to implementation details, provided the natural behaviours are replicated as described.

4.3.1 Search Environment

Figure 10 shows an indicative screenshot of the simulation environment. The search area comprises a two-dimensional representation of an area measuring 400 x 300 units; within this are two arbitrary non-central 'defended locations', shown by two small black crosses, used for those simulations where targets follow deterministic routes. In this scenario four targets (represented by blue dots when not detected and red once detected) are being sought by five agents whose search area is bounded by the large white rectangle and sensor range and field of view is shown by the green angled lines when not detecting a target, turning blue when detecting. The agents do not know of the targets' existence unless they have been detected and can only be certain of their presence whilst they are being detected.

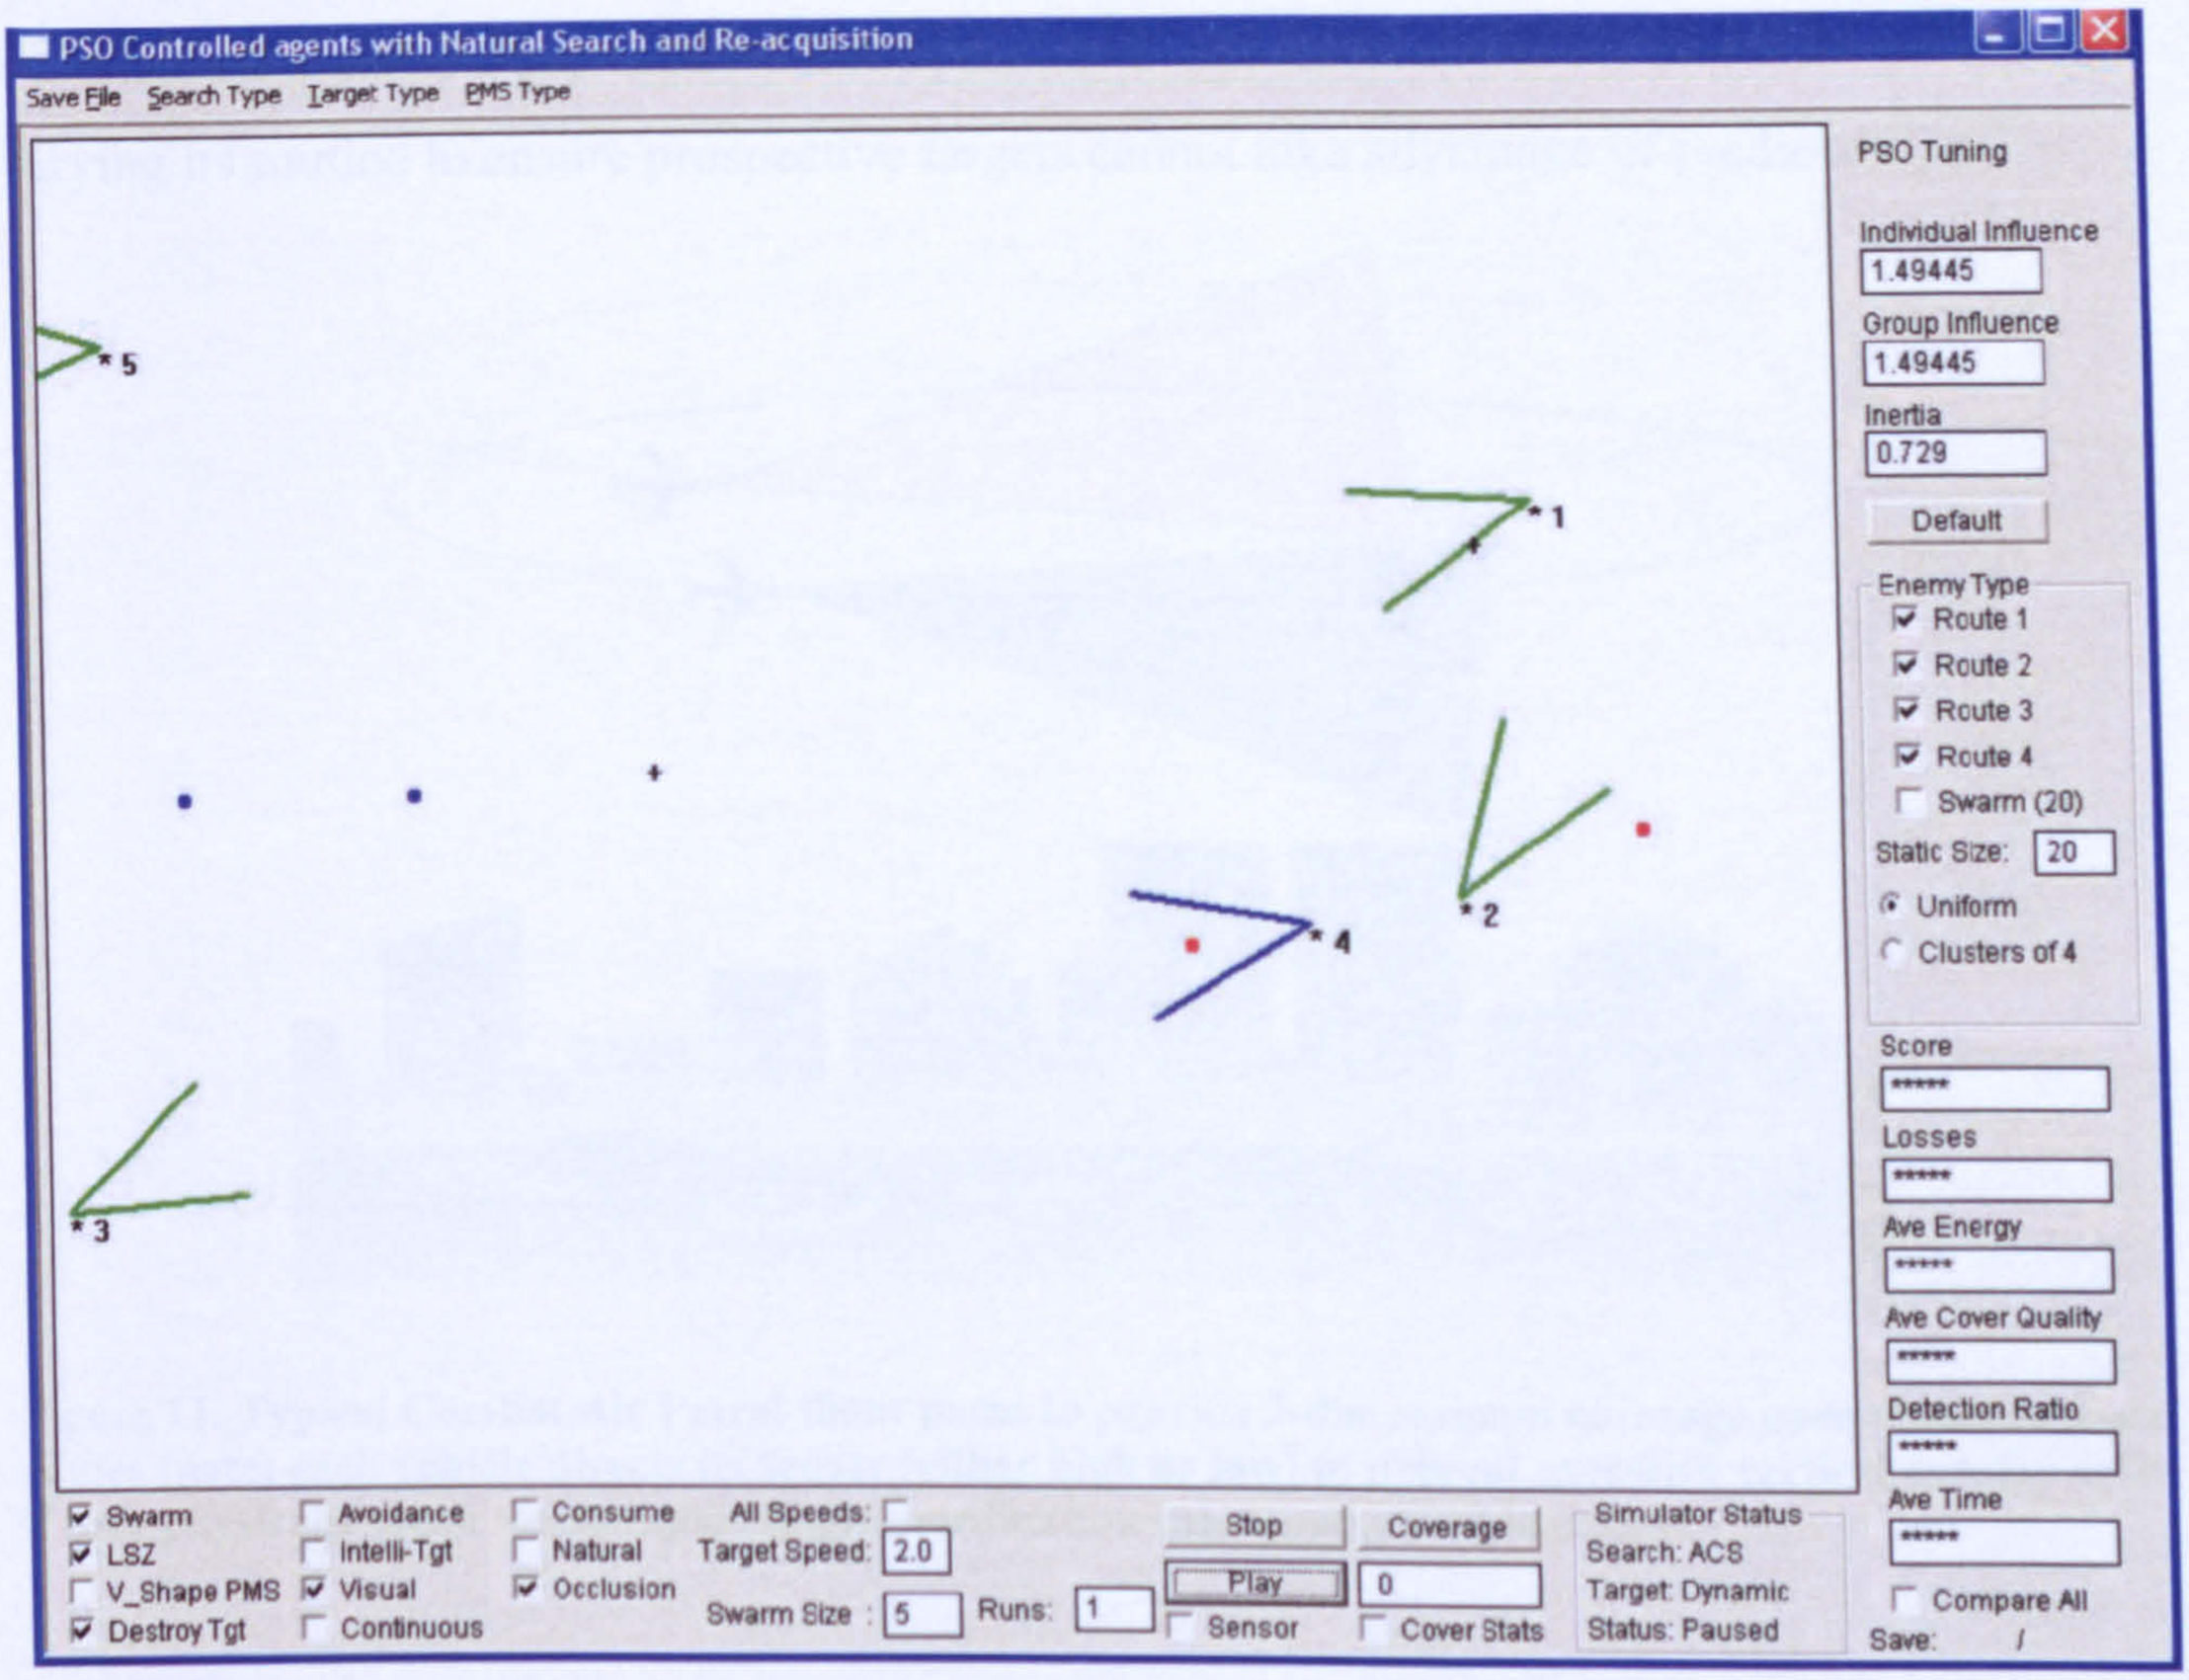

Figure 10. Simulation environment.

The search environment was deliberately featureless to ensure that the capabilities of the search strategies were measured with respect only to target movement characteristics. A significant simplification is achieved by the use of a two dimensional search space, which renders both implementation and, more importantly, analysis and visualisation tractable. Whilst extension to the vertical plane (altitude for air vehicles, depth for submersibles) is certainly an area worthy of further investigation, it should also be noted that such search is not always present in the natural world, nor in many applications of autonomous vehicles (particularly ground based vehicles). Animals often search within a set of vertical bounds (fish, for example, forage at different levels) and search vehicles operate such that the vehicle maintains a given height, or depth, and the sensor provides some coverage in the vertical plane. Therefore, autonomous vehicles may operate in the same manner, implementing the movement patterns essentially in two dimensions. Indeed, such principles are already applied in typical Combat Air Patrol flight patterns (see Figure 11), where, to achieve sensor

coverage of a given three dimensional space, multiple interceptors fly elliptical paths, each varying its routine to ensure prospective targets cannot take advantage of predictability.

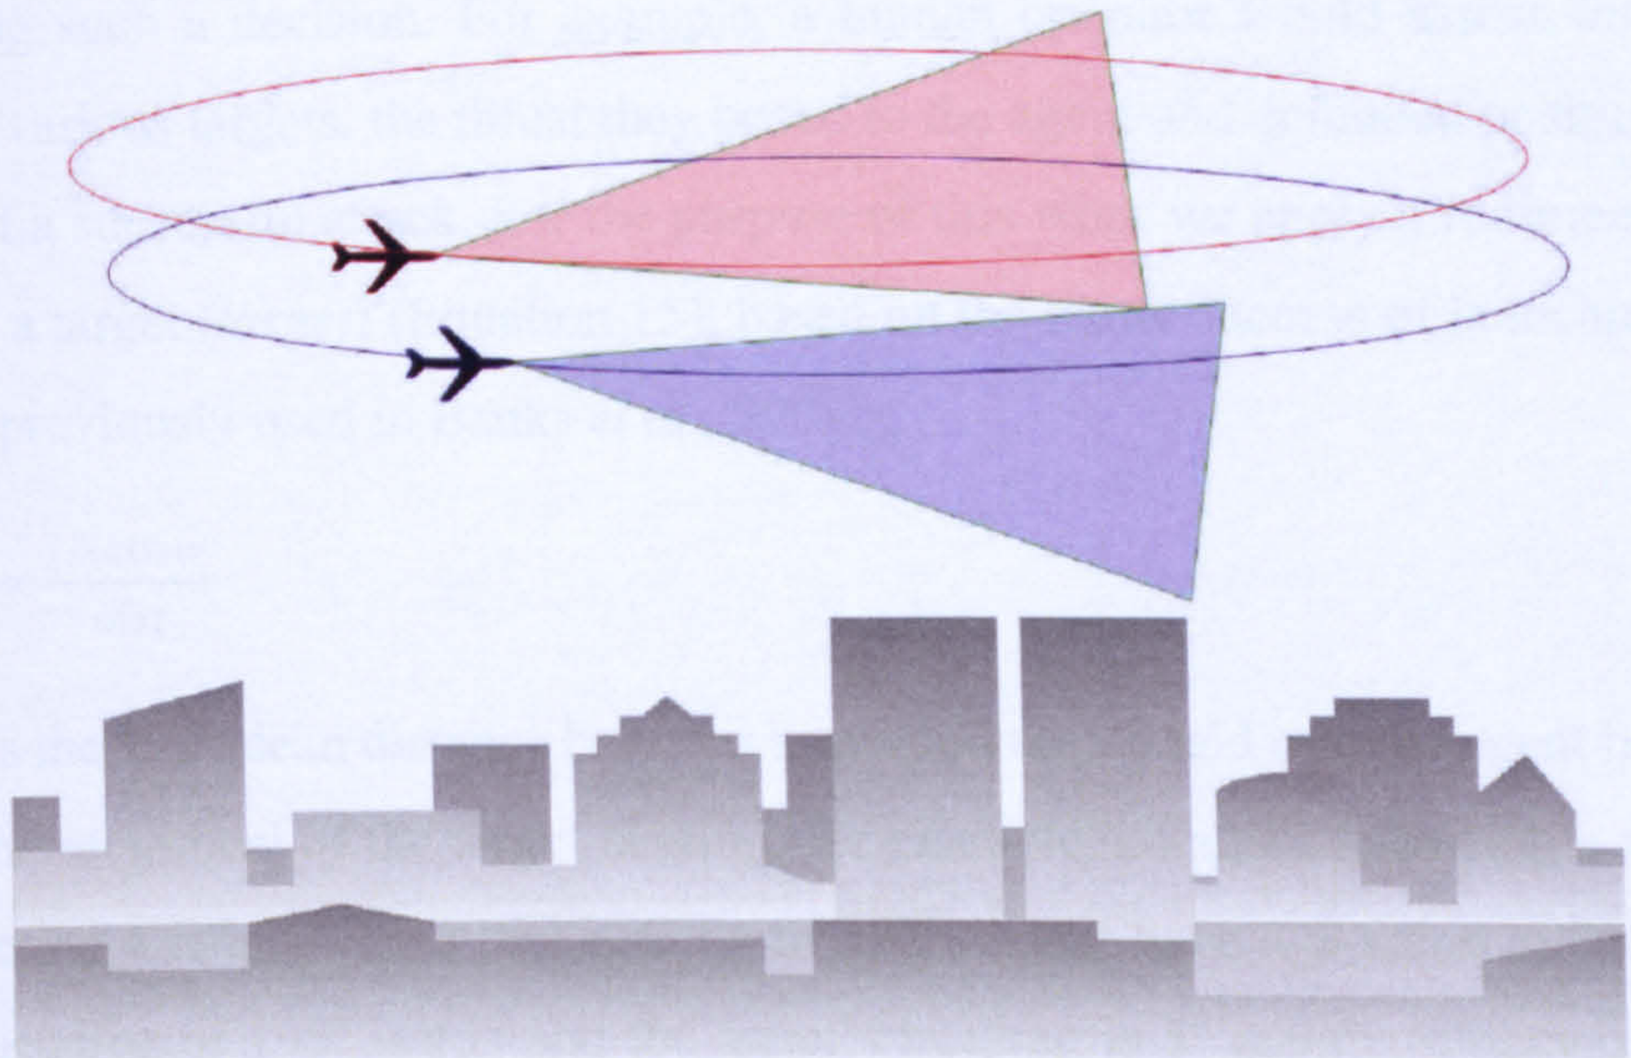

Figure 11. Typical Combat Air Patrol flight paths to provide 3-dimensional coverage using 2-dimensional routes (note: each vehicle directs its sensor (either high or low) to prevent excessive vertical overlap since if they physically move too far apart flight coordination can become problematic) .

4.3.2 Agent Characteristics

All movement in the physical world is constrained by the characteristics of the object and the environment through which it moves. In the natural world, search strategies are often physiologically limited, although some, such as fruit fly movements (where Tammero and Dickinson (2002) report turn rates of 90° in less than 100ms), easily outperform current vehicular technology. In this simulation, generic limitations were imposed on agent movement based on the author's experience of aircraft flight envelopes (note: both agent and target speeds were scaled to make simulation timescales viable). They do not aspire to replicate any given vehicle type, but represent an element of realism with a view to the range target applications.

Agents travel at a constant speed of 0.8 units per step (unless the search strategy dictates otherwise), have a sensor range of 40 units, a field of view (FOV) angle of 40° and maximum turn rates of 10° per step. Agents destroy targets if the targets are within 10 units and inside the sensor FOV.

Where an agent simultaneously detects multiple targets it needs to decide which one to pursue. In a given application there may be several factors that need to be taken into account when making such a decision. For example, a human operator would assess *inter alia* the value of the various targets, the threat they posed to the agent and defended positions, and the likelihood of a successful attack. For the purpose of this work we apply a rudimentary fitness function for a target (*target*) (Equation 15), based on the likely success of launching a guided weapon (as previously used in Banks *et al.*, 2008b).

$$f(tgt) = \frac{1 + \cos\psi}{dist} \tag{15}$$

where *dist* is the Euclidean distance between agent and target, and ψ is the agent heading with respect to the reciprocal of the target bearing; for example, an agent following a real heading of 10° detecting a target with a real bearing of 195° would have a ψ value of 355° since the reciprocal heading of 195° is 15° and the agent's heading is 5° counter-clockwise from that. The target, given a distance of 10 units, would therefore have a fitness of 0.1996. Compare that with a scenario where the target was the same distance away but on a bearing 280°, its reciprocal would be 100° giving a ψ of 270° and a lower fitness value of 0.1. This fitness function means that a target directly ahead of an agent would have maximum fitness whereas a target directly behind the agent would have minimum, because the chances of launching a successful attack on the former target would be greater than the latter. In reality a target may attempt an evasion tactic such as a 'notch' manoeuvre where the target turns perpendicular to the agent. Such a manoeuvre, where employed successfully, denies a radar return and the target effectively becomes occluded. Where agents are searching for a target following detection loss and a target is detected the agent will assume it has located the original target it was searching for, *i.e.* the agents are assumed not to be intelligent enough to distinguish individual targets and as such they assume any targets detected to be the lost target. This will only impact behaviour where several targets are located in the same vicinity and the agent could detect a different, previously undetected, target leaving the original target to continue (although other agents could still attempt to intercept it).

4.3.3 Target Characteristics

Targets move with the same characteristics as agents with the exception of speed, which varies for experimental purposes. Speeds are measured relative to maximum agent speed (*e.g.* a target speed of 0.5 represents half maximum agent speed, that is, 0.4 units per step).

Depending on the experiment, targets may be static, dynamic with randomised movements using some of the same movement patterns applied by the agents, or dynamic but following predetermined routes. Static targets were either uniformly distributed across the search area or clustered. The distribution of clustered targets was achieved by initialising groups of 4 targets with random uniform distribution across the search area. Each cluster was formed by placing an initial target in the search area and then locating 3 other targets (once more with a random uniform distribution) within 40 units of it. This cluster size ensured that a recognisable cluster existed, whilst it was possible that, in extreme cases, the targets could be located at a distance of twice the agent sensor range apart.

For those dynamic experiments where targets follow predetermined routes to the two 'defended locations' shown in Figure 12, these were selected such that any defensive system that favoured one particular area, especially the centre, of the search space would not gain advantage. The routes, which are unknown to the agents, were designed such that the attackers spent minimal time in the defended area and exited at or near the entry point, with routes generally avoiding defended positions where possible (the exception being route 4, which at one point is close to target position 1, to provide variation). Targets were assumed to have successfully attacked the defended locations if they moved to with 2.5 units.

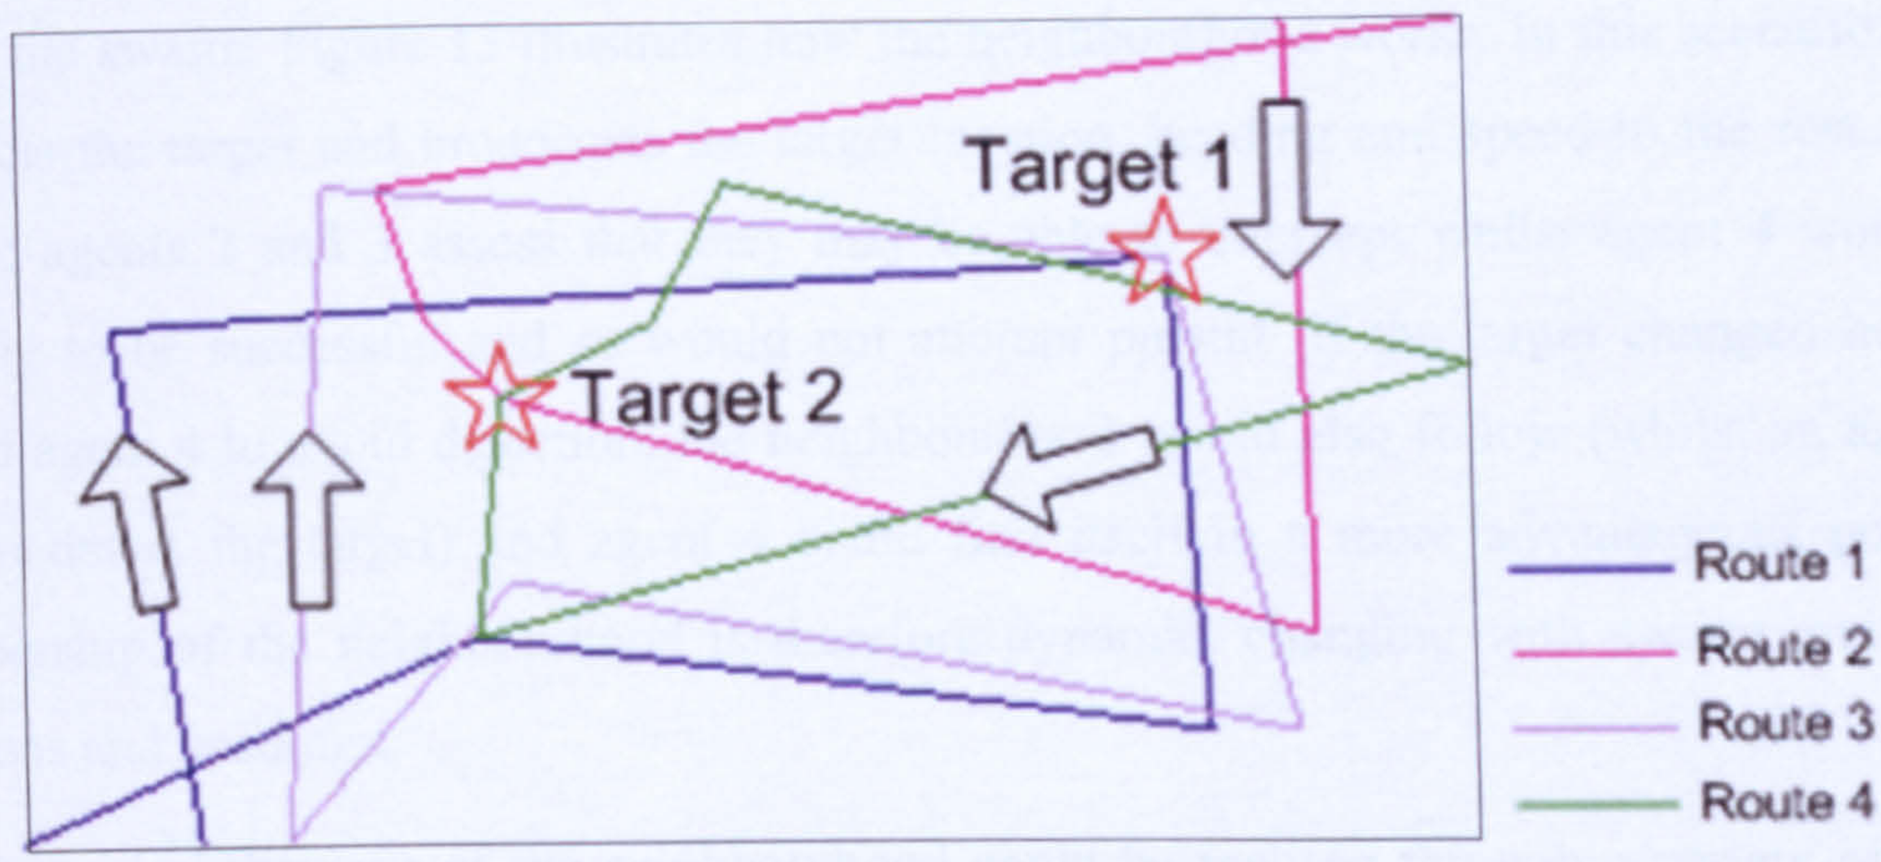

Figure 12. Map of deterministic target routes around search area indicating attackers' targeted locations (target 1 is attacked via routes 1 and 3 whilst target 2 is attacked via routes 2 and 4).

4.3.4 LSZ Neighbourhood

In the author's previous work on PSO based guidance, it was established that cooperation amongst agents could be detrimental by drawing agents towards detected targets that they could not possibly intercept, and thereby moving them away from other areas of the search space and allowing other targets to remain undetected. To address this, and promote productive swarm behaviour, a neighbourhood function termed the Launch Success Zone (LSZ) was introduced. The term was associated with the neighbourhood since it was based on a similar zone employed in aerial combat situations to describe the area in which an intercepting aircraft can successfully deliver a missile attack against an enemy. The term is retained here since its basic rationale remains and also to distinguish it from other neighbourhoods in operation.

When an agents detects a target, this neighbourhood determines which other members of the swarm, if any, might have reasonable prospects of intercepting the target. Those agents outside of the neighbourhood continue with their own search behaviours, whilst those within the neighbourhood cooperate in an attempt to intercept the target. It is this mechanism that is primarily responsible for the ability to detect and intercept targets moving at speeds faster than the agents.

Although developed for a PSO based control system rather than the natural search strategies explored in this work, the neighbourhood function is retained to enable fruitful cooperation

within the swarm. Figure 13 illustrates how the neighbourhood works. In this scenario agent 1 detects the target and broadcasts the target location, heading and speed to the rest of the swarm; agents 2 and 3 assess that they may be able to intercept whilst agent 4 would be unlikely to be successful and so would not attempt pursuit. If the target changed heading toward agent 4 to avoid detection, the neighbourhood would also follow (whilst an agent is able to detect the target) and agent 4 could find itself in a more advantageous position. Membership of the neighbourhood is therefore dynamic, changing with swarm and target locations and headings.

The dynamic behaviour of the neighbourhood could be realised through a variety of linear and non-linear functions, but, since it is the search behaviours that provide the focus for this work, in depth exploration of the neighbourhood function is avoided. Hence, a simple but effective linear function (Equation 16) was applied to control the length of the conical neighbourhood with a fixed angle (θ).

$$length\ (l) = r \cdot c \cdot s \tag{16}$$

Where r is the sensor range, c is a constant and s is the relative target speed. The angle, θ, and constant, c, are tuned to provide good performance (see Section 5.2.1, page 104). An agent may find itself in multiple neighbourhoods, in which case the selected target is resolved following Equation 15.

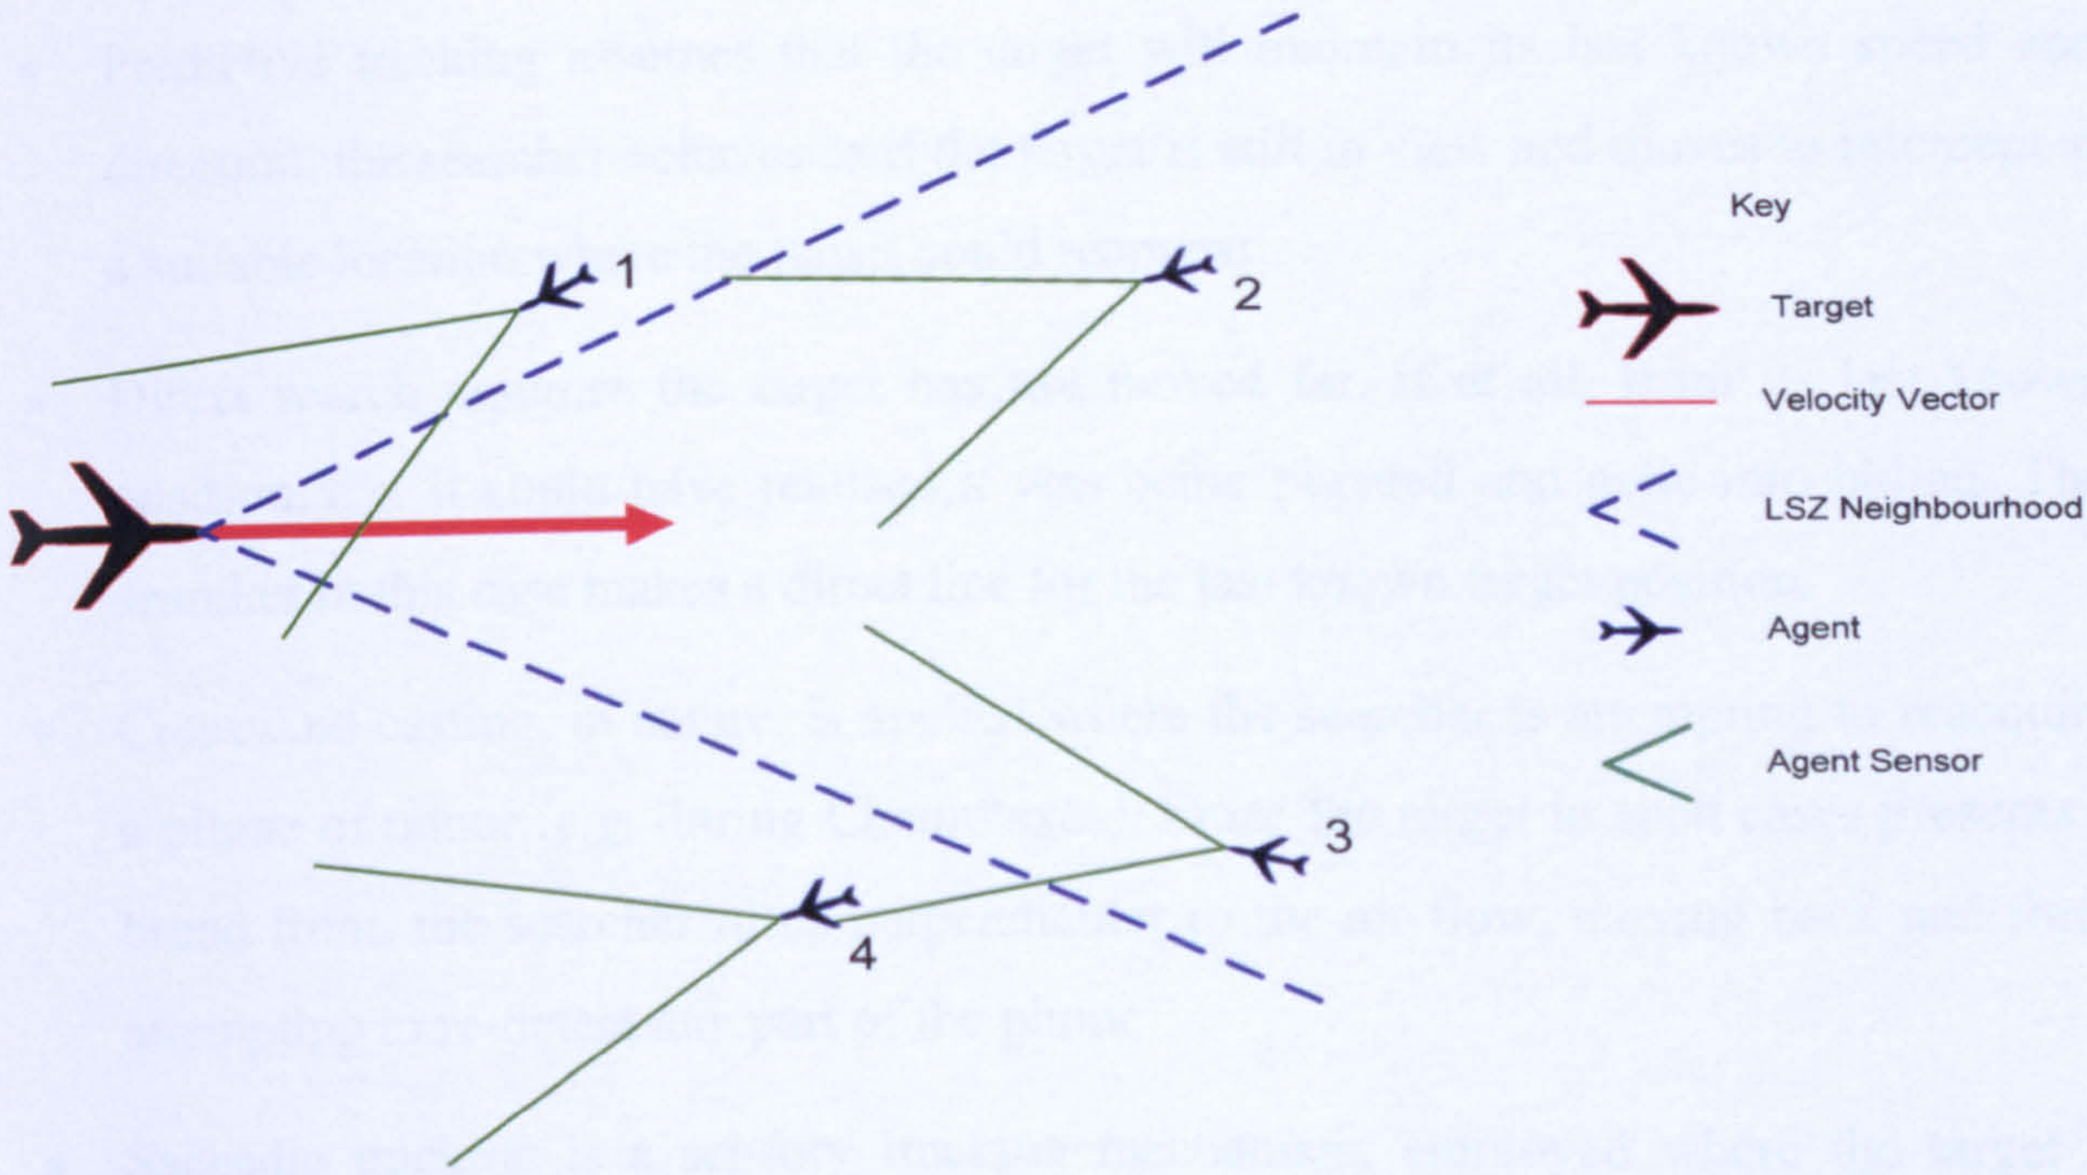

Figure 13 Example LSZ neighbourhood scenario.

4.3.5 Predictive Memory System Strategies

The author's previous work, which provided an initial investigation into the application of particle swarms to autonomous aerial vehicles guidance (Banks 2005), identified two main areas for further development to enhance the performance of swarms of vehicles defending an area against intruders. Both of these involved search activities: the first being the approach to initially locating targets (prior to the particle swarm based cooperation) and the second being the swarm behaviour where a target escaped detection. In the latter case, the existence of the target was immediately forgotten and the defending vehicles resumed initial search behaviour. This situation is undesirable since the information gleaned about the target during detection could be used to improve the likelihood of re-detection. As discussed in Section 3.10, nature has evolved strategies that utilise the information known about the last known location and behaviour of lost prey items to improve the chances of reacquisition. Such behaviours have developed because they can improve overall foraging efficiency since, if a prey item is lost, the initial effort in locating and pursuing it is wasted. Strategies from nature considered for this work are predictive tracking, direct search, crosswind casting and saccadic tracking. In summary:

- Predictive tracking assumes that the target will maintain its last known speed and direction, the searcher behaves as if the target is still in view and moves to intercept at a suitable location where the target could reappear.
- Direct search assumes the target has not moved far, if at all, from its last known position, *e.g.* it could have realised it was being pursued and gone into hiding. The searcher in this case makes a direct line for the last known target position.
- Crosswind casting, in nature, is applied where the searcher is attempting to reacquire a plume of odour (*e.g.* during Chemotaxes). Since the target in such cases presents a broad front, the searcher turns perpendicular to the air flow, moving back and forth attempting to re-detect any part of the plume.
- Saccadic tracking is a sensory tracking mechanism, employed where the target is moving faster that the searcher can track (due to limitations of visual sensing, where focus cannot be maintained). In such scenarios, the searcher shifts visual attention to a point just ahead of the target attempting to briefly re-acquire it as it passes the focal area before performing another saccadic shift. If a shift is performed and the target does not appear at the expected point, the focus moves back along the predicted track until either the target is reacquired or the search is given up.

The strategies implemented in this research that deal with reacquisition search, after loss of target detection, are termed Predictive Memory System (PMS) strategies, since they attempt to predict the most likely location of a lost target, utilising a short-term memory and simple trigonometry. The PMS, regardless of strategy, has several elements:

- *Target location prediction.* Targets are either predicted to have remained in the vicinity of the last known location or to have continued on their last known heading and speed. Both scenarios model behaviours observed in nature: the former is included to assess whether it could be of value against targets that have a high turn rate (and hence prediction could be less reliable); conversely, the latter is incorporated to assess whether it could add value where targets are more predictable.
- *Reacquisition attempt decision.* All agents, on noting that they are no longer receiving information regarding a previously communicated target, have to decide whether they should attempt to intercept the target based on a predicted position.

Such a decision is either based on membership of a given neighbourhood or the ability to reach a predicted target location within a given time period (*i.e.* no later than when the target is predicted to pass through).

- *When to give up.* Once an agent decides to search for a lost target it should only continue to do so for a finite period of time, otherwise it may be wasting energy seeking a target that is no longer present or has deviated significantly from its predicted behaviour. Even where agents do not initially decide to engage in a reacquisition search they should still maintain the target's predicted position in memory for a period of time in case the agent finds itself in such a position that to change to a reacquisition search may be profitable. In this implementation a single giving up time is used because the likelihood of redetection diminishes at the same rate (*e.g.* an agent finding itself entering a neighbourhood after 50% of the giving up time has passed will only attempt reacquisition for the remaining 50%).

The specific PMS strategies are now considered, starting with those where the decision to attempt reacquisition is based on membership of a neighbourhood.

Similar to the LSZ neighbourhood described earlier, neighbourhoods are implemented for the occlusion and communications failure experiments that consider target reacquisition (associated algorithms are described in Section 4.4.2). When a target is lost, the neighbourhood determines which agents may be associated with attempts to reacquire the target. The neighbourhood is referred to as the Predictive Memory System (PMS) neighbourhood. Two neighbourhood shapes were implemented: circular, for predictive tracking and direct search; and V shaped for predictive tracking and crosswind casting (hereafter, the two predictive tracking variants are termed predictive-C and predictive-V, respectively, to avoid confusion).

Neighbourhoods generally follow the underlying metaphor on which the reacquisition strategy is based. A circular pattern only was applied to direct search since that approach considers that the target is still located somewhere in the vicinity of its last known location; that being the case, any agent within a reasonable distance may be able to reacquire it. Similarly, crosswind casting applies only a V shaped neighbourhood since the underlying metaphor is based on tracking odour plumes and assumes that they continue to travel in the same direction, in which case agents located behind the target would not expect to be able to

redetect since the target would continue to travel away from them. Predictive tracking also expects the target to maintain its speed and direction; however, it was also implemented in a circular neighbourhood to reflect an 'optimistic' tracker that may continue to track the target in the hope that it would change direction or slow down.

Consider the scenario depicted in Figure 14. When target detection is lost (at point A) the agents continue to predict its path. If, at any time until the GUT (Giving Up Time) has expired, an agent finds itself in the neighbourhood it attempts to re-detect the target by moving toward its predicted actual location. No lead was included, since to do so would begin to move away from the predictive search metaphor into that of a saccadic search. At point B, agent 3 finds itself in such a position and, since predictive-V PMS is in operation the agent moves toward the predicted location.

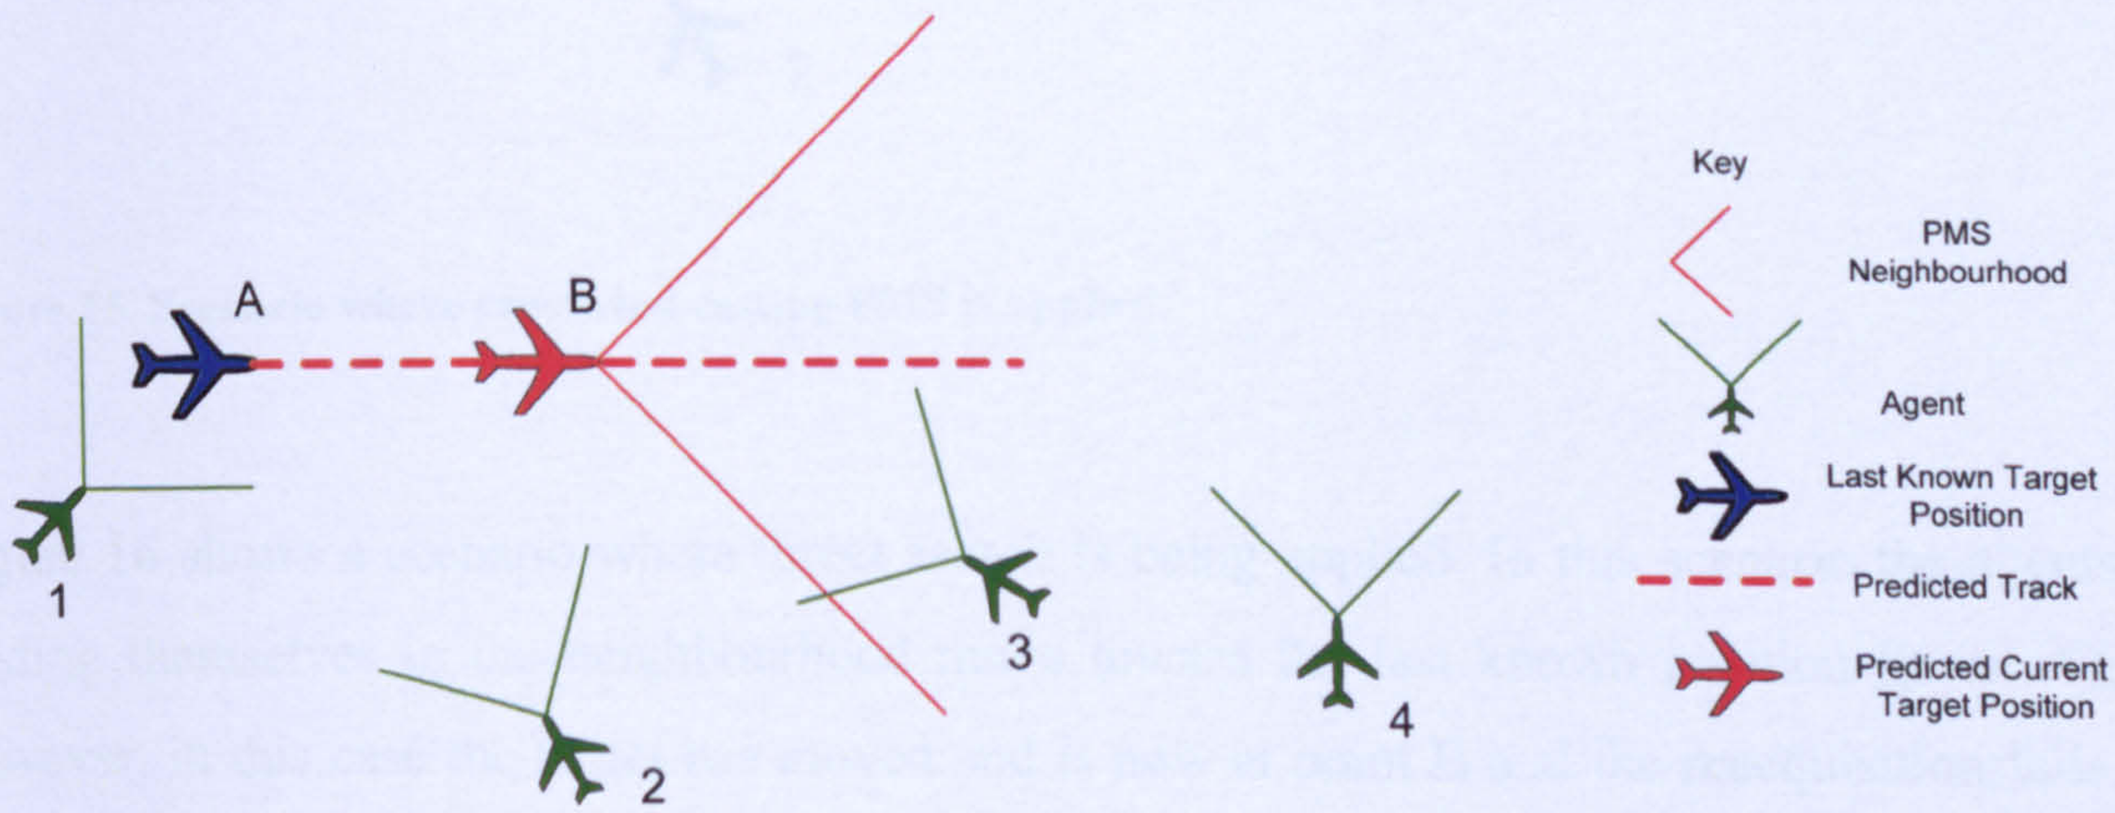

Figure 14. Scenario where predictive-V PMS is applied.

In the next example, Figure 15, a similar situation exists except crosswind casting is in operation. By the time the target, assuming it follows its predicted path, reaches point B agent 3 has been inside the neighbourhood for some time, has crossed the predicted path and the target will pass behind it; Agent 4, however has only just joined the neighbourhood and may well be in a better position to intercept. An obvious solution, to the problem of agents passing in front of a target and not being in a suitable position to re-detect, would be to allow the agents to slow down such that they do not pass the predicted path before the target. However, this was not implemented because it would move too far from the natural behaviour. Moreover, it may not be bad behaviour to operate as implemented because the

target could deviate from its path and even if it does not, the agent could reacquire when it turns back toward the predicted path to continue casting.

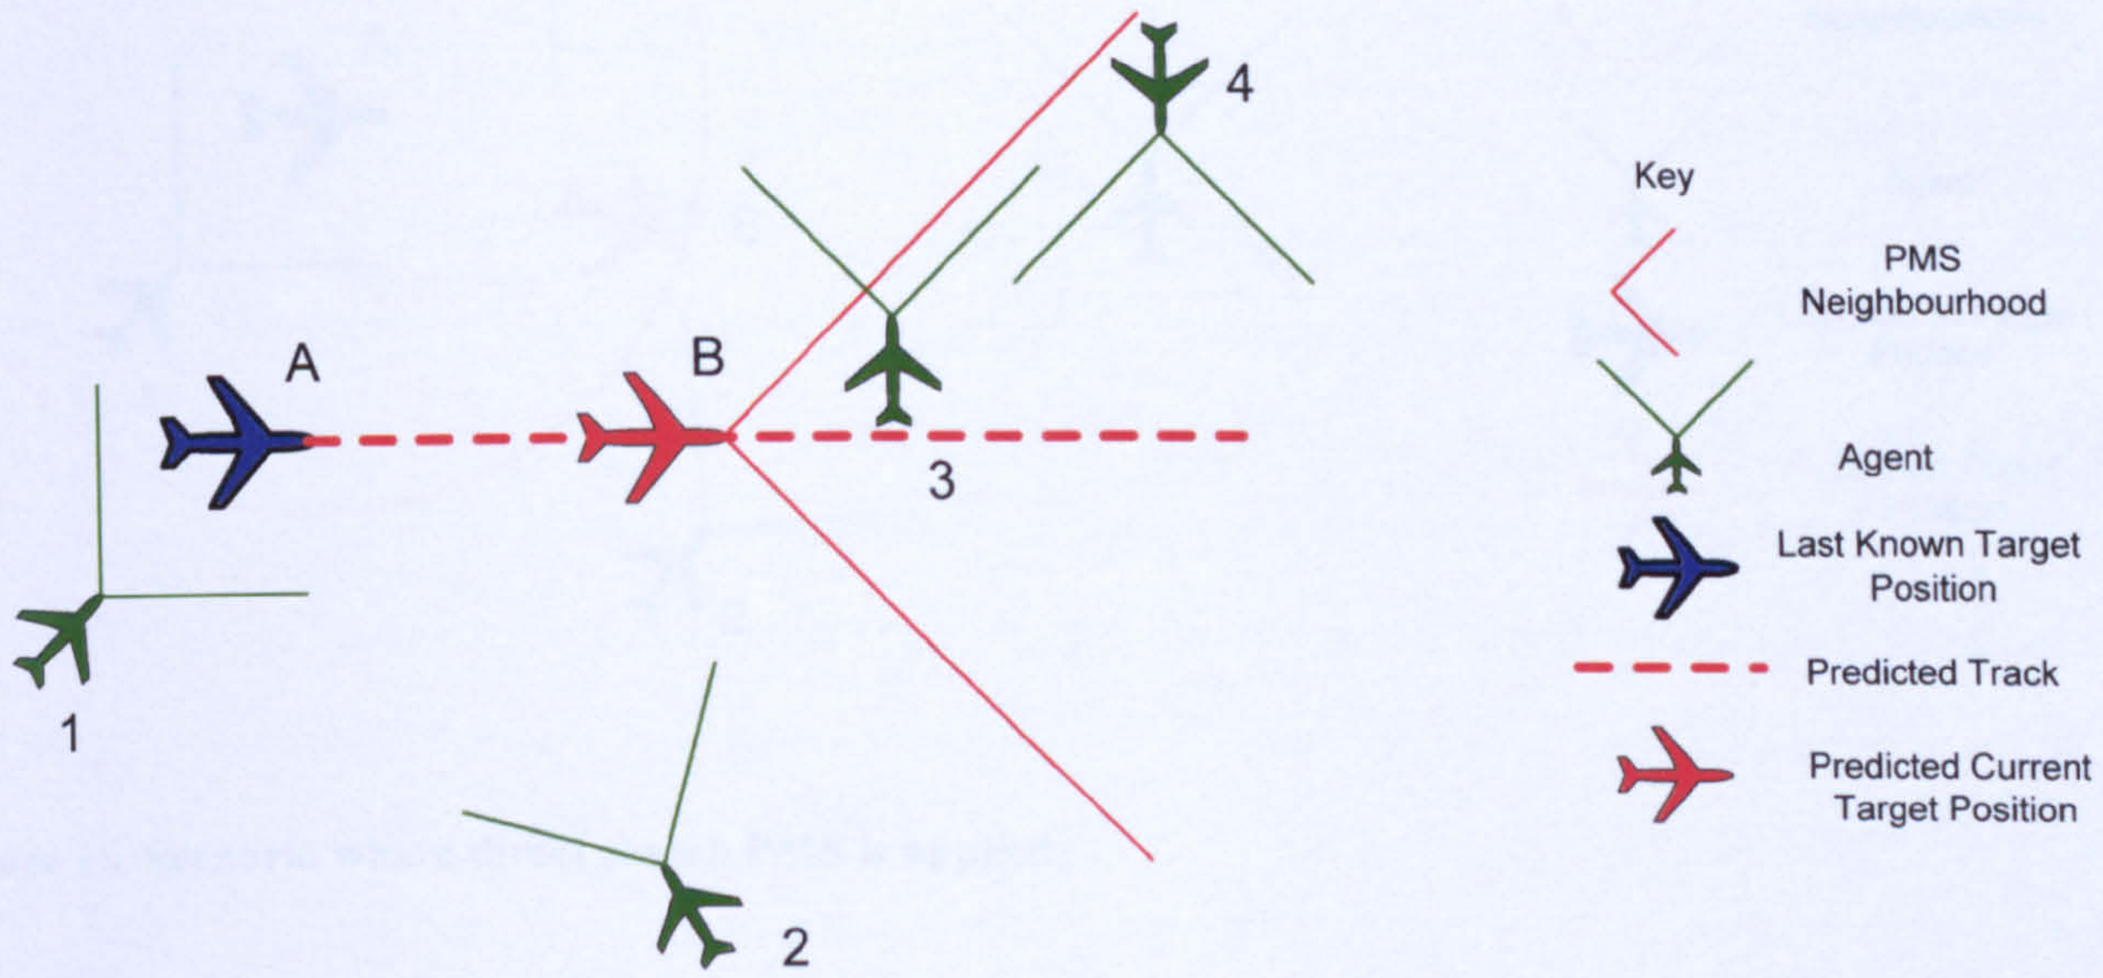

Figure 15. Scenario where crosswind casting PMS is applied.

Figure 16 shows a scenario where direct search is being applied. In this scenario the agents finding themselves in the neighbourhood move toward the last known position (point A); however, in this case the target has moved and is now at point B and the reacquisition fails. Had predictive-C been the PMS in operation, the locus of the neighbourhood would have been at a position just above point B and agents 2 and 4 may have moved to a potentially more profitable location. This does not suggest a predictive neighbourhood is always better, rather that it may have been in this situation.

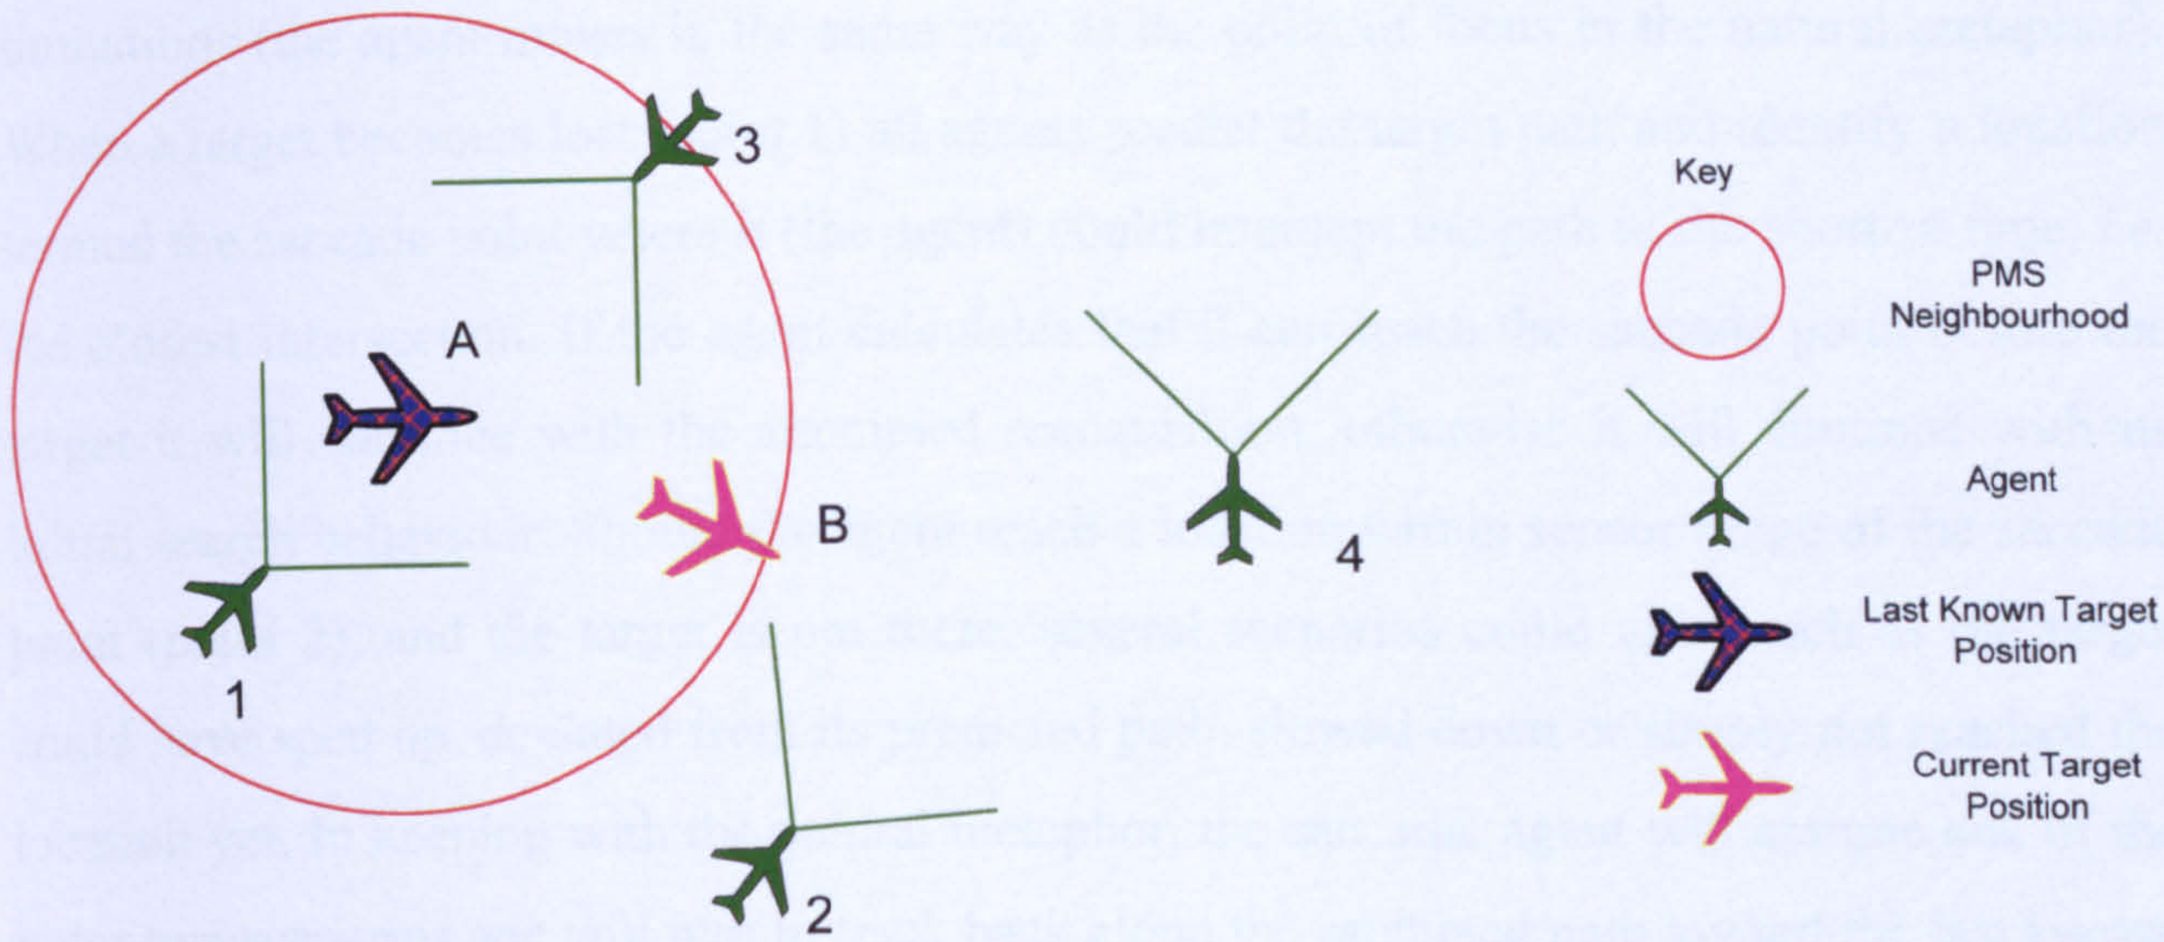

Figure 16. Scenario where direct search PMS is applied.

This optimism, suggested by agents using the predictive-C circular neighbourhood, was not conferred on the crosswind casting strategy since the agent turns at right angles to the target movement and therefore if the agent were behind the target to begin with it would become increasingly unlikely to redetect.

The linear function of Equation 16 (page 65) was applied to both neighbourhood shapes. A minimum neighbourhood size was also included for both shapes and set such that where a target has a relative speed of less than equality the neighbourhood was set to that of an equal speed relationship. This was introduced because an agent losing detection of a slow target could have found itself outside the neighbourhood straight away, precluding a reacquisition search; for example, if the target was lost at the edges of sensor range then, due to a low relative target speed, the neighbourhood could be smaller than the agent's sensor range. In such cases, with the minimum neighbourhood, the agent would switch back to initial search without attempting reacquisition. For the circular neighbourhood Equation 16 determines the radius of the neighbourhood in relation to the relative target speed and sensor range. Parameters of the neighbourhood are tuned for best performance (see Section 5.2.2, page 105).

Saccadic PMS (Figure 17) does not apply a neighbourhood but imitates the sensory behaviour employed where a predator is unable to continuously detect a target due to sensory

limitations (the agent moves in the same way as the point of focus in the natural metaphor). When a target becomes lost (point 1) all agents predict the target path and identify a location termed the saccade point where it (the agent) could intercept the path in the shortest time, *i.e.* the closest intersection. If the agent calculates that it can reach the saccade point before the target it will continue with the attempted reacquisition, otherwise it will continue with its initial search behaviour. Should the agent reach a location within sensor range of the saccade point (point 2), and the target is not there, several scenarios could exist such as the target could have sped up, deviated from its predicted path, slowed down or simply not reached the location yet. In keeping with the natural metaphor, the saccadic agent will assume one of the latter two scenarios and will turn to track back along the predicted path toward the last known position. Like direct search, the PMS is ceased if either the PMS GUT is exceeded or the last known position reached.

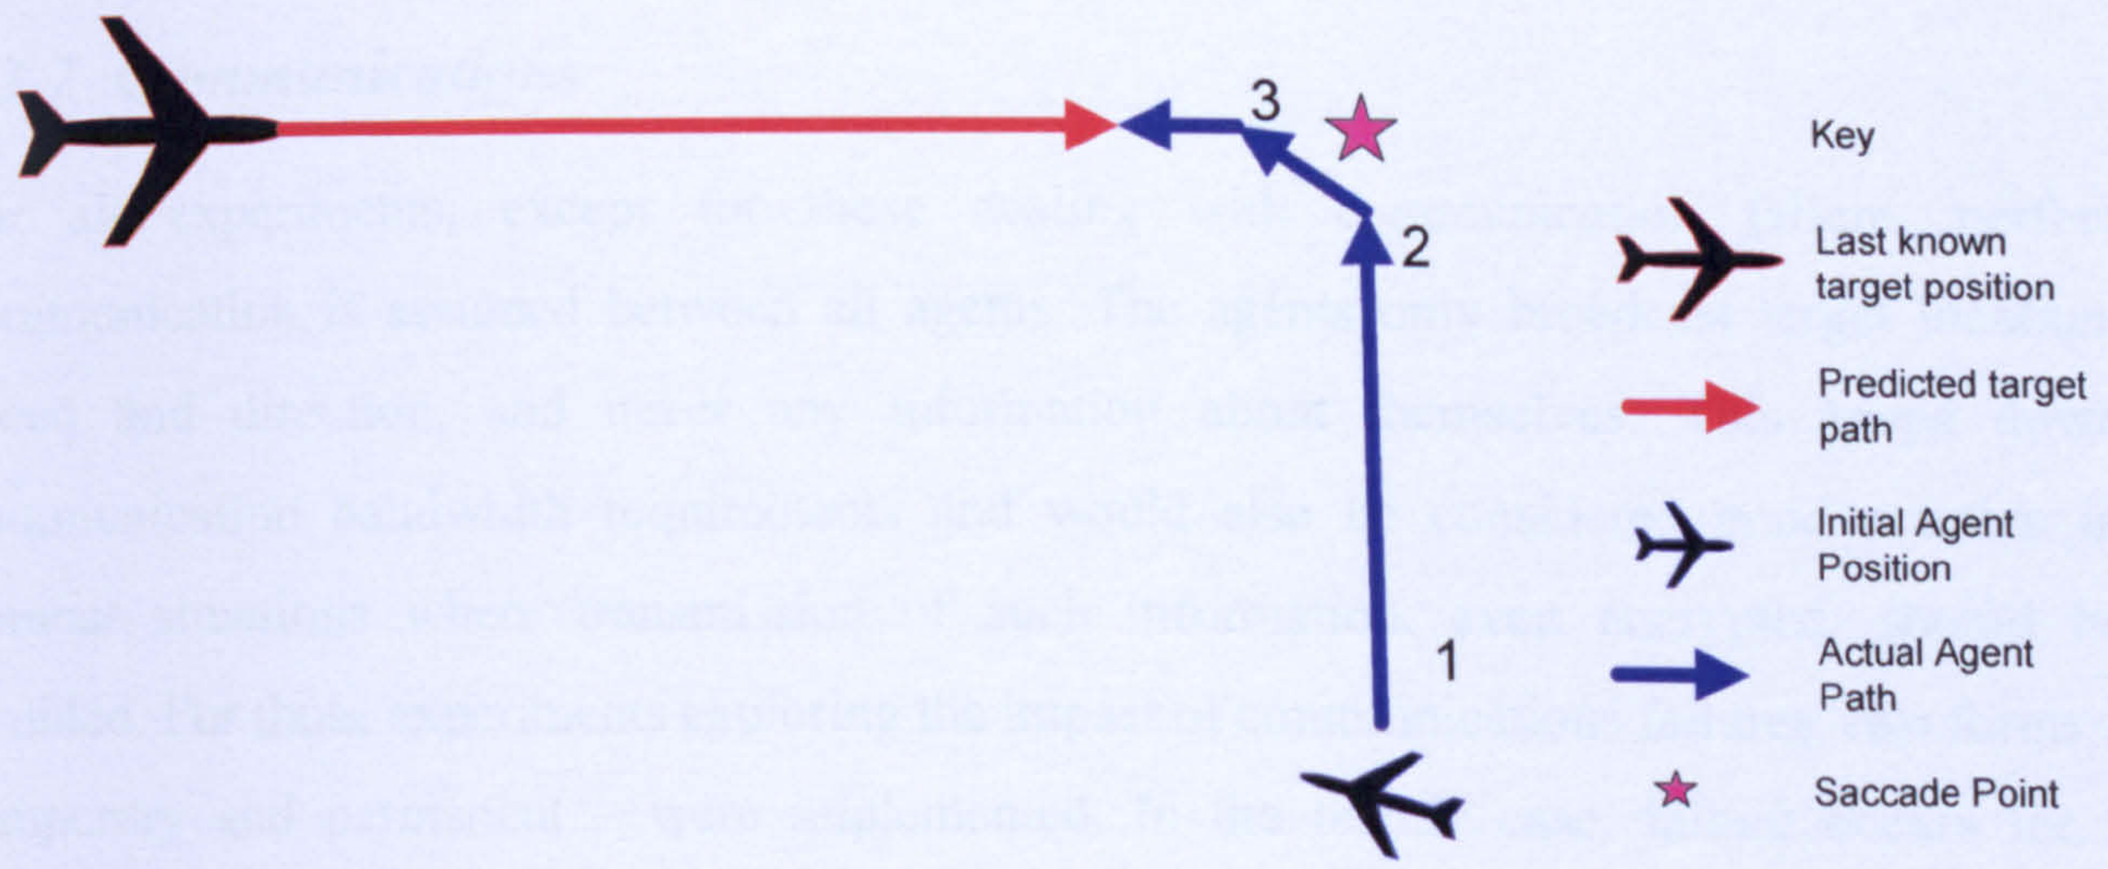

Figure 17. Example path of an agent intercepting an occluded target using a saccadic PMS.

4.3.6 Occlusion

Since the simulation environment is featureless, it was necessary to implement an occlusion mechanism to simulate the target becoming occluded from a detecting agent, thereby facilitating evaluation of the PMS in scenarios where targets are lost whilst still within sensory range.

Two indicative occlusion types were identified. The first, (hereafter termed permanent occlusion) simulates a target becoming occluded from the system due to an agent failure (such as sensor failure). When permanent occlusion is applied, it always occurs upon the first detection by an agent and the target remains occluded to the initial detecting agent but the target remains detectable by other agents. The second type, temporary occlusion, simulates where targets may become occluded through environmental features or target evasion manoeuvres such as 'notching' (where an aircraft turns at right angles to detecting radar). In this scenario the target occlusion is not certain and occurs randomly with 5% probability on each stepwise move. Redetection is possible by the original detecting agent, once again randomly, except with a 1% probability on each stepwise move. These figures were selected to offer a challenge to the system, and therefore allow the merits of the different PMS strategies to be evaluated. As with permanent occlusion, other agents could detect the target whilst it remains occluded to the agent that first detected it.

4.3.7 Communications

For all experiments, except for those dealing with communication failure, perfect communication is assumed between all agents. The agents only broadcast target location, speed and direction, and never any information about themselves. This keeps down communication bandwidth requirements and would also be considered good practice in combat situations where transmission of such information, even encrypted, should be avoided. For those experiments exploring the impact of communications failures, two forms – temporary and permanent – were implemented. In the former case, failure occurs for a random period between 10 and 20 steps following which good communications were maintained for the same period, before again failing. This time period was selected to provide a reasonable mix of availability; longer periods of outage would yields results more akin to having smaller swarm sizes, whilst shorter periods would be less effective at testing the predictive memory system. For permanent failures, the communication system of affected agents was unserviceable and could neither transmit nor receive information, although they could still detect and potentially destroy any targets within sensor range. In both types of failure, agents that were not in receipt of target information continued to behave as if no target information existed, *i.e.* they either continued with their initial search behaviour or applied a PMS strategy as appropriate.

4.4 Search Algorithms

The implementation of the search algorithms and the reacquisition of targets form the main emphasis of this section. However, a brief discussion concerning the implementation of the administrative hierarchy, for decision making, intercept behaviours and Giving Up Time (GUT), is also included to highlight important considerations made during the implementation.

4.4.1 Initial Search

The strategies applied by agents in the initial search phase (*i.e.* behaviours applied prior to the detection of a target, or when a target has escaped detection and the agents involved have given up their pursuit and returned to their search function) are described in this section. These strategies comprise three random walks (Brownian motion, Lévy paths and straight line search,), Area Concentrated Search (ACS), saltatory search, and two species specific movement patterns (bacterial and fruit fly). The illustrative search patterns formed by each initial search strategy shown in this section were taken directly from the simulation environment. The shaded line shows the path of a single agent over a given duration. The period of time spent in each location is indicted by its shade, darker squares indicating longer periods. Note that this illustrates agent location, rather than sensor coverage.

First, the basic random walks of Brownian motion, Lévy paths and straight line searches are considered, having the same basic formulation but differing in terms of leg lengths and turns rates. As discussed in Section 3.2.6 (page 32), ACS is a two-phase search, implemented here as a combination of straight line, for the extensive search phase, and Brownian motion, for the intensive search phase; the former being used until a target was detected and the latter until the giving up time (GUT) was exceeded (see Section 5.2.3, page 108, for tuning). Due to its trivial nature, the implementation of ACS is not discussed further. The remaining searches, whose implementation description follows that of random walks, also include elements of random walks and where this occurs any variations are highlighted.

Each natural search pattern was implemented such that it reflected the natural strategy it was inspired by. All random walks were implemented in the same manner (see Figure 18). To model the leg length probabilities described by Equation 7 (page 28) the leg lengths were set by truncating (removing fractions of a unit) the result of $\sqrt[-\mu]{rand}$, where rand is a uniformly

distributed random number in the range [0,1] and μ sets the path tortuousness, *i.e.* $\mu = 1$ produces longer leg lengths than $\mu = 2$ giving the latter a higher turn rate. The truncation is implementation specific and was necessary because, in this implementation, distance travelled was measured as time spent on a given heading rather than actual distance travelled; this simplification was possible since constant agent and target speed meant that the distance calculations were unnecessary. In implementations where actual distances are used the truncation would not be necessary. As an exception, for straight line search only, the random numbers were drawn from the range $0.0005 < \text{rand} \geq 1$ in order to constrain, for pragmatic reasons, the leg lengths, since values less than 0.0005 result in excessively large legs. The three random walks are characterised by differing values of μ, with $\mu = 1$ for straight line search, $\mu = 2$ for Lévy paths, and $\mu = 3$ for Brownian motion. Additionally, in keeping with the literature (Bell, 1991), the mean turn angle of straight line searching was restricted to 35°; whilst for others there was no such restriction. The raw leg lengths, for all three random walks, were then scaled up to suit the search area by a factor of 10 in keeping with agent sensory and speed capabilities.

```
if distance travelled < leg length
then
    maintain heading
    increment distance travelled
else
    if movement = Brownian
    then
        leg length := truncate(10 · ⁻³√rand)
        change heading (range ± [0°, 180°])
    else if movement = Lévy
    then
        leg length := truncate(10 · ⁻²√rand)
        change heading (range ± [0°, 180°])
    else
        leg length := truncate(10 · ⁻¹√rand)
        change heading (mean change ± [35°])
    end if
end if
```

Figure 18. Pseudo-code for random walks.

As can be seen in Figure 19 (left and right), the resulting paths were similar for Brownian motion and Lévy paths, with Brownian motion also being used for the intensive phase of ACS. A typical path of an agent applying straight line search is shown in Figure 20 (left); this behaviour was further utilised in the extensive phase of ACS.

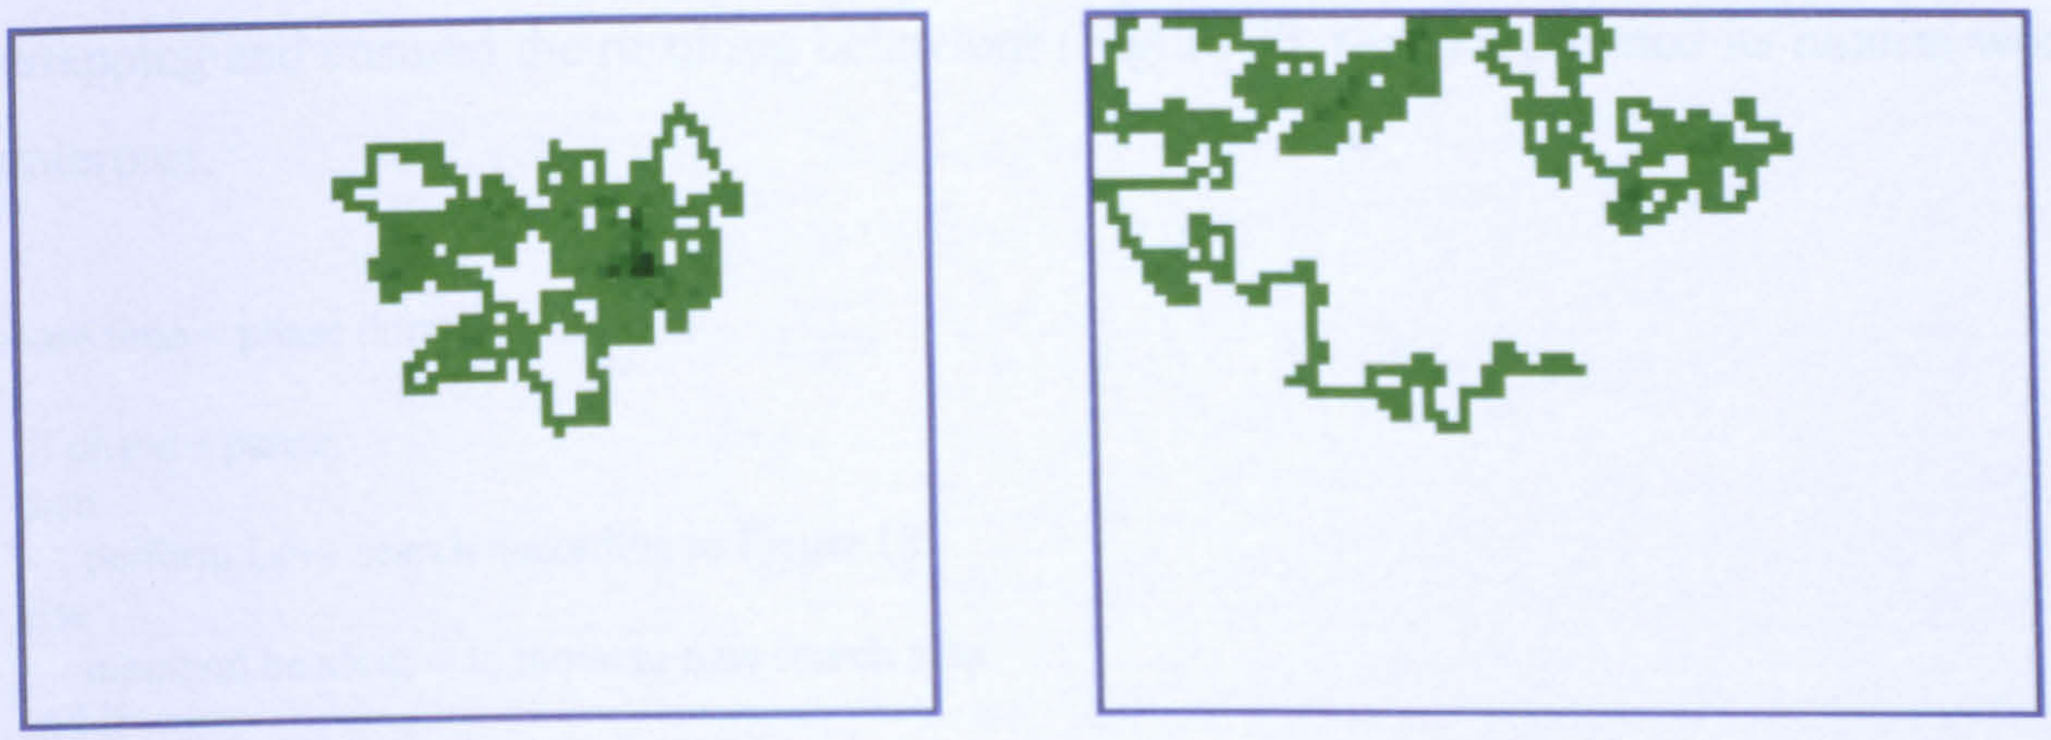

Figure 19. Typical movement pattern produced by, left: Brownian movement search and ACS (intensive phase); right: Lévy path search.

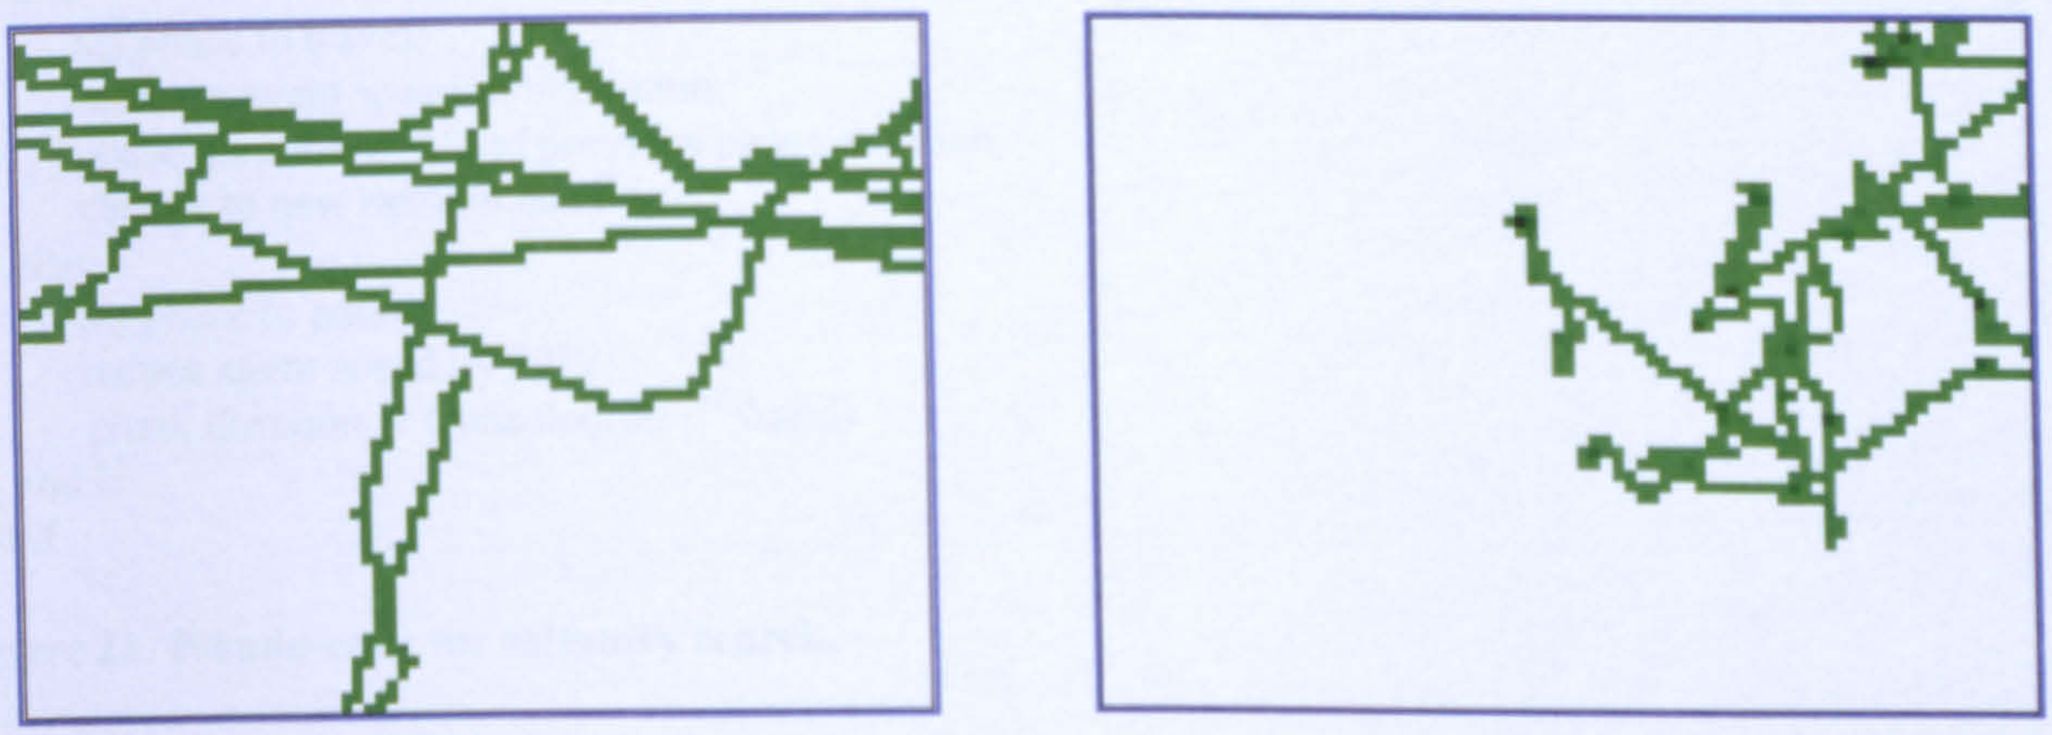

Figure 20. Typical movement pattern produced by, left: straight line search and ACS (extensive phase); right: saltatory search (note: dark points on route highlight where the agent paused and searched intensely).

Saltatory search applied Lévy path search during the intensive search phase, *i.e.* a random walk with $\mu = 2$, but also slowed the agent down by 50%. A modification to the parameters provided in the literature was the reduction in the GUT of the intensive search. This was made in acknowledgement of the environment, the targets were not hidden and to linger for too long (where there was plainly no target) would not have mimicked natural predators. In applications where targets could hide (for example, in dense vegetation) GUT times would be expected to be extended accordingly. The intensive search phase duration was calculated in the same manner as Lévy search except the leg length was extended by a scaling factor of 100 (rather than 10). A further modification to the indications from the literature was the relationship between the intensive and extensive search phases, which was extended from a 2/3 power to a 2/3 product. This was necessary to prevent the intensive search phases from

overlapping and ensured the resulting behaviour (Figure 20, right) replicated its natural world counterpart.

```
if phase time < phase duration
then
    if phase = pause
    then
        perform Lévy search according to Figure 18
    else
        maintain heading -- to move to new search area
    end if
    increment phase time
else
    if phase = pause
    then
        set phase to travel
        increase agent speed to maximum
        phase duration := 2/3 of previous pause duration
        change to new random heading
    else
        set phase to pause
        reduce agent speed by 50%
        phase duration := truncate(100 · ²√rand)
    end if
end if
```

Figure 21. Pseudo-code for saltatory search.

In bacterial search, where no targets are being detected, the agent is slowed down by 50% and a high a turn rate (as per bacterial tumbling behaviour) is maintained alternating with Brownian motion (as per bacterial swimming behaviour) for equal periods of time (see Figure 22). Figure 23 (left), illustrates the limiting effect this has on the explorative aspect of the search. When a target is detected, the whole group speeds up and switches to an unequal turn and movement pattern (ratio of 2:1). This mimics the behaviour of bacteria in the presence of a positive nutrient gradient, as illustrated in Figure 23 (right). A further attempt to mimic the bacterial response to a positive nutrient gradient is that the entire swarm of agents respond to the detection of the target. In the real world each bacterium responds independently, nonetheless, since all of the bacteria encounter the nutrient the effect is of a coordinated response. This cooperative implementation deviates from the individual response since only one agent needs to encounter a target for the swarm to switch mode, however the overall effect remains – the individuals switch modes in response to target presence. One could envisage a need for more localised responses when searching a wide area. Therefore, the

bacterial search was explored using various neighbourhood sizes, to identify whether there is any advantage in this experiment environment to providing a more local switch between behaviours. Results (which are included in Appendix 1) show that there is no benefit to be gained, here, with a local mechanism, and a simple global switch from 'tumbling' to 'swimming' modes when any agent detects a target works well.

```
if no targets detected by any of the agents
then
    slow to half normal speed
    if in tumble mode
    then
        change the heading randomly -- tumble phase
    else
        apply Brownian motion search  -- swim phase
    end if
    increment count
    if count >= duration
    then
        reset count
        toggle tumble mode
        if in tumble mode
        then
            duration := random duration (range [1,20])
        end if
    end if
else -- Chemotaxes (increase speed and swim length)
    move at normal speed
    if in tumble mode  -- tumble phase
    then
        change the heading randomly
    else
        apply Brownian motion search -- swim phase
    end if
    increment count
    if tumble count >= tumble duration
    then
        reset count
        toggle tumble mode
        if in tumble mode
        then
            duration := random duration (range [1,20])
        else
            duration := double the duration -- to maintain 2:1 ratio
        end if
    end if
end if
```

Figure 22. Pseudo-code for bacterial search.

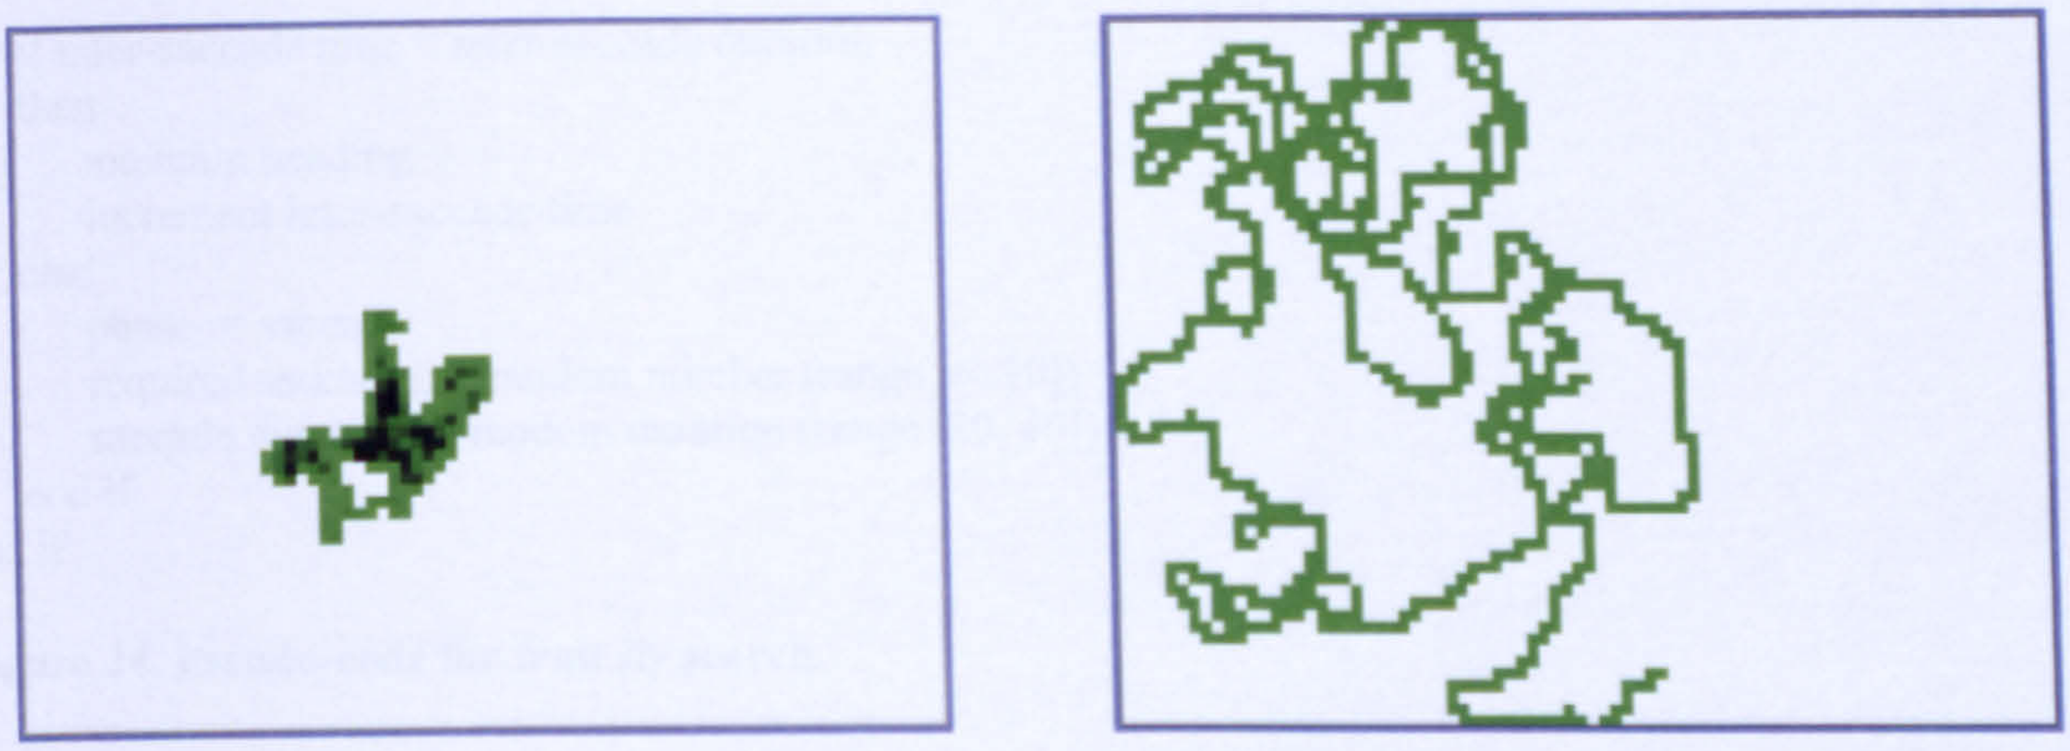

Figure 23. Typical movement pattern produced by bacterial search, left: no targets detected; right: targets detected.

Fruit fly search (Figure 24) utilised representative movement patterns, *i.e.* overall patterns produced mimicked those presented in the literature (for example, Tammero and Dickinson 2002), since the systems and stimuli to switch between phases of flight, such as the visual feedback from the environment described in the literature (Frye *et al.*, 2003), are absent in this application. To reproduce the required behaviour the agent applies a repeated two phase movement consisting of a long, straight, movement followed by a series of saccades; the saccades being short legs separated by sharp angular turns of 100° ± 20° (note: these turn angles were relaxed from those given in the literature to allow for more realistic vehicular movement; however, the ratio of inter-saccade leg lengths to saccade leg lengths and the number of saccades prior to moving to another area were maintained as described by Frye *et al.*, (2003). The resulting movement pattern is shown in Figure 25 (left).

```
if phase = saccade
then
    if saccade count < required saccades
    then
        if saccade time < saccade duration
        then
            maintain heading
            increment saccade time
        else
            change heading (range ± [80°, 120°])
            saccade duration := random duration (range [20, 40])
            increment saccade count
        end if
    else
        phase := inter-saccade
        inter-saccade duration := random duration (range [250, 350])
else
```

```
    if inter-saccade time < inter-saccade duration
    then
        maintain heading
        increment inter-saccade time
    else
        phase := saccade
        required saccades := random number (range [6, 10])
        saccade duration := random duration (range [20, 40])
    end if
end if
```

Figure 24. Pseudo-code for fruit fly search.

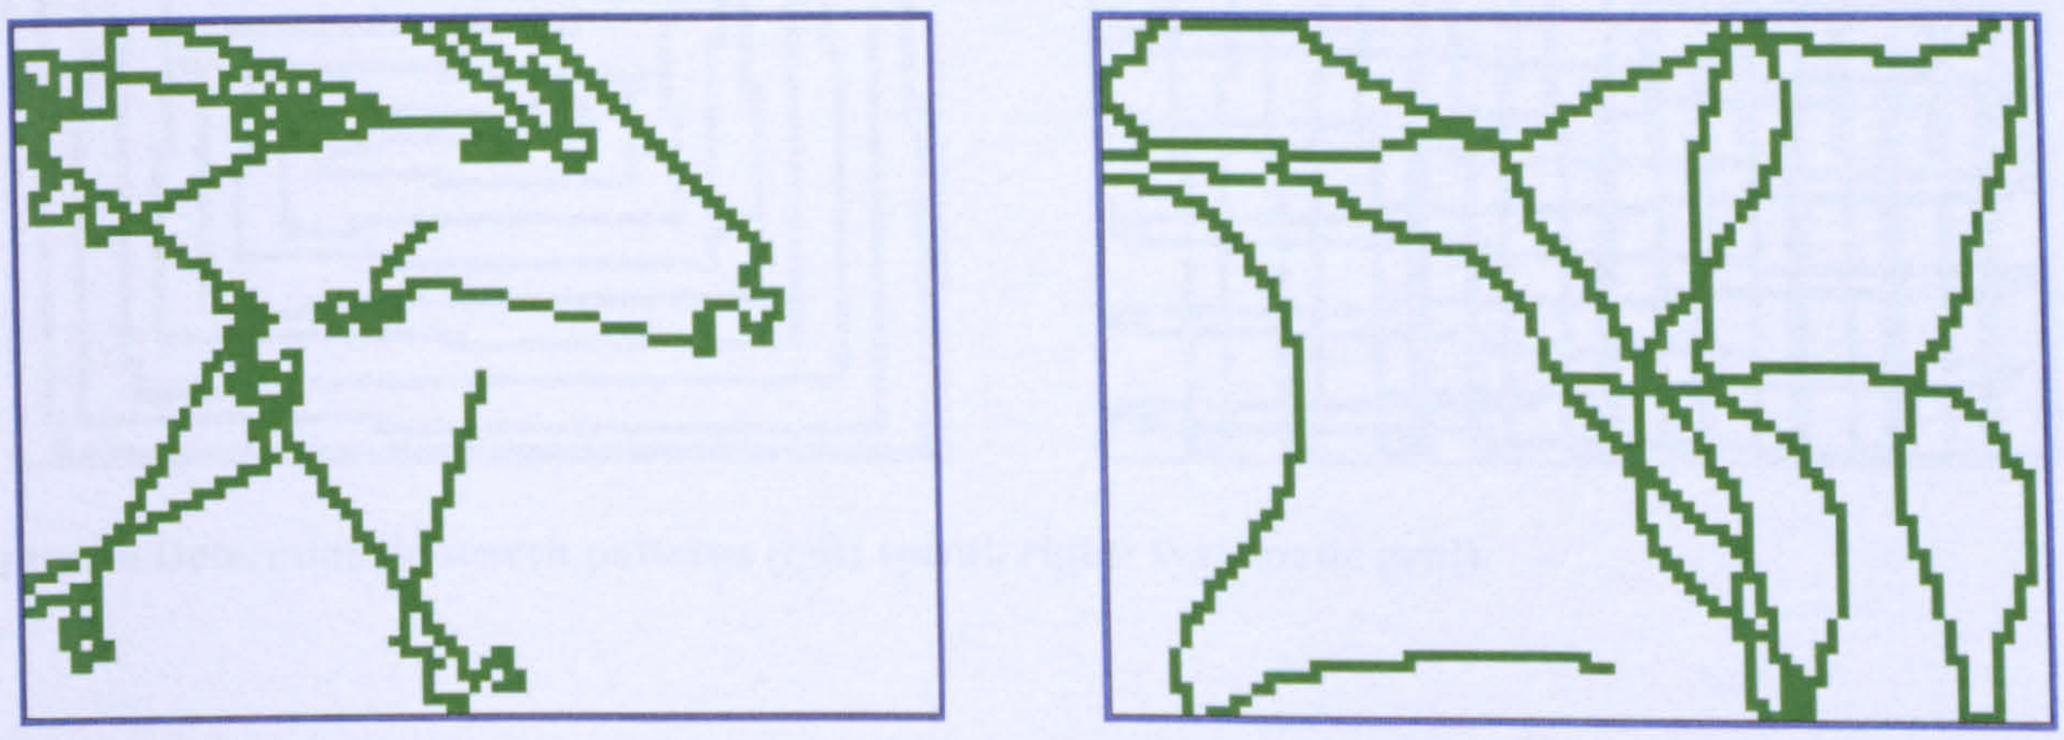

Figure 25. Typical movement pattern produced by, left: fruit fly search; right: random search.

The random movement search pattern was implemented to provide a relatively stable movement whilst maintaining sensor coverage. This was achieved by setting a uniformly distributed stochastic leg length in the range [1, 20] prior to turning with a uniformly distributed random heading change in the range [-20°, 20°]; Figure 25 (right) shows a typical movement path (the implementation of this is trivial and so not included as pseudo-code). These settings were selected to provide movement that did not replicate natural random walks or deterministic paths yet still allowed the agents to range across the search area.

Both deterministic patterns (spiral and systematic grid) were implemented by presenting the agents with different outbound and return routes, as illustrated in Figure 26. Both routes begin in the top left hand corner of the search space and follow a series of waypoints to a final destination; in the systematic pattern this is the bottom right of the space, and in the spiral it is the centre. Both strategies then take an alternative route back to the start. Agent positions are randomly initialised across the routes. To ensure complete space coverage it was

necessary for the sensor coverage to overlap. This was essential when searching for static targets and was identified as a weakness of the strategies: if an error in the pattern, or an obstacle, were to render an area occluded from the agents' field of view, then a target at that location could remain undetected.

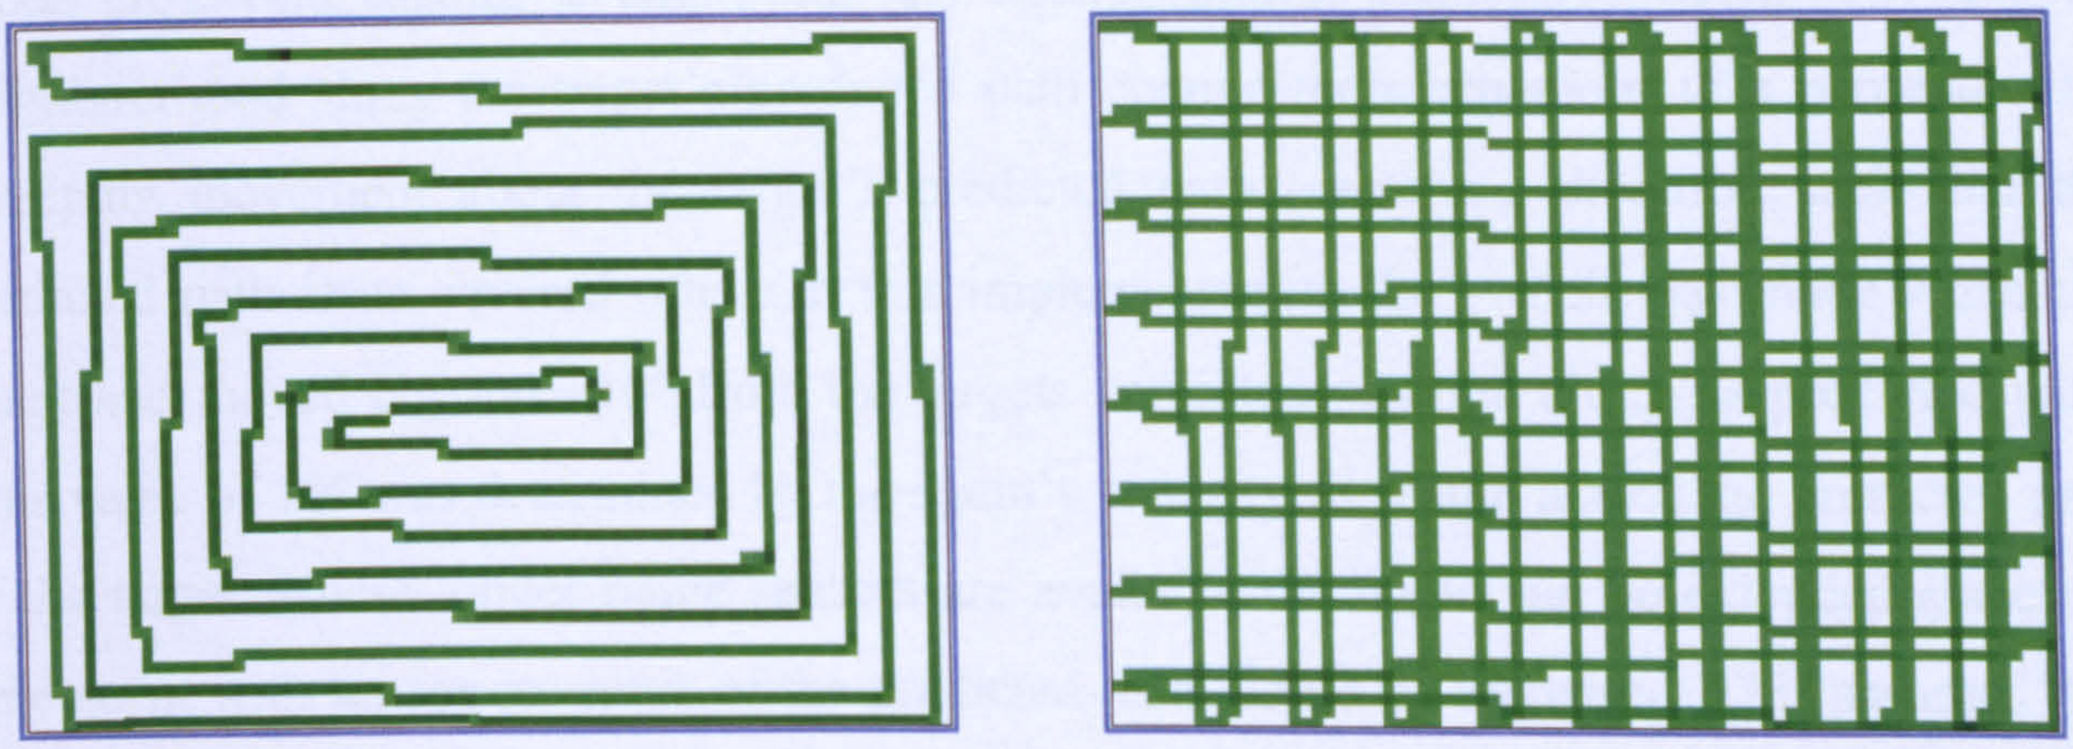

Figure 26 Deterministic search patterns (left: spiral, right: systematic grid).

4.4.2 Reacquisition

Five target reacquisition strategies were implemented, based on the behaviour of animals when a target becomes occluded, as identified in the background material (Section 3.10). Strictly speaking, only four actual strategies were employed since two of the five, predictive-V and predictive-C, are based on the same predictive tracking behaviour, but with different neighbourhood shapes (as discussed in Section 4.3.5). This section begins with a discussion of the four neighbourhood based strategies before looking at the implementation of the non-neighbourhood based saccadic search.

In this implementation, each agent holds lists of currently detected targets, both their own and those reported by other agents. Periodically (*i.e.* following each stepwise move), the list for each agent is updated and compared with the previous version of the list. Where a difference occurs, the agent adds the missing target(s) to its occluded target list. When an agent is not currently detecting, or is unable to cooperate with any other agents that are detecting targets, it then attempts reacquisition based on its PMS as described in Section 4.3.5. The algorithms for predictive-V, predictive-C, crosswind, and direct strategies are similar and the pseudo code is presented in Figure 27. During predictive tracking, the agents record the target's last

known position, heading and speed; from this information the agent predicts the target location using simple speed and time calculations. Agents finding themselves within the PMS neighbourhood (circular or V shaped) then set course for the predicted target position. They continue to attempt reacquisition until they either find themselves outside the neighbourhood or the PMS GUT expires (this giving up strategy also applies to crosswind casting).

When crosswind casting is employed, any agents finding themselves within a V shaped neighbourhood along the target's predicted path change their behaviour to a perpendicular sweeping movement about the target's predicted path, switching direction such that the predicted path is as covered often. In this implementation, the switch was made when the agent had moved beyond ±10° from the targets predicted position along its predicted path. This value of 10° was determined by the agent's capacity to search across the predicted path of the target, where longer range sensors are available the angle can be extended since the time spent with sensor coverage of the predicted path would be increased. This assumes that the agent would wish to spend as little time as possible located on, and perpendicular to, the path of the target because it could be vulnerable to counter attack.

Direct search does not predict the target course but assumes that there is just as much chance of re-detecting the target at the location it was last seen at; hence, agents employing this strategy that find themselves within the circular PMS neighbourhood, change course for the last known target location. Once again, if the PMS GUT is exceeded the search is ceased; however, since the neighbourhood does not change, agents may also cease the direct search once they have ensured that the target is not at its last known location.

Finally, saccadic PMS is implemented based on animal eye movements searching for fast moving lost prey items. In this implementation, a saccadic style prediction is made of the closest point the occluded target would pass by the agent's location at the time of initial occlusion. If the agent can move to that position before the target, given its last known speed, then it will attempt an interception. A more detailed description of the various PMS implementations follows.

In the neighbourhood scheme based PMS, the agent first assesses whether it (the agent) is within the neighbourhood (which is PMS dependent) associated with a lost target. In the case of multiple lost targets, it then considers the fitness of its own position based on the Euclidean distance to the target positions in use (either predicted or last known, depending on

the reacquisition strategy). Once the decision has been made to attempt to re-acquire a target, the agent either sets course for it at its last known position for direct search or its predicted position for predictive searches, or, in the case of crosswind casting, starts perpendicular sweeping movements with respect to the target's predicted heading.

It should be noted that, in all cases where the agent heads toward the target position, that is either the predicted position during reacquisition (*e.g.* not when heading for a saccade point or during crosswind casting in PMS) or actual position for detected and known targets (Sections 4.4.3 and 4.4.4 respectively), the agent does not move directly toward it, instead the agent's next position is set by difference between the target location and the agent's current location. This is a feature of the previous particle swarming research and hence retained in this implementation, the effect of the swarming on the results is examined in Section 5.5 (page 137).

```
fittest_target_fitness := 99999.9    -- i.e. value > maximum distance
for each occluded target
    if direct search
    then
        set target position to last known target position
    else
        set target position to predicted target position
    end if
    target_distance := Euclidean distance from the target position
    target fitness := Euclidean distance from the target position
    if PMS = predictive-V or PMS = crosswind casting
    then
        if target_speed <= agent_speed
        then
            neighbourhood_length := sensor_range · c
        else
            neighbourhood_length := sensor_range · c · (target_speed /agent_speed)
        end if
        if target_heading < | relative_target_bearing  - 40°| and
            target_distance < neighbourhood_length
        then
            target is in the neighbourhood
            if target_fitness < fittest_target_fitness
            then
                required_agent_position := target_position
                required_target_heading := target_heading
                relative_target_bearing  := bearing of target from agent
            end if
        else
            target is not in the neighbourhood
        end if
    else if PMS = predictive-C or direct
    then
```

```
    if target_distance < neighbourhood_length
    then
        target is in the neighbourhood
        if target_fitness < fittest_target_fitness
        then
            required_agent_position := target_position – current_agent_position
        end if
    else
        target is not in the neighbourhood
    end if
end if
end loop
if fittest_target_fitness < 99999.9   -- i.e. agent is in a neighbourhood
then
    if PMS = crosswind casting
    then
        if relative_target_bearing  < 100 and then relative_target_bearing  > 80
        then
            agent heading := required_target_heading - 90°
        else if relative_target_bearing  < 280 and then relative_target_bearing  > 260
        then
            agent heading := required_target_heading - 90°
        else if relative_target_bearing  < 180
        then
            agent heading := required_target_heading - 90°
        else
            agent heading := required_target_heading + 90°
        end if
    else
        required_agent_position := target_position – current_agent_position
    end if
end if
```

Figure 27. Pseudo-code of non-saccadic PMS algorithm.

Where agents employ a saccadic PMS, the initial action on discovering target occlusion is to identify the closest point along the target's predicted path to their current position and record it as the saccade point (see pseudo-code in Figure 28). If the intercepting agent calculates that it could be within sensor range of the saccade point before the target, then it sets course for the saccade point, otherwise it continues with its initial search behaviour. If, on reaching sensor range of the saccade point, the agent has not detected a target, it alters course for the last known target position travelling back along the predicted path. This reflects the natural metaphor described in Section 3.10, where the animal searches back along the saccade in the absence of prey. Because the real-world saccadic tracking is based on a single target, it was not possible to implement a nature inspired solution to multiple occlusions. A common sense approach was therefore applied, where, when multiple targets are occluded concurrently, the agent selects the target with the closest saccade point. If the multiple occlusion occurs when

the agent is attempting a reacquisition and has gone beyond the saccade point of its current target, it moves on to the attempt to re-acquire the newly occluded target only if the new saccade point is closer than the last known position of the current target.

```
for each occluded target
    if target is newly occluded -- i.e. first iteration
    then
        calculate saccade point
        if saccade point reachable before target passes by
        then
            target_saccade_defined := true
            target_saccade_reached := false
            target_fitness := Euclidean distance to saccade point
        else
            target_saccade_defined := false
            target_saccade_reached := false
            remove target from occluded target list
            resume initial search
        end if
    else if target_saccade_defined
    then
        if distance to saccade point < sensor_range
        then
            target_saccade_defined := false
            target_saccade_reached := true
            target_fitness := Euclidean distance to target
        end if
    else if target_saccade_reached
    then
        if distance to last known target position < sensor_range
        then
            target_saccade_reached := false
            remove target from occluded target list
            resume initial search
        end if
    end if
end loop
fittest_target_fitness := 99999.9    -- i.e. value > maximum distance
for each occluded target
    if target_fitness < fittest_target_fitness
    then
        required_agent_position := target_position
    end if
end loop
```

Figure 28. Pseudo-code of saccadic intercept algorithm.

4.4.3 Detected Target Interception

Once a target has been detected, an agent should attempt to get close enough to engage and destroy it as quickly as possible. However, where an agent is detecting multiple targets it must decide which would be the most appropriate to pursue. Such decisions are application specific. In this military inspired model, the criterion is based on the fact that a missile is most effective against targets that are moving in line with the agent (Figure 29).

```
fittest_target_fitness := 0.0
for each target (target)
    target_distance := target distance from agent
    -- find the reciprocal of target bearing with respect to agent heading
    rec_relative_target_bearing := 180 + | target_bearing – agent_heading |
    -- apply the fitness function
    target_fitness := 1.0 + COS(rec_relative_target_bearing)/target_distance
    if target_fitness > fittest_target_fitness
    then
        required_agent_position := target_position – current_agent_position
    end if
end loop
```

Figure 29. Pseudo-code of detecting agents' target interception algorithm.

4.4.4 Known Target Interception

Assuming they are not already detecting a target, each agent calculates whether they have membership of the LSZ neighbourhood for any known targets (*i.e.* targets being detected by other agents, as described in Section 4.3.4); if they are in the LSZ, the target is pursued. Where an agent finds itself located in multiple LSZs, the fitness of all the relevant targets is calculated as indicated in the pseudo-code and the fittest pursued. Figure 30 provides the relevant pseudo-code.

```
fittest_target_fitness := 0.0
for each target (target)
    target_distance := target distance from agent
    if target_distance < LSZ length (l)
    then
        -- find the bearing of the agent from the target using
        -- the target location relative to agent location
        agent_bearing := target_bearing + 180°
        if |agent_bearing - target_heading|  < LSZ angle (θ)
        then
            rec_relative_target_bearing := 180 + | target_bearing – agent_heading |
            -- apply the fitness function
            target_fitness := 1.0 + COS(rec_relative_target_bearing)/target_distance
            if target_fitness > fittest_target_fitness
            then
                required_agent_position := target_position – current_agent_position
            end if
        end if
    end if
end loop
```

Figure 30. Pseudo-code of known target interception algorithm.

4.4.5 Behaviour Hierarchy

The algorithms described above were constructed within an administrative behaviour hierarchy (see Figure 31) in which priority for target pursuit was determined by certainty of knowledge. Thus, targets being detected by the agent were given priority over those detected by other agents. Similarly, where an agent finds itself within one or more LSZ neighbourhoods due to other agents detecting targets, these are pursued in preference to any occluded targets. As an example, given a situation where an agent is detecting two targets, the agent should decide which one to pursue, this is decided by the fitness function described earlier and if that target then becomes occluded the agent should either try to reacquire it or pursue the other target still being detected. By applying the hierarchy the agent would change its focus to the remaining target being detected. The hierarchy implemented for this work is based on a simple 'bird in the hand' principle; however, it would not be difficult to envisage the need for more complex decision making mechanisms (such as one based on Andrews *et al.*, 2007), where, for example, the agent may need to consider other factors such as whether sufficient fuel remained or whether the target was of sufficient value for the agent to risk sacrificing itself in its pursuit and engagement.

```
if detected targets > 0
then
    remove known targets from memory
    calculate movement as described in 4.4.3
    broadcast all detected target information to the swarm
else if known targets > 0
then
    calculate movement as described in 4.4.4
else if PMS targets > 0
    calculate movement as described in 4.4.2
else
    search for a target using one of the initial
    search strategies described in 4.4.1
end if
```

Figure 31. Pseudo-code of the administrative behaviour hierarchy.

4.4.6 Giving Up Time

During any type of search, an autonomous agent should be able to decide when following a given strategy is proving sub-optimal, for example, due to diminishing returns or increased exposure to danger, and then change some element of their activity. This could be as simple as changing the time or location of the search, or indeed the search strategy itself. To examine issues such as strategy switching due to search success are beyond the scope of this work. However, some of the aspects of the research have intrinsic strategy switching that requires a mechanism that ensures that any required switching occurs in a timely manner. For the purposes of a simple experimental environment, a system based on time (see Section 4.5.2 for a description of the time metric) was developed. This consisted of a timer that was set at the start of a given agent behaviour and reset, along with the agent behaviour, once a specified time had elapsed. Aspects of agent behaviour that required such a mechanism were: the return, from intensive search, to extensive searching in ACS (the intensive search being entered after target destruction); and the switch to the initial search strategy from a reacquisition search. These 'giving up' times are tuned for best performance. Saltatory search also has a two-phase nature but is not tuned here because the strategy would only require tuning for operations in an environment that contained features where targets could hide (for example, areas of dense vegetation); for this work, the agent switches to intensive phase for sufficient time to thoroughly search the agent's vicinity. Whilst, due to the stochastic nature of the movement, complete coverage during intensive search could not be ensured, given its sensory capability which, in the simulation environment was assumed to be perfect, an

adequate period was allocated such that full coverage of the locale would be expected. This duration was selected following observations of agent behaviour and actual values used are provided in Figure 21 (page 77). In more demanding environments, where exhaustive local search may not be practical, work may be required to construct a mechanism for a giving up time that can adjust for factors such as varying degrees of vegetation density.

4.5 Metrics

Satisfactory analysis of search performance requires a variety of metrics and visualisations. For those scenarios where agents are defending the two fixed positions in the search space, a losses metric was designed to reveal the effectiveness of the defence. Whilst useful in this context, it does not allow performance in the other search scenarios, where there are no specific defended locations, to be evaluated. This led to consideration of time (to detect and/or intercept) as a suitable measure of performance. When relative performances were compared in the initial search scenario using both losses and time, it was evident that time was indicative of losses (see Figure 32). A notable feature of the comparison is the results of the spiral search, this occurs because agents turn more frequently toward the middle of the search area and so the chance of the targets being detected increases further into the scenario (agents always begin at the edges of the search area for deterministic attacks).

Current thinking is, however, that it is not just the effectiveness of a defensive system in detecting and removing intruders that makes the system effective; it is also the appearance of defending the area (as confirmed by the authors' personal experience over the course of 22 years service in the Royal Air Force). This then led to consideration of the coverage of the search space, being both essential for the appearance of a well defended space, as well as for the ability to detect targets throughout the space. The coverage metric allows quantifiable comparison of each strategy's defensive posture. Finally, given the various turn rates involved in the search patterns, it was evident that a metric was required that would illustrate whether performance gains were at the cost of energy efficiency.

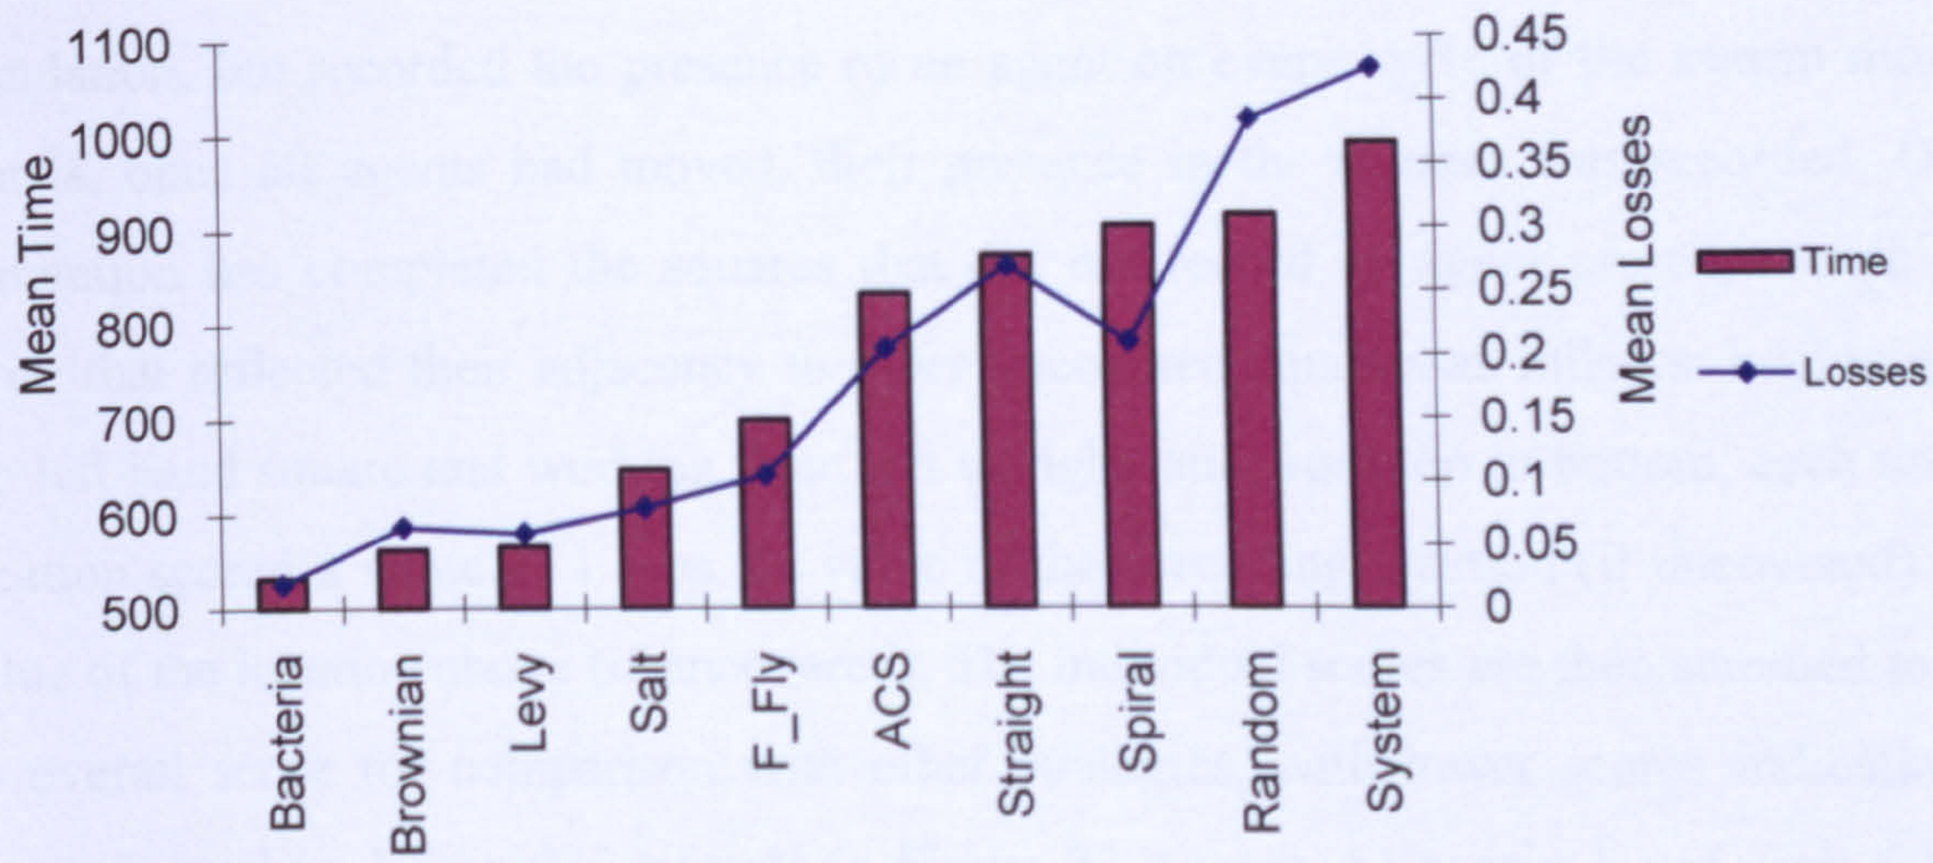

Figure 32. Comparison of mean time and losses metric for 20 agents acting independently against 4 targets travelling at maximum agent speed.

4.5.1 Losses

As described in Section 4.3.1, two non-central locations were identified as defended positions within the search space. Each time an intruder got close enough (within 2.5 units) to either of the defended positions to release a weapon, the losses metric was incremented by 1.

4.5.2 Time

The time, measured in search iterations, from the when targets appear in the search area to when the last one is located and/or destroyed or has left the area (depending on the specific scenario). Chronological time and CPU time are not suitable mechanisms, since available resources would not permit isolating software processes such that operating system background tasks cannot be guaranteed not to interfere during experimentation.

4.5.3 Coverage

The coverage metric was designed to quantify contiguous areas of uncovered space – large contiguous areas of search space are considered undesirable because of the potential for exploitation by observant intruders. To assess effective agent coverage across the search area, the number of uncovered areas is assessed with respect to their proximity to other uncovered areas. For the purposes of this metric, the search area is divided in squares of area 25 units^2

(*i.e.* the search area contained 80 x 60 such squares). These had no influence on the simulation, but recorded the presence of an agent on every cycle of the swarm movement; that is, once all agents had moved, their presence in the squares was recorded. Once the simulation has completed the squares that did not record an agent presence were given a score that reflected their adjacency to other uncovered squares as follows: beginning at the top left hand square and working from left to right, and from top to bottom, each uncovered location scored a value of 1 plus the value of the preceding location (if uncovered) plus the value of the location above (if uncovered). The individual scores are then summed to provide an overall score for comparison with other strategies, with lower scores indicating better coverage quality. Using the example in Figure 33, square A1 scores 1 and each subsequent square on that row scores 1 plus the value in the previous square since there are no adjacent squares above (so the total score for that row is 10). Square A2 scores 1 plus the value in the adjacent square above (A1) only since there is no adjacent square to its left but B2 scores 1 plus values of the adjacent squares above and to the left (B2 and A1 respectively); however, C2 scores nothing and so D2 scores only 1 + the value of D1.

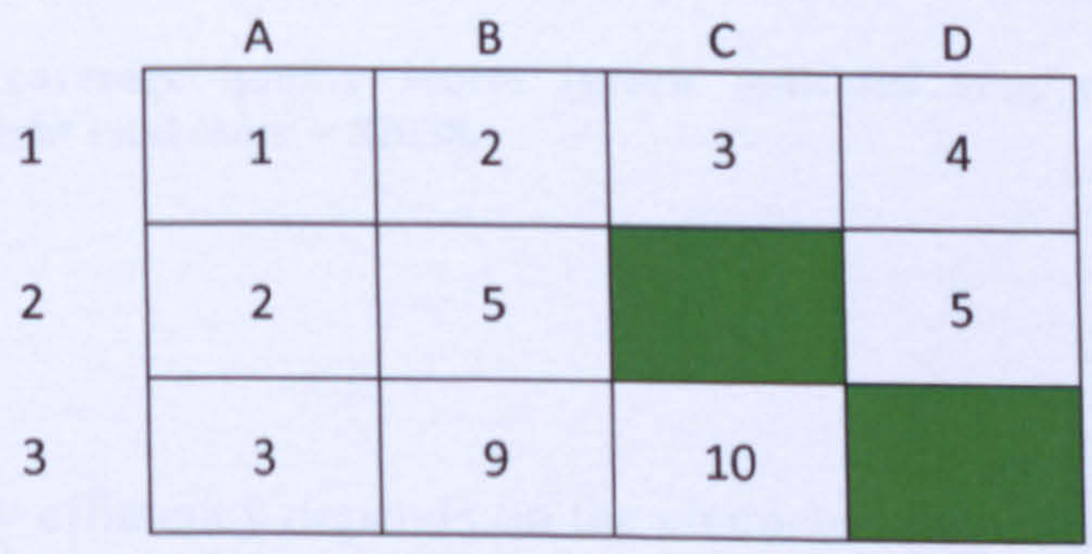

Figure 33. Coverage metric calculation example (green indicates area covered by an agent).

For a larger scale example, consider the two search areas depicted in Figure 34, both agents covered the same amount of search space (18 out of 100 squares) but the quality of coverage on the left would be considered superior because it contains less contiguous uncovered space; *i.e.* the right hand area has a very large uncovered area that would be easier to exploit. This is reflected in the coverage scores of 11,094 (left) and 82,038 (right).

1	2	3		1	2	3	4	5	6
2	5	9			3	7	12	18	25
3	9	19	20		4	12	25	44	70
4	14	34	55		5	18	44	89	160
5	20	55	111			19	64	154	315
6	27	83	195	196				155	471
7	35	119	315	512	513	514		156	628
8	44	164	480	993				157	786
9	54	219				1	2	160	947
10	65	285		1	2	4	7	168	1116

1	2	3		1	2	3	4	5	6
2	5	9						6	13
3	9	19	20	21	22	23			14
4	14	34	55	77	100	124	125		
5	20	55	111	189	290	415	541	542	
6	27	83	195	385	676	1092	1634	2177	
7	35	119	315	701	1378	2471	4106		
8	44	164	480	1182	2561	5033	9140		1
9	54	219	700	1883	4445	9479			2
10	65	285	986	2870	7316	16796		1	4

Figure 34. Two example coverage quality scores (green indicates area covered by an agent), left total score = 11094 and right total score = 82038.

4.5.4 Efficiency

The assessment of energy efficiency depends on the characteristics of the vehicles involved. Since these vary significantly from application to application, indicative values for aircraft were used as an example. Two Royal Air Force navigators were interviewed and a consensus reached that given an energy consumption of 1 unit per stepwise movement at maximum speed, 1% would be added per 1° of turn, and energy could be reduced on a *pro rata* basis with speed. Therefore, an agent travelling for 100 steps performing a 10° turn at 75% of maximum speed would consume $(100 + (100 \times 0.1)) \times 0.75 = 82.5$ units.

4.6 Statistical Analysis

Statistical tests were conducted on the results to assess their significance. Kolmogorov-Smirnov normality tests (Kolmogorov, 1933, and Smirnov, 1939) were initially used to assess the normality of the distributions of results. These concluded, unsurprisingly, that the

results were not normally distributed but were positively skewed (Figure 35 and Figure 36 show representative histograms). Such distributions arise because where agents were located near to targets on initialisation those targets were located and/or destroyed quickly, but where agents were remote, relatively long periods of time could pass before encounters occurred. This was especially true for Brownian motion targets, since they would not range across the search area, decreasing the likelihood of encounter still further. Evidence of this can be seen when comparing the histograms in Figure 37. These show two specific strategy results from the same experiment as the Brownian motion histogram (Figure 35, left) The spiral strategy, because it ranges throughout the search area enhancing the likelihood of encountering the remote and non-ranging target, clearly performs better than the bacterial search, which, due to the reduction in agent speed and high turn rate when no targets are being detected, has some instances that take twice as long to complete the search. This non-normality of results meant that results could not be compared reliably using parametric methods and, therefore, for the experiments exploring the initial search strategies, Kruskal-Wallis mean rankings (Kruskal and Wallis, 1952) were used to compare performance. Whilst these rankings are valid for non-normally distributed data, they remove the magnitude of the difference between approaches. Therefore, means and related box plots, are used where necessary to provide context for the discussion, although these must be treated with caution.

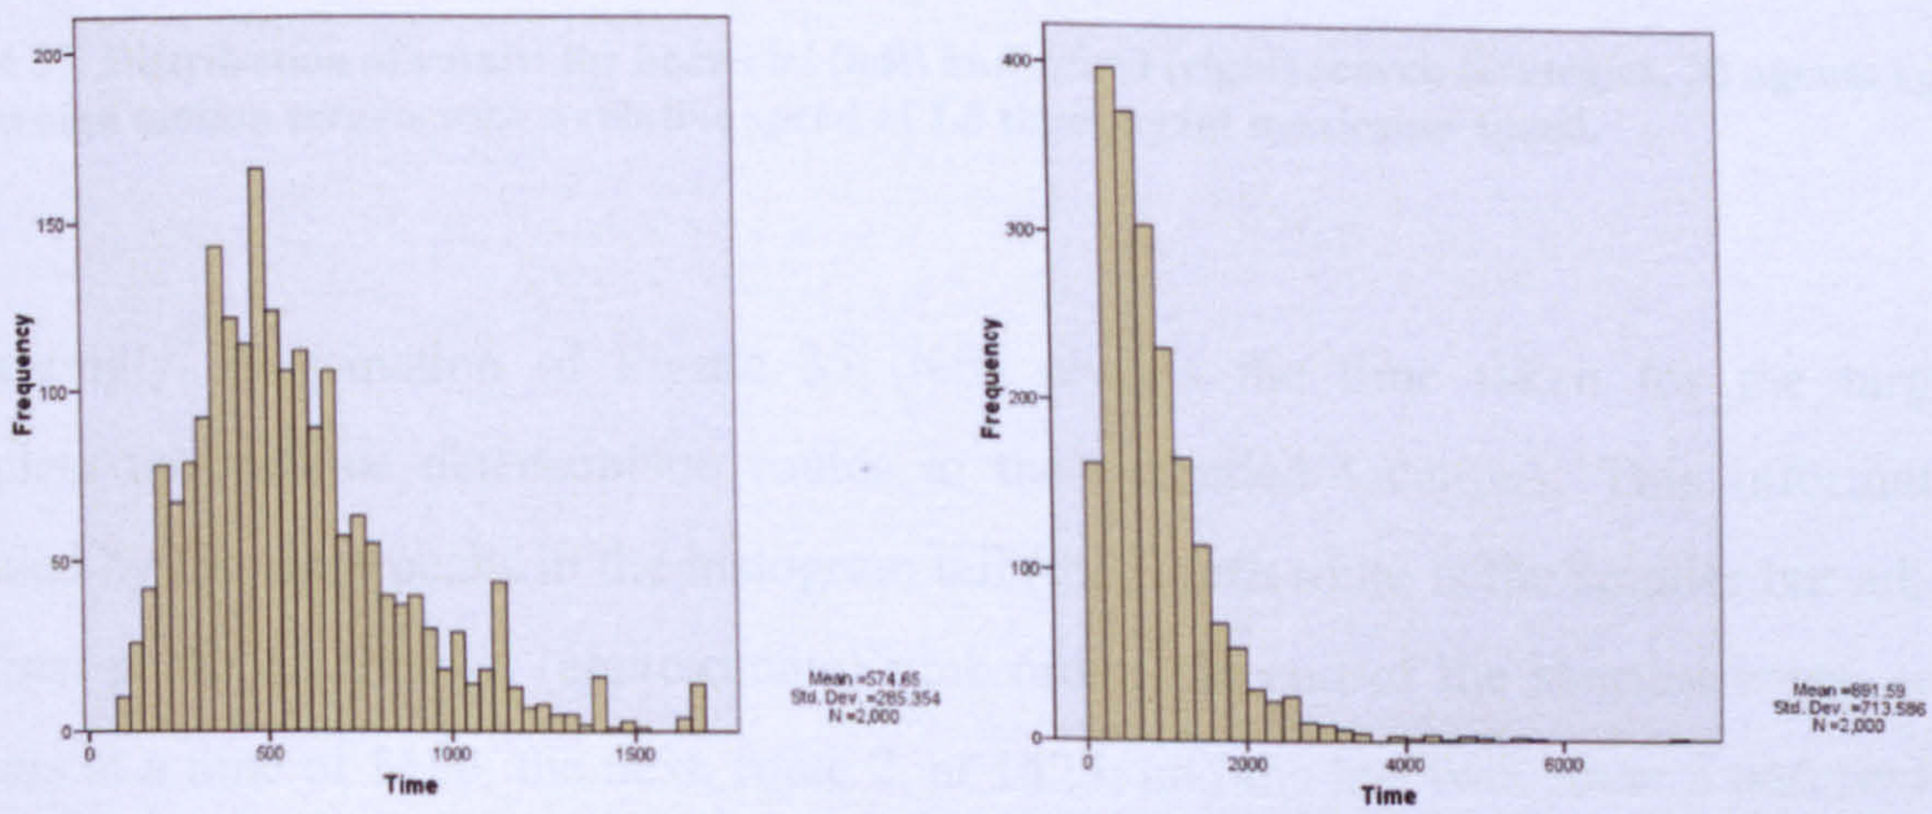

Figure 35. Distribution of results for all search strategies, 20 agents against 4 targets with a relative speed of 1.5 times agent maximum speed (left: target movement = deterministic, right: target movement = Brownian motion).

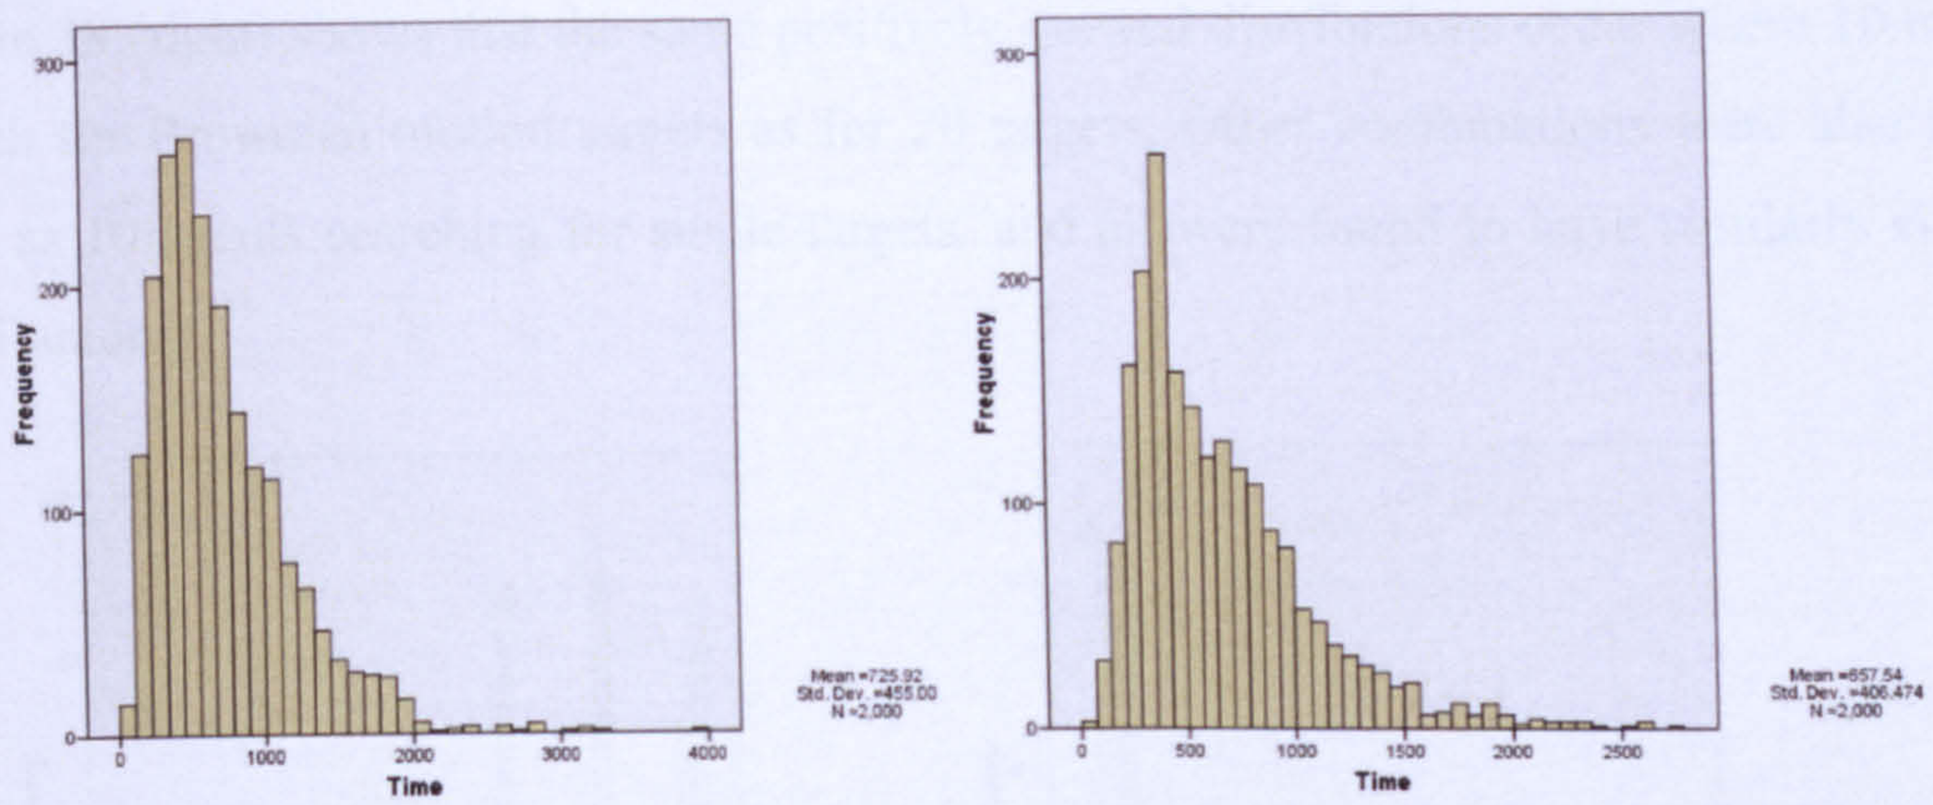

Figure 36. Distribution of results for all search strategies, 20 agents against 4 targets with a relative speed of 1.5 times agent maximum speed (left: target movement = random, right: target movement = fruit fly).

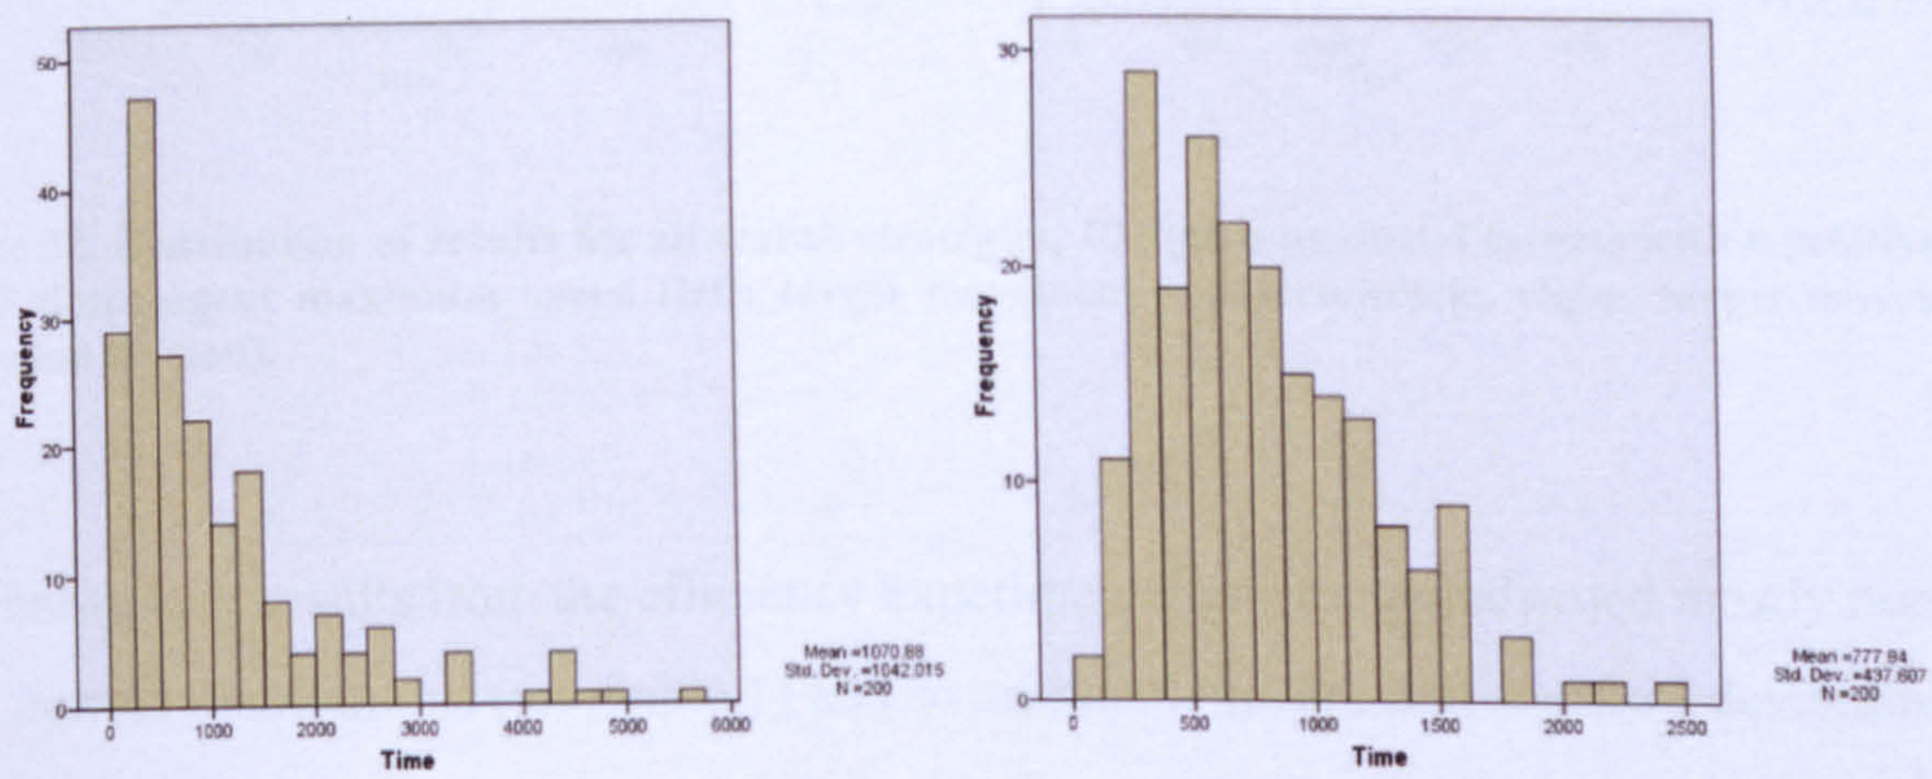

Figure 37. Distribution of results for bacterial (left) and spiral (right) search strategies, 20 agents against 4 Brownian motion targets with a relative speed of 1.5 times agent maximum speed.

Interestingly, examination of Figure 35 (left) reveals the time taken for the targets to complete the various deterministic routes to the defended locations. This information is revealed by the three peaks in the histogram tail (the fourth route is the smaller bar adjoining the final peak). In shortest (approximate) time order, the end of the shortest route, route 4, appears at a time of 1150; the next, route 2, at 1425; and the last two, route 3 and route 1, at 1650 and 1700 respectively. This phenomenon is exaggerated further where only 10 agents are present since there is more chance of the targets completing their routes (see Figure 38, left).

Figure 38 (right) shows that the same positively skewed distributions occur where 10 targets search for Brownian motion targets as for 20 targets. Other combinations were also tested such as 10 agents searching for single targets, and all were found to have similarly skewed distributions.

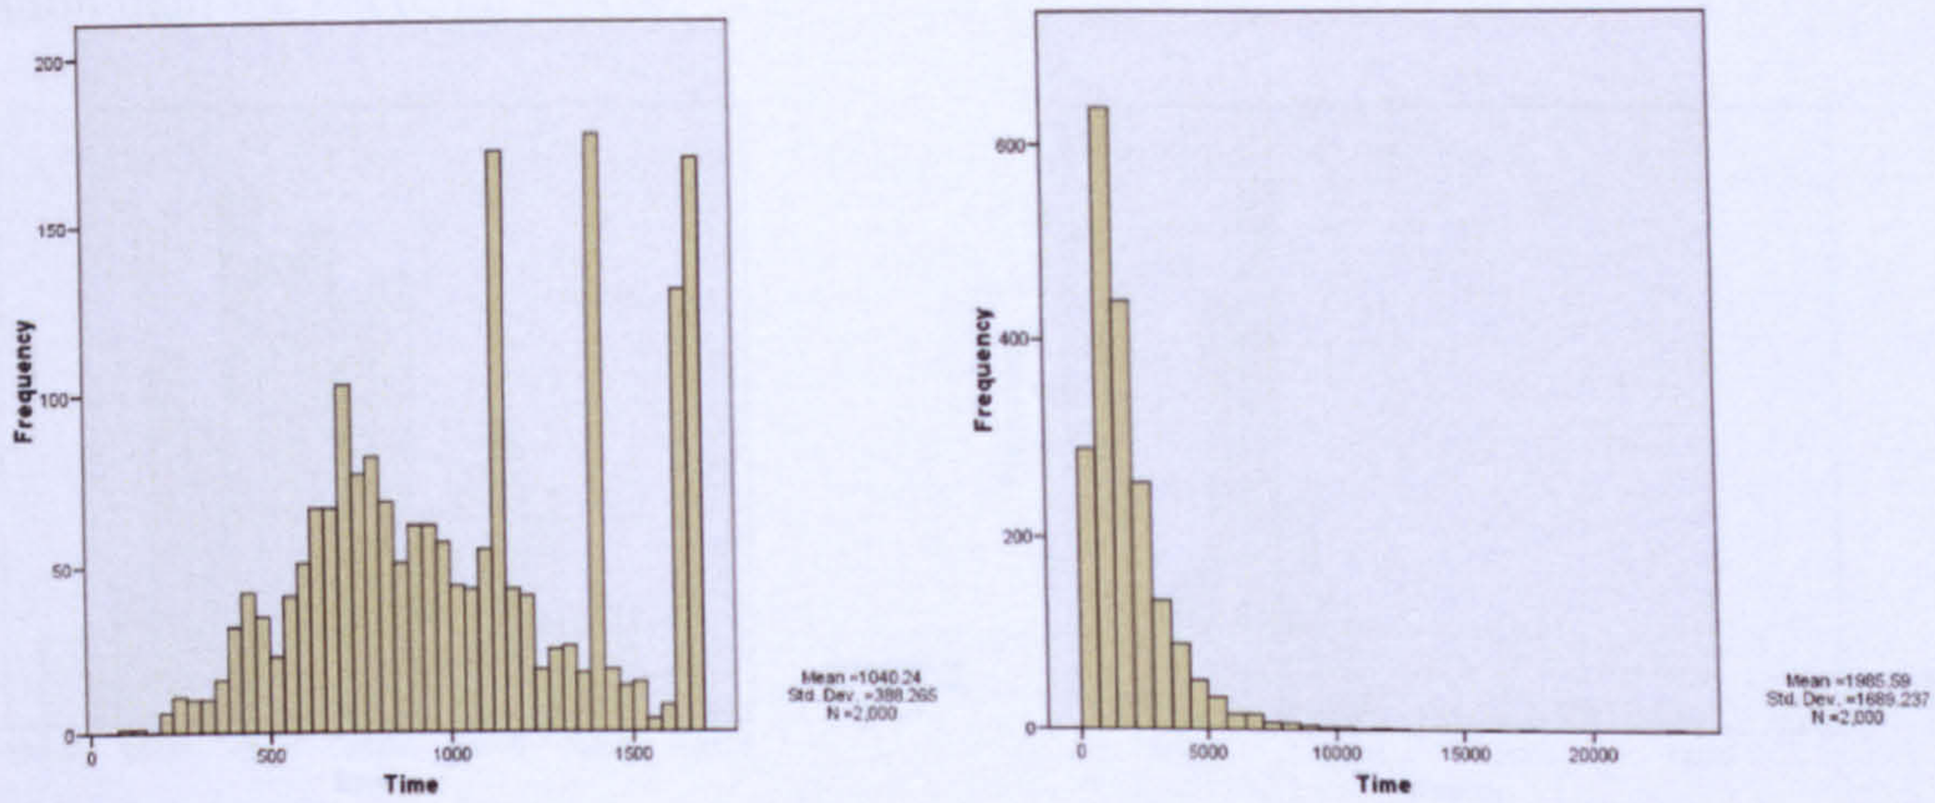

Figure 38. Distribution of results for all search strategies, 10 agents against 4 targets with a relative speed of 1.5 times agent maximum speed (left: target movement = deterministic, right: target movement = Brownian motion).

Normality tests results from the efficiency experimentation results indicated mostly normal or near normal distributions (see Table 1) and so arithmetic means and standard deviations were used for comparison.

	Search	Kolmogorov-Smirnov	
		Statistic	Significance
Energy	ACS	0.07	0.02
	Bacteria	0.25	0.00
	Brownian	0.05	0.20*
	F_Fly	0.05	0.20*
	Levy	0.06	0.05
	Random	0.06	0.05
	Salt	0.05	0.20*
	Spiral	0.09	0.00
	Straight	0.08	0.00
	System	0.05	0.20*
*This is a lower bound of the true significance.			

Table 1. Kolmogorov-Smirnov normality test results for efficiency experiment results.

The large difference observed in the test statistic for bacterial search in Table 1 is caused by the relatively high degree of variance in agent movement under no target conditions, *i.e.* the rapid switching between swimming and tumbling. This is illustrated in Figure 39 where the energy requirements of agents applying ACS has a higher mean with lower standard deviation than the bacterial search.

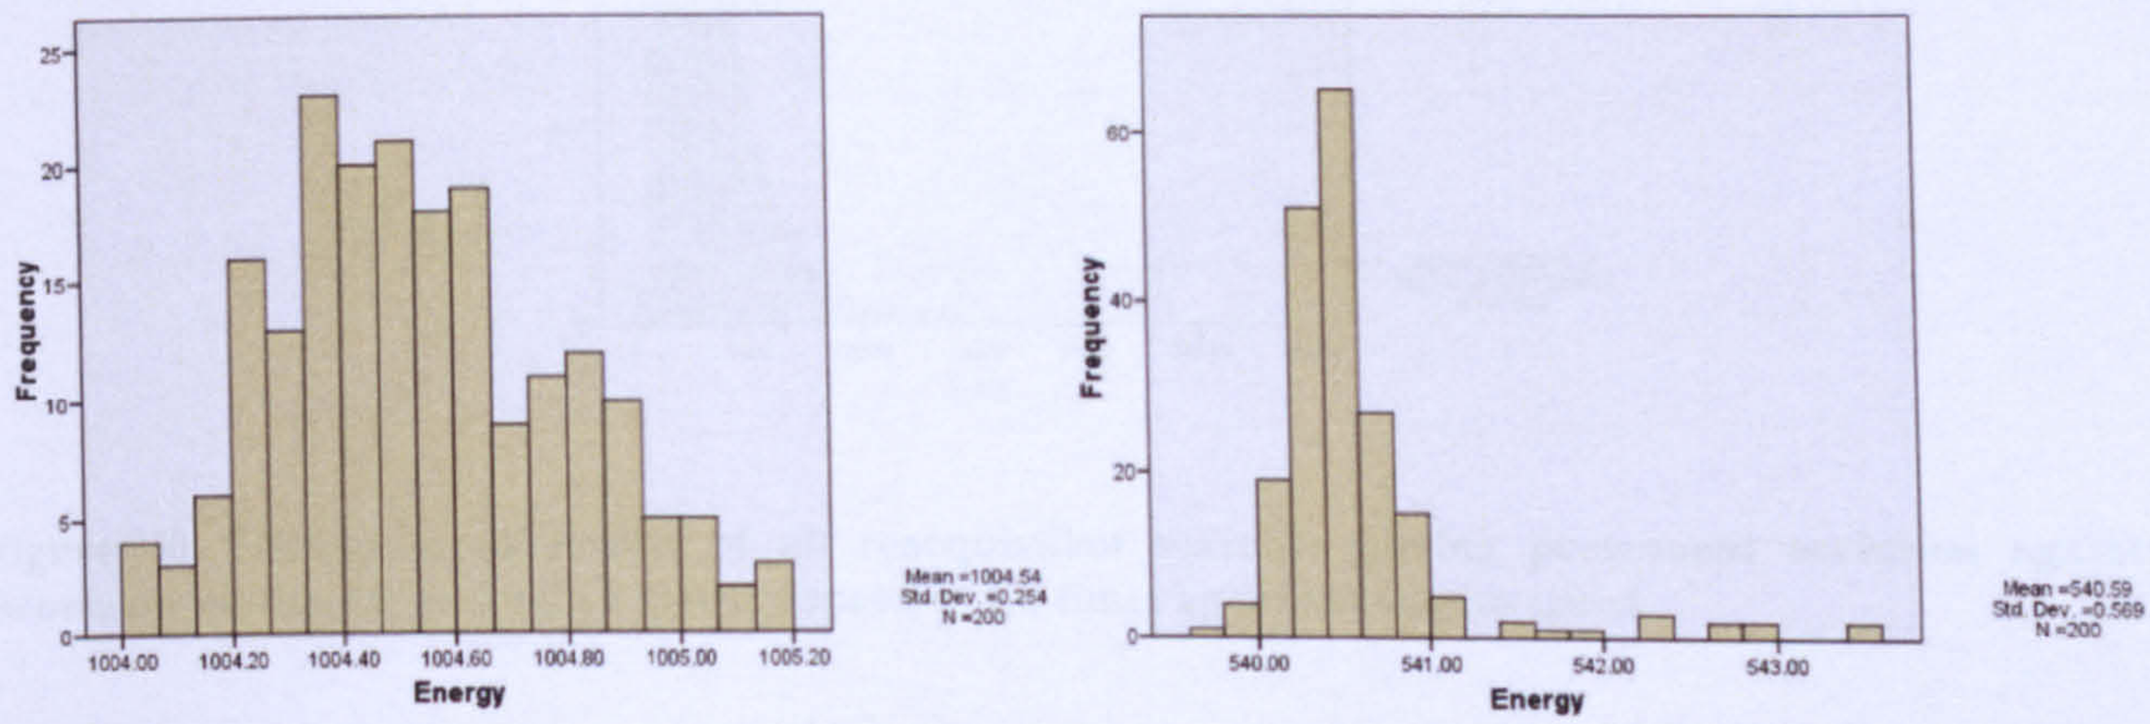

Figure 39. Frequency of agent energy requirements under no target conditions (left: ACS, right: bacterial movement).

For reacquisition searches (exploring PMS performance), arithmetic means were used despite non-normal results (Figure 40 shows indicative distribution). This was because the experimentation aimed to assess whether the various PMS could improve on system performance during occlusion or communications failures. Furthermore, because all initial search strategies were employed and results combined to provide general improvement trends, mean rankings were less meaningful since a worse performing PMS could attain some better ranking positions because of its initial search, thus obscuring results.

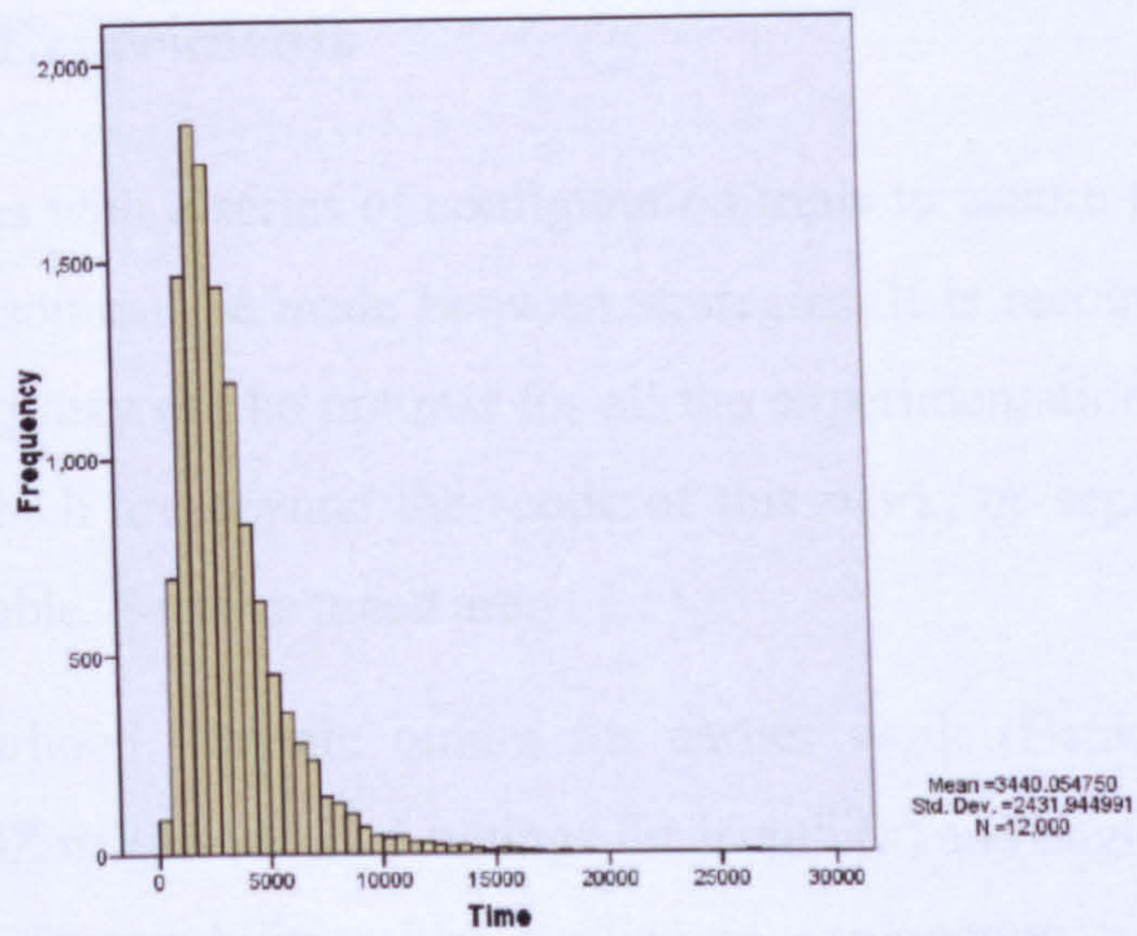

Figure 40. Histogram of results of all reacquisition searches during permanent occlusion against Brownian motion targets with a relative speed of 1.5 times agent maximum speed.

To evaluate whether results were statistically significant, Mann-Whitney U tests (Mann and Whitney, 1947) were used, as appropriate for non-normally distributed data. As evaluating the significance of each possible pair of strategies and configurations is both infeasible and unnecessary, samples from the best and worst performing strategies in each experiment are taken. For the initial search strategies, the top three were compared with the bottom three in turn (9 tests per experiment) and, for the PMS strategy experiments, comparisons between the top two and bottom two were made (4 tests per experiment). Generally, the mid-table performers did not perform significantly differently and were not of interest to the discussion. Results were deemed to be significant where there was less than 5% probability of the null hypothesis, *i.e.* that the results were drawn from the same underlying distribution. To reduce the combinatorial explosion still further, only one relative speed was selected for the dynamic target results and a mid-range number of failing agents from the communications failure scenarios. In the former case, the relative speed selected was decided by the point on the results table where the difference was largest since this was the main aspect used for discussion; in the latter case the scenario where 8 agents failed was used, again representing the mid-point in experimentation.

4.7 Schedule of Experiments

Experimentation begins with a series of configuration trials to ensure that settings are tuned such that fair comparison can be made between strategies. It is recognised that the settings provided by the tuning may not be optimal for all the experimentation but without adaptive learning strategies, which are beyond the scope of this work, or separate tuning for every scenario, this is inevitable. Settings tuned are:

- LSZ neighbourhood. Despite tuning for earlier work (Banks 2005, Banks *et al.* 2008b), the LSZ neighbourhood settings for length (c) and angle (θ) are re-assessed to ensure continued suitability when applied in conjunction with the natural search strategies.
- PMS neighbourhood. Settings for both circular and V shaped neighbourhoods are examined. In the former case the circle radius (c) is tuned and, for the latter, it is the length (c) and angle (θ).
- ACS GUT. Agent ACS Giving Up Time is adjusted to ensure that the periods of time spent in intensive search following the destruction of a target are suitable for the environment and any likely clustering.
- PMS GUT. Too long spent trying to re-acquire a target that has either moved away from the area or been destroyed by another agent could be a waste of resources. A range of times were therefore tested using indicative scenarios.

Following tuning, experiments to provide evidence of initial search strategy performance against static and dynamic targets are carried out as follows:

- Static, locate-only. Where agents are not required to destroy targets they can mark the location and continue searching without disruption to the search pattern. For this series of experiments, targets were randomly initialised in the search area (with both uniform and clustered distributions) and removed from the search as soon as agents detected them.
- Static, locate-and-destroy. Where target destruction is required, it would be unlikely that an agent would have the ability to do so as soon as the target was located, *i.e.* sensor range can be much greater than effective weapon range. So, to assess

performance where, following target location, agents have to get into weapon range to engage, the locate-only experiments are repeated except agents are required to get much closer to the target in order to destroy them, possibly disrupting their search path to do so.

- Dynamic, deterministic targets. Whilst static target searching can be a non-trivial task, searching for dynamic targets represents a much more difficult search problem. In the first series of dynamic target search experiments, targets (operating at a range of speeds) follow pre-determined paths in an attempt to attack locations defended by the agents who locate-and-destroy targets[1] at all times. Agent performance is compared using both the losses and time metric.

- Dynamic, non-deterministic targets. To compare search performance against a range of target movement types, targets (again operating at a range of speeds) are randomly initialised in the search area and then moved non-deterministically using behaviours ranging from high turn rate Brownian motion, through fruit fly movement to the more ranging random movement. In these scenarios the targets do not attempt to destroy specific defended locations and therefore only the time metric is used for comparison.

- Effect of swarming. This work builds on previous work that introduced particle swarm guidance for autonomous vehicles; the influence of the existing swarming mechanism on the performance of the search strategies is assessed, in a representative scenario, against each of the target types used in the dynamic target experiments.

Two further elements of the initial search patterns are then investigated: coverage and efficiency. These experiments are not intended to be exhaustive but to allow general insight into two major factors that could influence their selection for implementation.

- Coverage. Oftentimes, when defending a given area, the perception of strong defence is as important as the actual defensive capability. To assess how effectively the swarm of agents move about the search area, experiments are conducted where a number of agents are permitted to move about the search area for fixed periods of time.

---

[1] Note: Agents are required to locate-and-destroy all dynamic targets and so this element is not further included in experiment settings although it always applies.

Strategies are then compared using the coverage metric, which quantifies the contiguous areas of uncovered search space. Following these experiments, a qualitative view is taken by plotting indicative sensor coverage for each of the strategies, with the aim of further understanding performance, particularly during the search for static targets.

- Efficiency. A given search strategy may be more effective but incur such large energy requirements make it less attractive in practice. To examine whether any of the search behaviours perform better at the cost of additional energy requirements, a fixed size swarm is initialised and then agents are allowed to move for a fixed period in the absence of any targets. This target free scenario is designed to reflect the mode that searching swarms would be expected to spend the majority of the time in, looking for targets. The energy requirements are then recorded and compared.

Finally, experiments are conducted to assess whether the performance of the agents applying the initial search strategies can be further enhanced in scenarios where real world limitations occur on sensory capability and communications. PMS strategies are applied in an attempt to reacquire targets following loss of detection or communication. Two basic scenarios are examined (note: in the second series, only the best performing PMS is compared to ensure results held from the occlusion experiments):

- Occlusion. Performance of PMS strategies in scenarios where a target becomes occluded from the detecting agent either temporarily or permanently, whilst other agents can still detect the target.

- Communications failures. To assess performance during communications failure, the best performing strategy from the occlusion experiments was tested using two types of failure. First, agents may lose their ability to communicate their knowledge permanently, for example, following equipment failure. In the second scenario, agents lose communication temporarily; simulating situations where, for example, environmental features interfere or high levels of electro-magnetic noise interfere intermittently. All communications failures simultaneously affect both transmission and reception channels although agents may fail at different times.

5 Research Findings

This section begins, in Section 5.2, with the findings of the tuning experiments prior to presentation and discussion of main research findings[2] that commence, in Section 5.3.1, with strategy performance using the time metric to locate static targets. This is followed in Section 5.3.2 by the same experiments repeated with agents having to locate-and-destroy static targets. Section 5.4 investigates performance against dynamic targets, which is subdivided between deterministic and non-deterministic targets (actual results are presented in Appendix 2). A brief foray into the effect of swarming on the results is included in Section 5.5. Section 5.6 reports on coverage experiments and Section 5.7 reviews the energy requirements of the strategies to assess whether the more successful behaviours are higher energy consumers. Finally, Sections 5.8 to 5.10 present the results of PMS experimentation for scenarios where targets become occluded and agents experience communication system failures.

5.1 General Considerations

Research findings for all experiments begin with a more detailed description of experiment configurations than provided in the general Schedule in Section 4.7, followed by discussion of the results.

Experiment results are then provided graphically[3], due to the large quantity of data, with all graphs of dynamic target results shown as 2nd order polynomial trend lines to highlight result trends. Trend lines were used because noise in the actual plots (provided in Appendix 2) made it difficult to observe trends in search performance; a comparison between Figure 41 and Figure 42 indicates that such an approach is reasonable.

All experiments were repeated 200 times per search strategy (for example, in dynamic target experiments each experiment was repeated 200 times for each of the 10 search strategies over 20 relative target speeds giving 40,000 simulation runs for each experiment configuration). Although sample sizes are necessarily limited by computation effort, to ensure that the

---

[2] The statistical analysis methods used in the research findings are discussed in Section 4.6.

[3] Graphs are presented with different scales so that differences in strategy performance for a particular scenario are highlighted and, therefore, care should be exercised when comparing results between different experiments.

chosen sample size was sufficient for reliable comparison, representative experiments using a fixed configuration were also performed using larger sample sizes of 2000 runs per strategy (400,000 runs in total). This experiment used 10 agents, applying all initial search strategies, against a single target moving deterministically following route 1 (see page 64) at relative speeds of {0.1, 0.2, ... , 1.9, 2.0}. Swarming was enabled with LSZ neighbourhood. Results are reported for 2,000 and 200 runs per strategy for each relative target speed.

The results of the extended runs are shown in Figure 41 with the 2nd order polynomial trend lines shown in Figure 42, which can be compared with the same experimental set up, over 200 runs, results shown in Figure 43. Small differences can be observed, but these graphs indicate that a sample size of 200 is a sufficient basis on which to make comparative assessments of the various experiment results, and all results are reported for that number of runs.

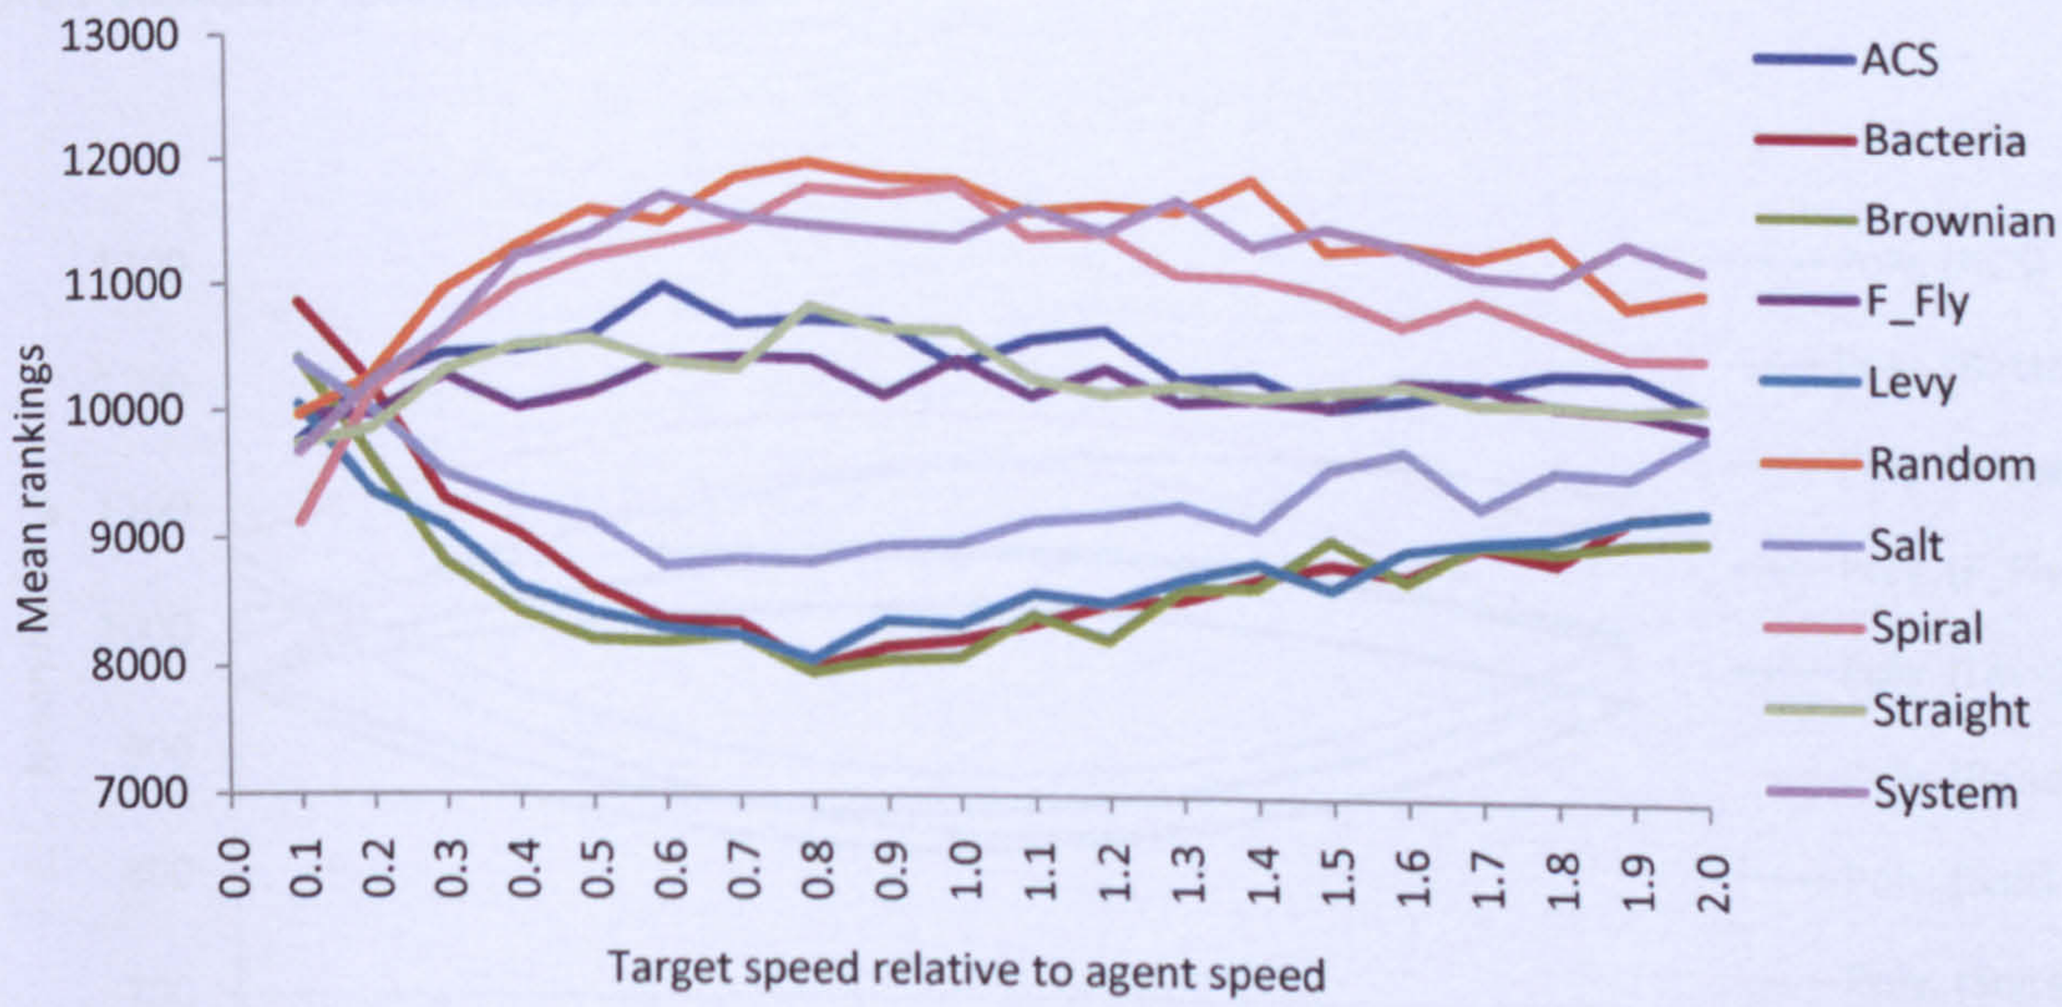

Figure 41. Actual mean rankings of 10 agents against 1 target moving deterministically (2000 runs per trial).

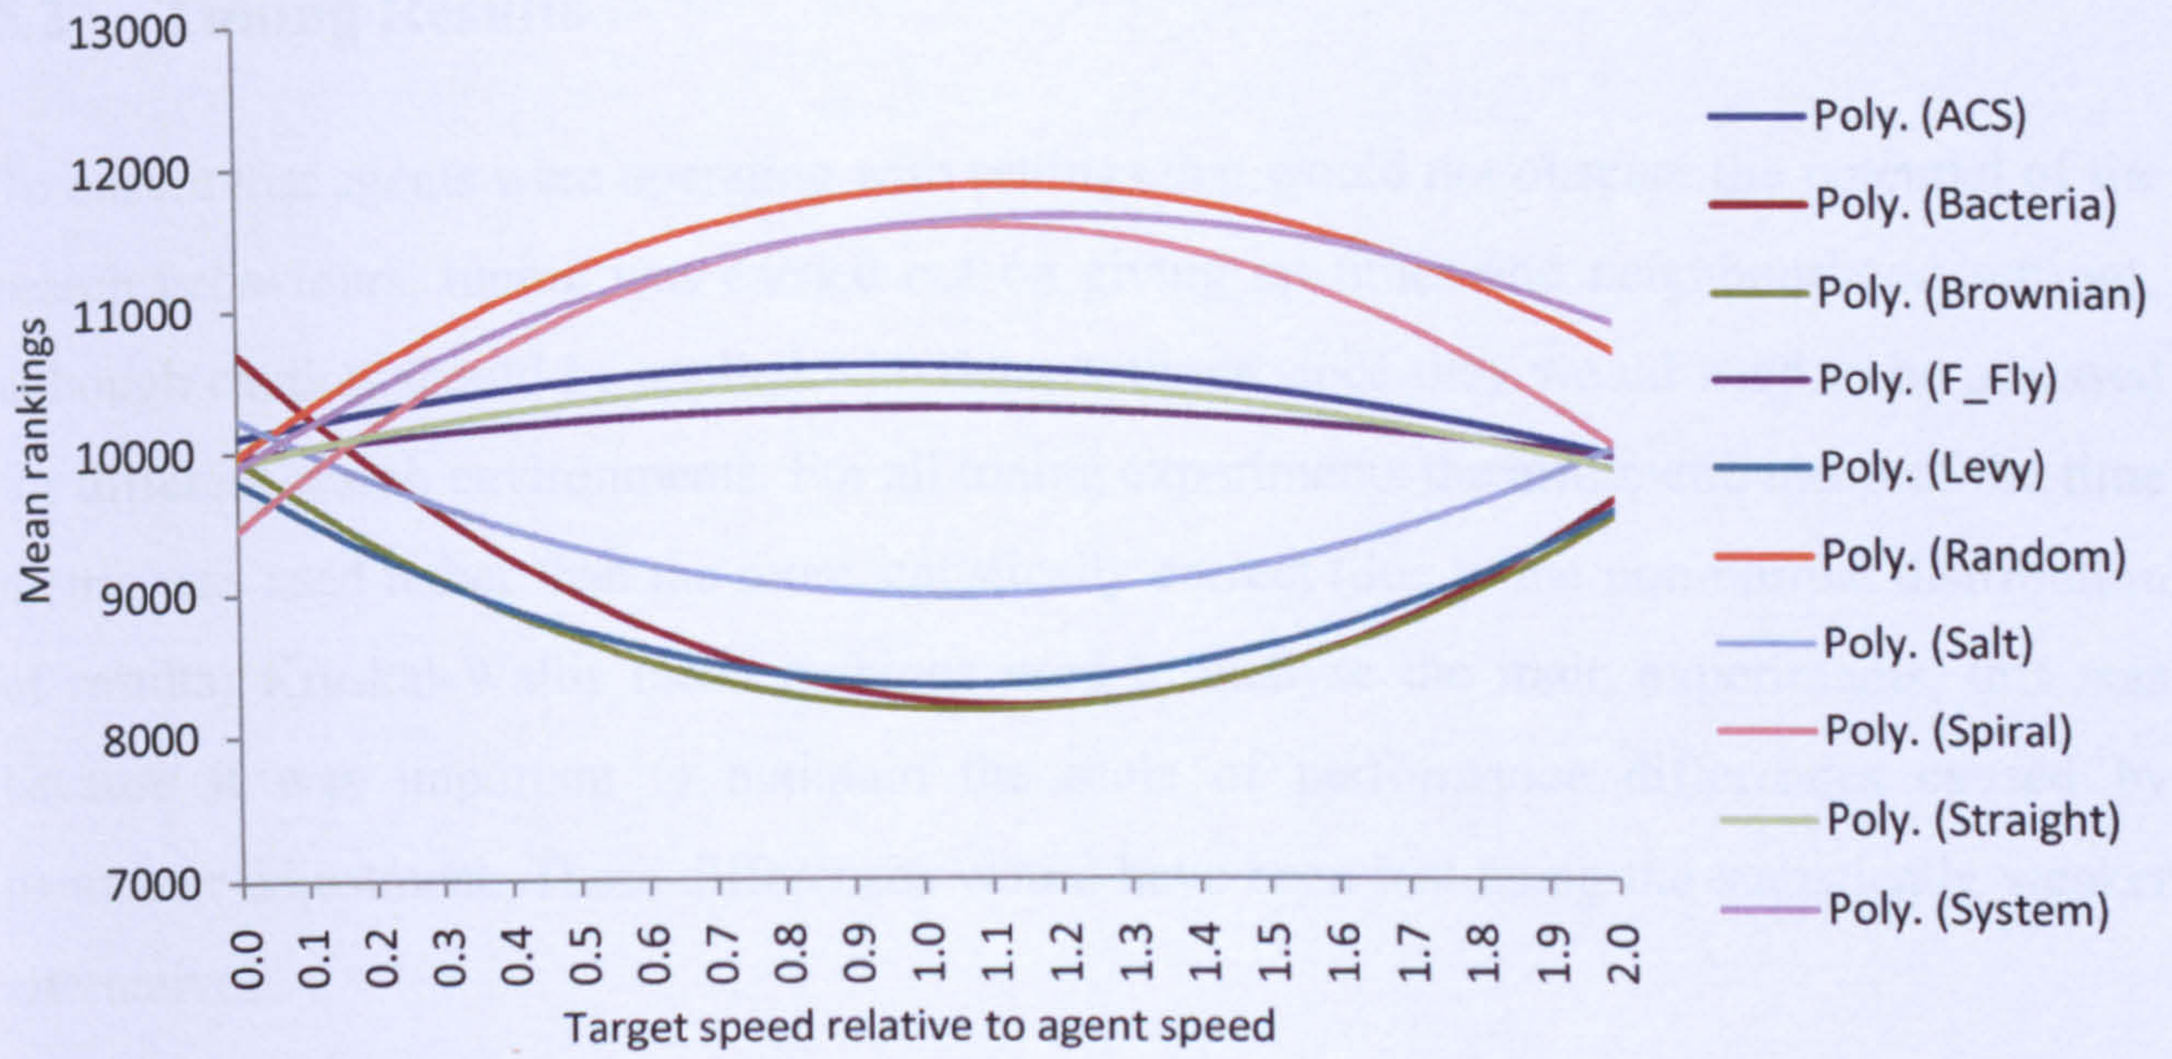

Figure 42. 2nd order polynomial trend lines of mean rankings of 10 agents against 1 target moving deterministically (2000 runs per trial).

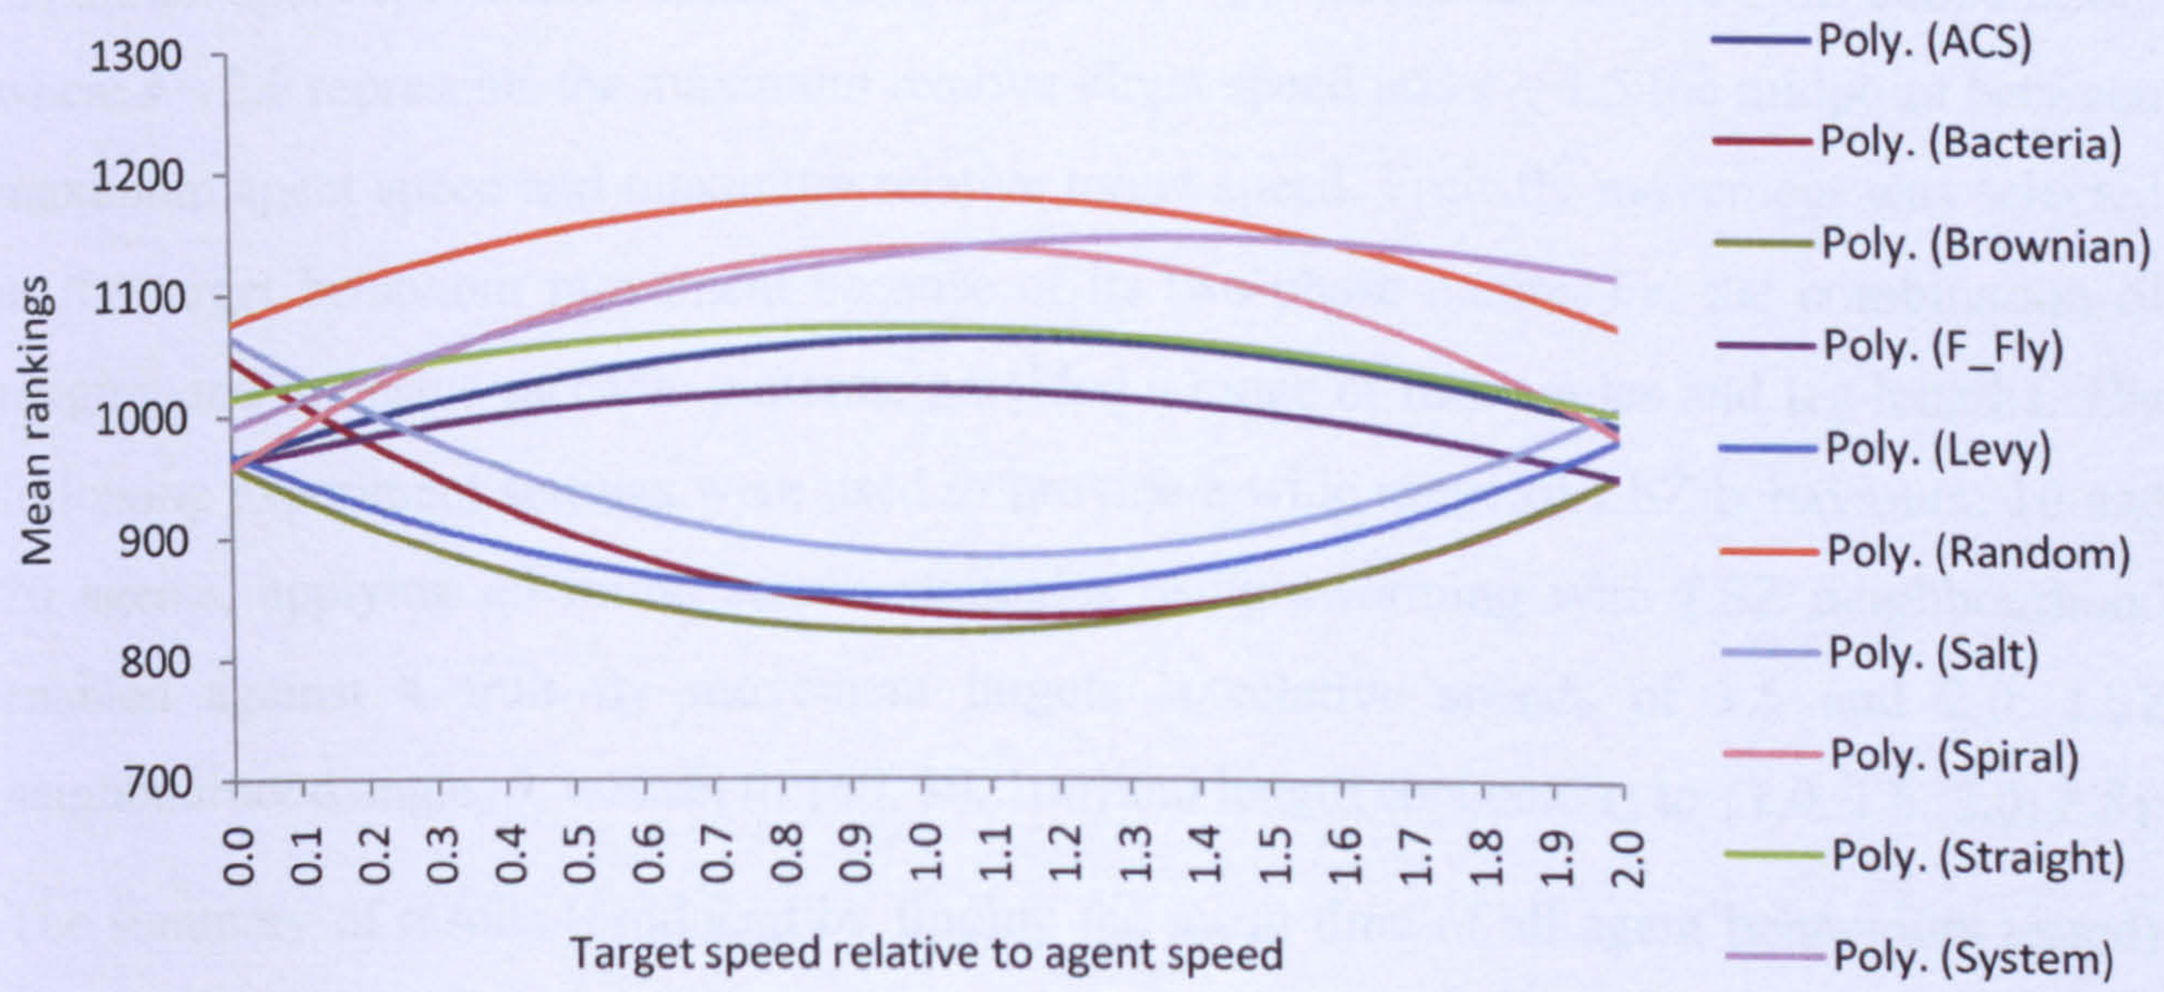

Figure 43. 2nd order polynomial trend lines of mean rankings of 10 agents against 1 target moving deterministically (200 runs per trial).

5.2 Tuning Results

To ensure that agents were operating with settings that would not obscure the potential of the search behaviours, tuning was carried out on giving up times and neighbourhood settings, although caution should be applied with these settings since they would need to be assessed for different search environments. For all tuning experiments the arithmetic mean of the time metric was used rather than the more statistically correct (due to the non-normal distribution of results) Kruskal-Wallis mean rankings used to analyse the main experiments, this was because it was important to maintain the scale of performance differences caused by parameter adjustment. These differences would have been lost using the statistically weaker alternatives.

5.2.1 LSZ Neighbourhood

Experimentation was performed to optimise the LSZ neighbourhood settings θ and c (see Equation 16, page 65). Two values of s (Equation 16) were selected, both faster than maximum agent speed since speeds of equality or below would not require LSZ cooperation, where $s = 2.0$ represents the maximum relative target speed and $s = 1.5$ the midpoint between maximum agent speed and maximum relative target speed. Fruit fly movement was selected as the target behaviour movement because of its two-phase nature, *i.e.* the combination of ranging and intensive saccadic patterns, provided a range of turn angles and leg lengths. The following experiment settings were used to provide a wide range of LSZ behaviours: 10 and 20 agents, applying all initial search strategies using swarming with LSZ neighbourhood enabled against 4 fruit fly movement targets at relative speeds of 1.5 and 2.0; LSZ neighbourhood angle, θ, was set to {60, 80, 100}and length constant, c, to {1.0, 1.5, 2.0, 2.5}.

The summary of results (produced by finding the mean time of all agent behaviours tested), shown in Figure 44, indicates that whilst the setting of $\theta = 60°$ and $c = 2.0$ was competitive for large swarms it was not for smaller groups; this is because with smaller groups there is insufficient opportunity for cooperation. This contrasted with $\theta = 100°$, where performance was relatively better with smaller swarms; in this case, where the swarm was larger, too much cooperation occurred, drawing in too many agents and leaving the search area exposed to multiple targets. Notably, all experiments with short LSZs were poor performers regardless of θ and swarm size, due to the ranging aspect of target movement often placing the target beyond the LSZ within a short period. The optimal balance was found to occur at $\theta = 80°$ and

$c = 2.0$ and, therefore, for all further experiments an agent is considered a member of an LSZ neighbourhood if it is ±40° from the target and the target distance is less than the sensor range multiplied by twice the ratio of target to agent speed.

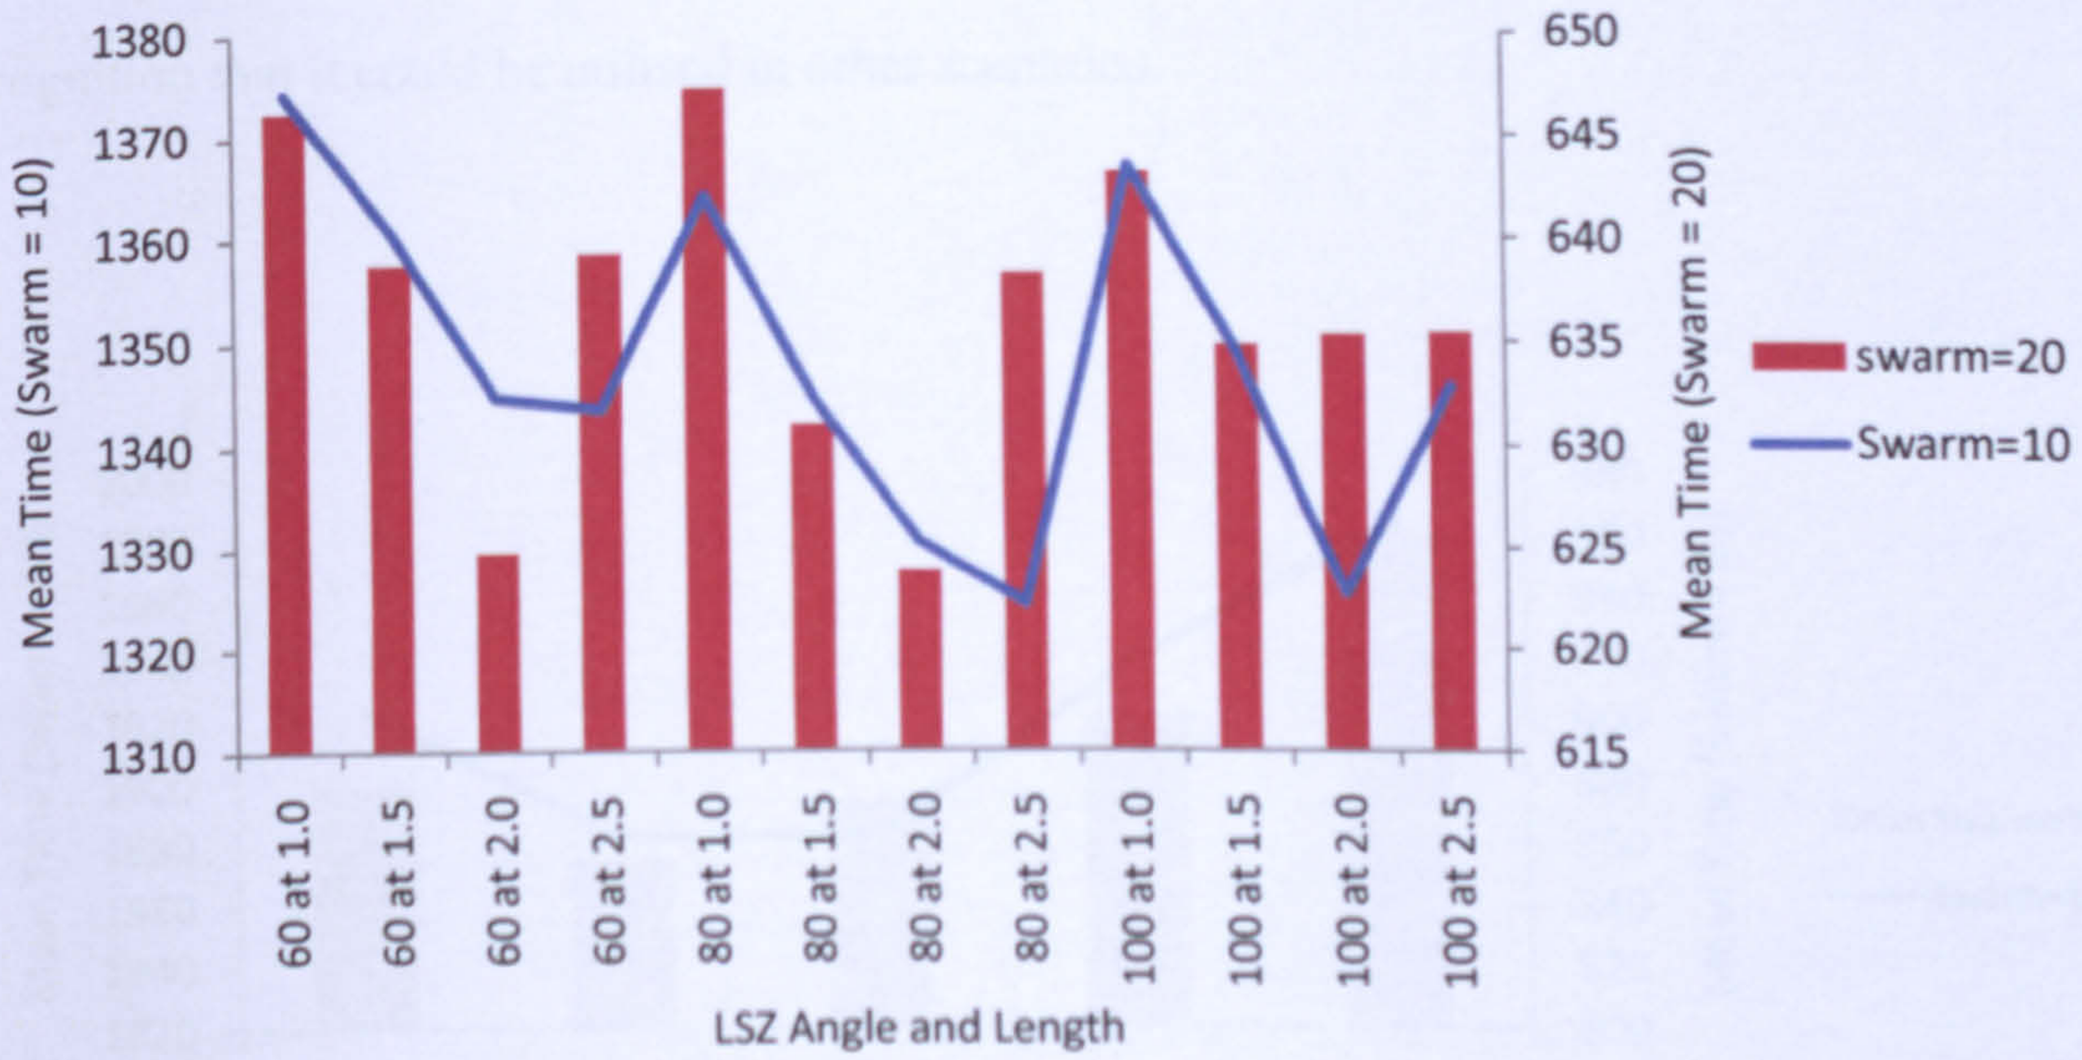

Figure 44. Summary of LSZ tuning results.

5.2.2 PMS Neighbourhood

As for LSZ neighbourhood settings, the PMS neighbourhood may vary between applications and therefore require investigation on a case by case basis. For the purposes of this work, experimentation was conducted to find reasonable sizes for circular and V shaped neighbourhoods (the same settings were applied to all PMS strategies employing each neighbourhood). The saccadic PMS does not have a neighbourhood shape and therefore does not require tuning. In the experiments, 10 and 20 agents applied all of the initial search strategies with LSZ neighbourhood swarming enabled against fruit fly movement targets at relative speeds of 1.5 and 2.0. The switch from initial search to the application of the PMS (predictive-C and predictive-V were used) was triggered by permanent occlusion. PMS neighbourhood settings tested were angle, θ, of {60, 80, 100, 120} (V shaped neighbourhood only) and length, c, of {0.5, 1.0, 1.5, 2.0, 2.5, 3.0 - where required}. The giving up time was

fixed (heuristically) at 50. Later tuning of this parameter found this value to be reasonable for both this and V shaped neighbourhoods against fruit fly targets (see Section 5.2.4).

The results shown in Figure 45 indicate that the constant c multiplier was not necessary for this implementation of the circular neighbouhood (the best performing setting of c being 1.0); however, c was retained in the implementation for consistency across neighbourhoods and in recognition that it could be utilised in other scenarios.

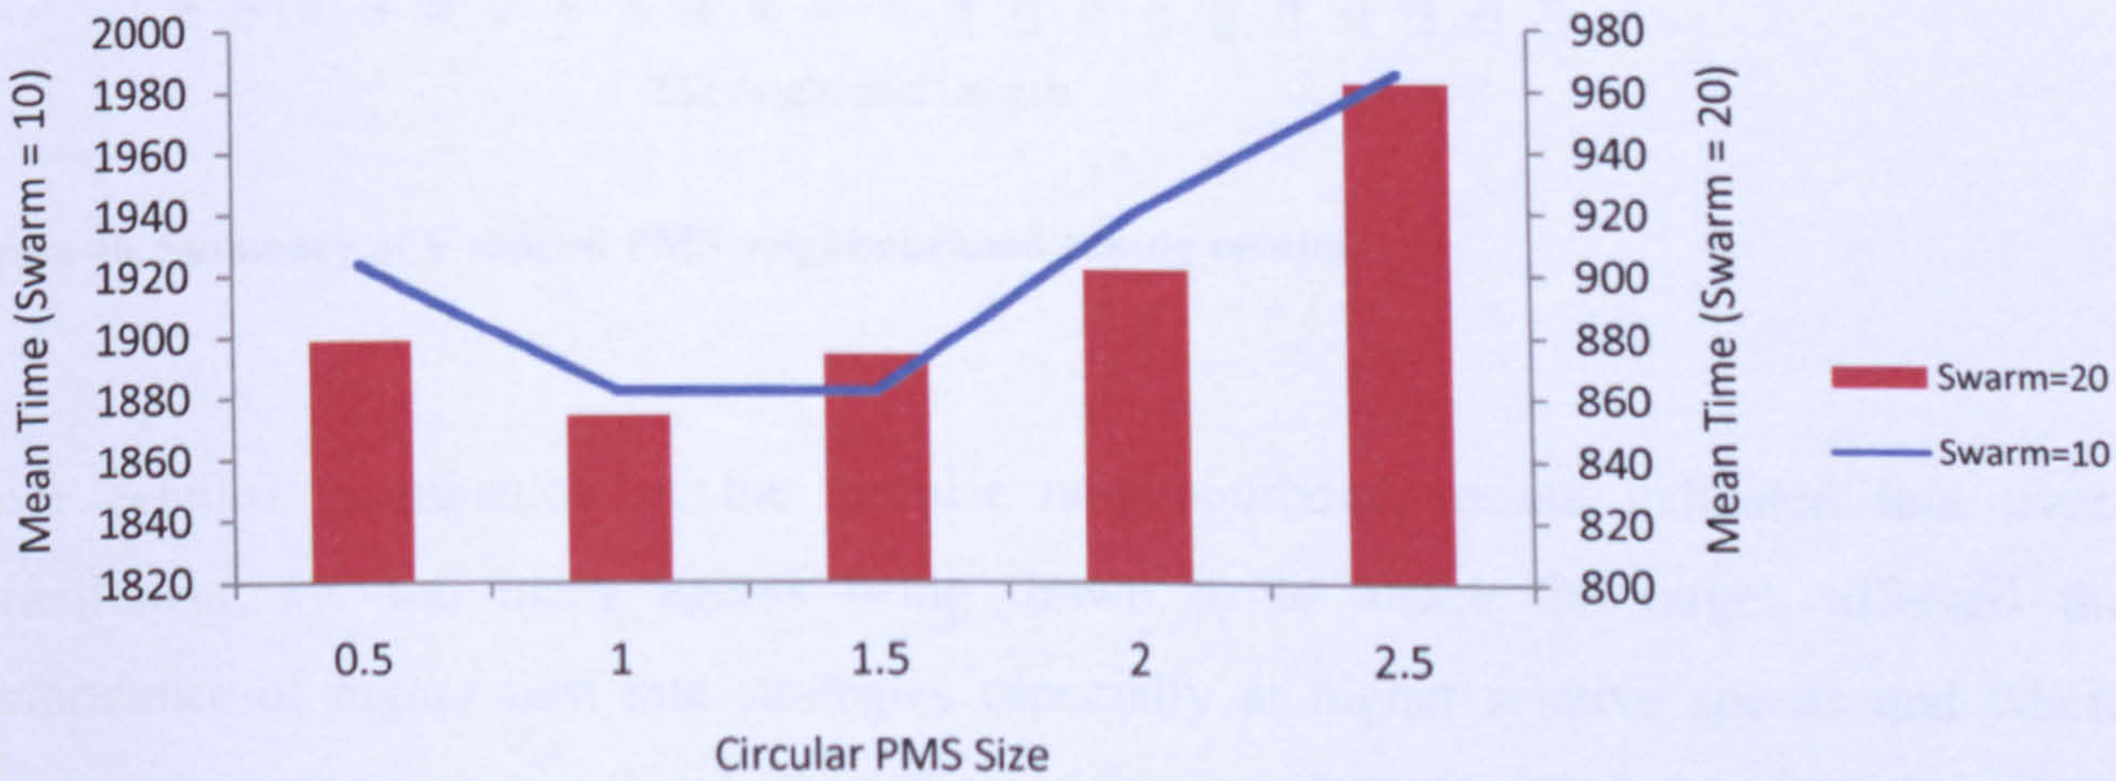

Figure 45. Summary of circular PMS neighbourhood tuning results.

The tuning results for V shaped neighbourhoods, presented in Figure 46, indicate that angles of $\theta = 80°$ provide large improvements in performance, although these were less notable at small and large values of c. The unusual reduction in performance where $\theta = 60° / c = 2.0$ and $\theta = 120° / c = 1.5$ were due to the two phase nature of the targets, with the settings on either side being suited to the ranging, inter-saccadic, movements and the high turn rate saccade phases respectively. Settings of $\theta = 80°$ and $c = 2.5$ were selected for use in all further experiments.

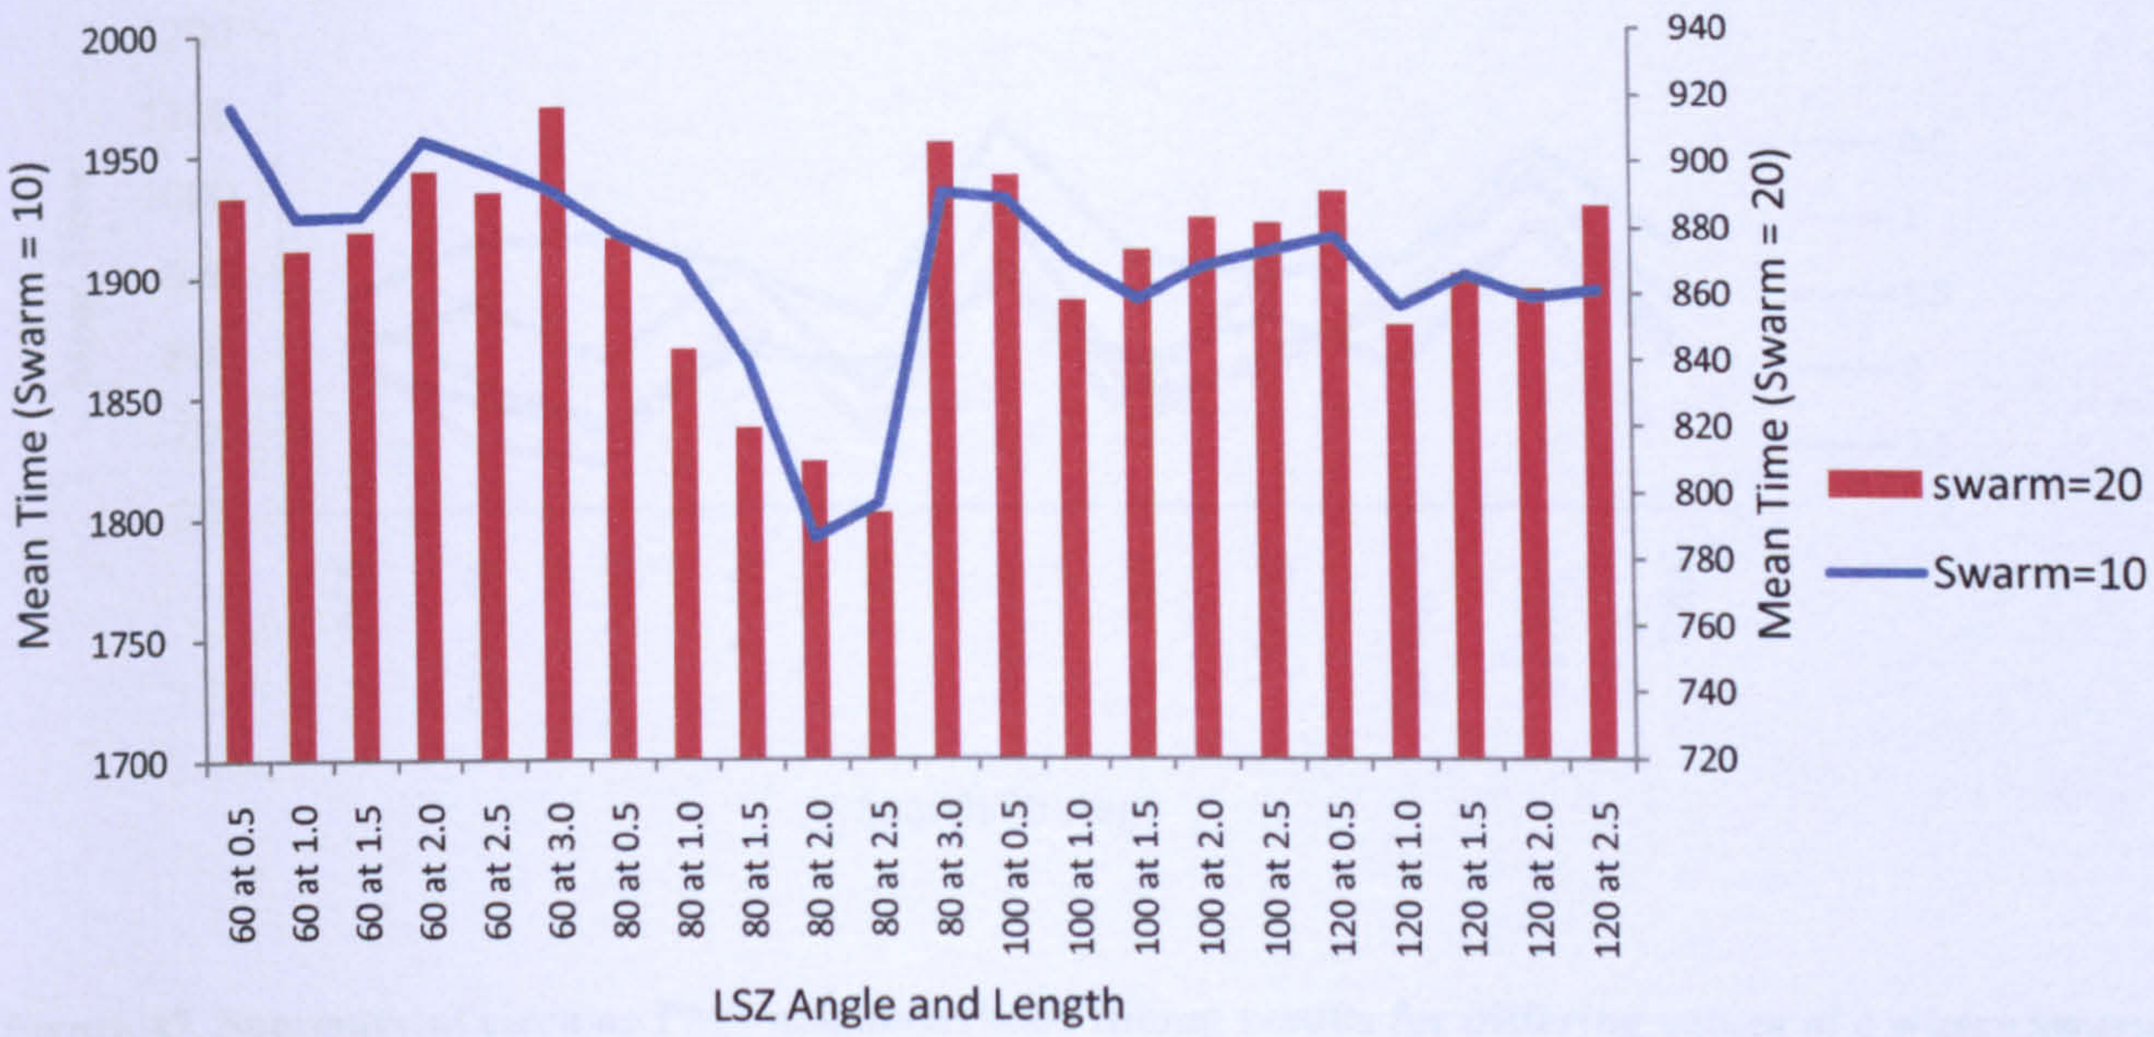

Figure 46. Summary of V shaped PMS neighbourhood tuning results.

More detailed examination of the circular neighbourhood results indicated that over-commitment, *i.e.* too many agents being drawn in to attack the target, affected the performance of higher turn rate strategies especially at higher relative speeds and where larger swarms are operating more than those that ranged farther. For example, Figure 47 shows that the difference in mean times between $c = 2.0$ and 2.5 for high turn rate strategies such as bacteria and saltatory are much larger than those for ranging types such as fruit fly and straight-line searches. This effect was caused by the large numbers of agents being drawn toward the target taking a relatively long time to re-disperse across the search area to locate any remaining targets.

Figure 47. Summary of circular PMS neighbourhood tuning results for differing values of *c* where swarm size = 20 and *s* = 2.0.

5.2.3 ACS Giving Up Time

ACS GUT in the natural world can depend on a wide variety of factors such as risk of predation and prey patchiness, and in real implementations such factors should play a part in the development of autonomous search strategies. In the non-threatening simulation environment the former does not require consideration although, given that, in the experimental environment, the goal is to complete the search and destroy task in the shortest time, it is important to consider the ACS GUT, since fruitless intensive search is inefficient and, conversely, insufficient intensive search in a patchy environment can also be detrimental to performance due to the missed opportunity of locating a nearby target.

To ensure good performance against both uniform and clustered targets, experimentation was performed with the goal of providing a balanced GUT that would not be too short for clustered targets nor too long for uniform targets. Further experiments against indicative dynamic (fruit fly) targets were performed to ensure that the chosen settings did not disadvantage the strategy when applied in a dynamic setting. For static target experiments, 5 and 10 agents searched for 20, uniformly distributed and clustered, targets whilst for dynamic targets 10 and 20 agents were pitted against fruit fly targets moving at relative speeds of {0.5, 1.0, 1.5}. For both static and dynamic experiments, the ACS strategy was applied with LSZ neighbourhood swarming enabled and giving up times of {100, 150, 200, 250, 300, 350, 400} were used.

Results of ACS GUT tuning, compared using the mean time, are indicated at Figure 48.

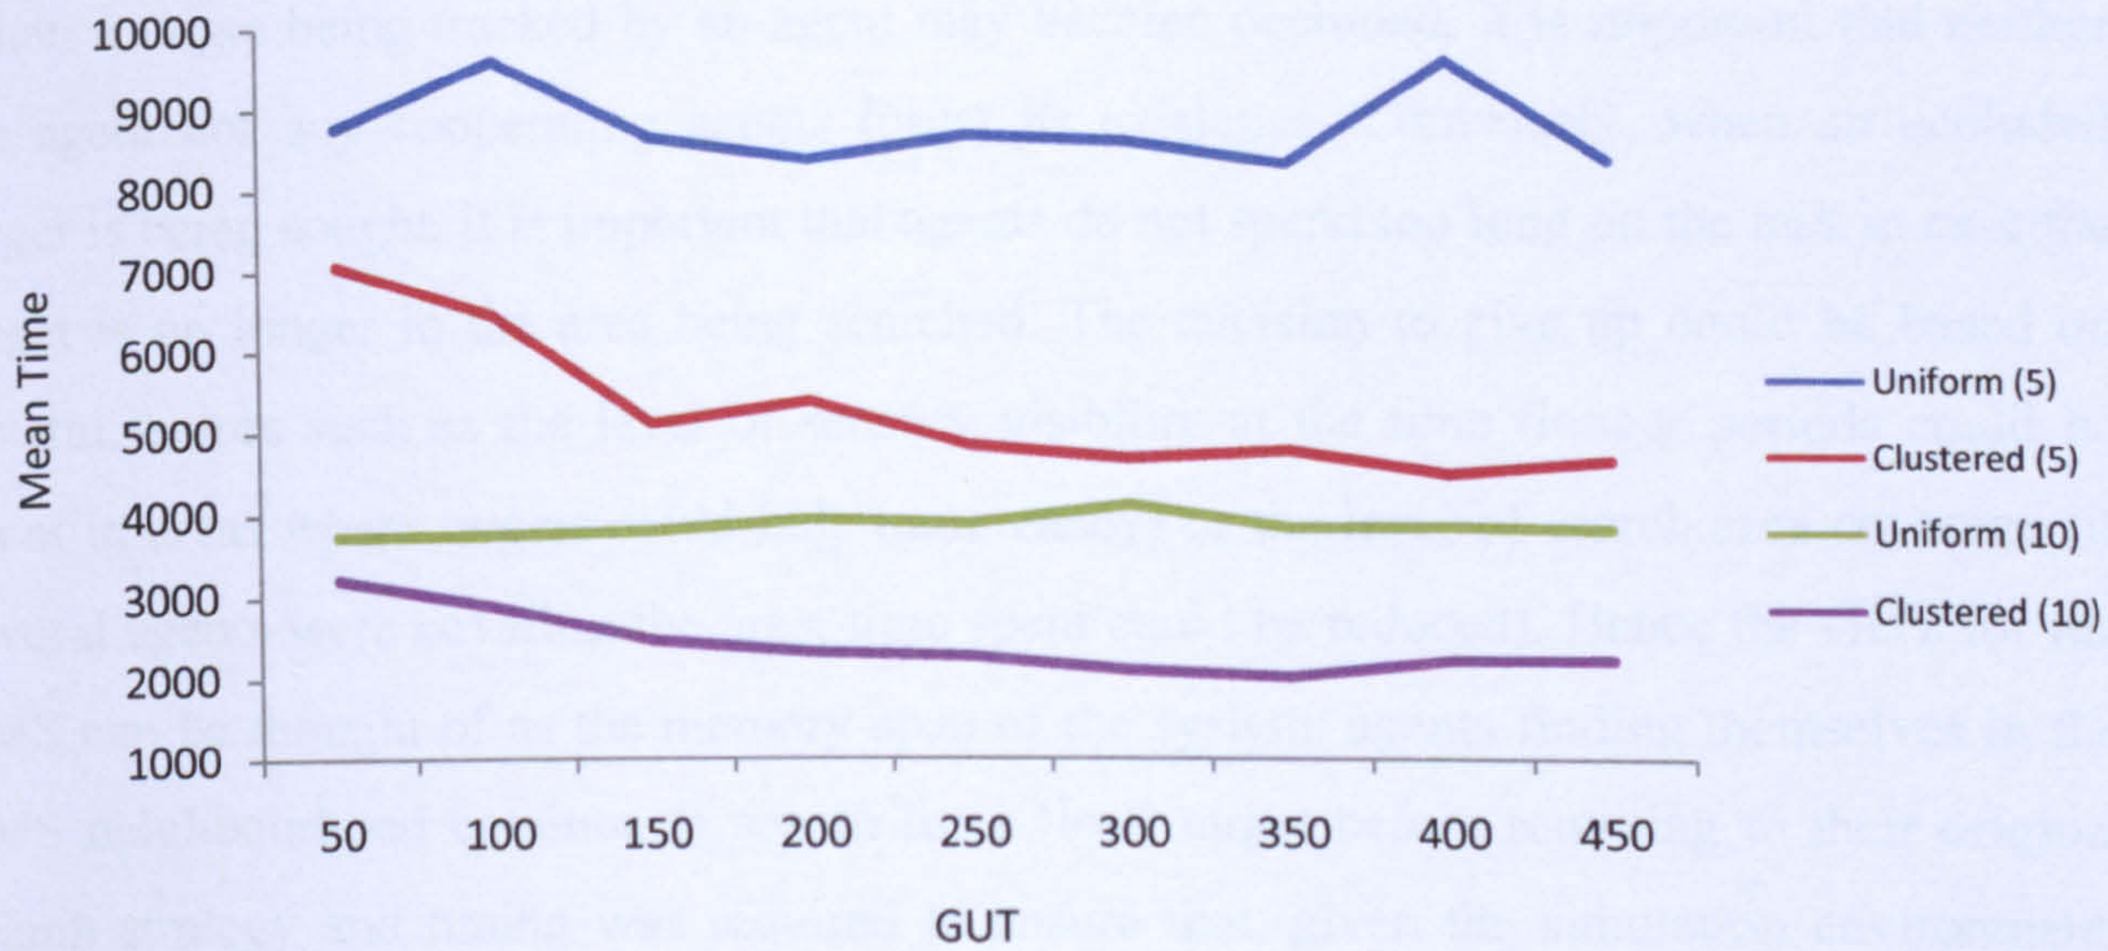

Figure 48. Effect of GUT on time taken to destroy 20 targets by agents employing an ACS strategy.

Based on these results a GUT of 350 was selected for the ACS strategy. The tuning experiments were repeated for dynamic (fruit fly movement) targets and differences in performance between the settings were found to be marginal (results are shown in Figure 49).

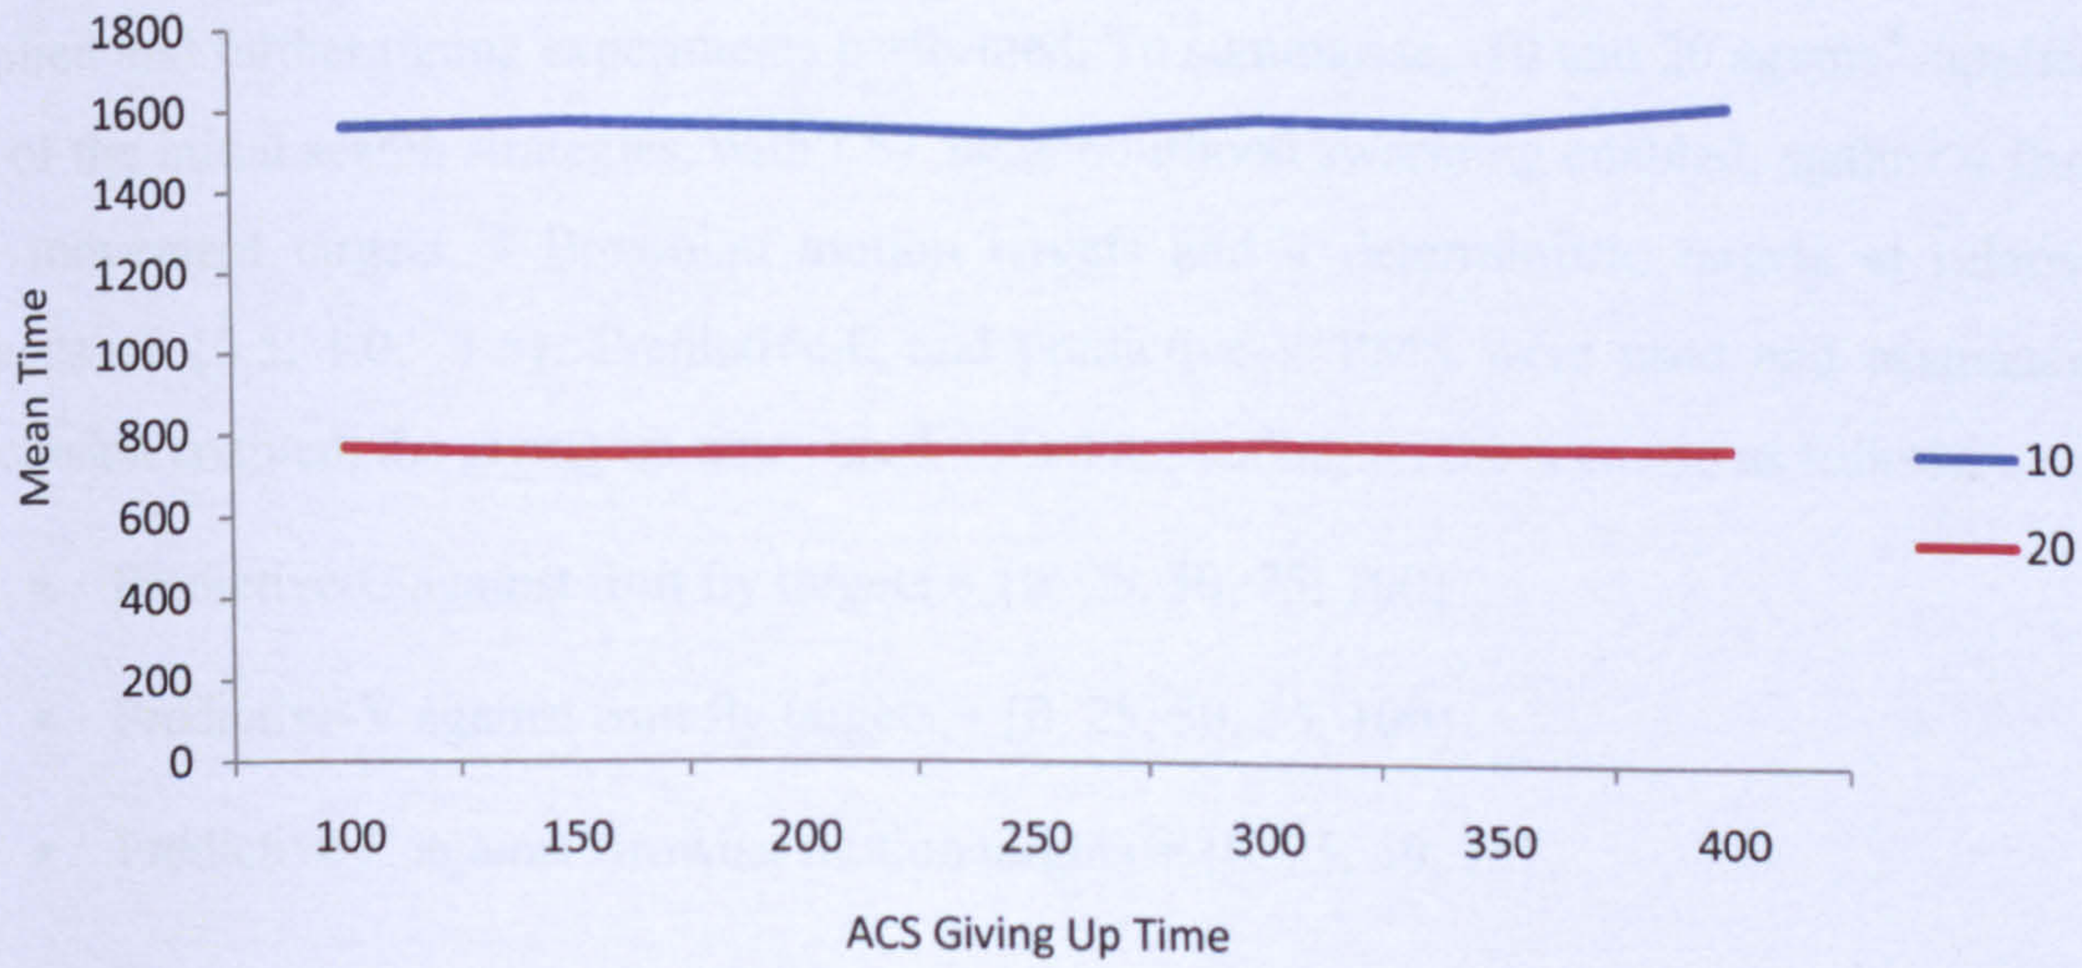

Figure 49. Effect of GUT on time taken to destroy 4 fruit fly targets by 10 and 20 agents employing an ACS strategy.

5.2.4 PMS Giving Up Time

When a target being tracked by an agent may become occluded, it is important that neither the agent nor any cooperating agents forget its existence. Conversely, when an occluded target is being sought, it is important that agents do not spend too long on the task in case the target is no longer in the area being searched. The decision to give up could be based on several factors such as the level of sensory visibility at the time (longer periods could be spent in areas where targets could hide more easily) or the level of search area coverage (if several agents were covering the area, time spent could be reduced). Hence the GUT for the PMS can be thought of as the memory span of the system, agents finding themselves in the PMS neighbourhood continue to search for a 'lost' target before returning to their original search strategy and tuning was required to ensure that, given the simulation environment, agents were spending a reasonable amount of time attempting to re-acquire lost targets.

Initially, experimentation was performed using fruit fly movement targets; however, this work was found to be unsuitable due to the unpredictable movement of the targets post-detection loss. A series of additional experiments were conducted to confirm the findings against less predictable targets. These experiments were conducted using 10 agents only (performance of 20 agent swarms mirrored that of 10 agent swarms in all other PMS GUT scenarios). To overcome problems of unpredictability, deterministic target behaviour was applied and further tuning experiments performed. To summarise, 10 and 20 agents[4] applied all of the initial search strategies, with LSZ neighbourhood swarming enabled, against 4 fruit fly movement targets, 4 Brownian motion targets and 4 deterministic targets at relative speeds of {0.5, 1.0, 1.5}. Predictive-C and predictive-V PMS were used and permanent occlusion enabled; the giving up times used varied depending on the scenario as follows:

- Predictive-C against fruit fly targets = {0, 25, 50, 75, 100}.
- Predictive-V against fruit fly targets = {0, 25, 50, 75, 100}.
- Predictive-C against Browian motion targets = {0, 25, 50, 75}.

---

[4] Note: only 10 agents were used against Brownian motion targets since those experiments were only used to confirm problems against unpredictable targets.

- Predictive-V against Brownian motion targets = {0, 25, 50, 75}.
- Predictive-C against deterministic targets = {0, 25, 50, 75, 100}.
- Predictive-V against deterministic targets = {50, 100, 150, 200, 250, 300, 350, 400}.

The results of circular and V shaped neighbourhood PMS GUT tuning using fruit fly targets can be seen in Figure 50 and Figure 51 respectively. From these surprisingly similar results, it can be seen that it is difficult to tune giving up time using random movement driven targets, even targets that have periods of stable movement such as that provided through fruit fly emulation. Given the level of unpredictability, the shorter times of 25 and 50 were shown to perform marginally better with the former being better suited to smaller swarms. This was due to the reduced time spent looking for targets that had made direction changes away from the predicted path and thus reduced likely of valuable assets being unnecessarily drawn toward non-existent targets.

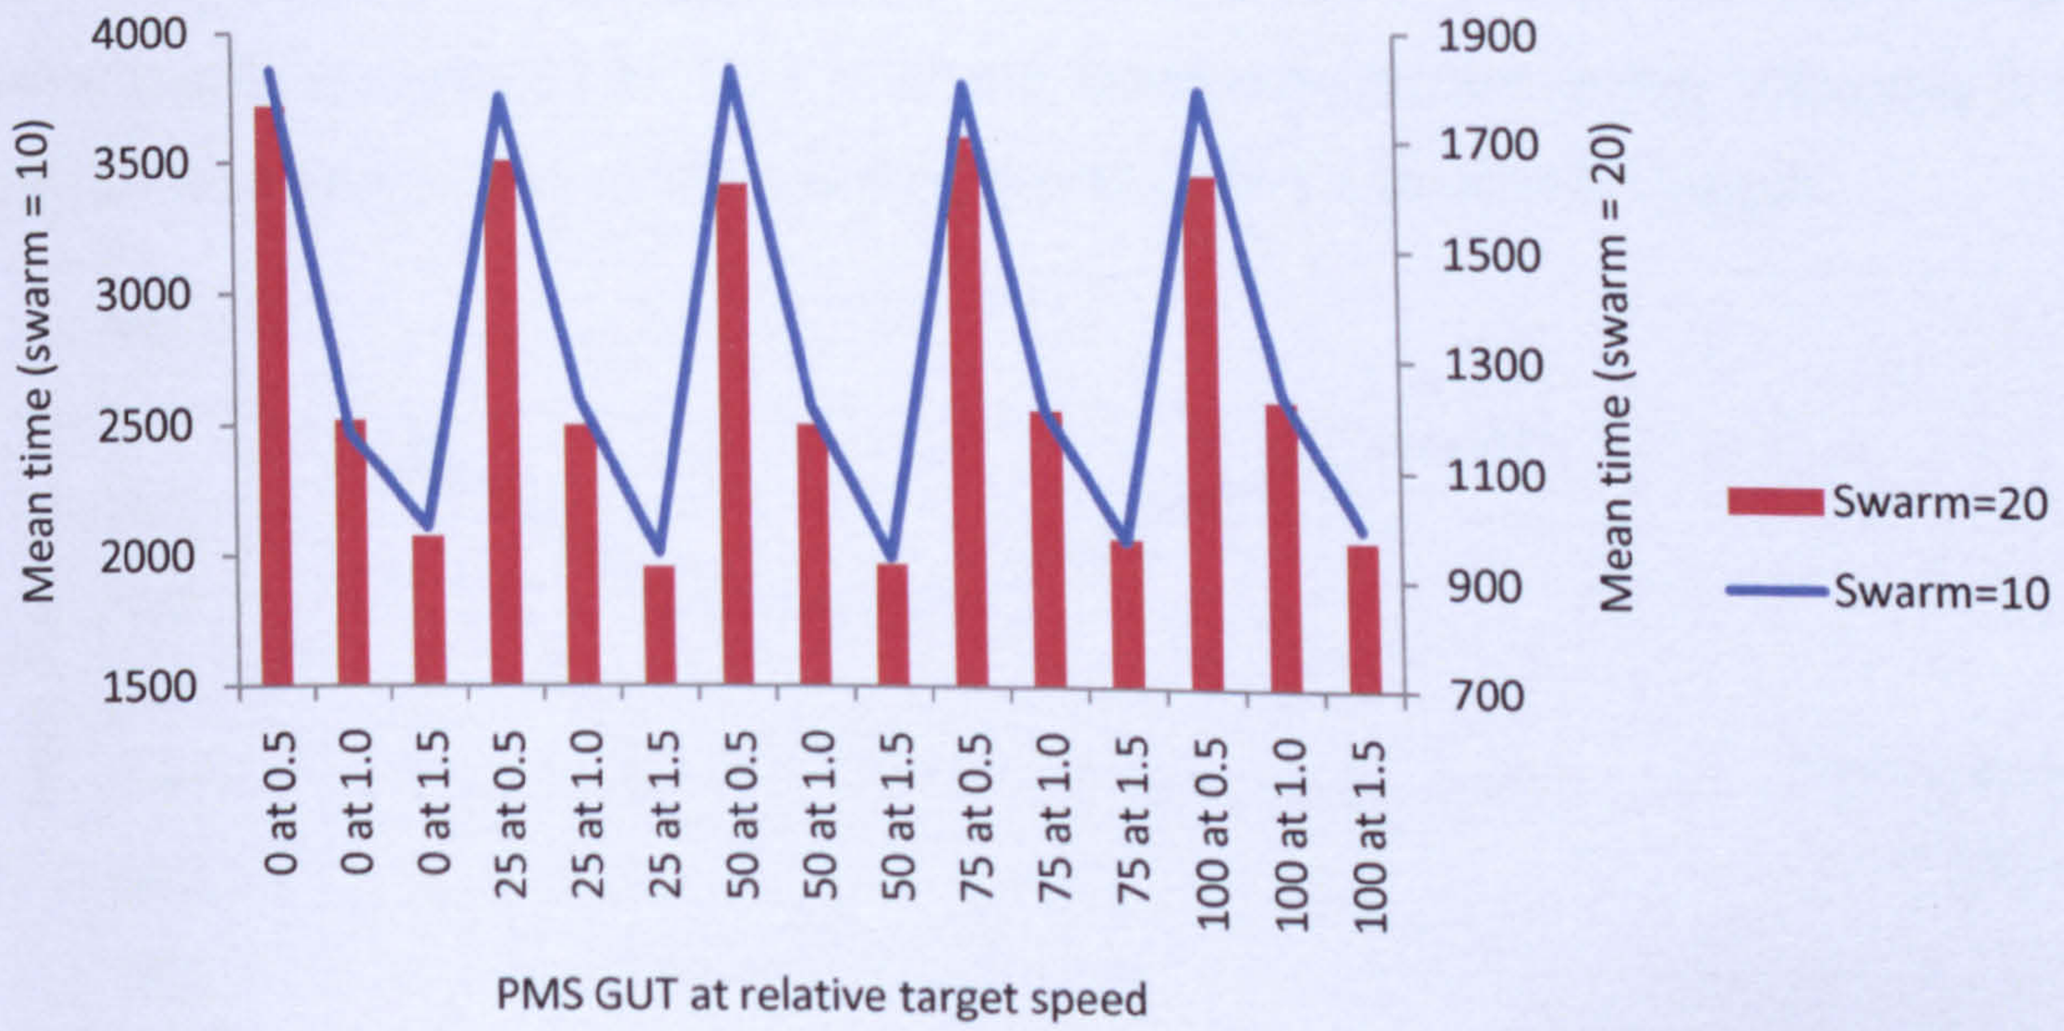

Figure 50. Summary of predictive-C PMS tuning results against 4 fruit fly targets.

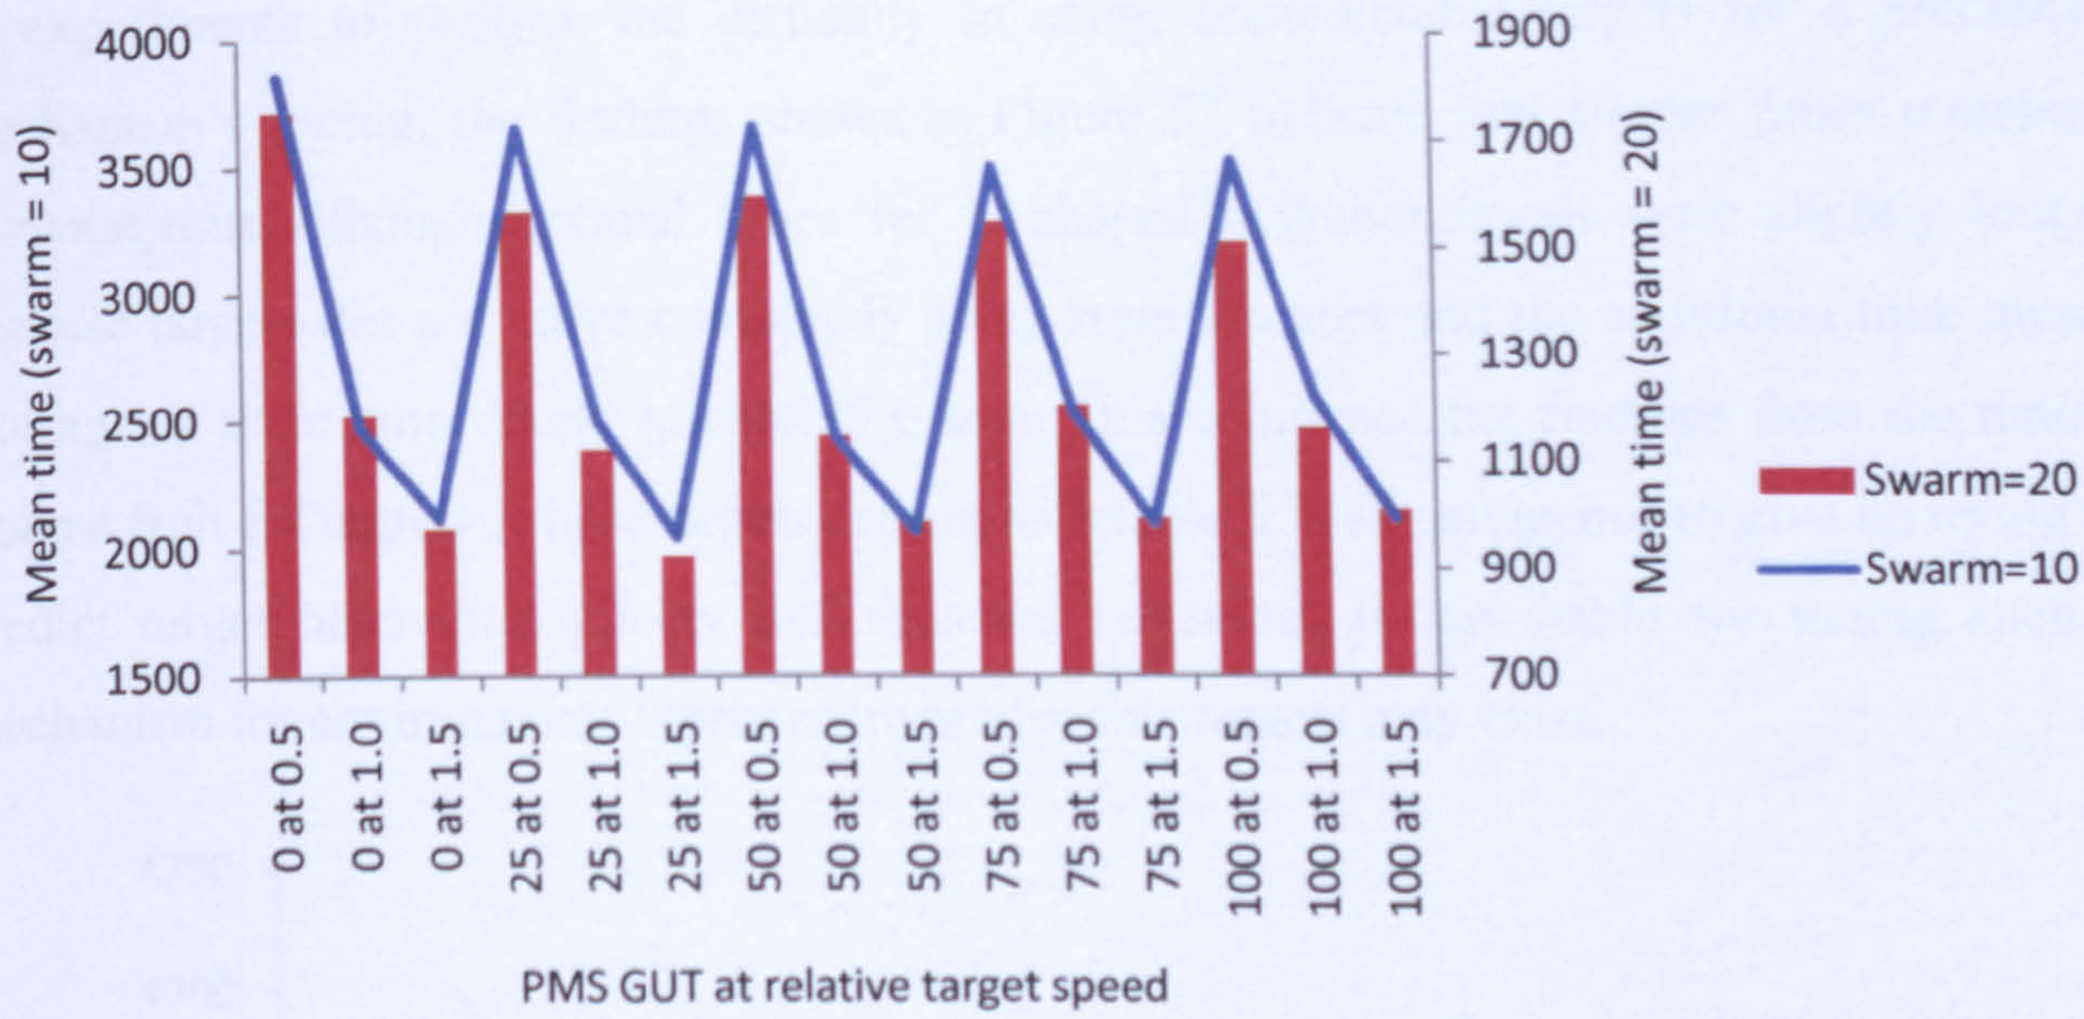

Figure 51. Summary of Predictive-V PMS tuning against 4 fruit fly targets.

When comparing the overall performance, *i.e.* the mean of both swarm sizes and all three relative speeds, in Figure 52 the GUT of 25 was found to be the best setting, indicating it was better for the swarm to give up after shorter periods against unpredictable targets.

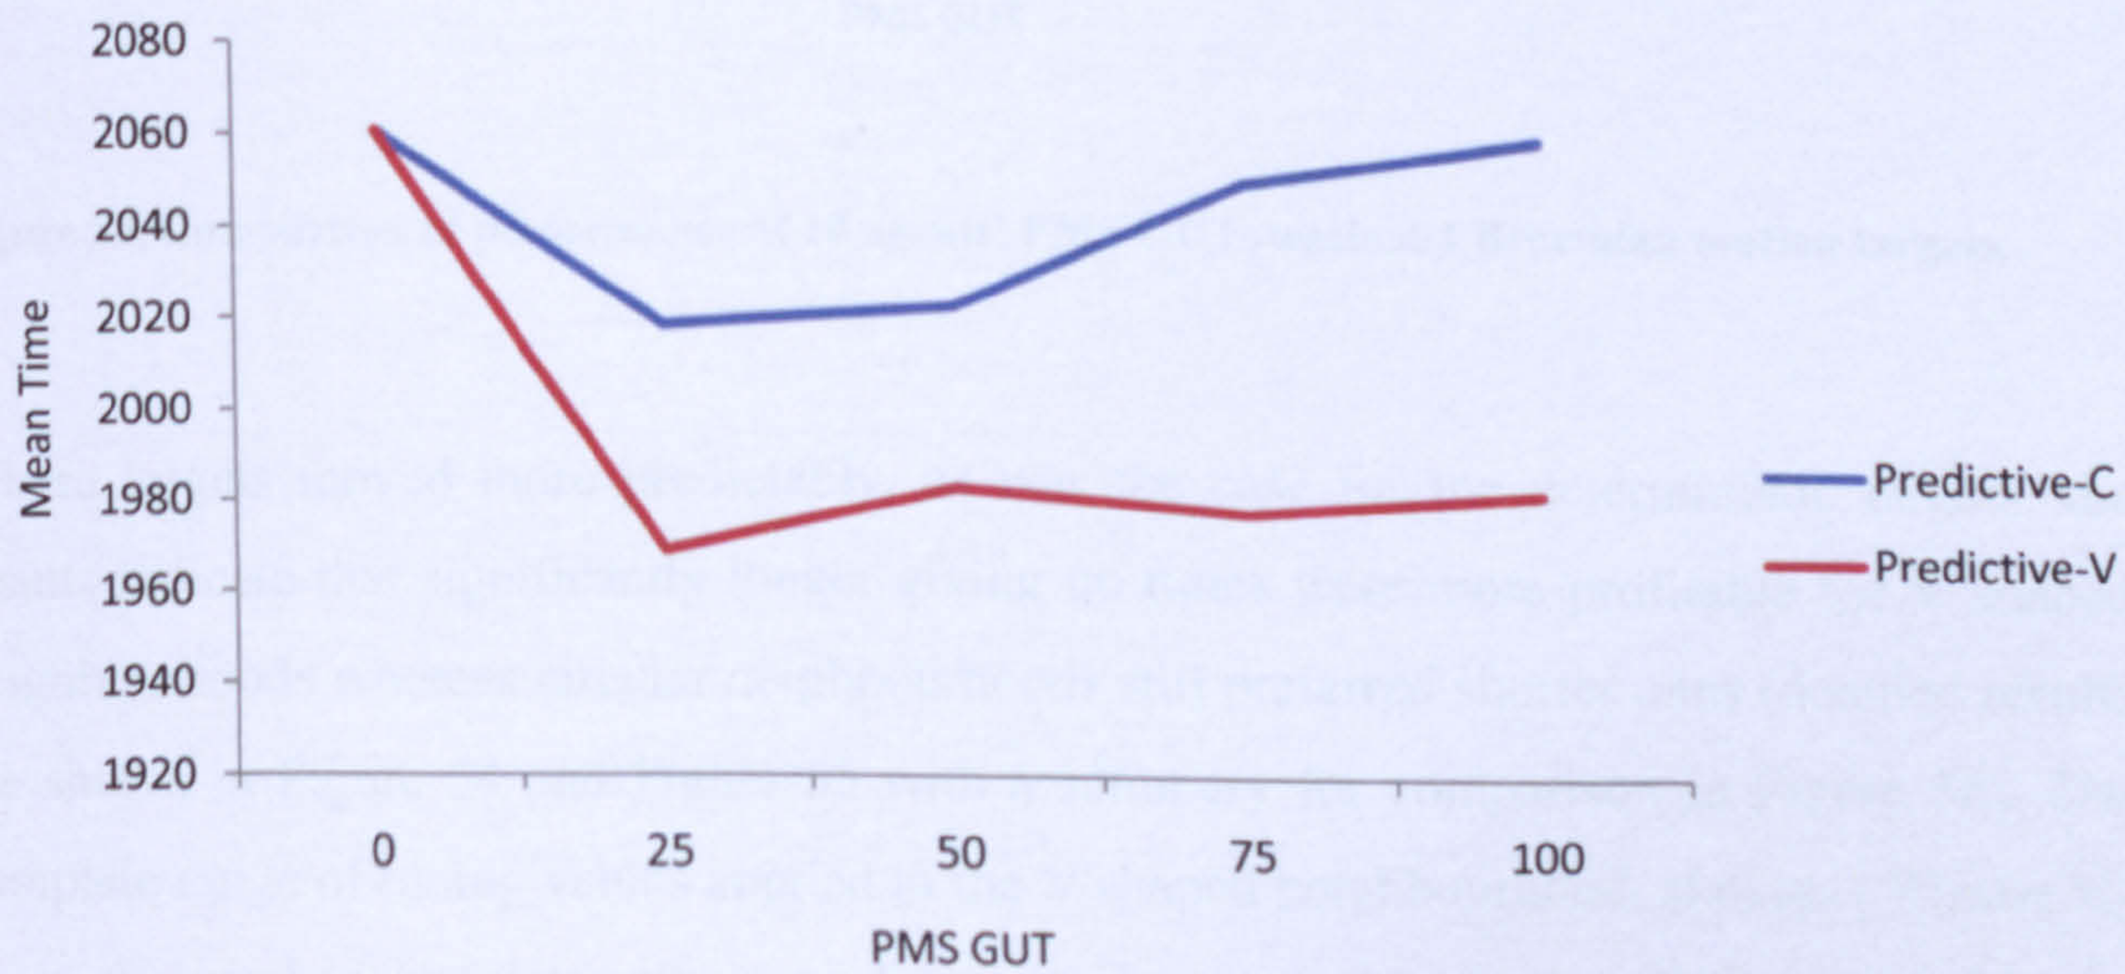

Figure 52. Comparison of overall performance of PMS GUTs against 4 fruit fly targets.

In experiments to confirm the difficulty in using unpredictable targets for a predictive mechanism's tuning, the findings shown at Figure 53 indicate that shorter times remained advantageous although optimal times for V shaped neighbourhoods were slightly longer because targets did not move completely away from the area and the additional time meant that agents were more likely to reach the area. This confirmed the findings from the tuning against fruit fly targets: where targets are unpredictable it is advantageous to give up trying to predict target behaviour quickly and thus the behaviour is unsuitable for tuning such a mechanism for environments where more predictable targets may exist.

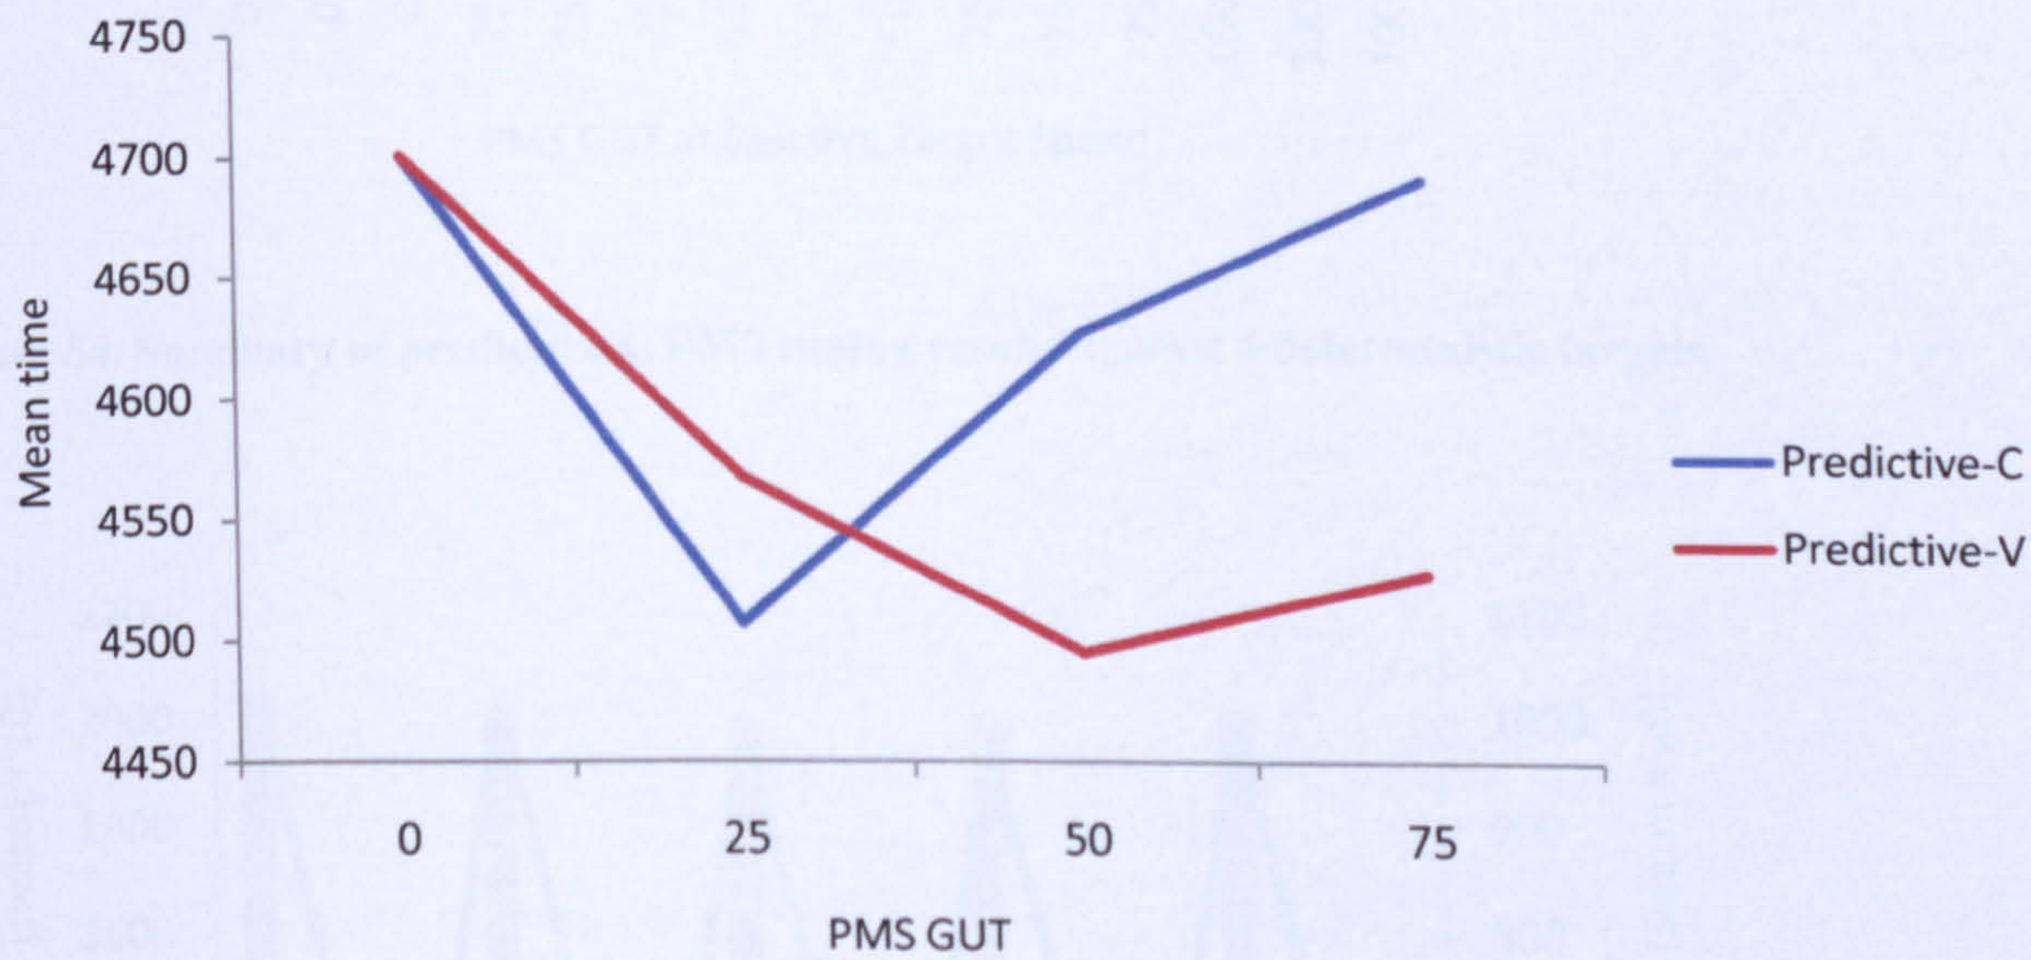

Figure 53. Comparison of performance of 10 agents' PMS GUTs against 4 Brownian motion targets.

Where targets moved more predictably, as was the case for the deterministic targets, the results indicate that significantly longer giving up times were more profitable for V shaped neighbourhoods whereas circular neighbourhoods still preferred shorter ones (detailed results are shown in Figure 54 and Figure 55 with a summary for comparison in Figure 56). The complete range of tuning values applied to the V shaped neighbourhood, shown in Figure 57, show the marked improvement in performance between the shorter GUT as used in the circular neighbourhood and those used as the final settings.

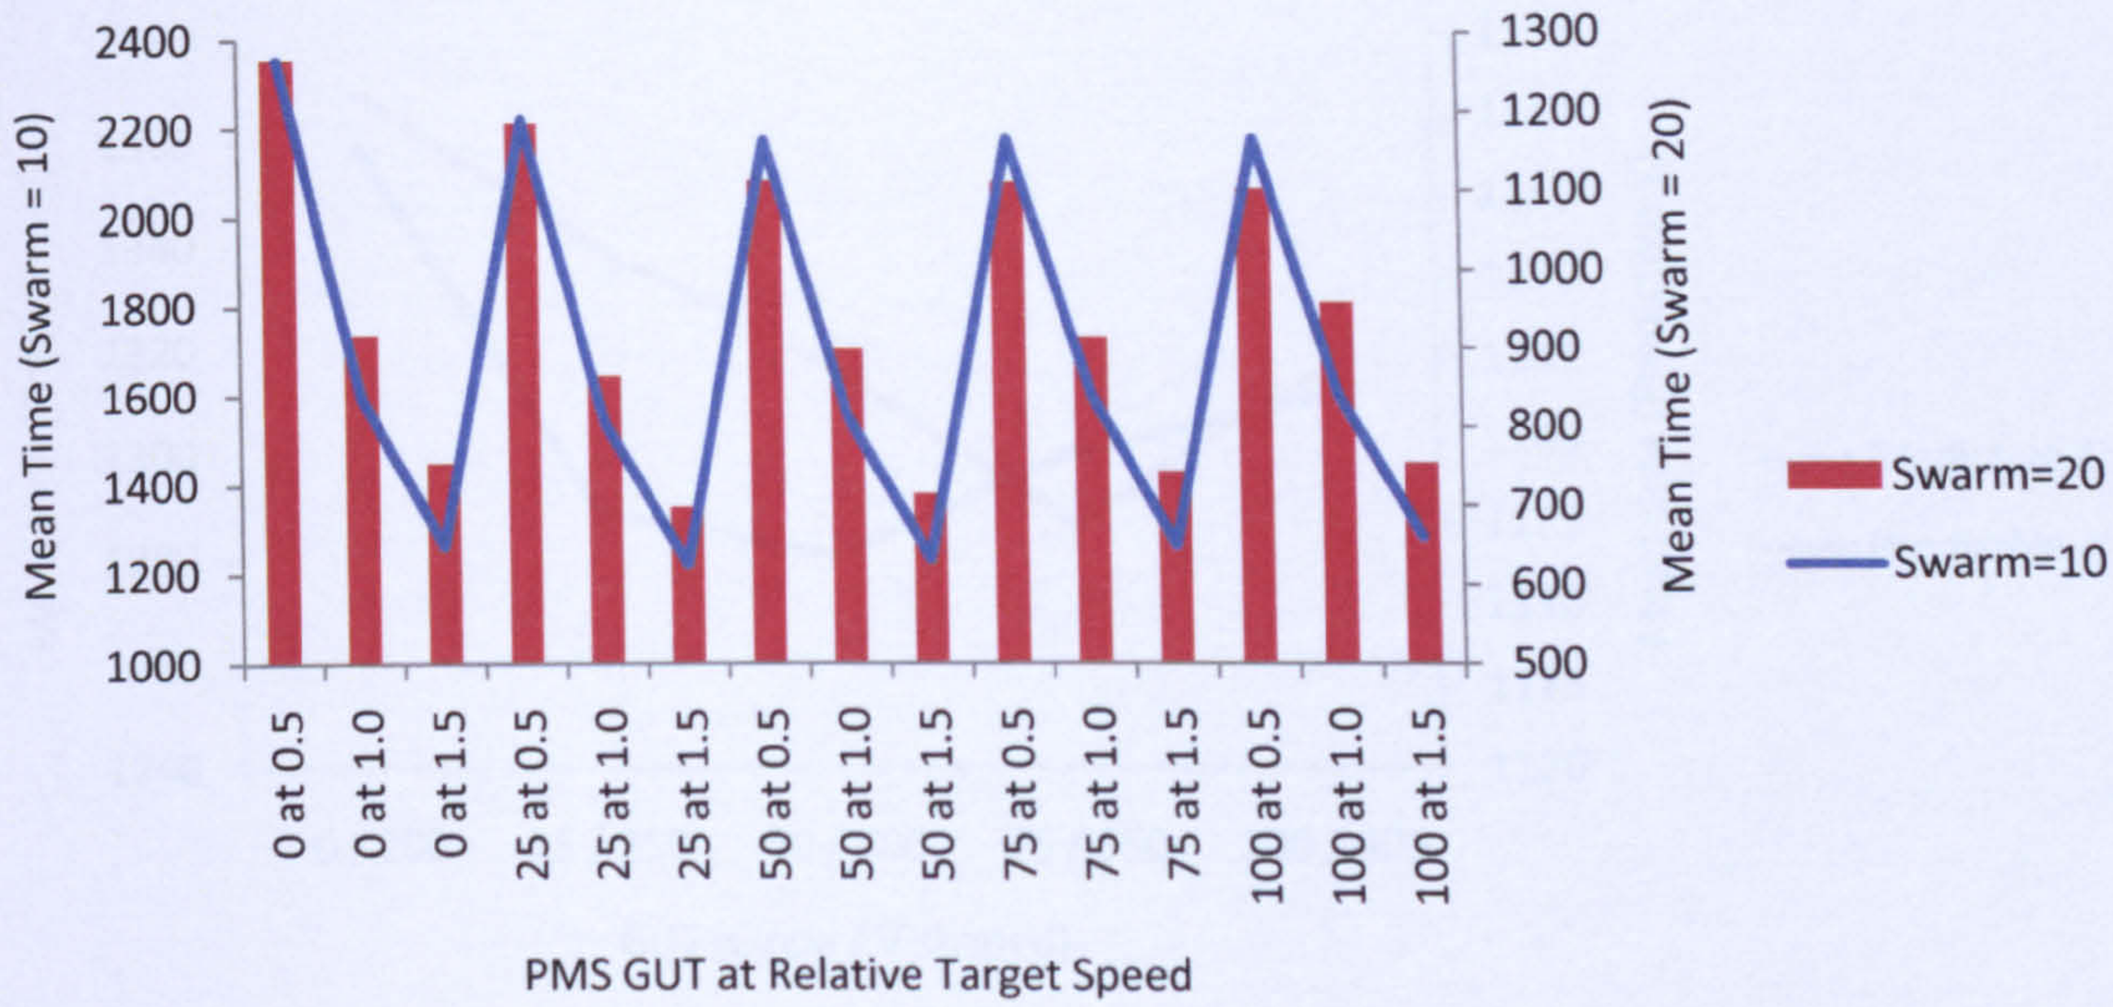

Figure 54. Summary of predictive-C PMS tuning results against 4 deterministic targets.

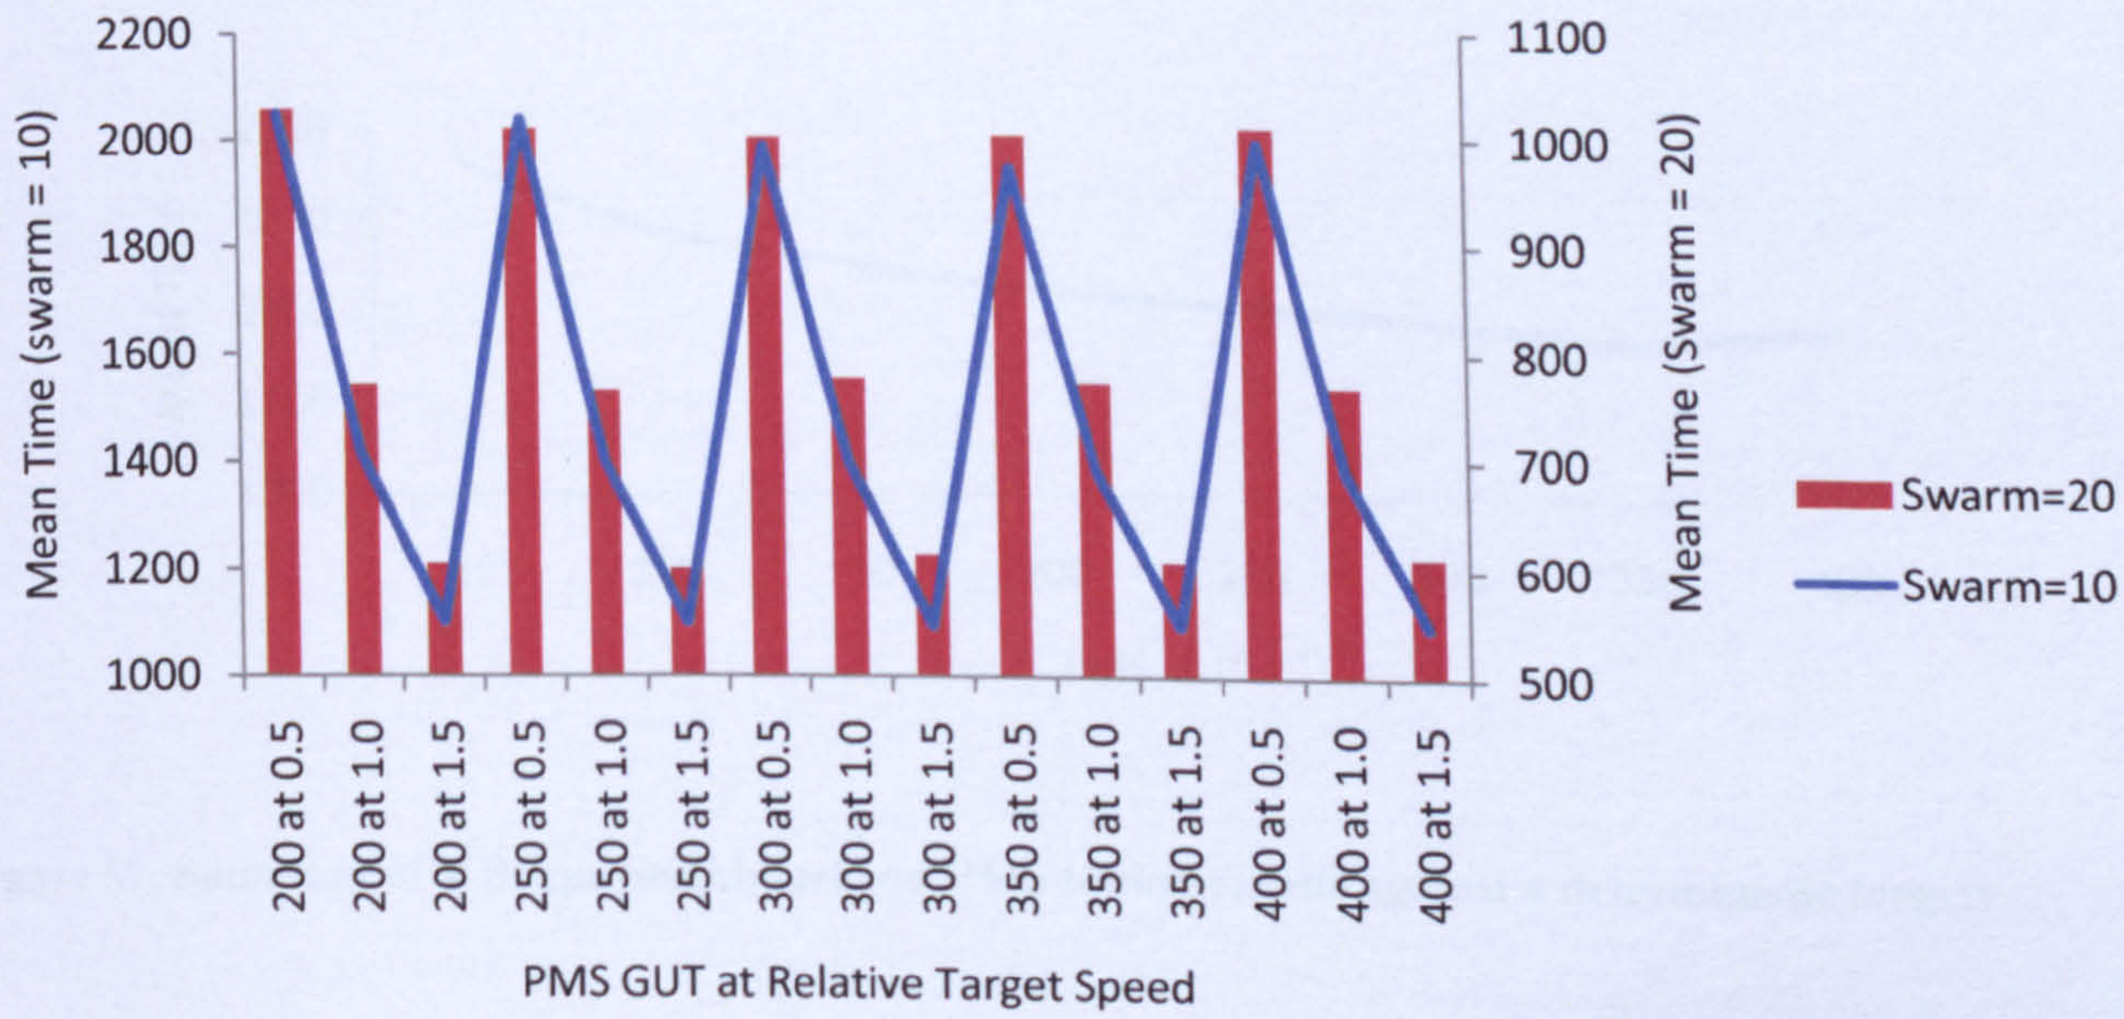

Figure 55. Summary of predictive-V PMS tuning results against 4 deterministic targets.

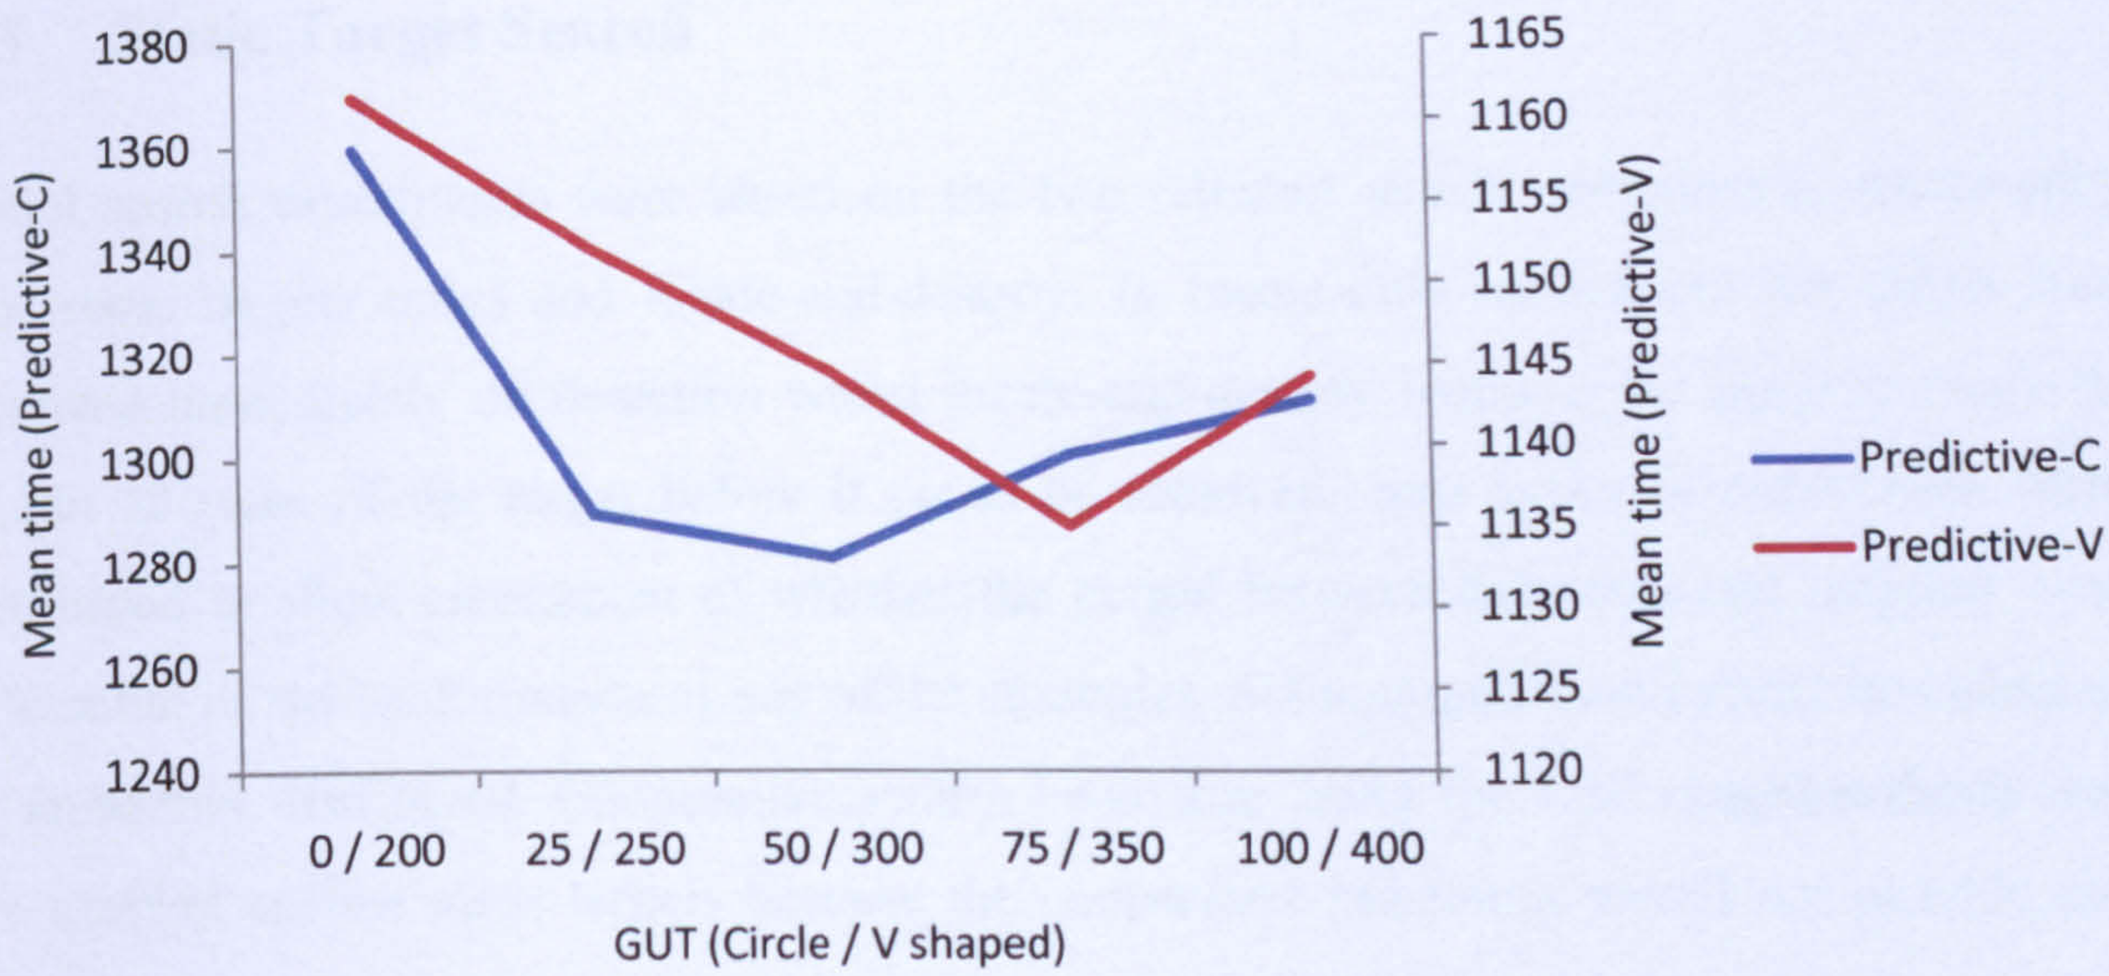

Figure 56. Comparison of overall performance of PMS GUTs against deterministic targets.

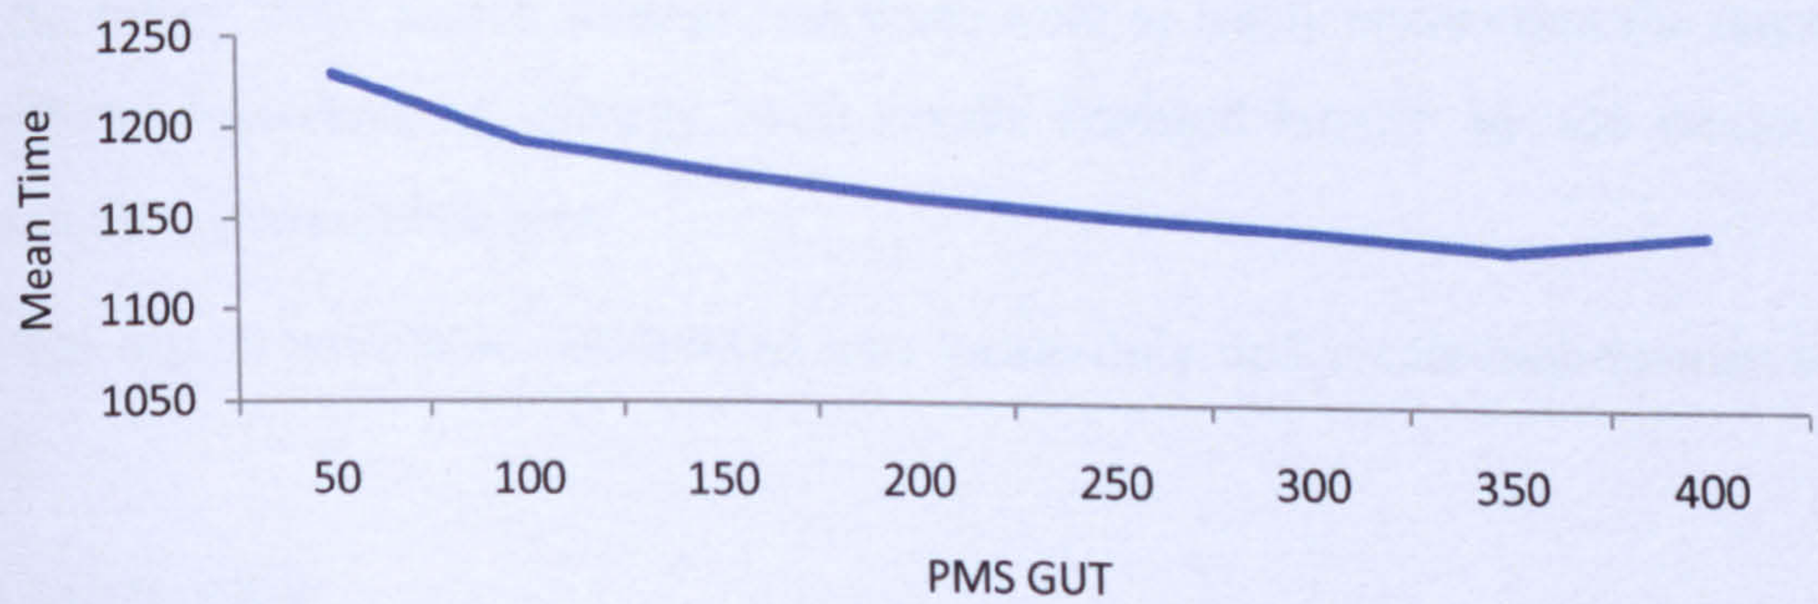

Figure 57. Summary of V Shape neighbourhood PMS tuning results against 4 deterministic targets.

Final settings for the PMS GUT were selected from the deterministic target results (*i.e.* 50 and 350 for the circular and V shaped neighbourhoods respectively) since the results indicated that the PMS would be more useful against predictable targets.

5.3 Static Target Search

Initial search experiments were based on the two relevant search components, locate-only (for static targets only) and locate-and-destroy. In locate-only assessment the target was removed immediately on detection whilst locate-and-destroy required the agent to move to within 10 units of the target before it could be removed; both types of experiment were performed to allow assessment of whether the period between detection and 'capture' was influential on the performance of any of the strategies. Static targets could either be clustered or uniformly distributed. Cooperative swarm behaviour using the LSZ neighbourhood was not enabled against static targets because the cooperative behaviour would not provide any benefits.

Experiments using static targets were configured with small and large groups of agents (5 and 20 respectively) tasked with detecting (locate-only) or detecting and 'destroying' (locate-and-destroy) low and high numbers of targets (4 and 20 respectively) that were either uniformly distributed or arranged in clusters of 4. Larger agent group sizes were tested but were found to mask the effect of the search strategy, *i.e.* there were so many agents that the targets were being detected regardless of strategy, with results dictated largely by the random initial positions of the agents and targets.

Static target results have been subdivided into locate-only and locate-and-destroy scenarios for clarity.

5.3.1 Locate-only

Against uniformly distributed targets (see Figure 58[5]), strategies that ranged more extensively outperformed those with high turn rates; this was most significant where small numbers of agents were searching for large numbers of targets because of scenarios where the targets and agents were not initialised in proximity to each other the high turn rate agents took longer to move to the regions containing targets, a problem that was less likely where agent numbers were large enough to increase the likelihood of agents being well dispersed across the search area.

---

[5] In figures where the key indicates the relationship between agents and targets the number of agents are presented first, therefore a key of '5 Against 20' represents where 5 agents operated against 20 targets.

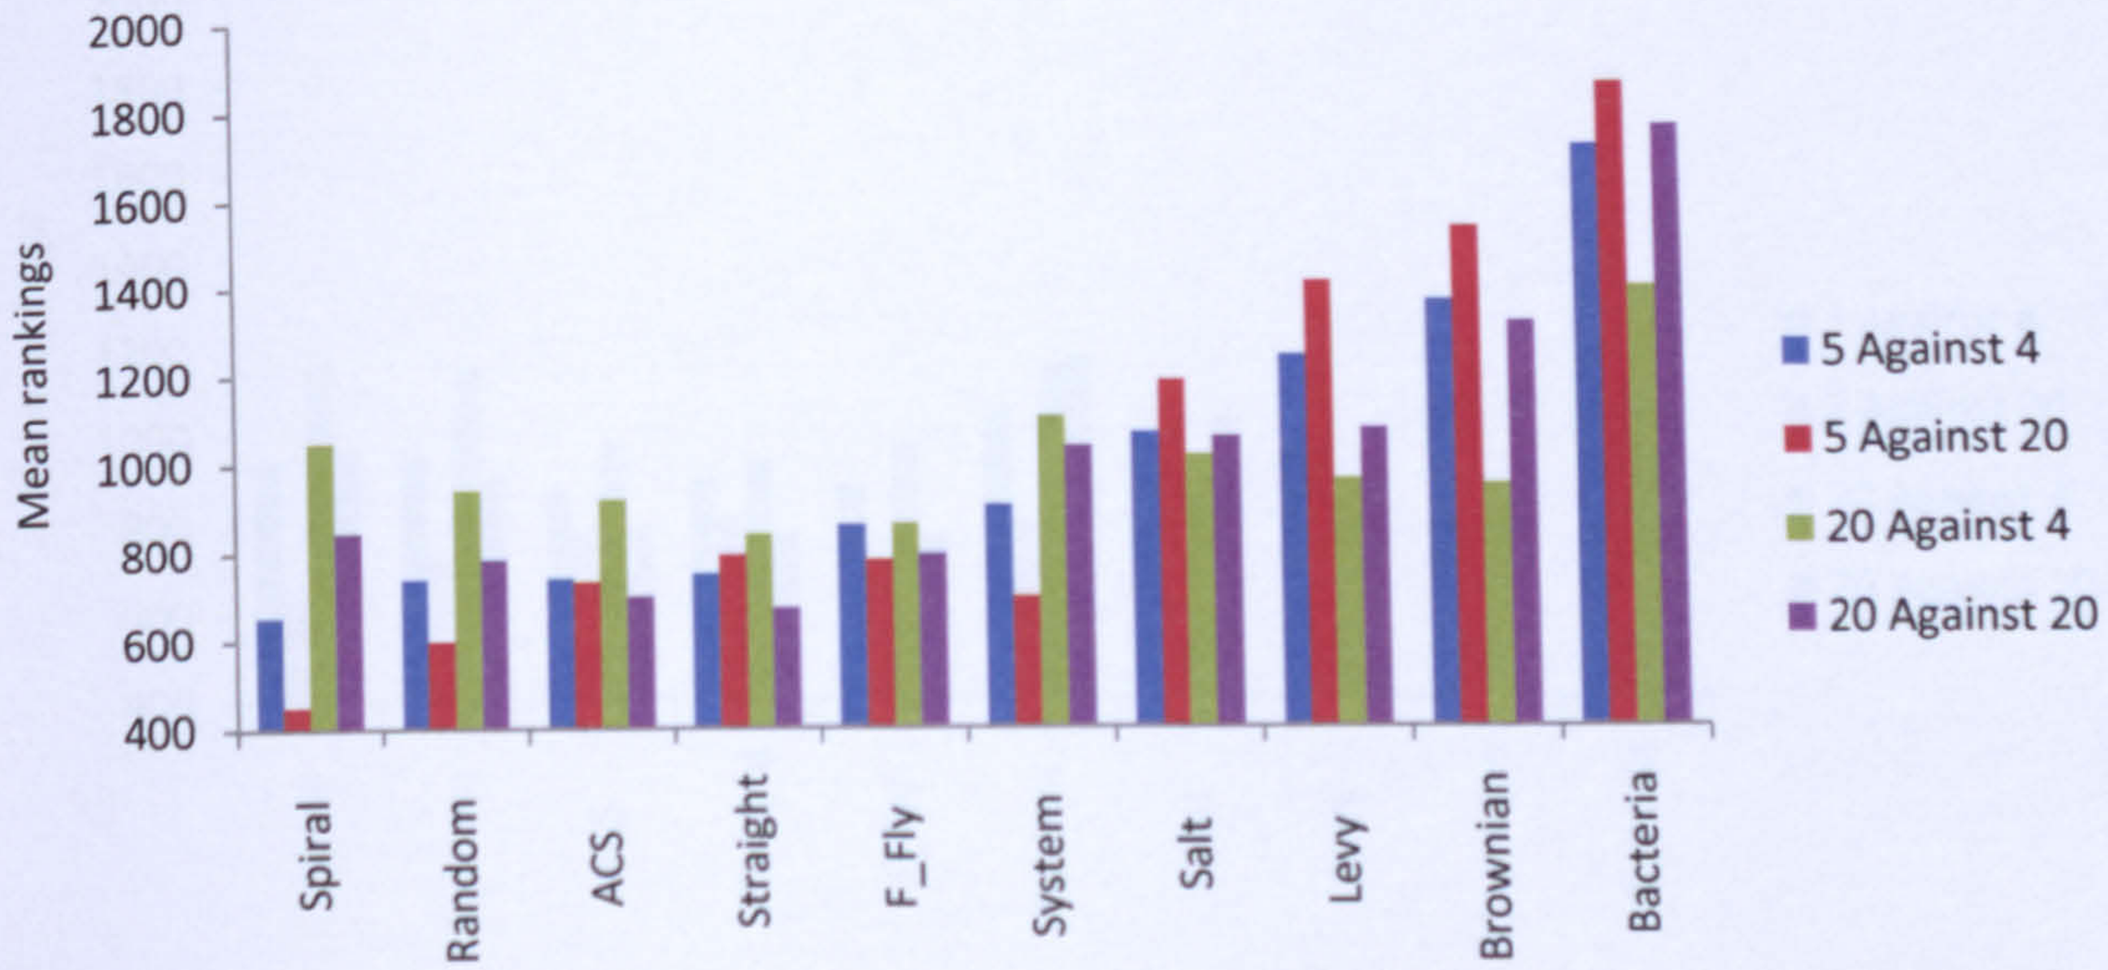

Figure 58. Summary of results for strategies in locate-only search against static uniformly distributed targets.

Where targets were clustered and agent density was low, the findings described above generally held (see Figure 59), although, somewhat surprisingly, the expected improvement in comparative performance of ACS did not occur; this was due to the targets being removed immediately upon detection. Under normal conditions, *i.e.* when the ACS driven agent discovers a target it will track it until destruction, after which it will enter an intensive search phase. In this scenario, locate-only experimentation, the agent will often be located some way from the cluster centre (due to the range of sensory perception) when entering the intensive search phase and therefore may not reach other targets in the cluster. Also notable in this series of experiments was the contrast between the spiral strategy's excellent performance when few agents were seeking large numbers of targets and its poor performance where agent numbers were high and target numbers low. The former being due to the efficient manner in which small numbers of agents eventually cover the whole search area whilst the latter was a result of the agents rarely turning within a cluster, so targets were passed by despite agents discovering others nearby.

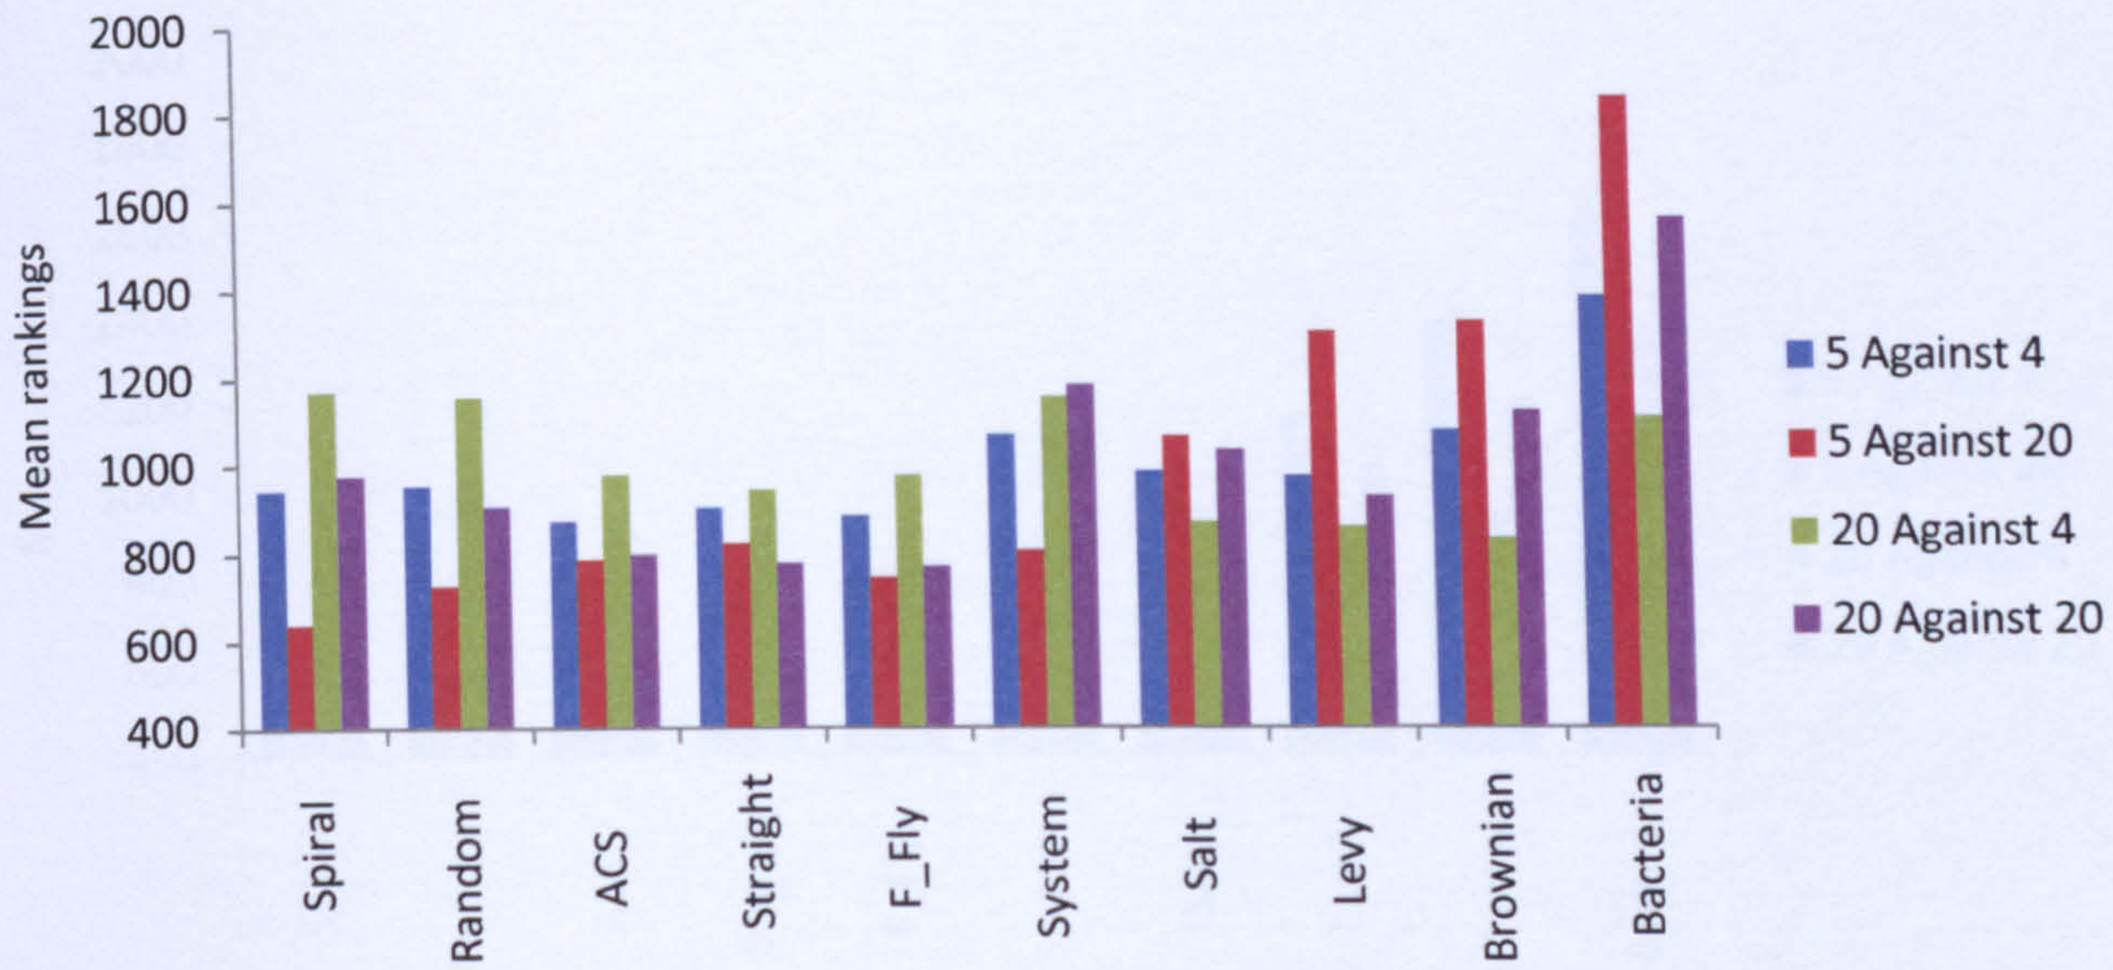

Figure 59. Summary of results for strategies in locate-only search against static clustered targets.

5.3.2 Locate-and-destroy

Generally, the results for uniformly distributed targets held when agents were required to move close enough to the target to destroy it (compare Figure 58 and Figure 60). Bacterial search may have been expected to improve its relative performance since its agents would switch to a wider ranging search during target detection (Figure 23). This failure to improve was relative to other strategies using mean rankings– its actual performance improved significantly; for example comparing the mean time (20 agents against 20 targets) the performance improved by 23.6% whilst all other strategies remained at similar levels (the next largest change was Brownian motion, with 6.5% improvement).

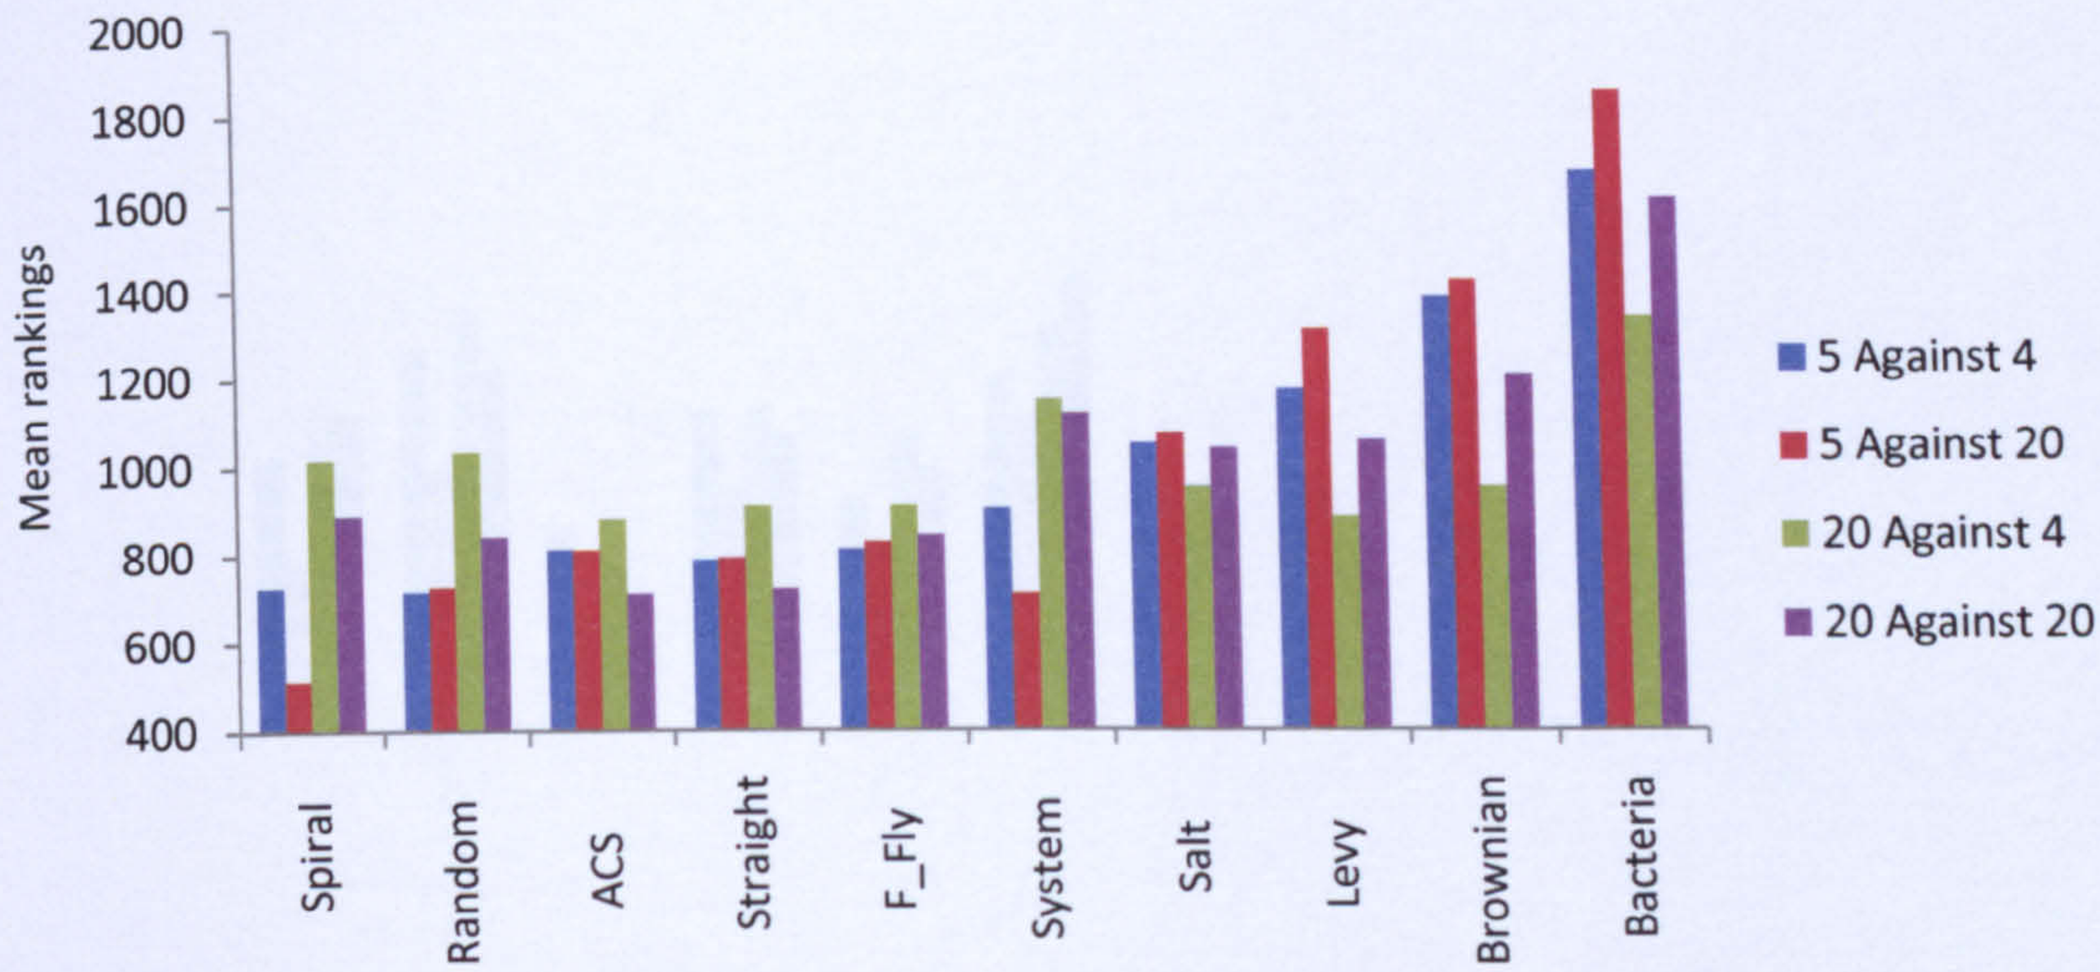

Figure 60. Summary of results for strategies in locate-and-destroy search against static uniformly distributed targets.

The improvement expected from the ACS strategy against clustered targets was evident when the agents were required to move close enough to destroy the target. This was because the switch to intensive search generally occurred closer to the cluster centre and, therefore, provided the agent had sufficient sensory acuity, increased the likelihood of other localised targets being detected. Where *a priori* knowledge exists regarding the level of likely clustering then strategic decisions can be made regarding the intensive search; for example, if clustering is wider than can be detected by normal sensor range then more speed could be maintained to allow the agents to range further whilst having the high turn rate required to cover the area. The two phase approach of bacterial search also paid dividends in the clustered environment; this time with sufficient improvement to markedly improve comparative performance provided sufficient agents were present to achieve an initial detection.

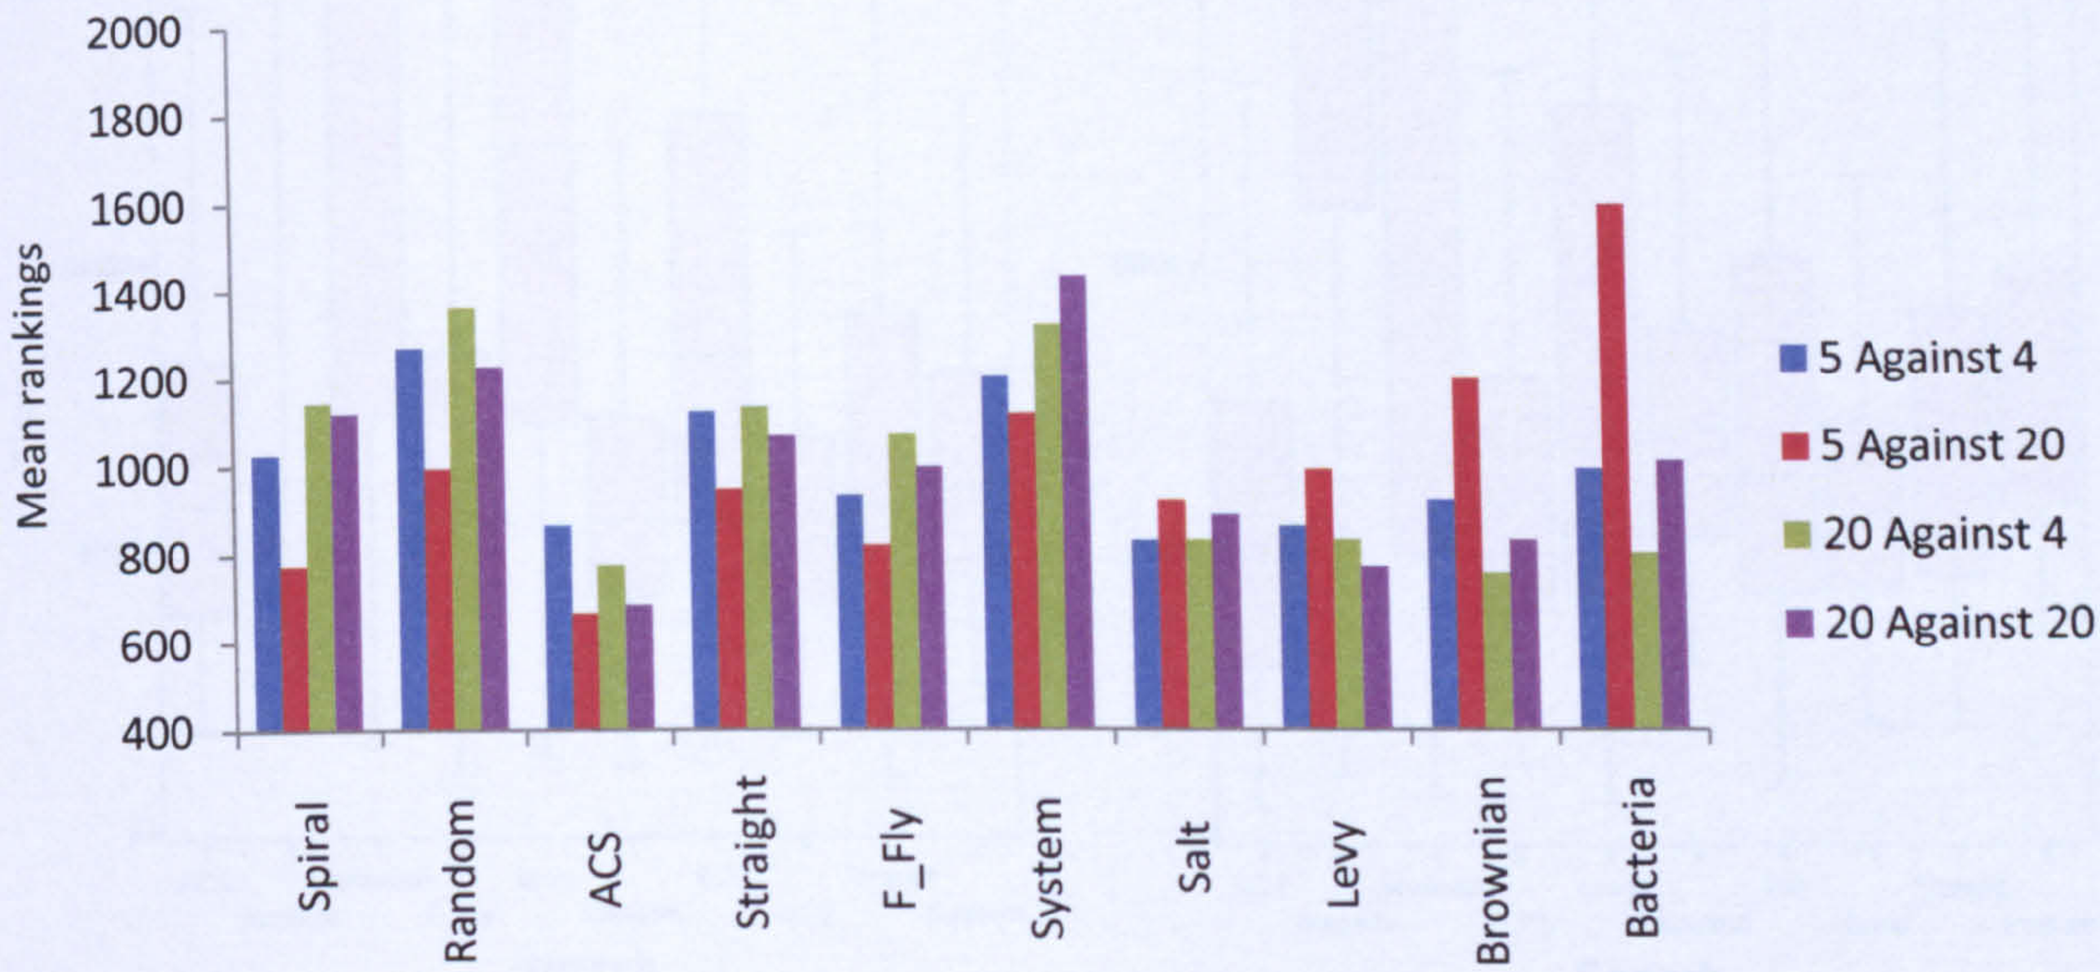

Figure 61. Summary of results for strategies in locate-and-destroy search against static clustered targets.

Closer examination of the actual results, Figure 62 and Figure 63, shows that whilst some strategies improved others performed less well when required to interrupt their search behaviour to get close enough to destroy the target. In particular, strategies with high turn rates (*e.g.* bacterial, Brownian motion and Lévy path) gained most because when one target was detected the agent was drawn closer to the other clustered targets faster than would have been the case following the usual, generally tortuous, path. The systematic grid pattern's performance suffered because following the move to destroy a target an agent would generally resume its search on a different part of the grid pattern, potentially taking it away from detecting the other targets in the cluster.

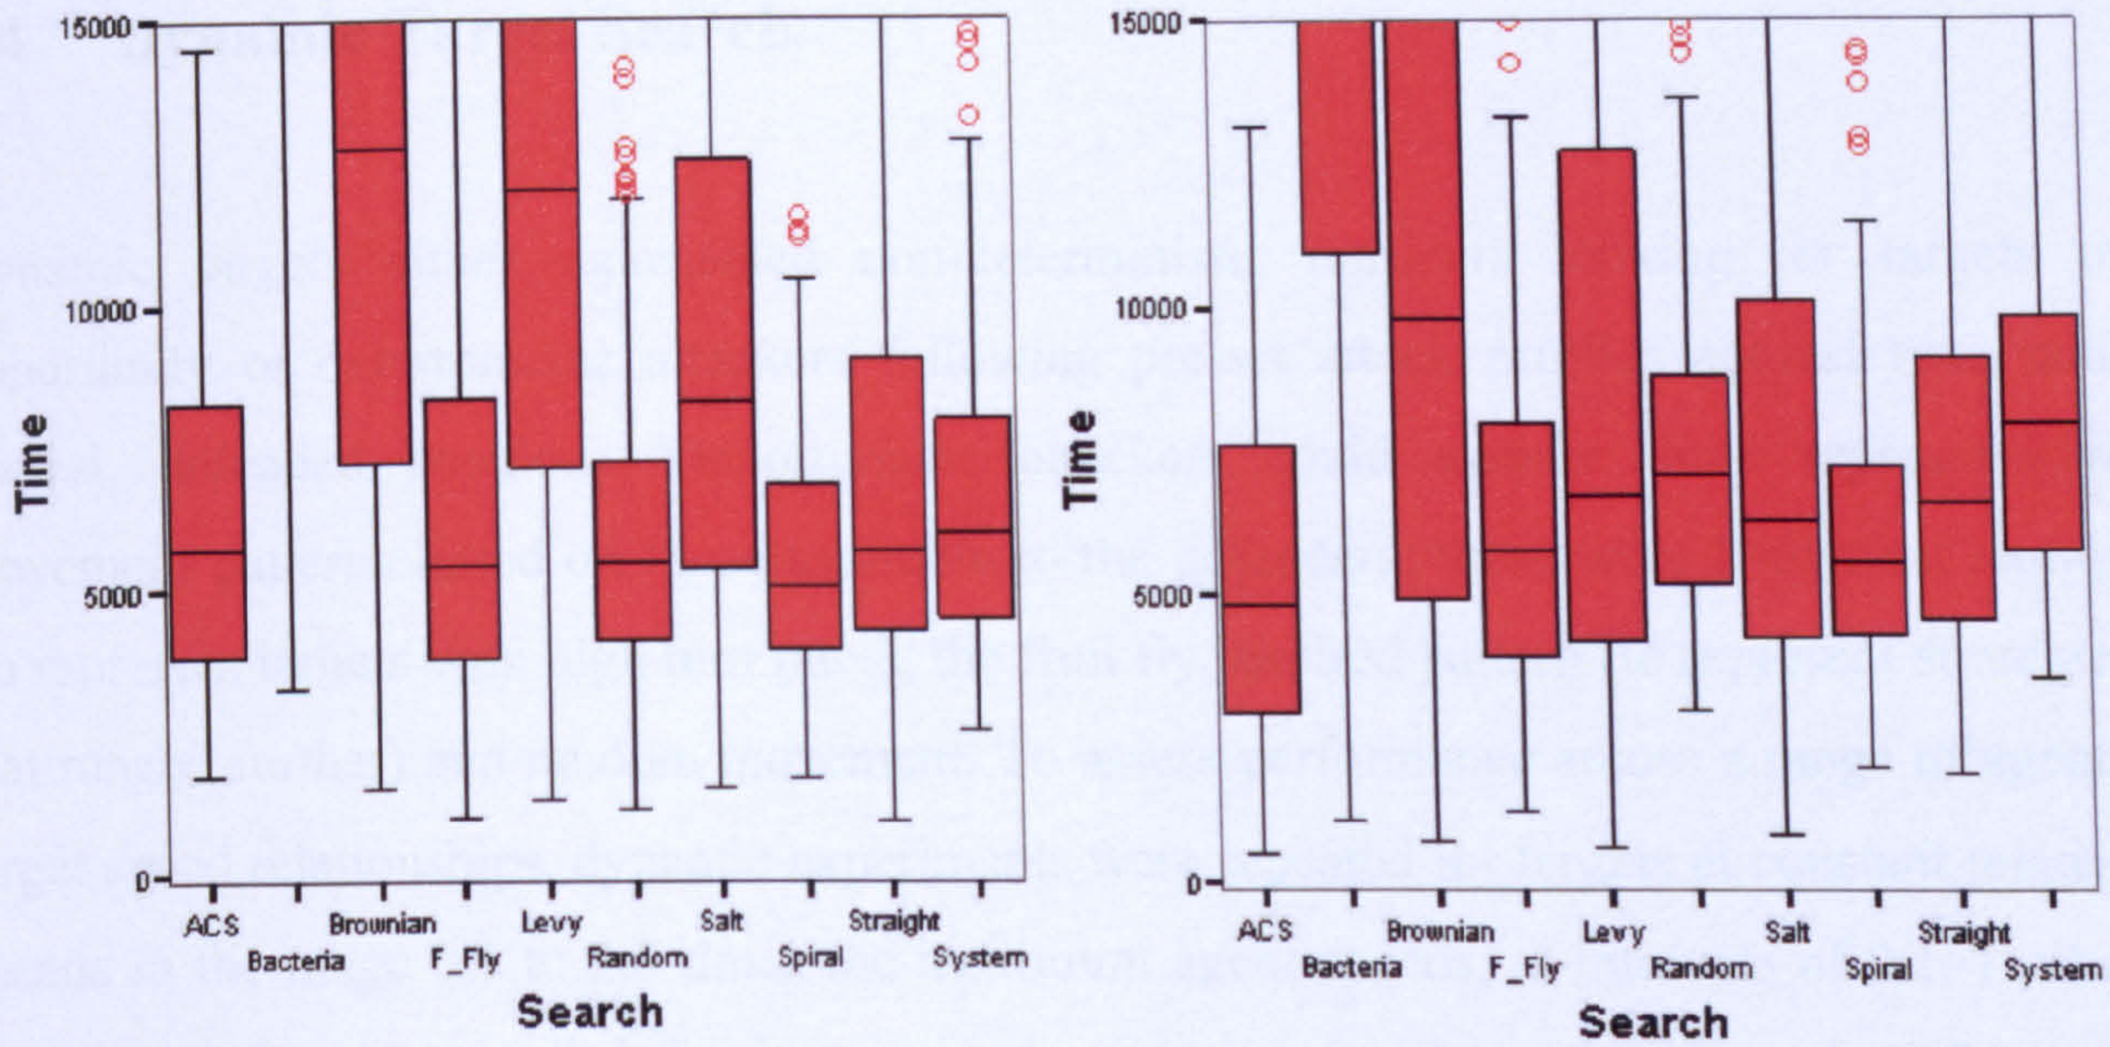

Figure 62. Comparison of actual results for 5 agents against 20 static clustered targets concentrating on the bottom half of the time scale, left: locate-only and right: locate-and-destroy (note: bottom half of the graphs only shown here because the effect discussed is hidden when shown full scale – see Figure 63 for full scale graphs).

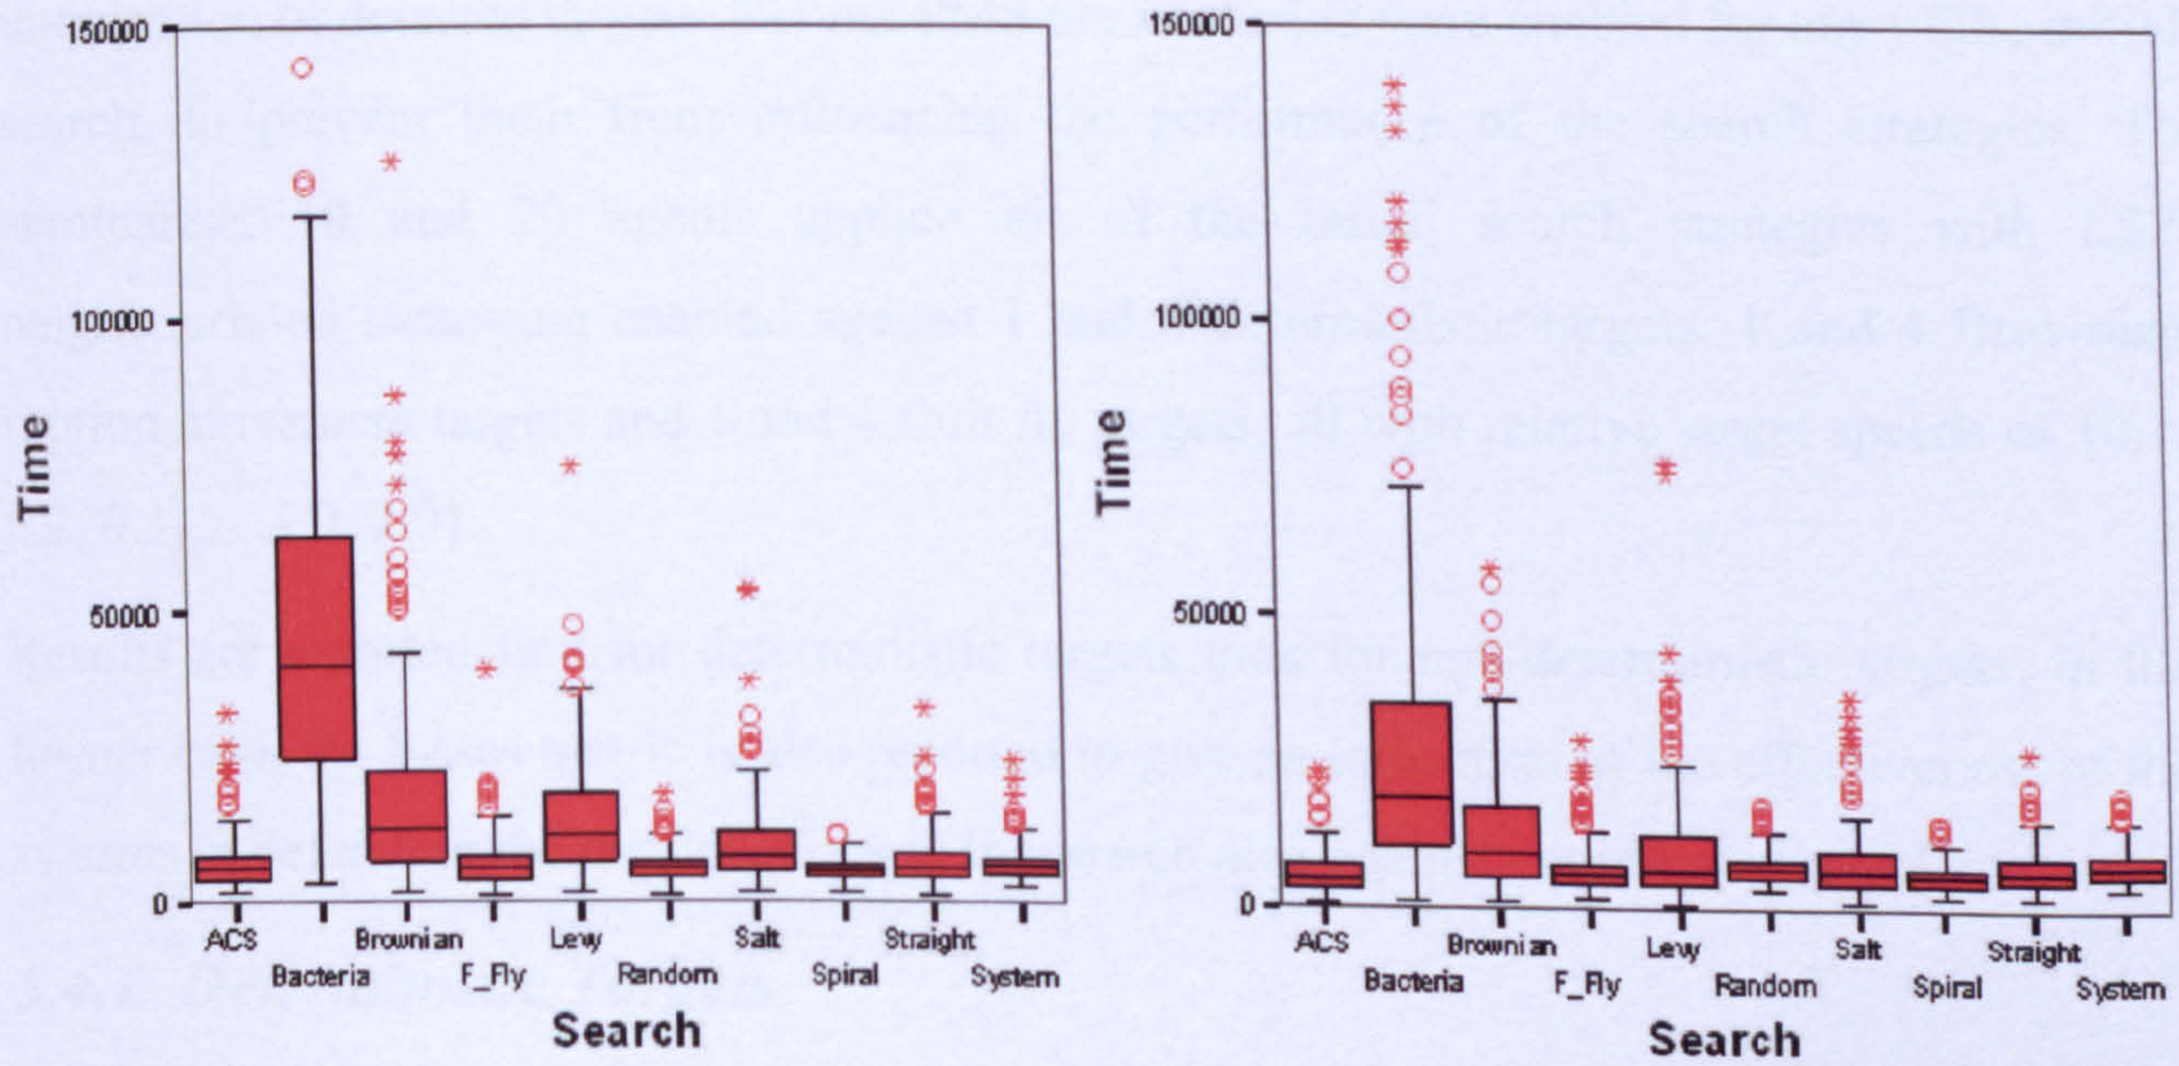

Figure 63. Comparison of actual results for 5 agents against 20 static clustered targets, left: locate-only and right: locate-and-destroy (note: these are full scale versions of Figure 62).

5.4 Dynamic Target Search

Dynamic targets either represented non-deterministic intruders looking for targets of opportunity or deterministic attackers following pre-set attack profiles against two, non-central, defended positions. Opportunistic attackers could assume three representative movement patterns based on those available to the defenders; these were Brownian motion (to represent targets with high turn rates), the fruit fly inspired pattern (to represent strategies that ranged further) and random movement. To assess performance across a range of agent–target speed relationships, dynamic experiments were repeated for targets at constant relative speeds in the range 0.1 to 2.0 times the maximum agent speeds, at intervals of 0.1. During dynamic testing, the small defensive group size was increased to 10 agents, since 5 agents was found to be too few for the large search space, and target encounters were more through random chance than engineered by the search strategies. Swarming with the LSZ neighbourhood was enabled for these experiments, allowing the agents to cooperate on the interception of detected targets. Neither PMS nor occlusion were enabled for any of the initial search, to prevent them from influencing the performance of the search strategies. To summarise, 10 and 20 agents applied all of the initial search strategies with LSZ neighbourhood swarming enabled against 1 and 4 deterministic targets, 1 and 4 Brownian motion movement targets and 1 and 4 fruit fly targets, all with relative target speeds of {0.1, 0.2, 0.3,... ,1.9, 2.0}.

Results are reported first for deterministic targets then for non-deterministic targets; in the former case, the losses metric is also reported to give an indication of the effectiveness of the swarms in defending the two locations in the search area against coordinated attack.

5.4.1 Deterministic Targets

In the dynamic environment where targets followed deterministic routes, the results presented in Figure 64 show that as the targets speed up the ranging search techniques become less effective whilst those with higher turn rates become more attractive. This was most notable with Brownian and bacterial making large gains in relative performance at the expense of spiral and random searches. All three of the highest turn rate strategies generally outperform the ranging strategies, except at the slowest target speeds, where the target movement is insufficient to bring the opposing individuals into close proximity. As target speed increases

beyond agent speed, the performance difference once again reduces as target–agent encounters become more inevitable through chance; although, the high turn rate strategies remained influential, as can be seen from the generally better performance at the fastest target speeds.

These trends are also apparent with more targets and larger defensive swarms (Figure 65, Figure 67 and Figure 68), although the differences in performance is more pronounced. Where more targets were present the difference was amplified as the influence of the strategy on performance rather than chance meetings between agent and target became more apparent.

The mean losses results presented in Figure 66 and Figure 69 confirm the time-based performance results. In the first of these, the random and systematic movement perform less competitively throughout the range of relative speeds whilst bacterial search, due to its limited ability to range about the search area, performed less well compared to the ranging strategies at low speeds but as the relative speeds increased it quickly became one of the most competitive due to its high turn rate. Interestingly, the high turn rate strategies that did not slow down managed to effectively defend the area even at low speed (although Figure 65 shows they took longer to destroy the intruders).

With higher agent numbers the low speed of the bacterial search was less problematic and, whilst the mean rankings of the time-based metric results (Figure 68) show that the strategy did not perform well at the low speeds, the losses indicate that, due to the swarm size, it was still effective at defending the locations shown in Figure 12, page 64.

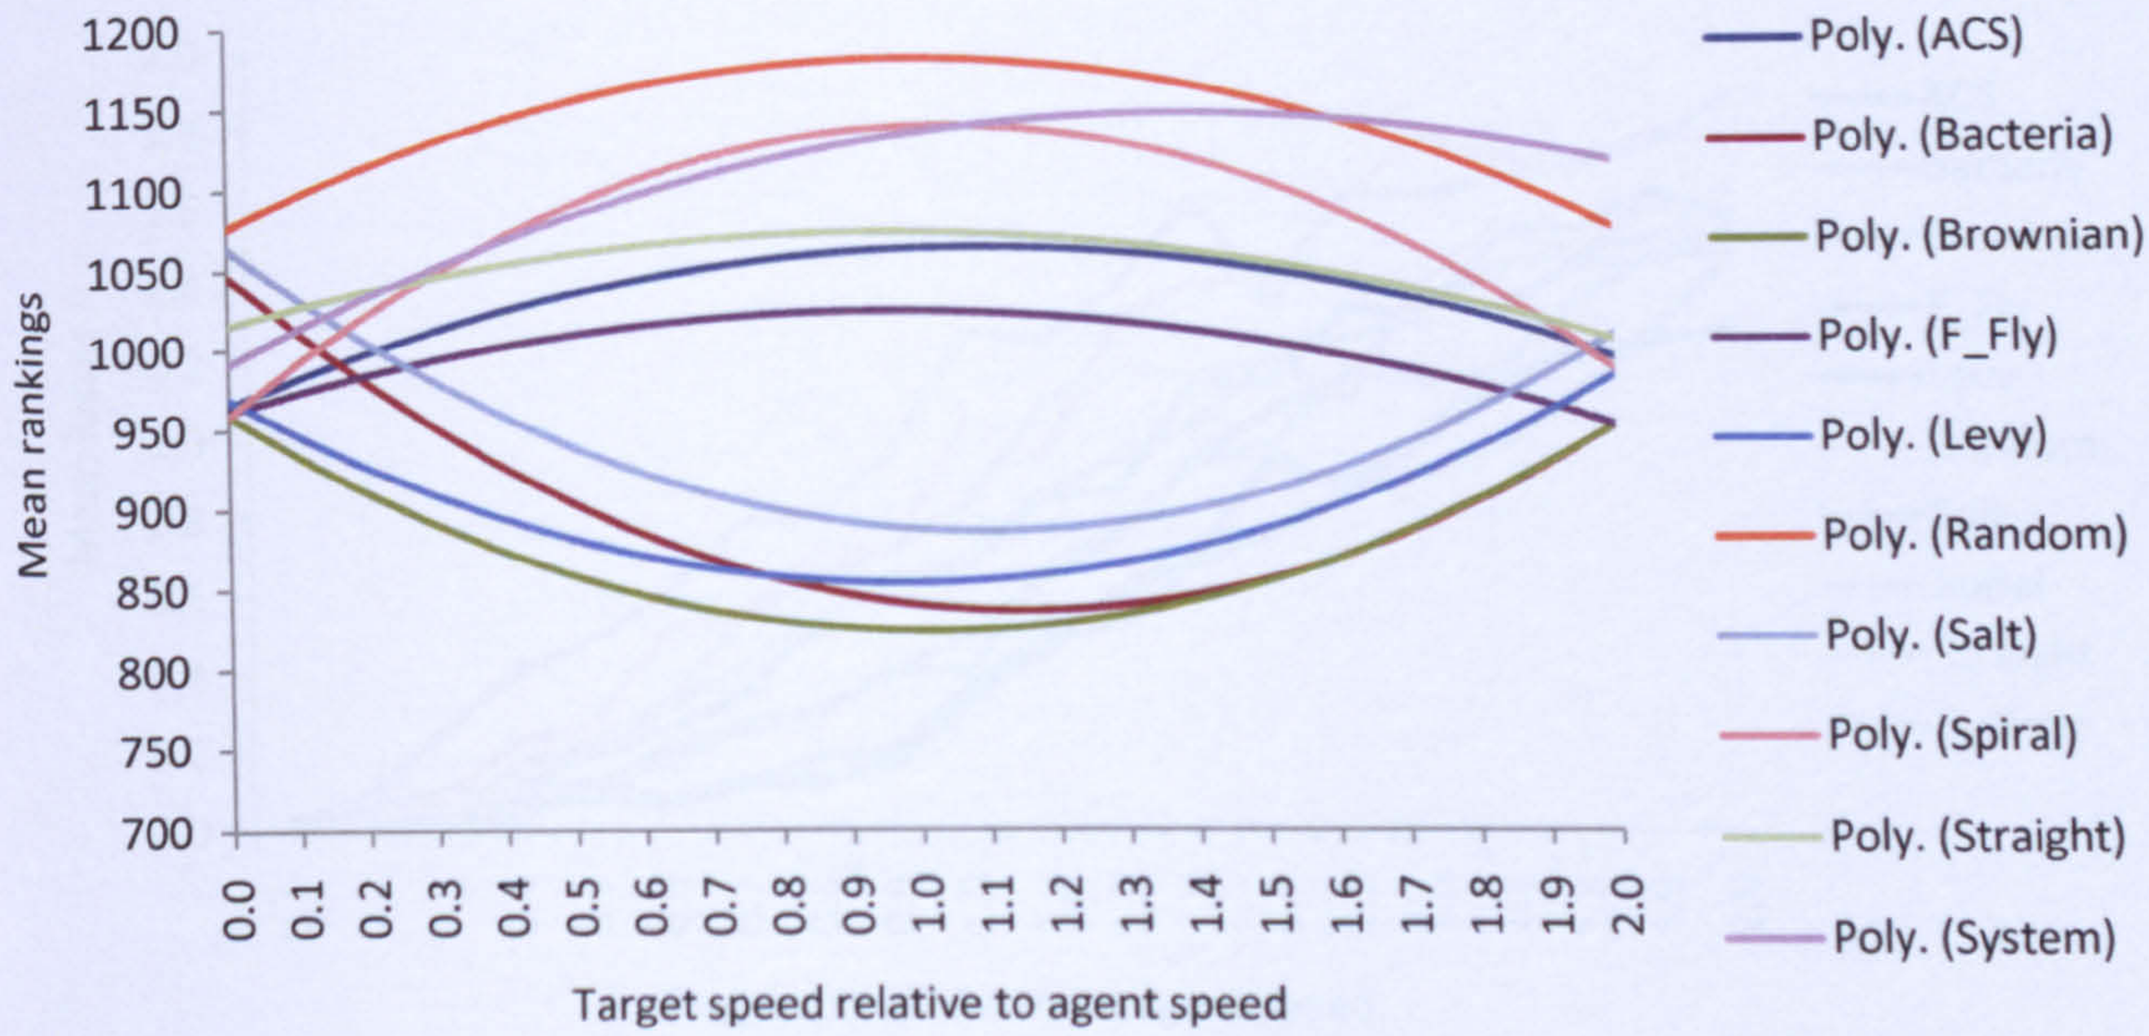

Figure 64. Summary of mean ranking results for 10 agents against 1 target moving deterministically.

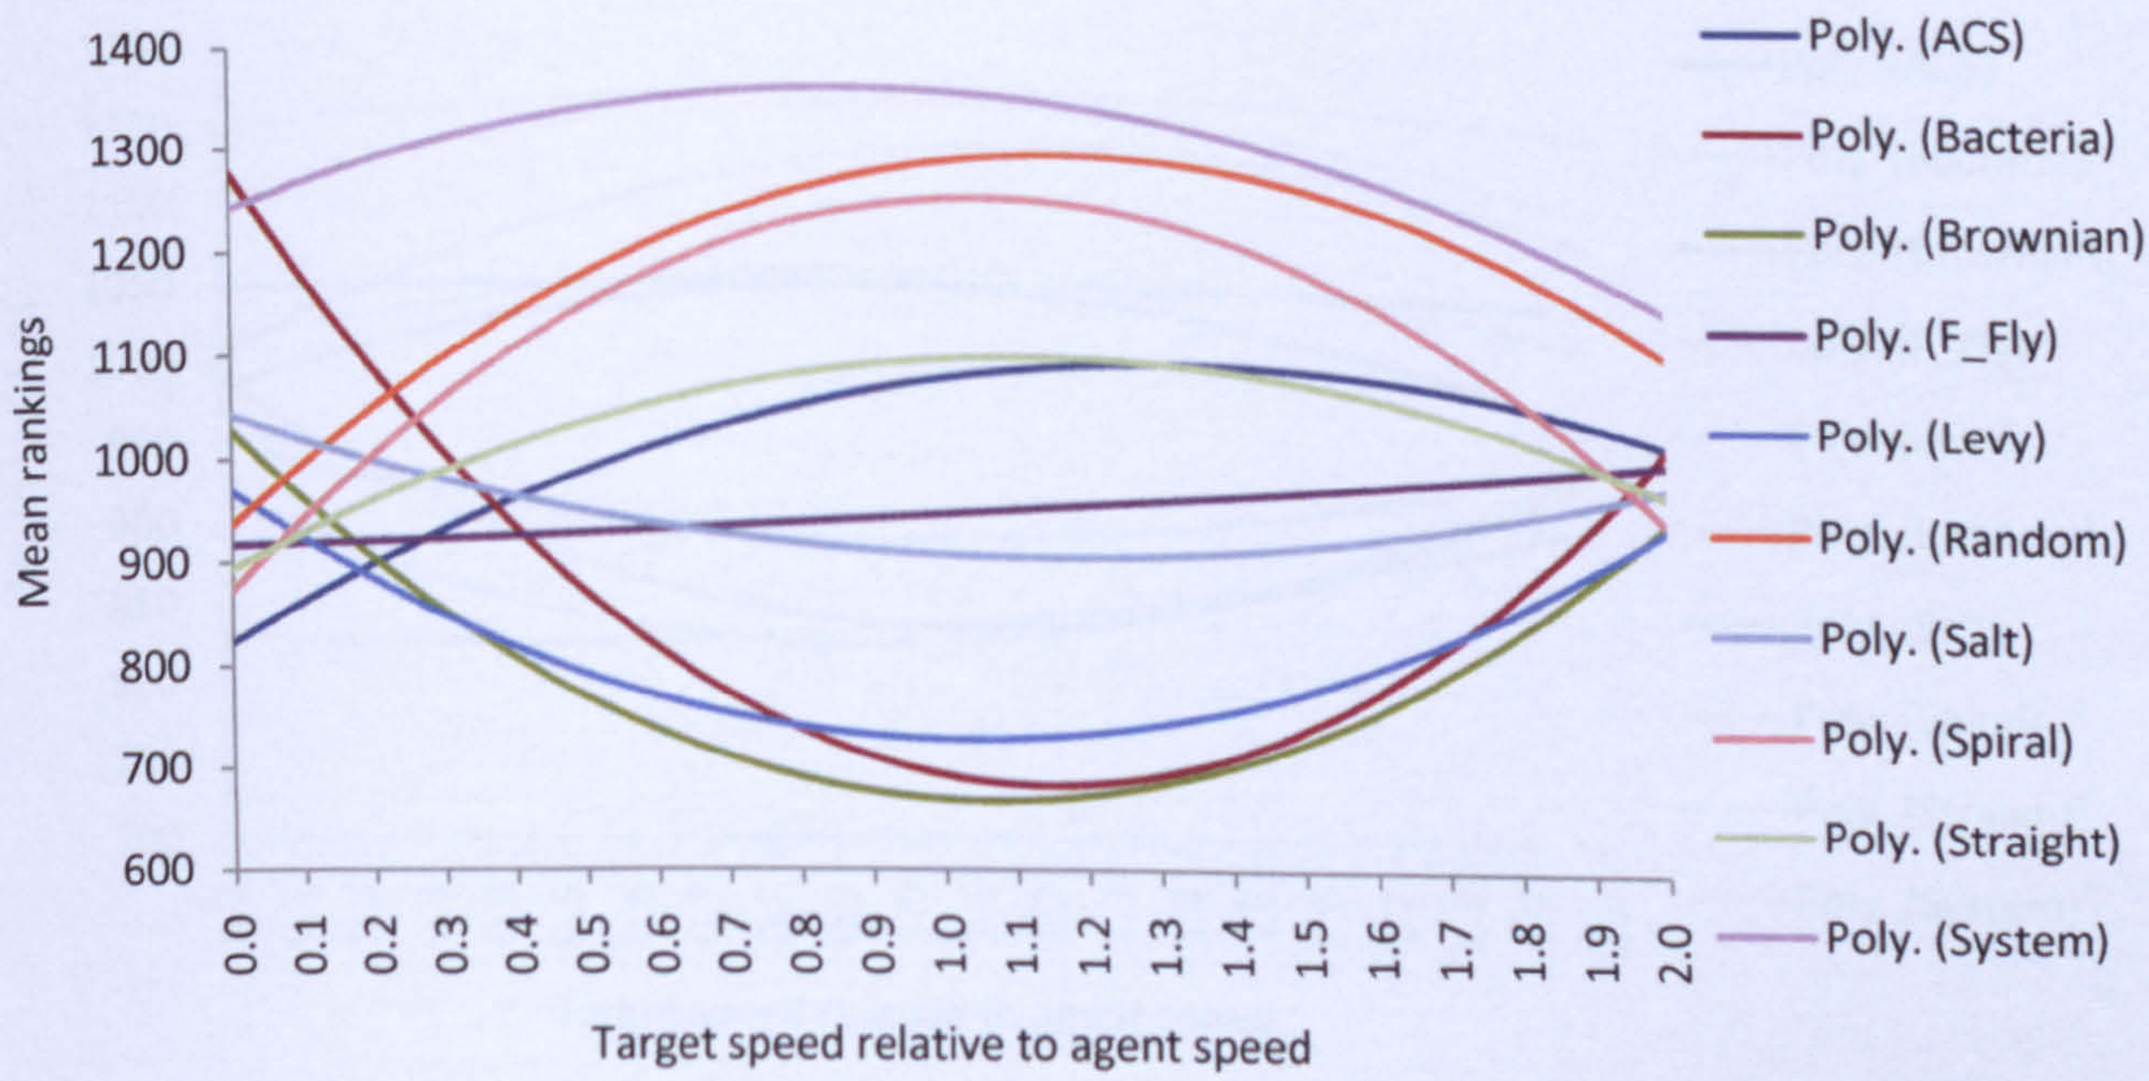

Figure 65. Summary of mean ranking results for 10 agents against 4 targets moving deterministically.

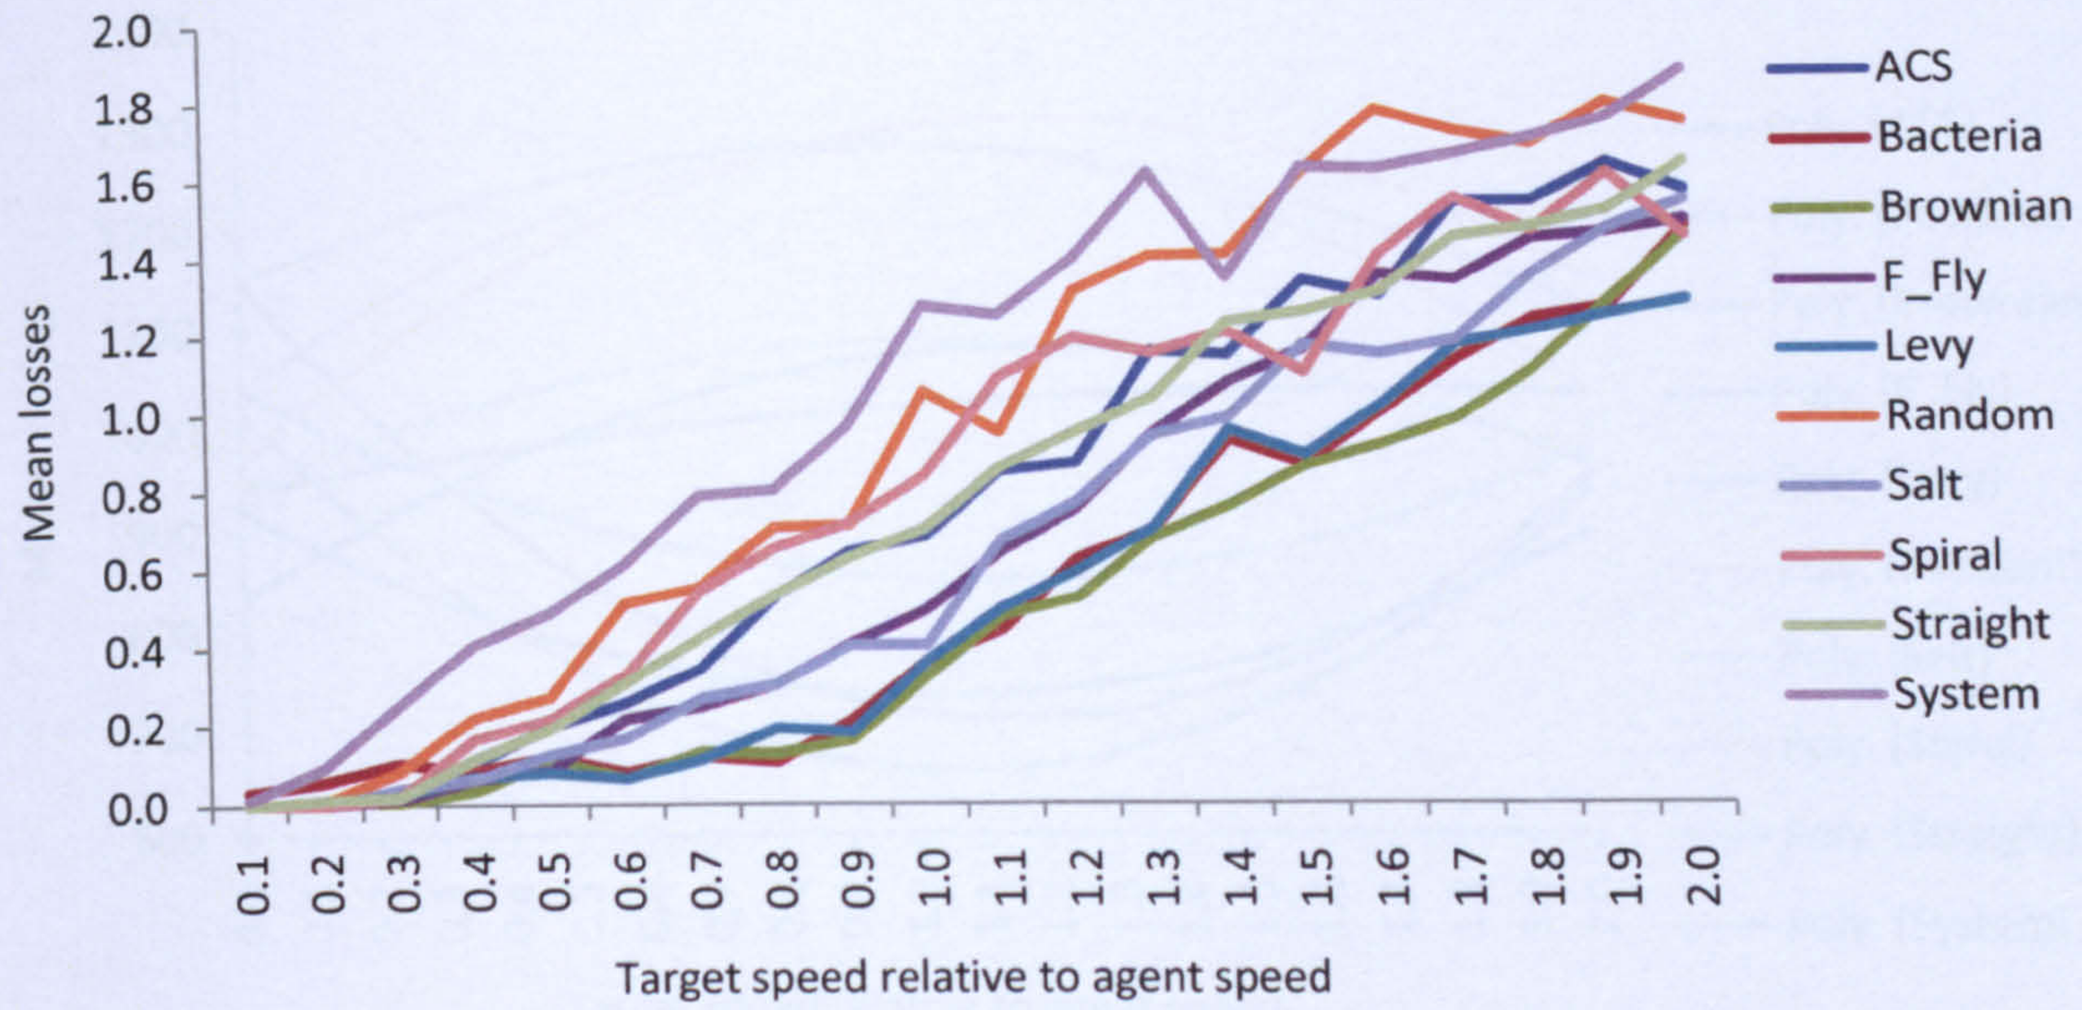

Figure 66. Summary of mean losses results for 10 agents against 4 targets moving deterministically.

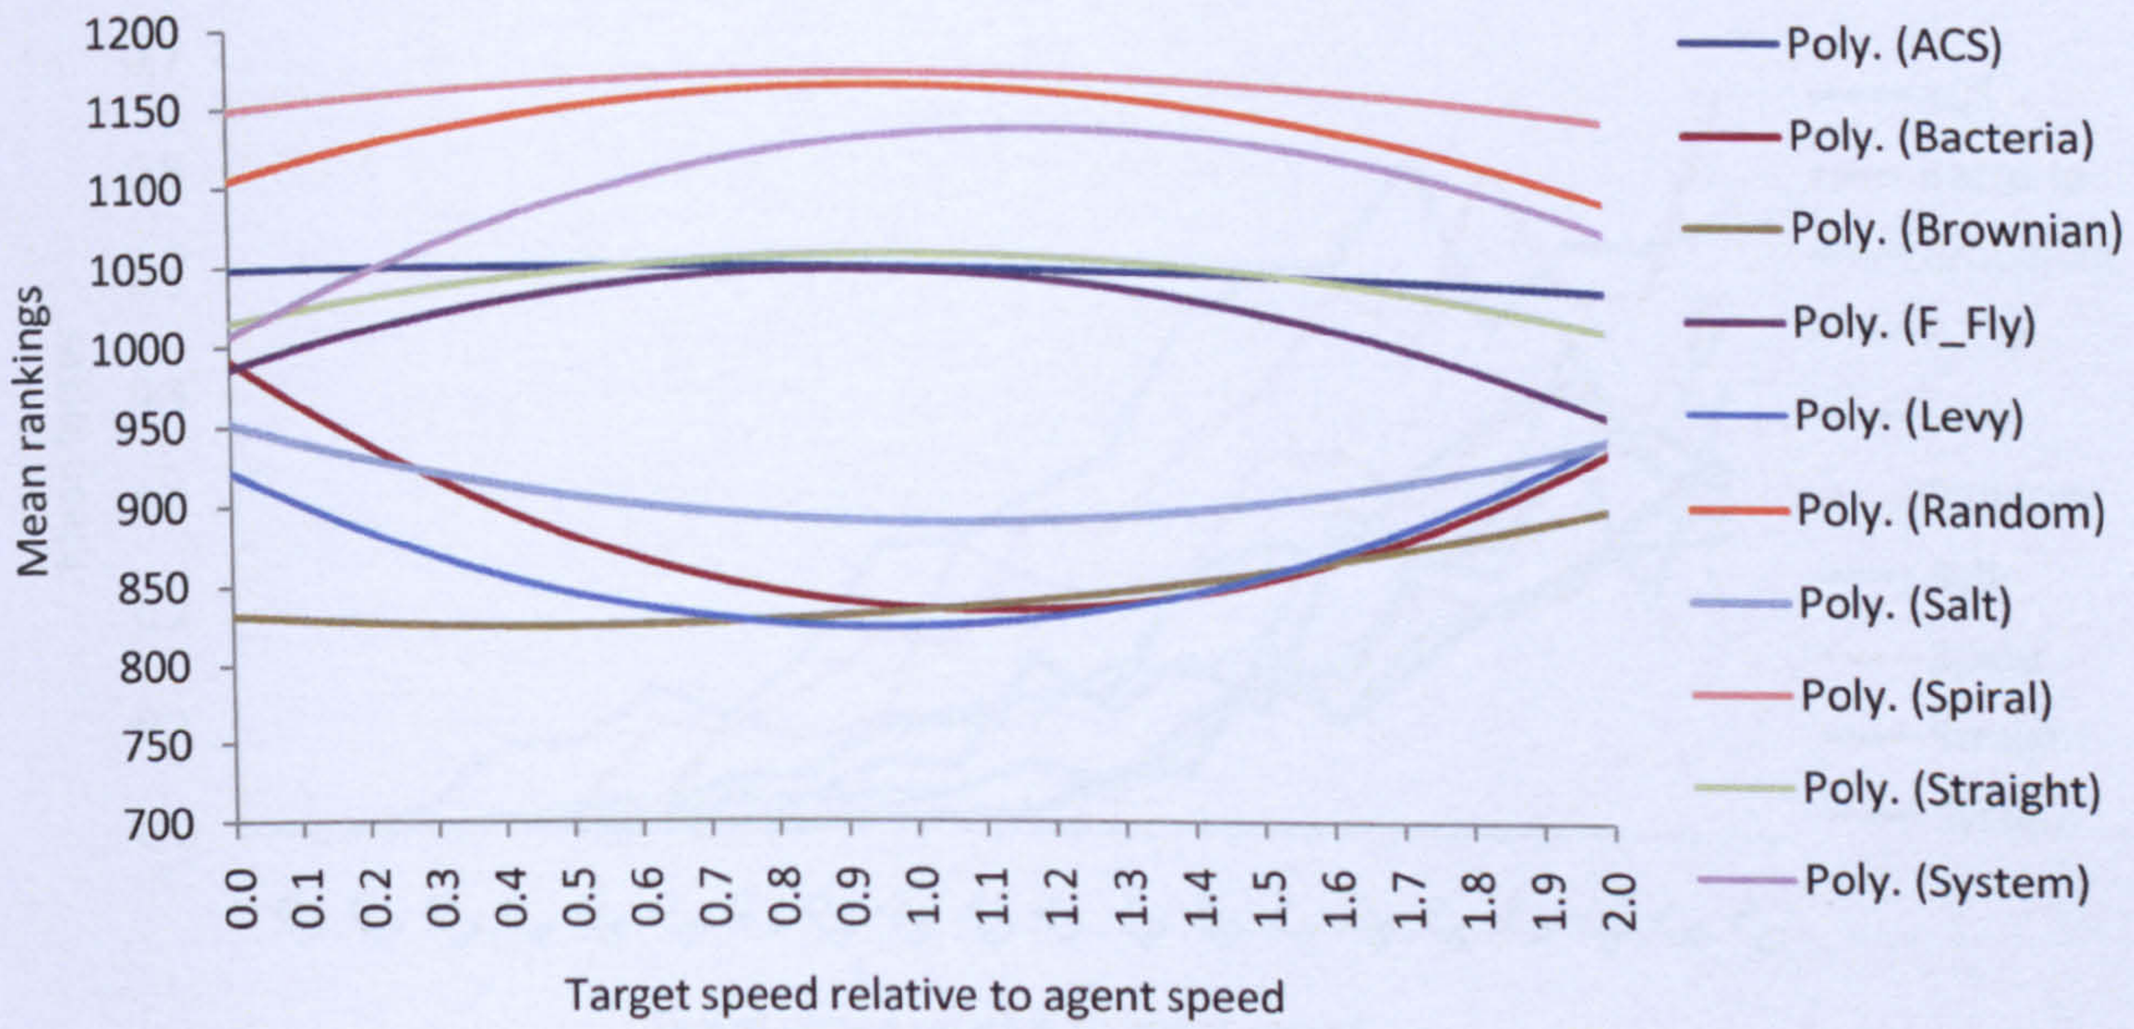

Figure 67. Summary of mean ranking results for 20 agents against 1 target moving deterministically.

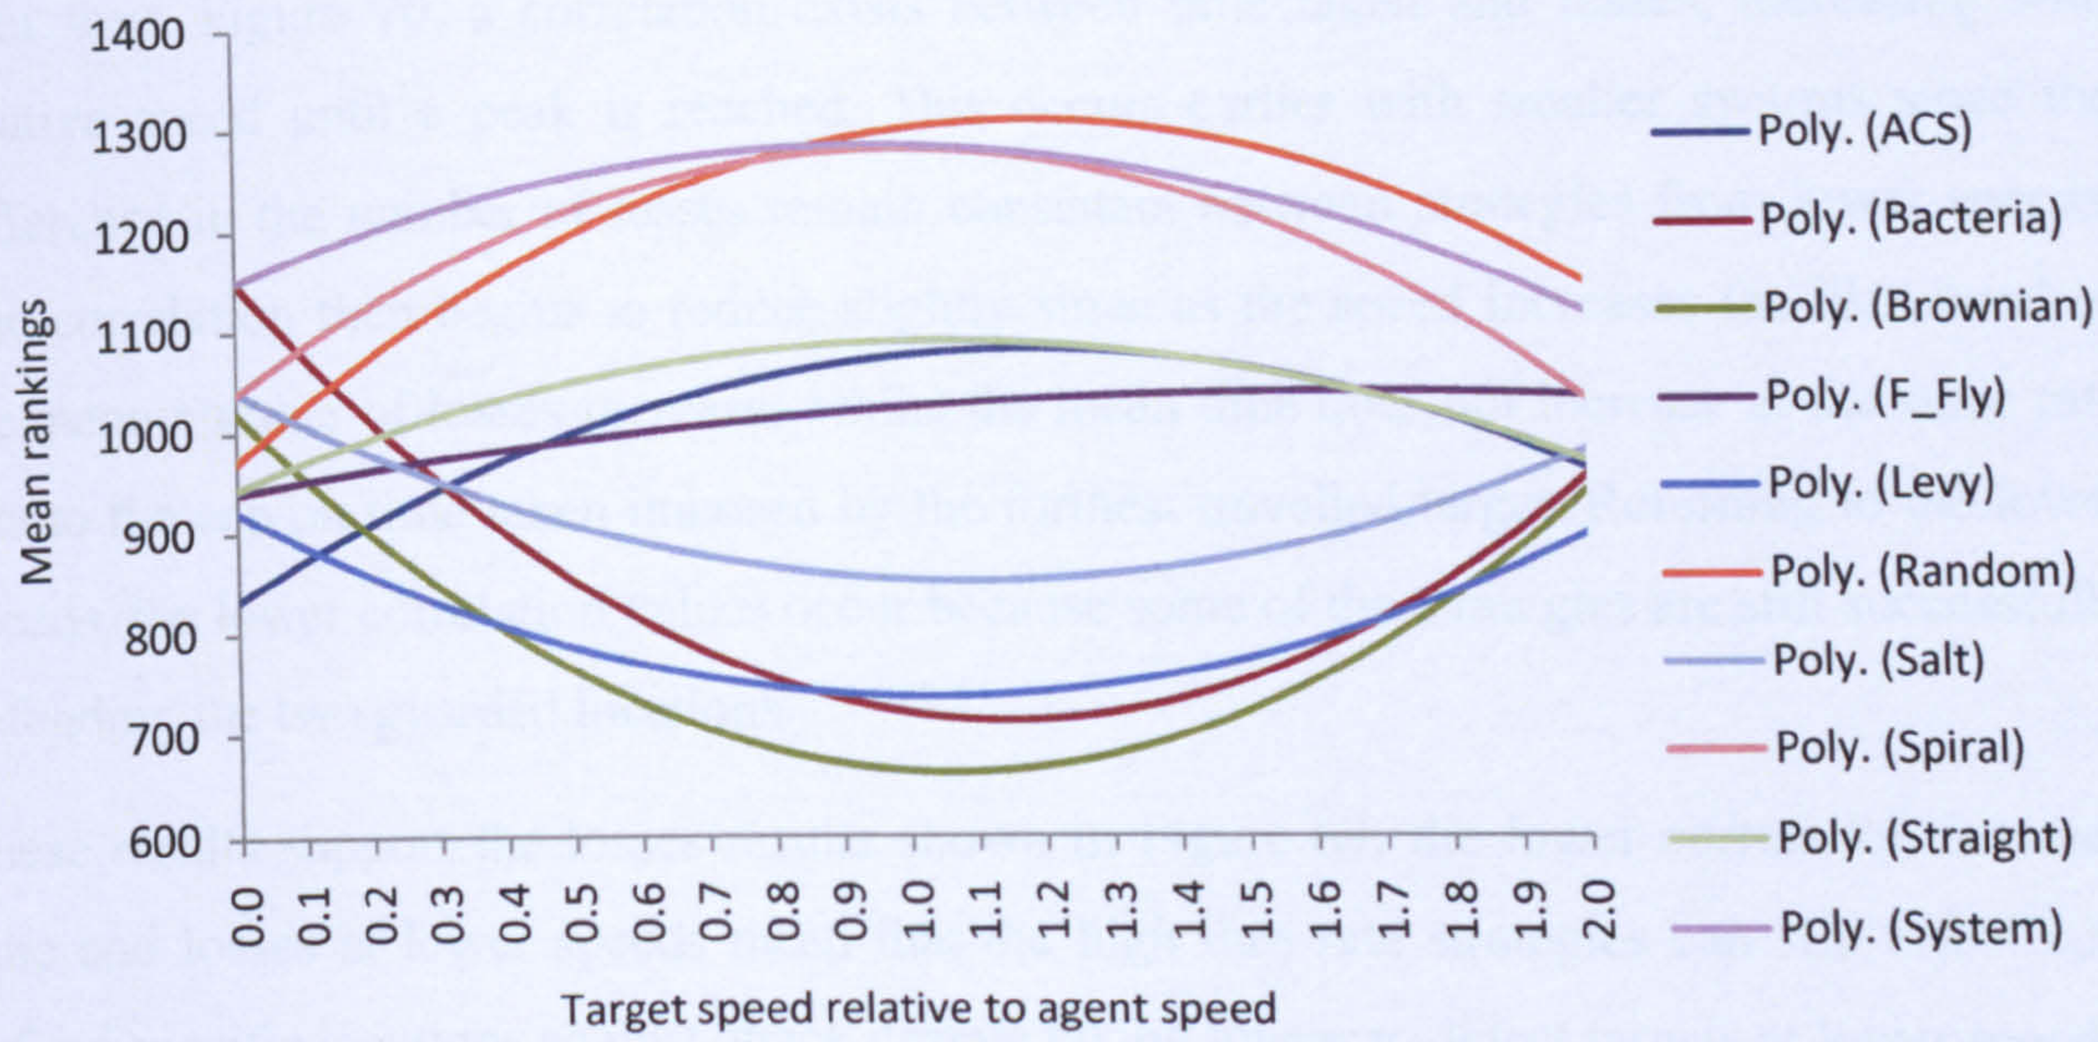

Figure 68. Summary of mean ranking results for 20 agents against 4 targets moving deterministically.

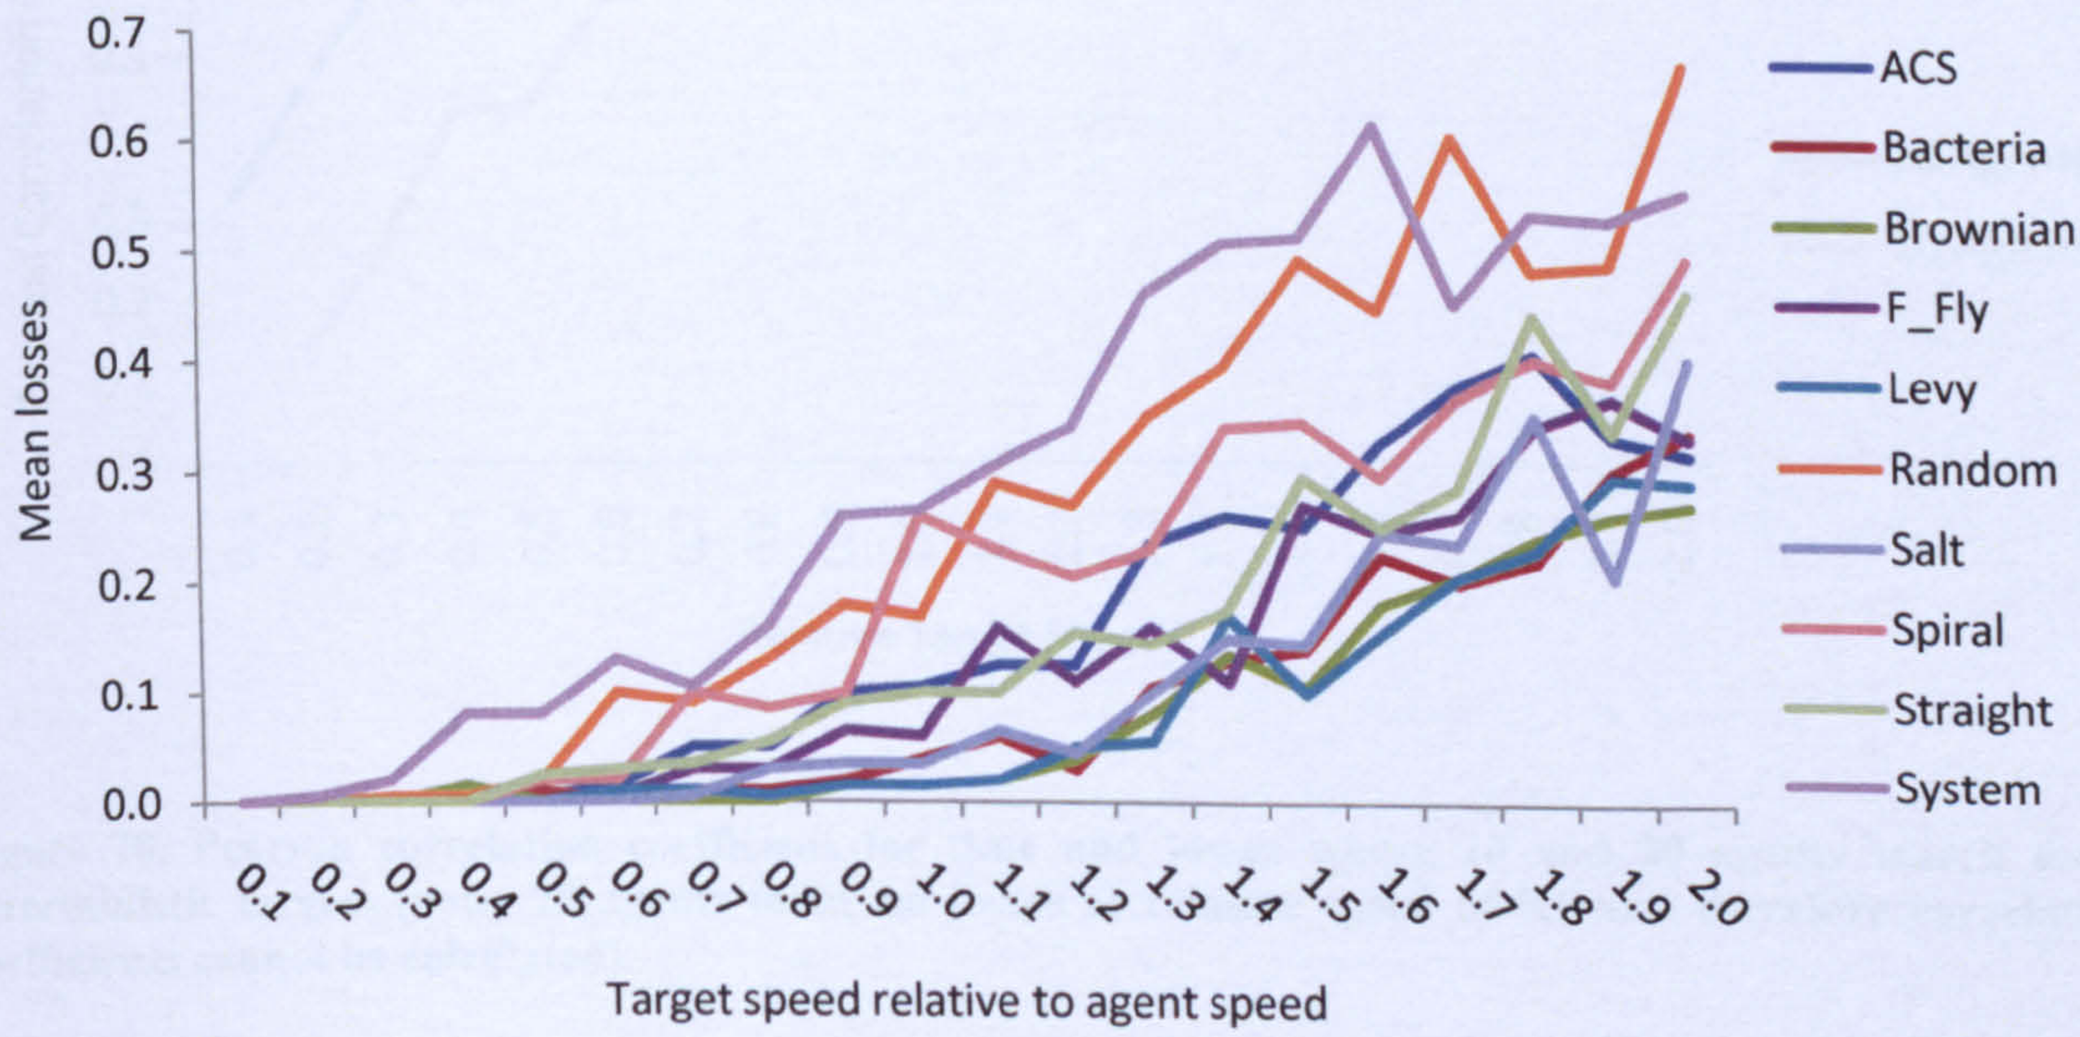

Figure 69. Summary of mean losses results for 20 agents against 4 targets moving deterministically.

Statistical analysis was also performed using Pearson's correlation tests (Pearson, 1896) to determine whether a relationship existed between time and the number of times targets successfully attacked the defended positions using deterministic attack profiles. As can be

seen from Figure 70, a correlation exists between time taken and losses, increasing with relative speed until a peak is reached. This occurs earlier with smaller swarms since the difference in the number of losses remain consistant between strategies from lower speeds. The correlation then begins to reduce slightly since as the speed increases the likelihood of the mean number of losses increases whilst the mean time does not increase at the same rate due to the cap on time taken imposed by the furthest travelled target. Returning to the lower speeds, the lower correlation values occur because some of the strategies are still successfully defending the two guarded locations.

These results support the losses results shown in Figure 69; the lower correlation between time and losses at lower speeds mean that the high turn rate strategies can still effectively defend specific locations against attack despite taking longer to detect targets at lower speeds.

Figure 70. Pearson correlation coefficient for time and losses where 10 and 20 agents search for 4 deterministic targets (note: 20 agents incur no losses at relative speed of 0.1 and therefore correlation coefficients cannot be calculated).

5.4.2 Non-deterministic Targets

Where the targets moved randomly (Figure 71), the increased performance of the high turn rate strategies was less pronounced, and advantages over ranging strategies did not occur until the targets were moving at just over half the speed of the agents, after which relative performance generally held steady regardless of speed. This effect was again amplified with more targets (Figure 72) because at lower speeds the agents had to wait longer to encounter

the targets. At the faster end of the relative speed scale the effect of chance meetings is evident from the re-converging of the strategies' performances although it should be noted that even though the ranging behaviour of the targets had increased such that encounters were partially a result of chance, the likelihood of detection was still increased in strategies such as Brownian, bacterial and Lévy searches due to their high turn rates. This last point is further evidenced where the swarm size is increased (Figure 73). In this scenario, the high turn rate strategies perform better from quite low speeds since the systemic coverage of the area is sufficiently high that targets do not have to travel far before an encounter occurs, although this did not occur at such low speeds where more targets were present (Figure 74) since there was an increased likelihood of at least one of the targets being in an area sparsely populated by agents.

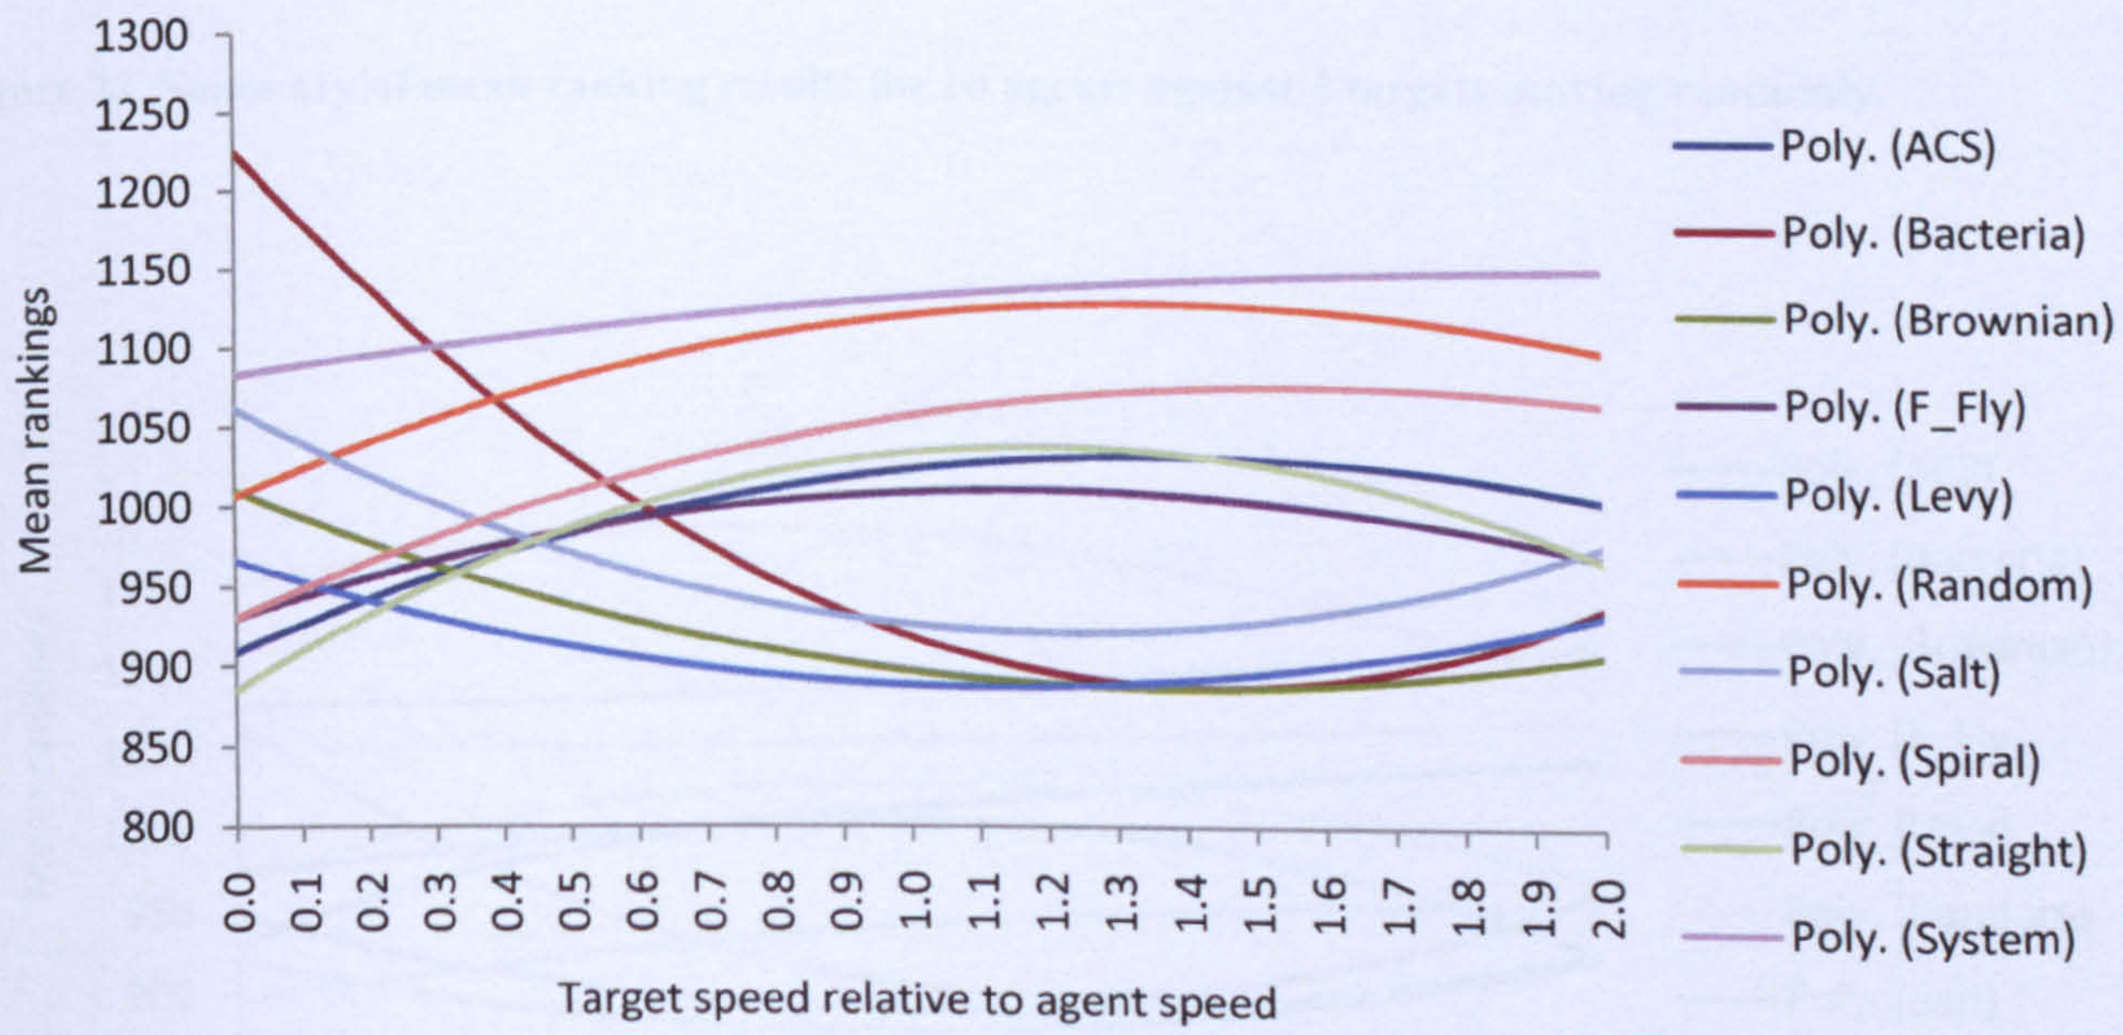

Figure 71. Summary of mean ranking results for 10 agents against 1 target moving randomly.

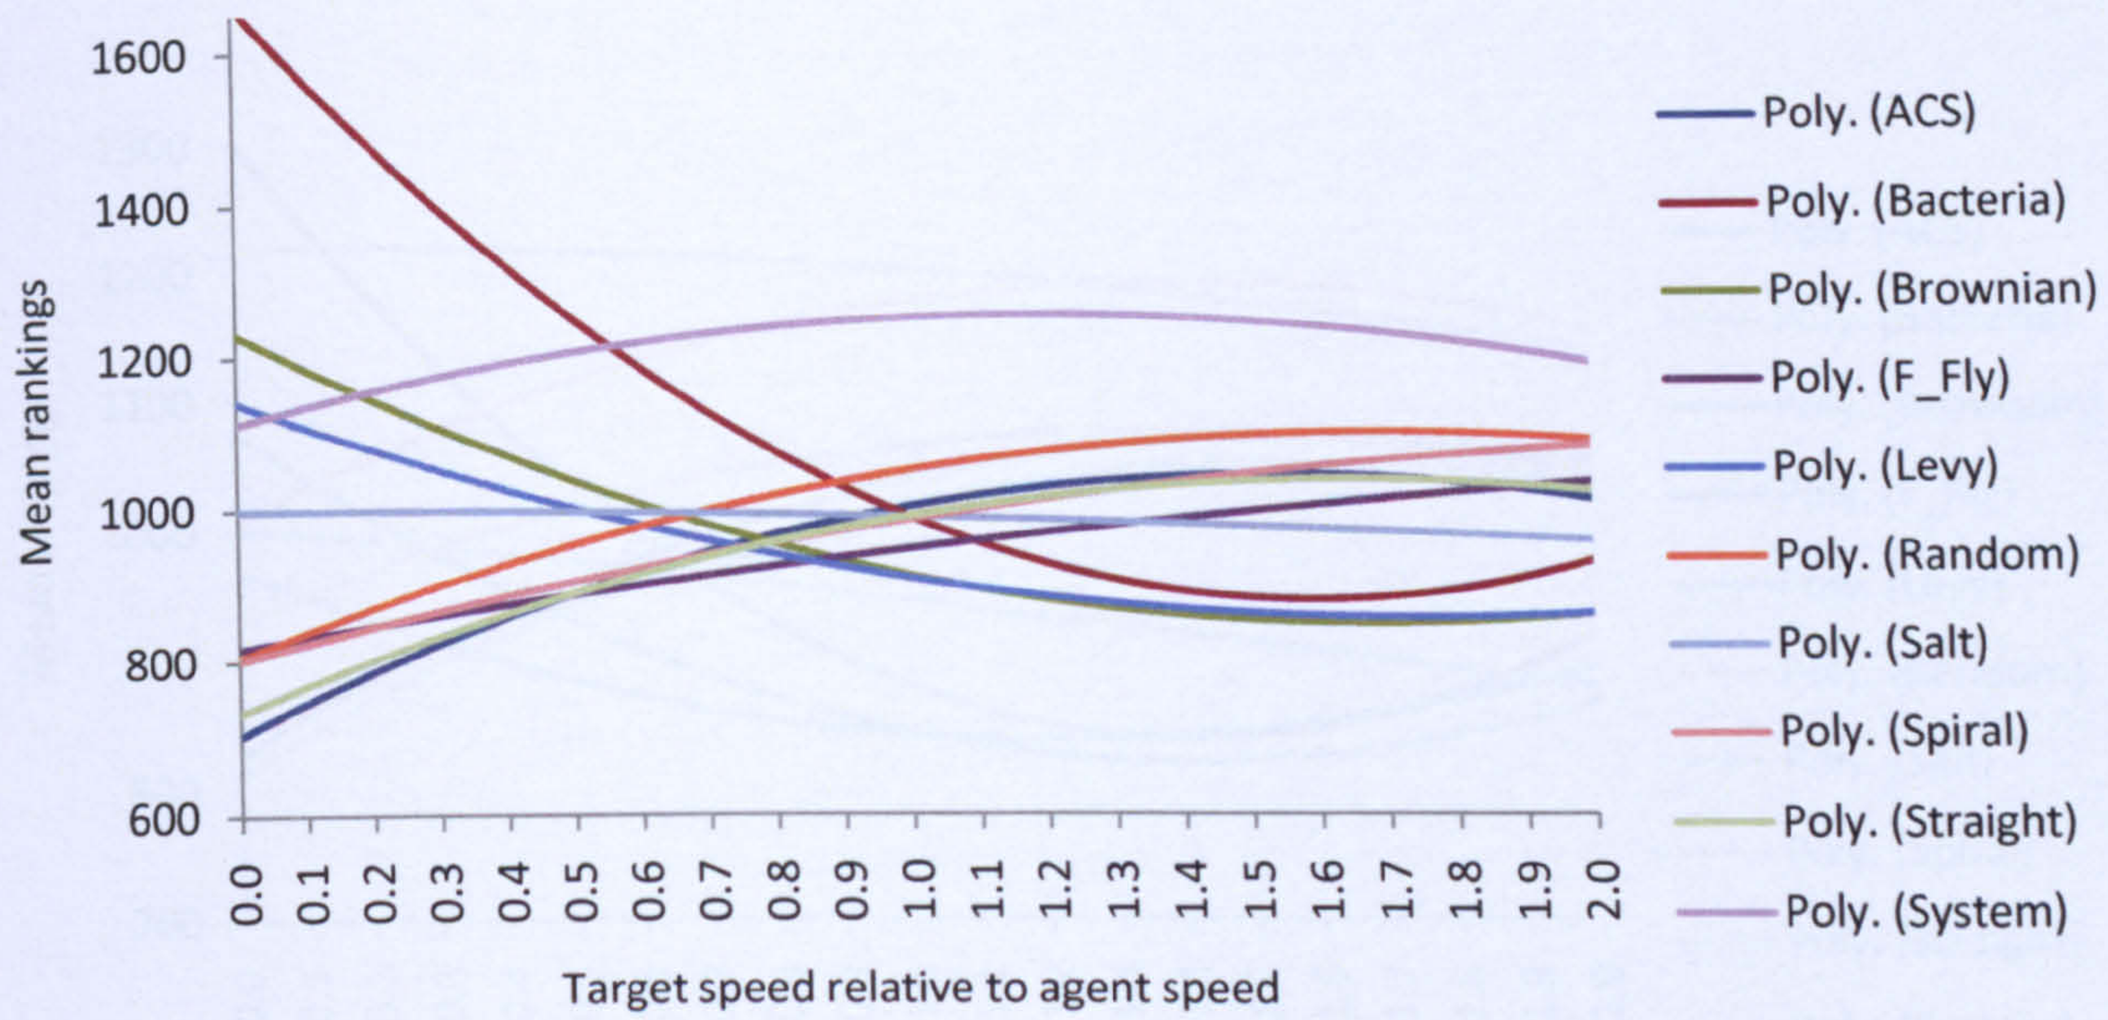

Figure 72. Summary of mean ranking results for 10 agents against 4 targets moving randomly.

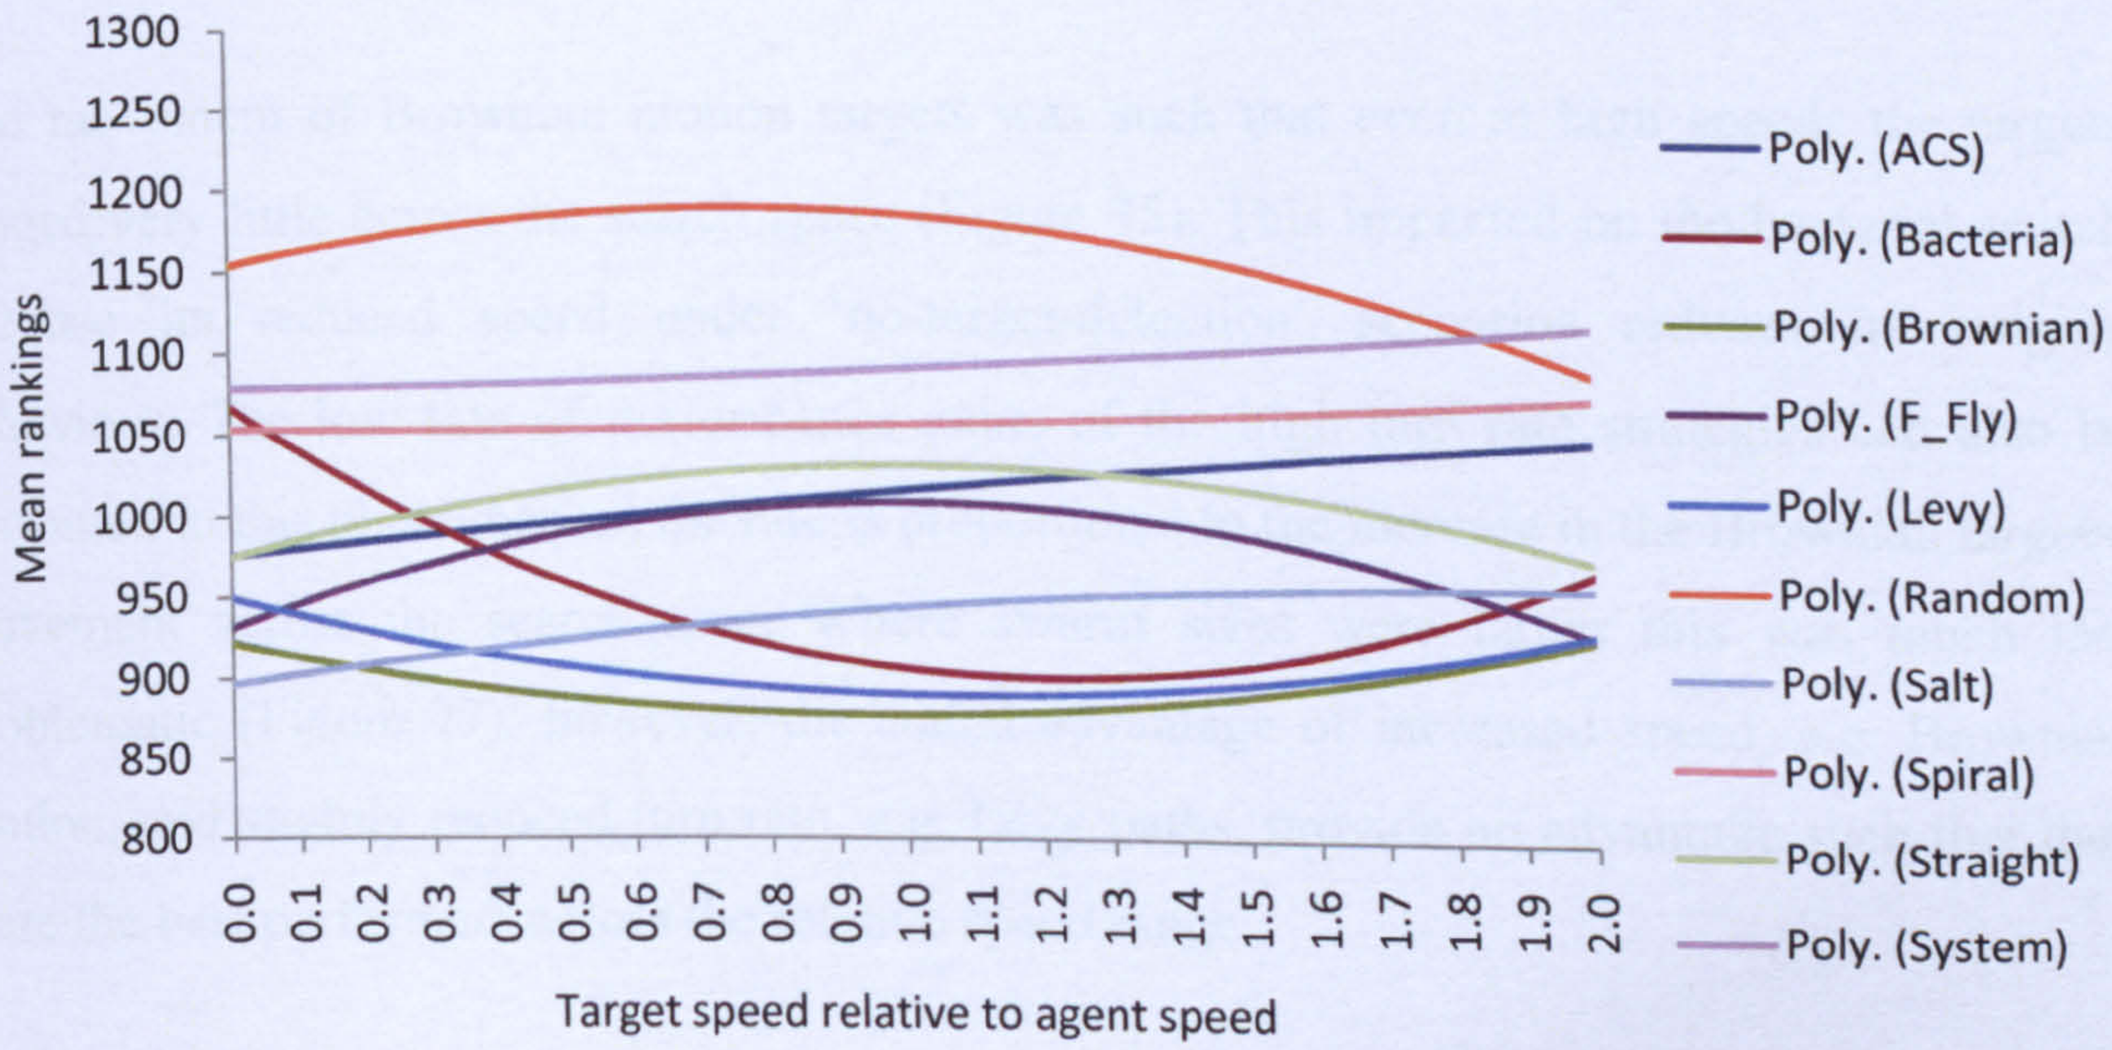

Figure 73. Summary of mean ranking results for 20 agents against 1 target moving randomly.

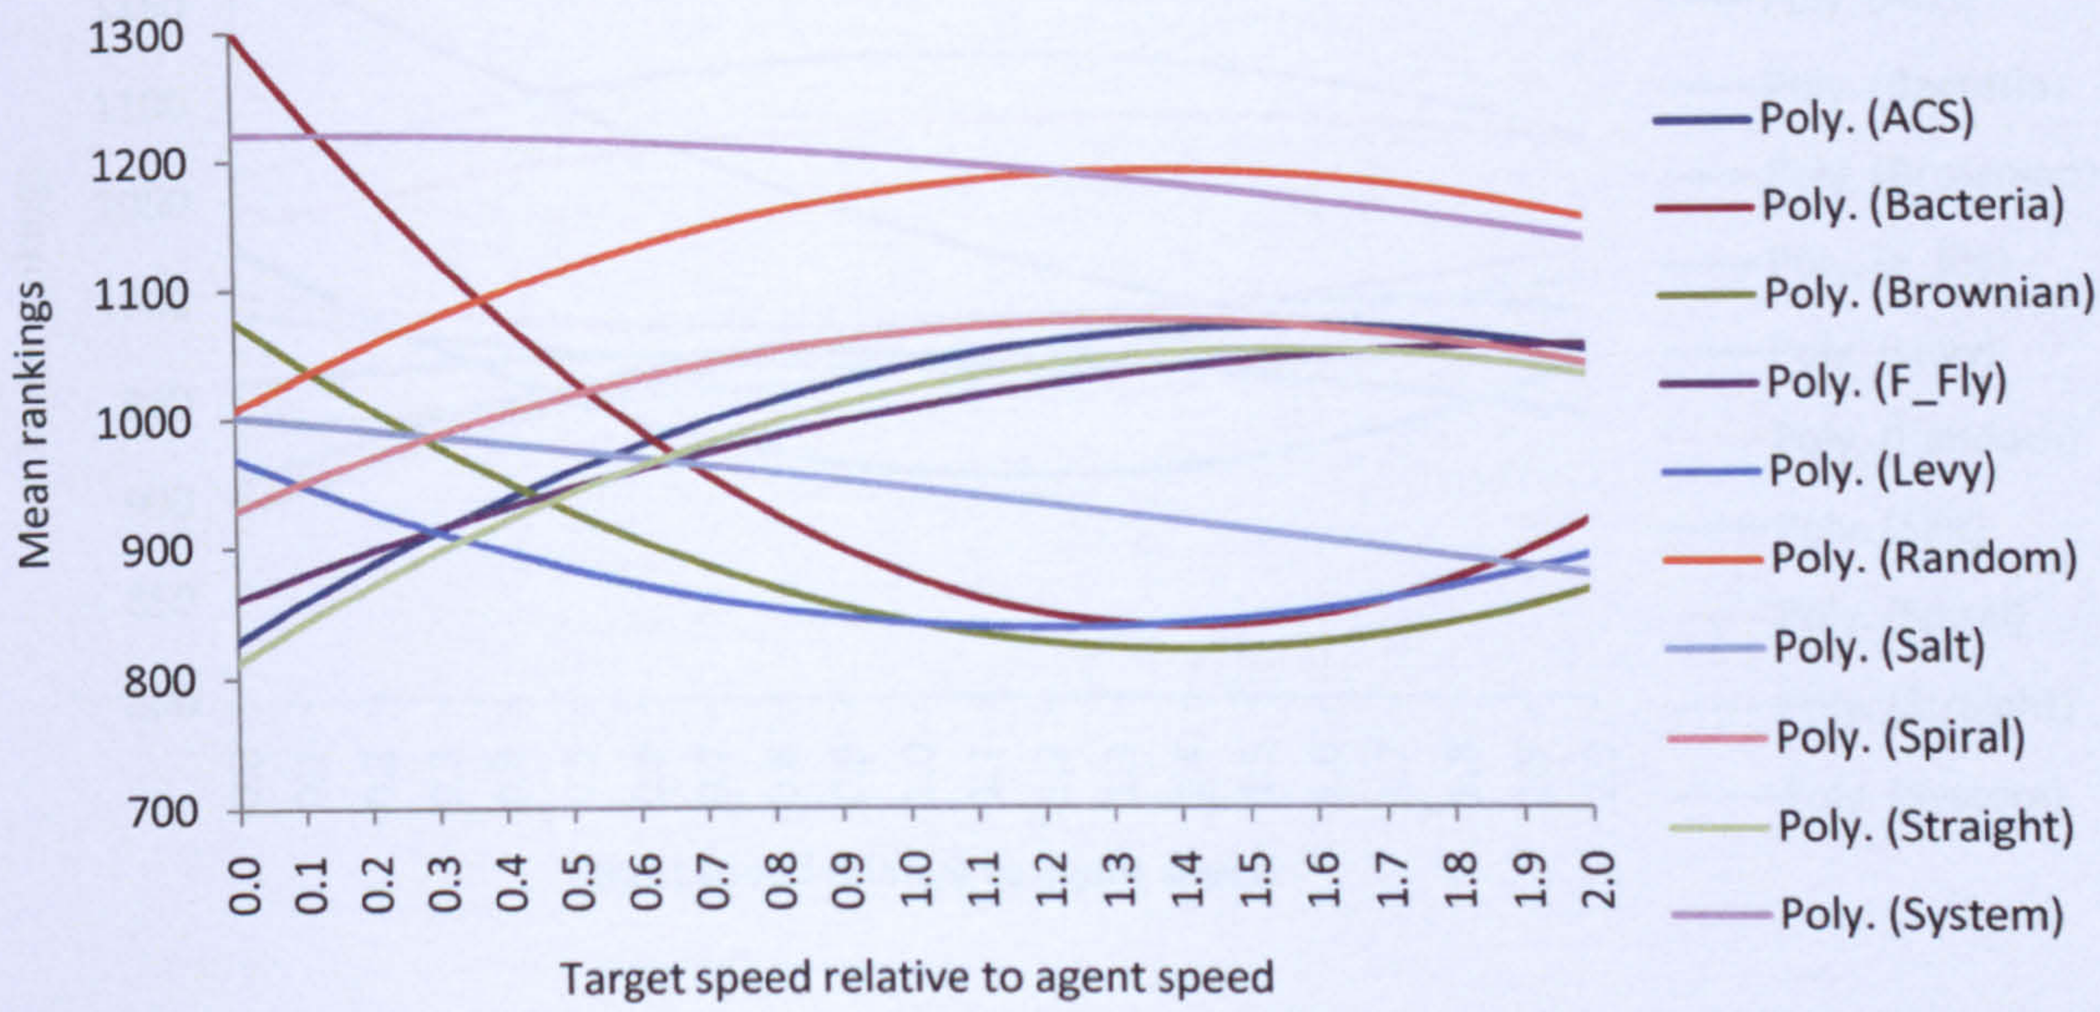

Figure 74. Summary of mean ranking results for 20 agents against 4 targets moving randomly.

The movement of Brownian motion targets was such that even at high speeds the targets ranged very little across the search space (Figure 75). This impacted on the bacterial search because its reduced speed under 'no-target-detection' scenarios reduces its ranging behaviour. The low rate of performance gains of the high turn rate strategies can also be attributed to this phenomenon; the rate is proportional to the increase in the Brownian targets' movement across the search area. Where swarm sizes were larger this was much less problematic (Figure 77); however, the added advantage of increased speed, *e.g.* Brownian motion, and slightly reduced turn rate, *e.g.* Lévy paths, provide an advantage such that they were the best performers across the relative speed range.

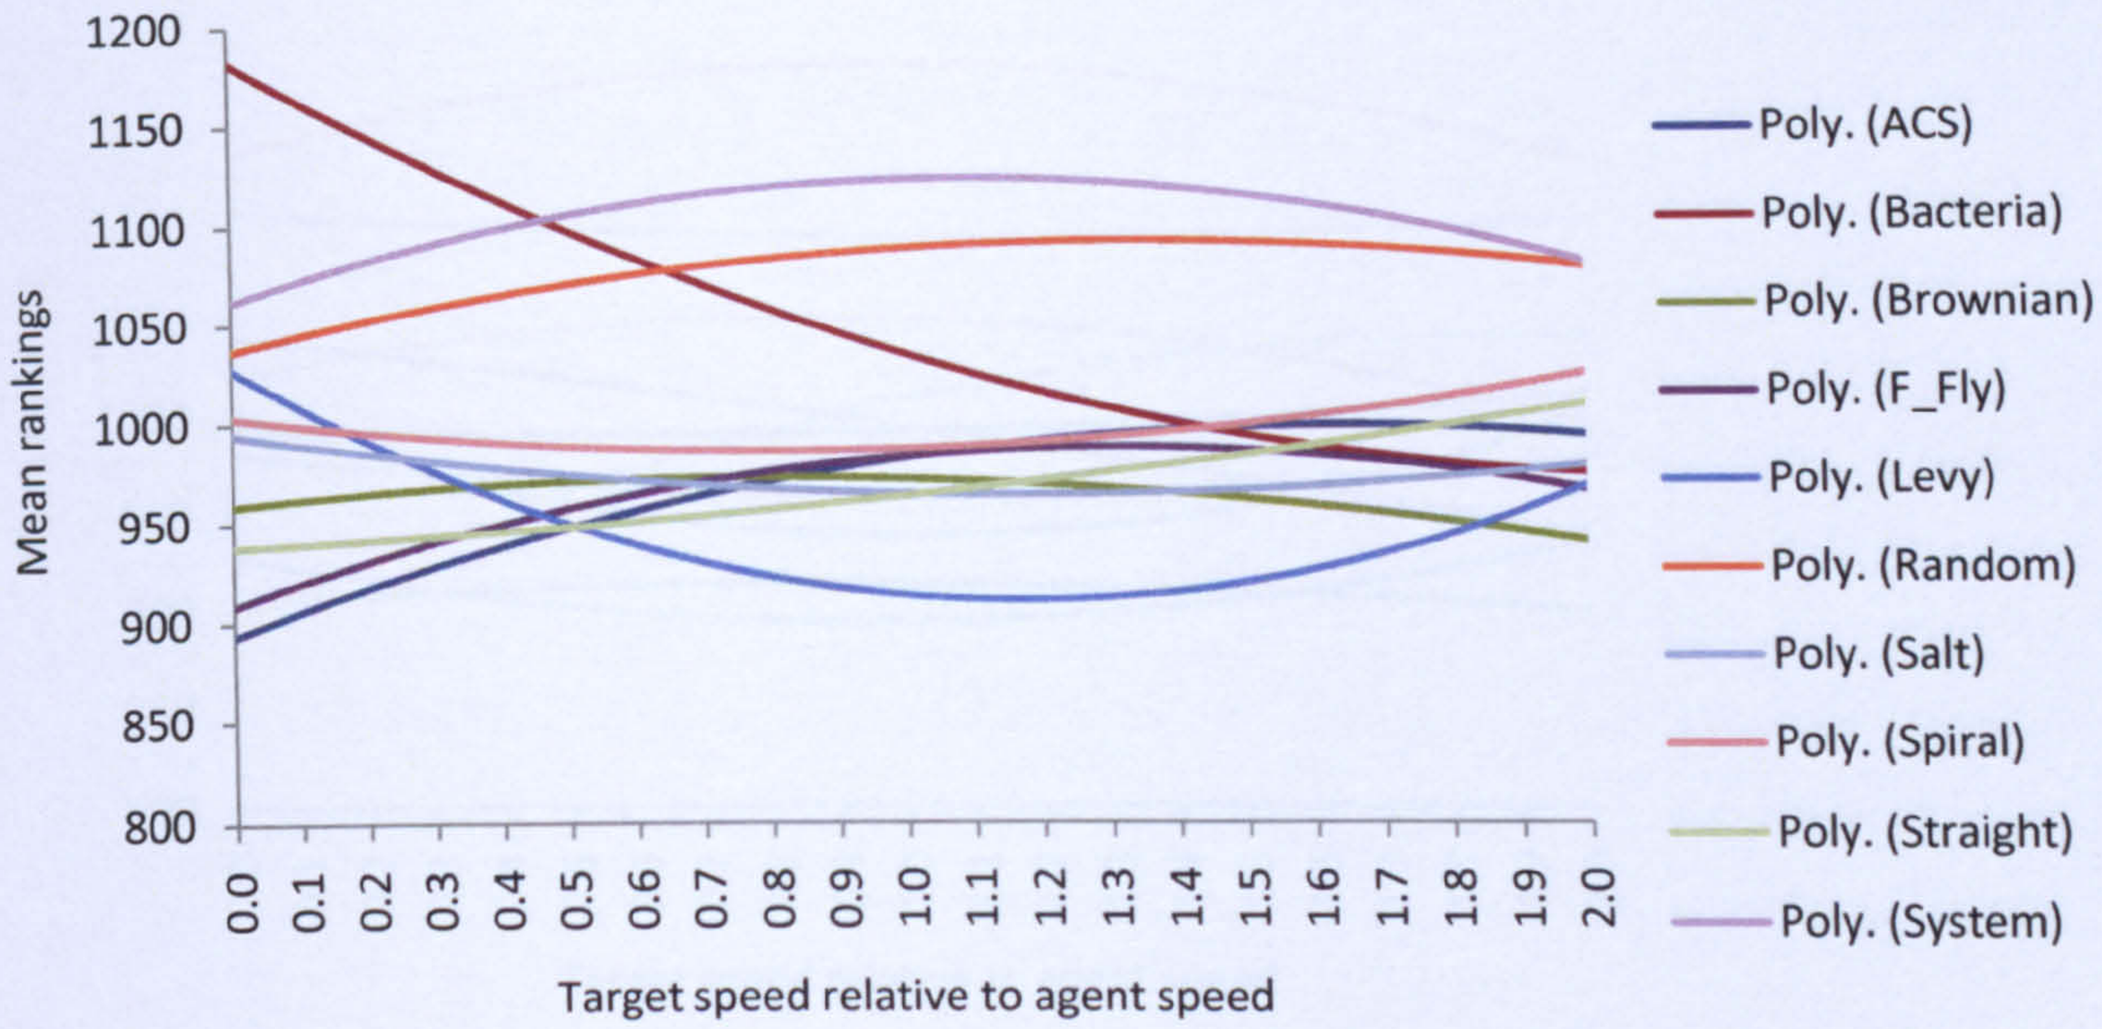

Figure 75. Summary of mean ranking results for 10 agents against 1 Brownian motion target.

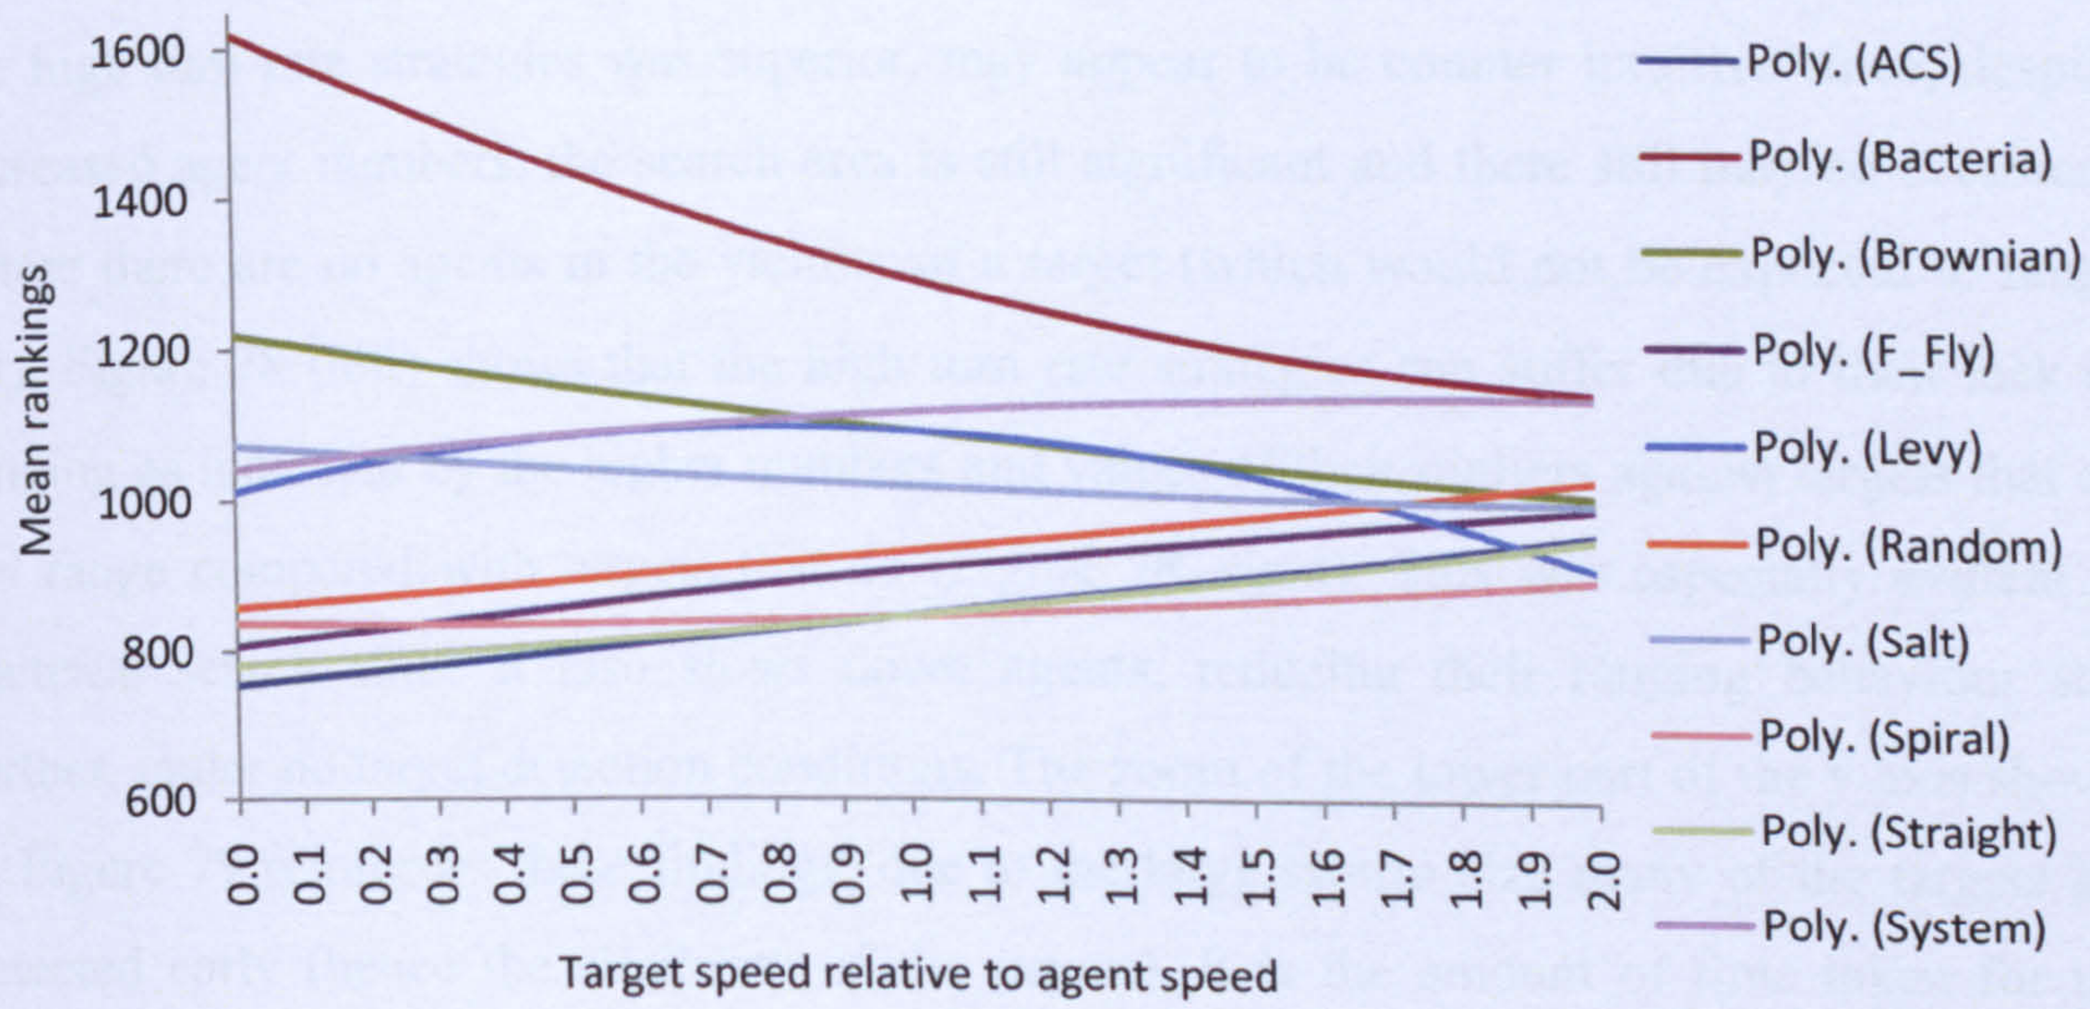

Figure 76. Summary of mean ranking results for 10 agents against 4 Brownian motion targets.

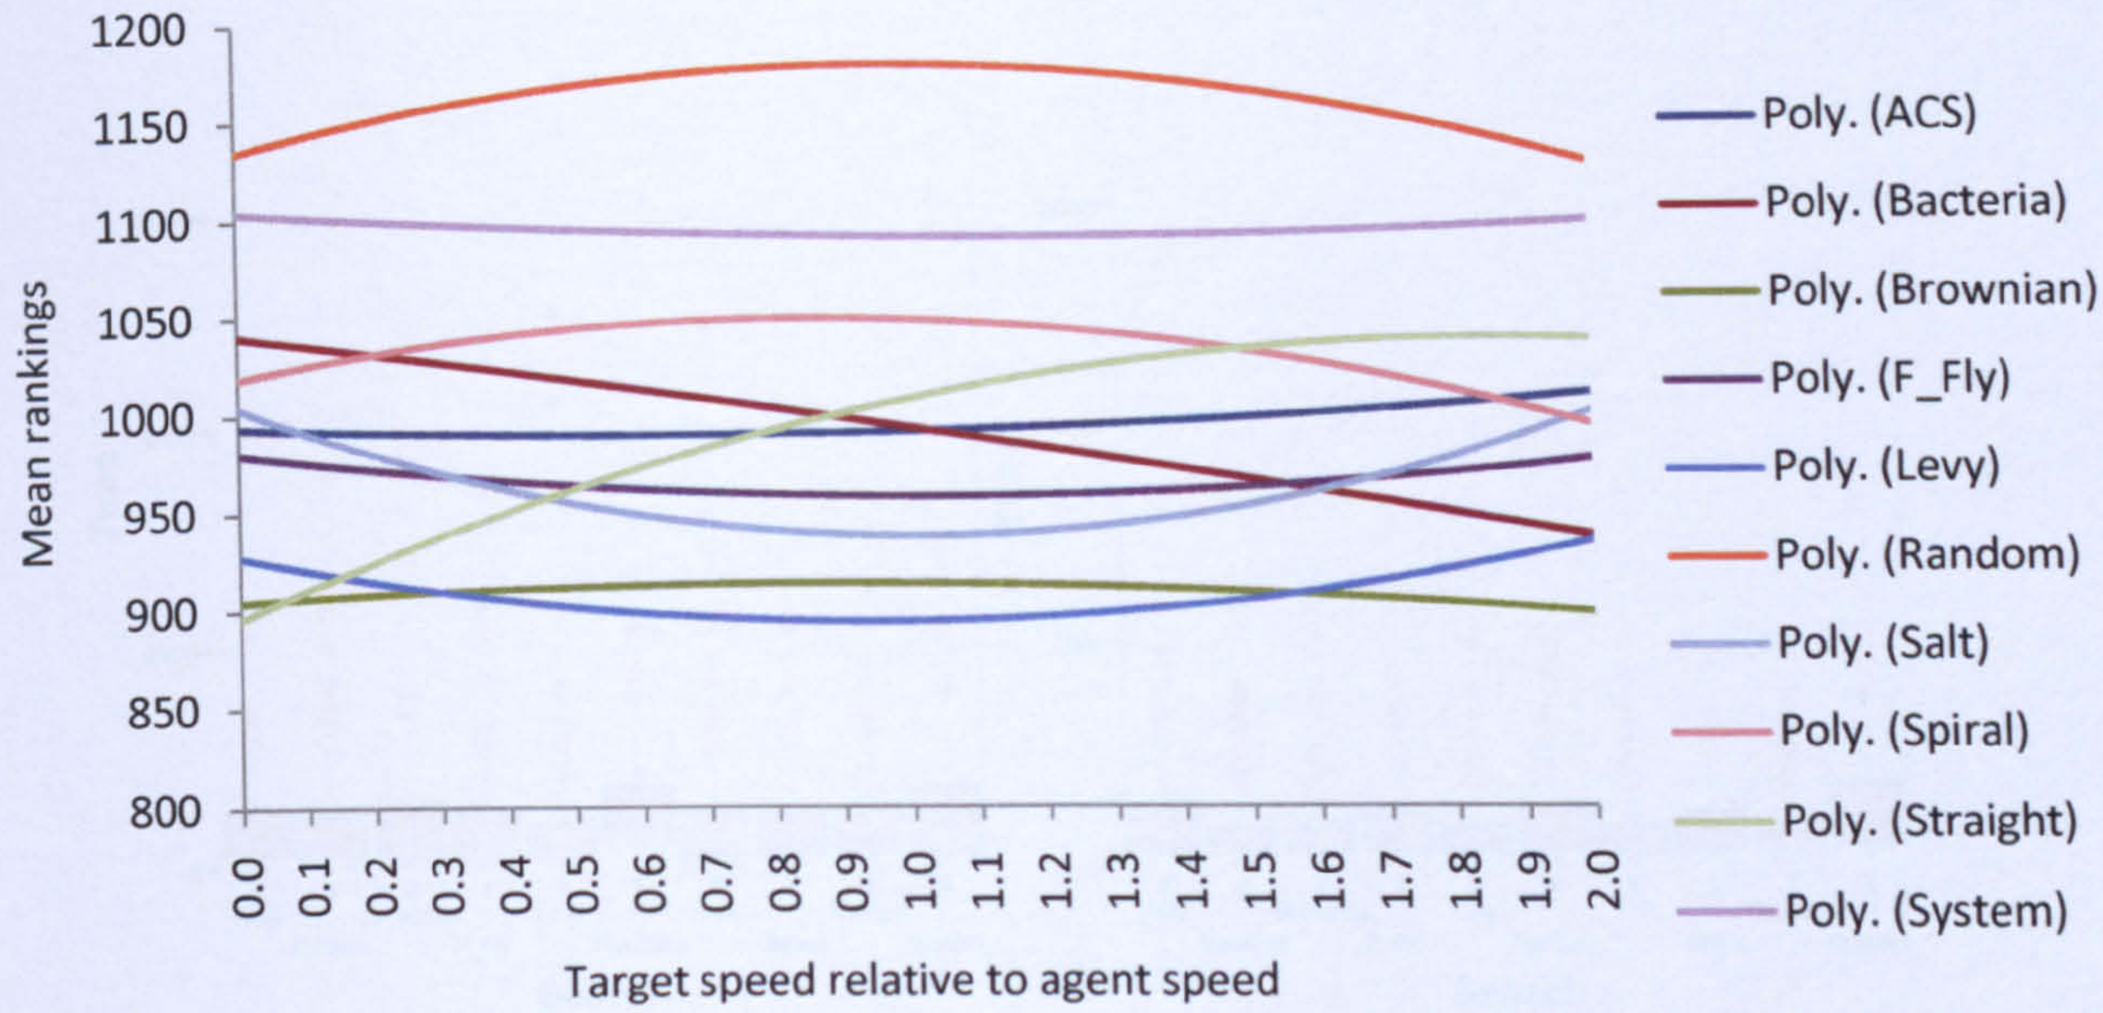

Figure 77. Summary of mean ranking results for 20 agents against 1 Brownian motion target.

The results shown in Figure 77, whilst indicating that the general comparative performance of the high turn rate strategies was superior, may appear to be counter intuitive since, despite increased agent numbers, the search area is still significant and there still may be occasions where there are no agents in the vicinity of a target (which would not be expected to range far). Figure 78 (left) shows that the high turn rate strategies can suffer due to their lack of ranging as indicated by the higher numbers and values of their outliers against targets that do not range compared with targets that do (Figure 78, right). This was especially evident in bacterial search since it also slows down agents, reducing their ranging behaviour still further, under no target detection conditions. The zoom of the lower part of the y axis shown in Figure 79 reinforces these findings; due to the large swarm size many of the targets are detected early (hence the similarity of the means), it is the amount of time taken for the targets, which are not close to any agents, and the agents to move together that makes the difference in the strategies' actual performance.

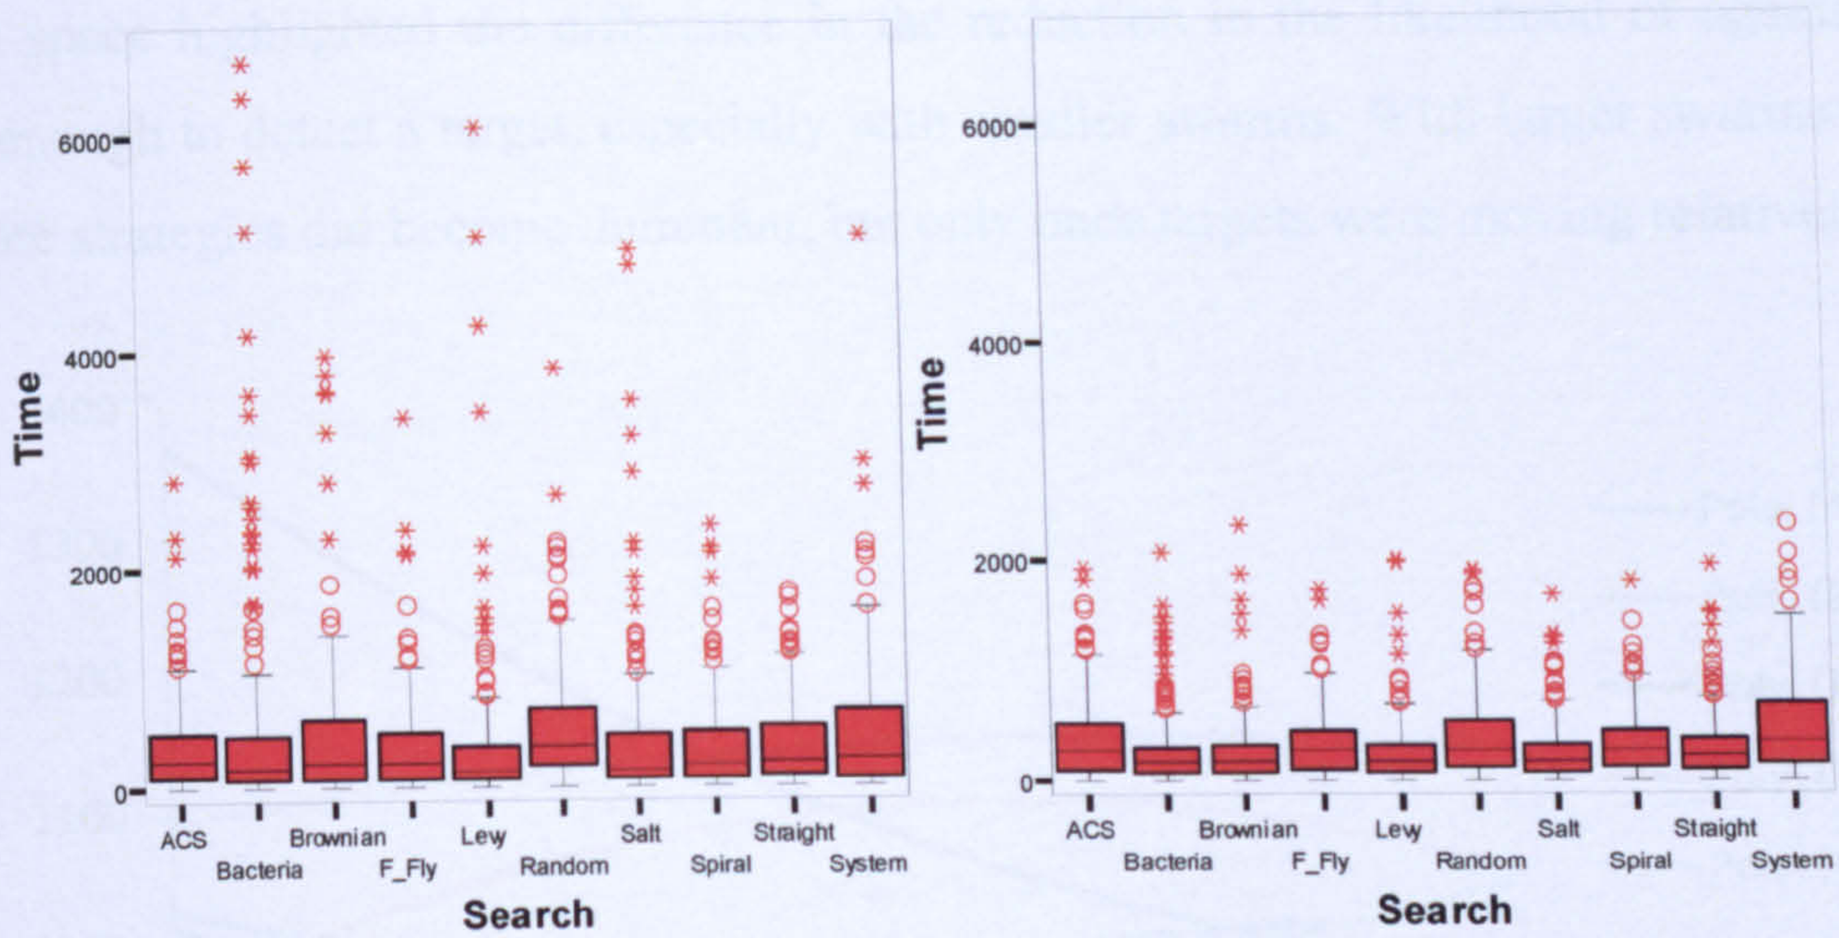

Figure 78. Actual results of 20 agents against 1 target travelling at maximum agent speed (left: Brownian motion, right: fruit fly movement).

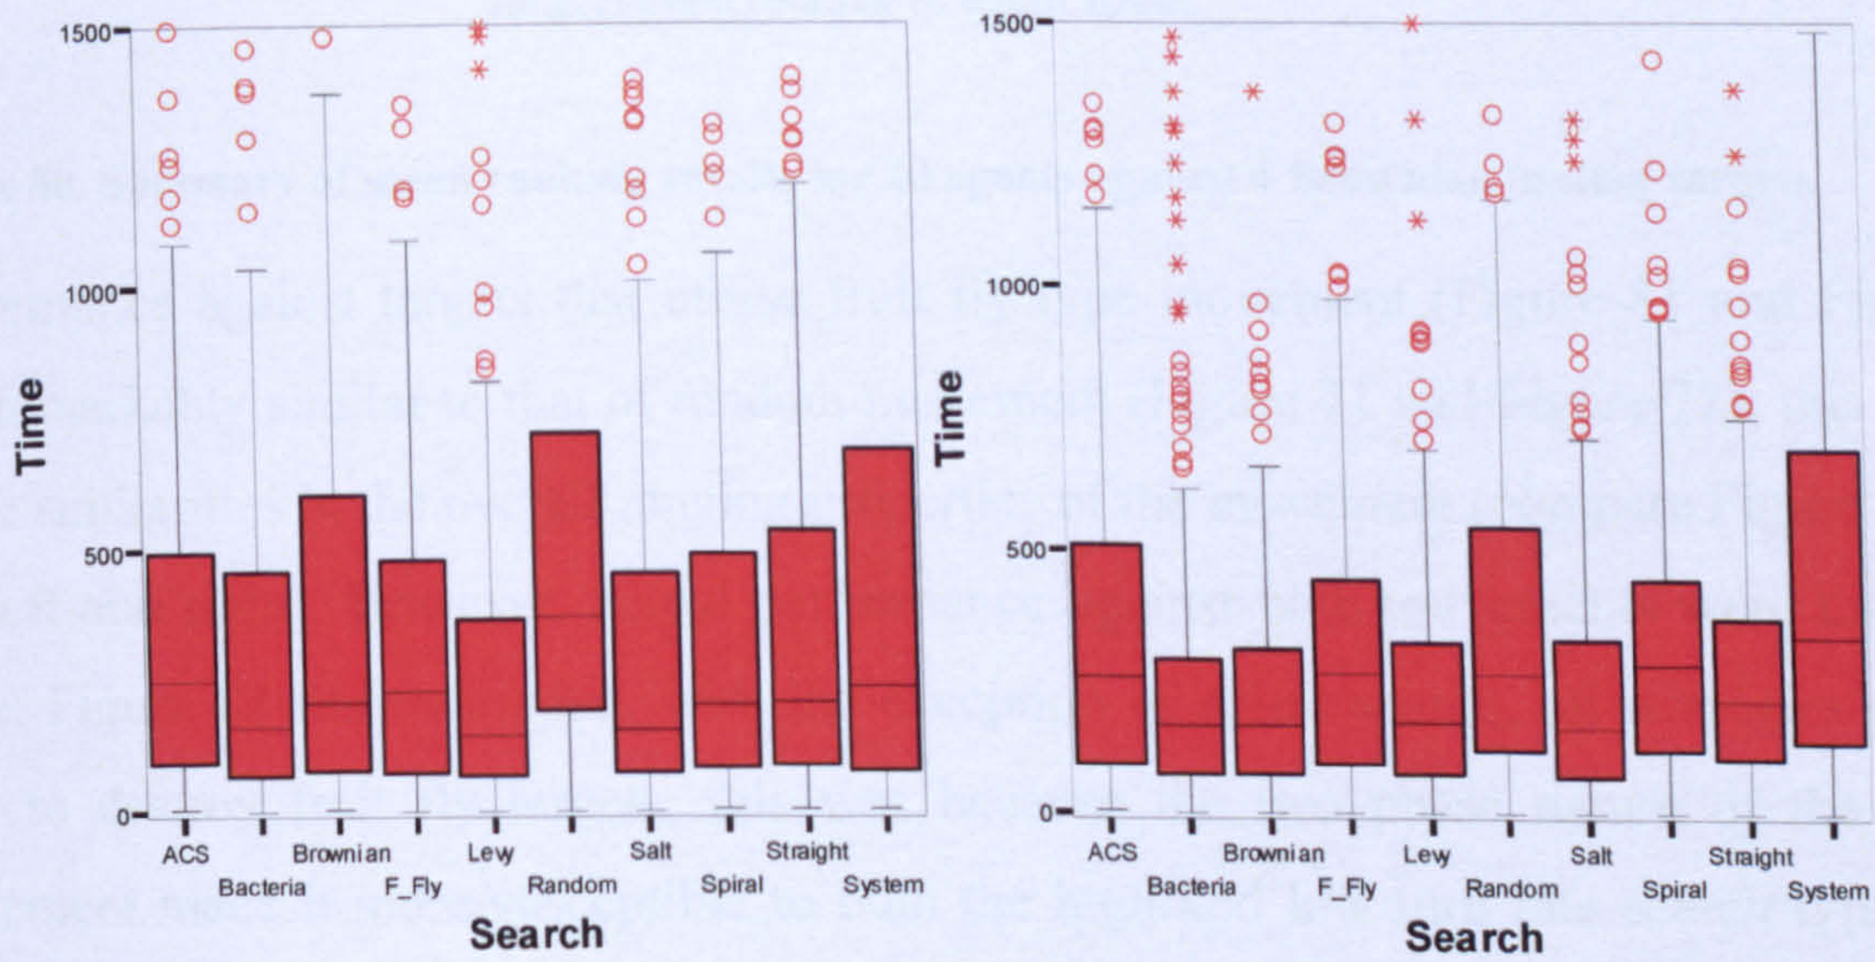

Figure 79. Lower axis zoom of actual results of 20 agents against 1 target travelling at maximum agent speed presented in Figure 78 (left: Brownian motion, right: fruit fly movement).

Where more targets were present (Figure 76 and Figure 80) the lack of movement across the search space highlighted the difference in the reduction in the likelihood of agents getting close enough to detect a target, especially with smaller swarms. With larger swarms the high turn rate strategies did become dominant, but only once targets were moving relatively fast.

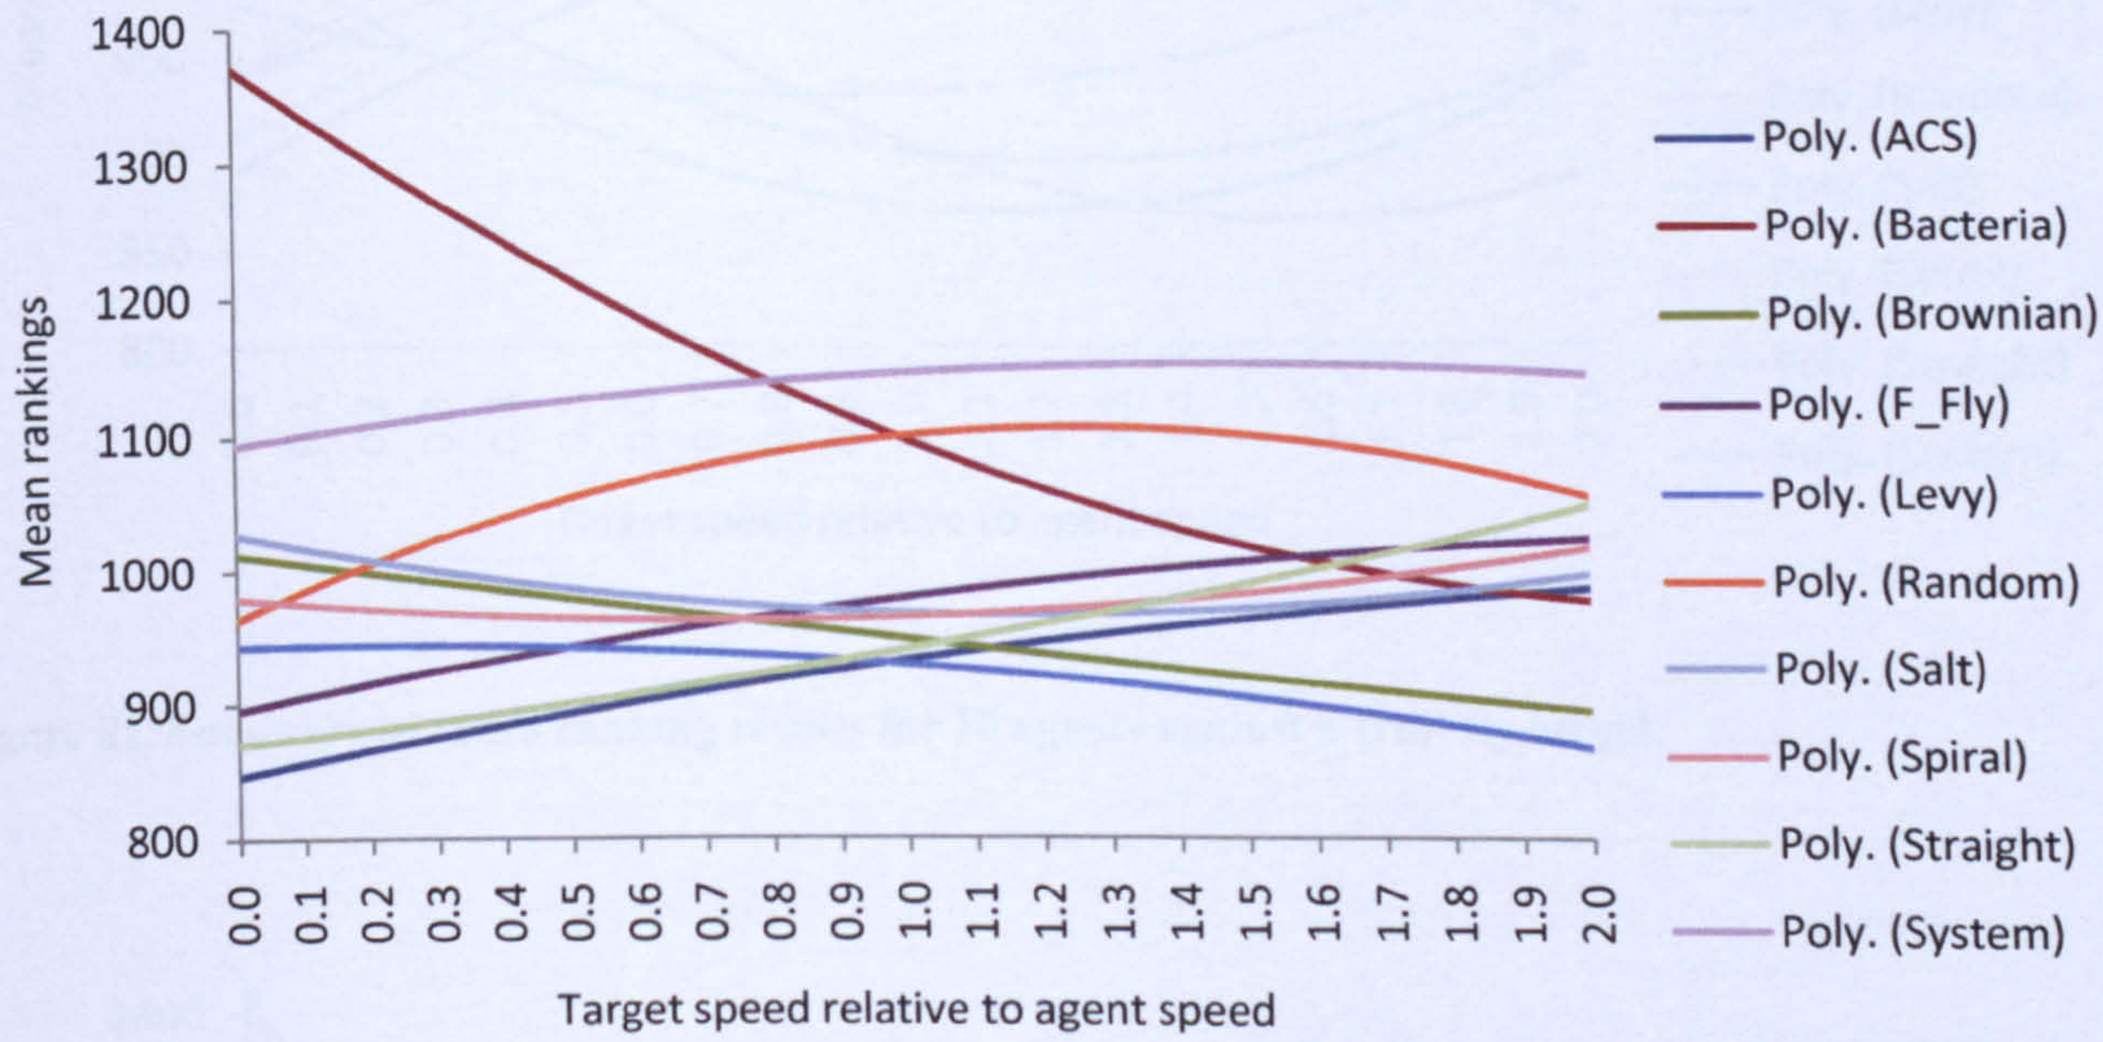

Figure 80. Summary of mean ranking results for 20 agents against 4 Brownian motion targets.

Performance against targets that utilise fruit fly type movement (Figure 81 and Figure 82) was remarkably similar to that of random movement (Figure 71 and Figure 72); this was due to the similarities in the overall ranging properties of the movement (compare Figure 25, page 80, left and right). However, actual performance against both movements were the not the same. Figure 83 illustrates that, with the exception of spiral search, all strategies took less time to destroy fruit fly targets; this was because the two phase nature of the fruit fly movement made it more susceptible to both the high and low turn rate search types (inter-saccade and saccade phases respectively).

Also notable from the actual results of the fruit fly target movement was the change in result consistency compared to those of Brownian motion (Figure 78, left and right). This may be attributed to the target's increased movement about the search space reducing the chances of it remaining outside agent sensor range.

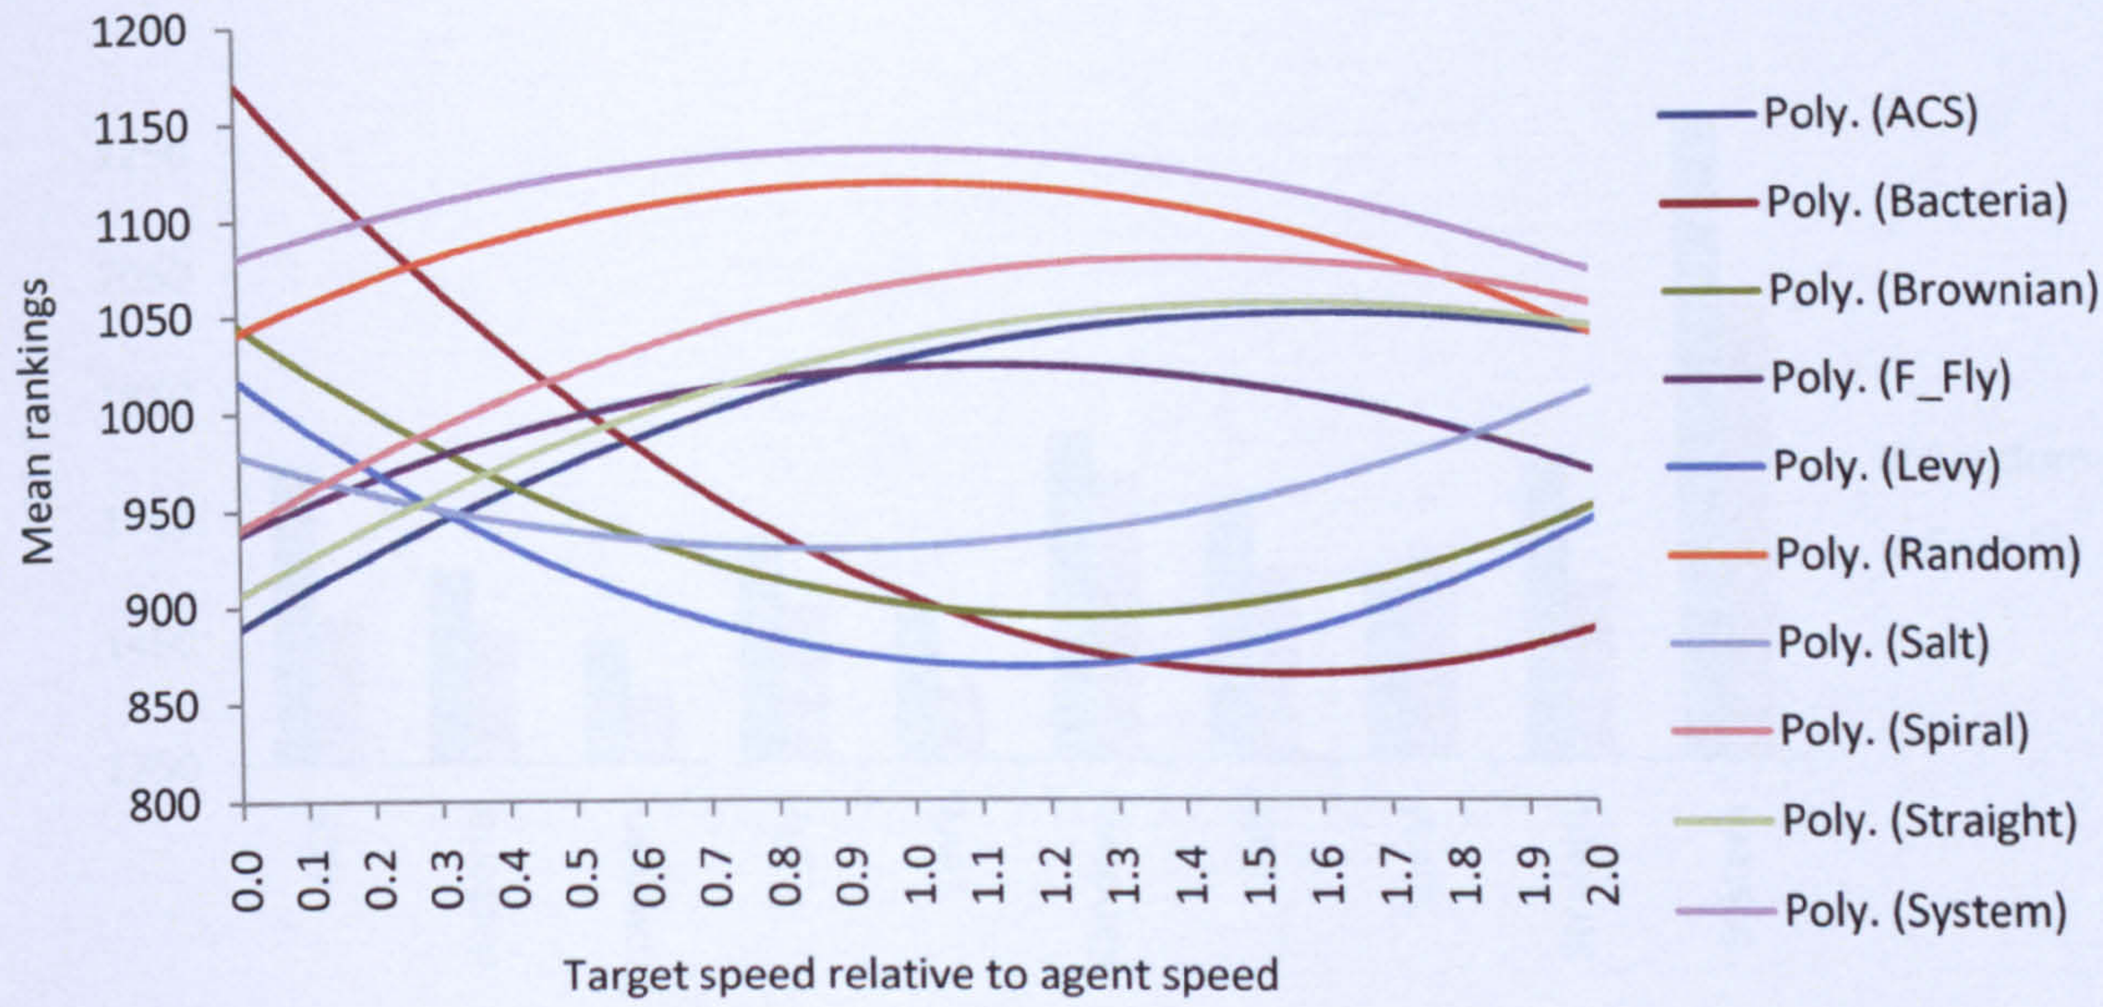

Figure 81. Summary of mean ranking results for 10 agents against 1 fruit fly target.

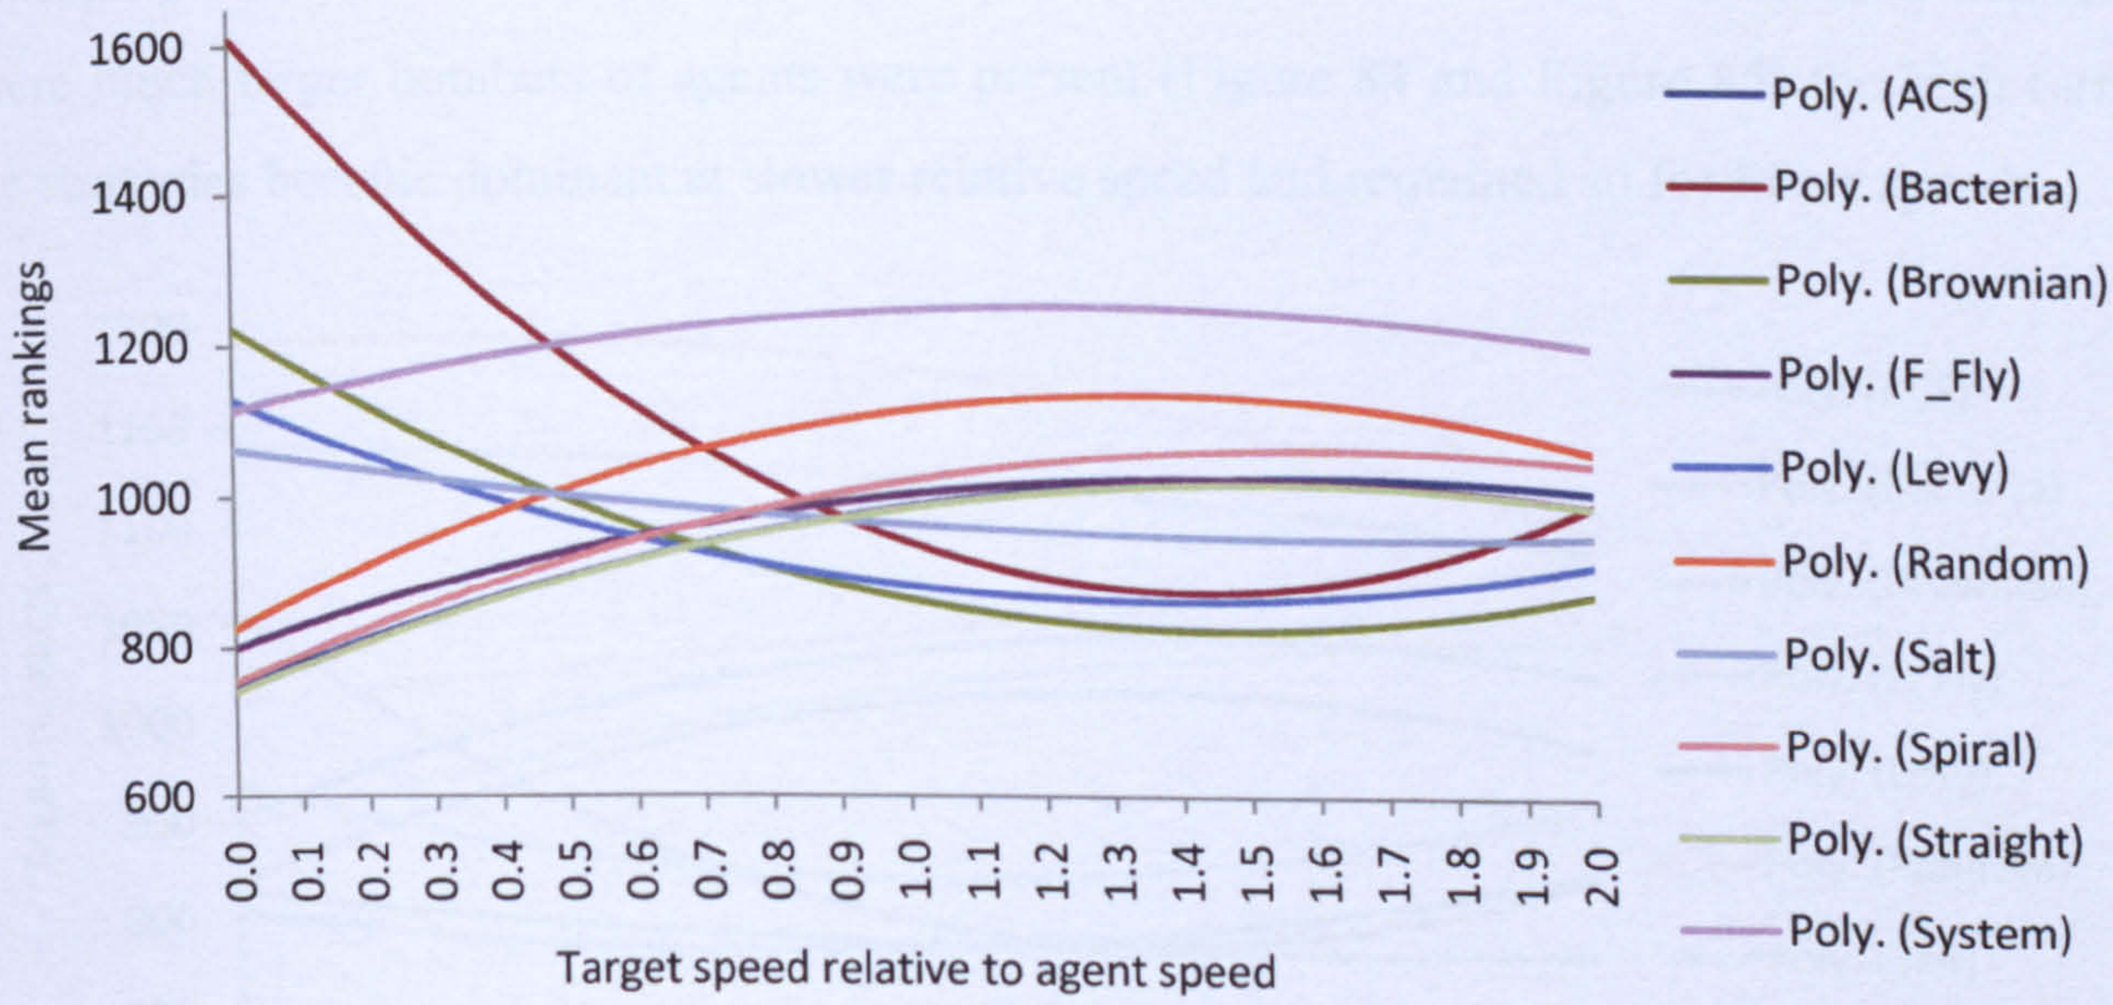

Figure 82. Summary of mean ranking results for 10 agents against 4 fruit fly targets.

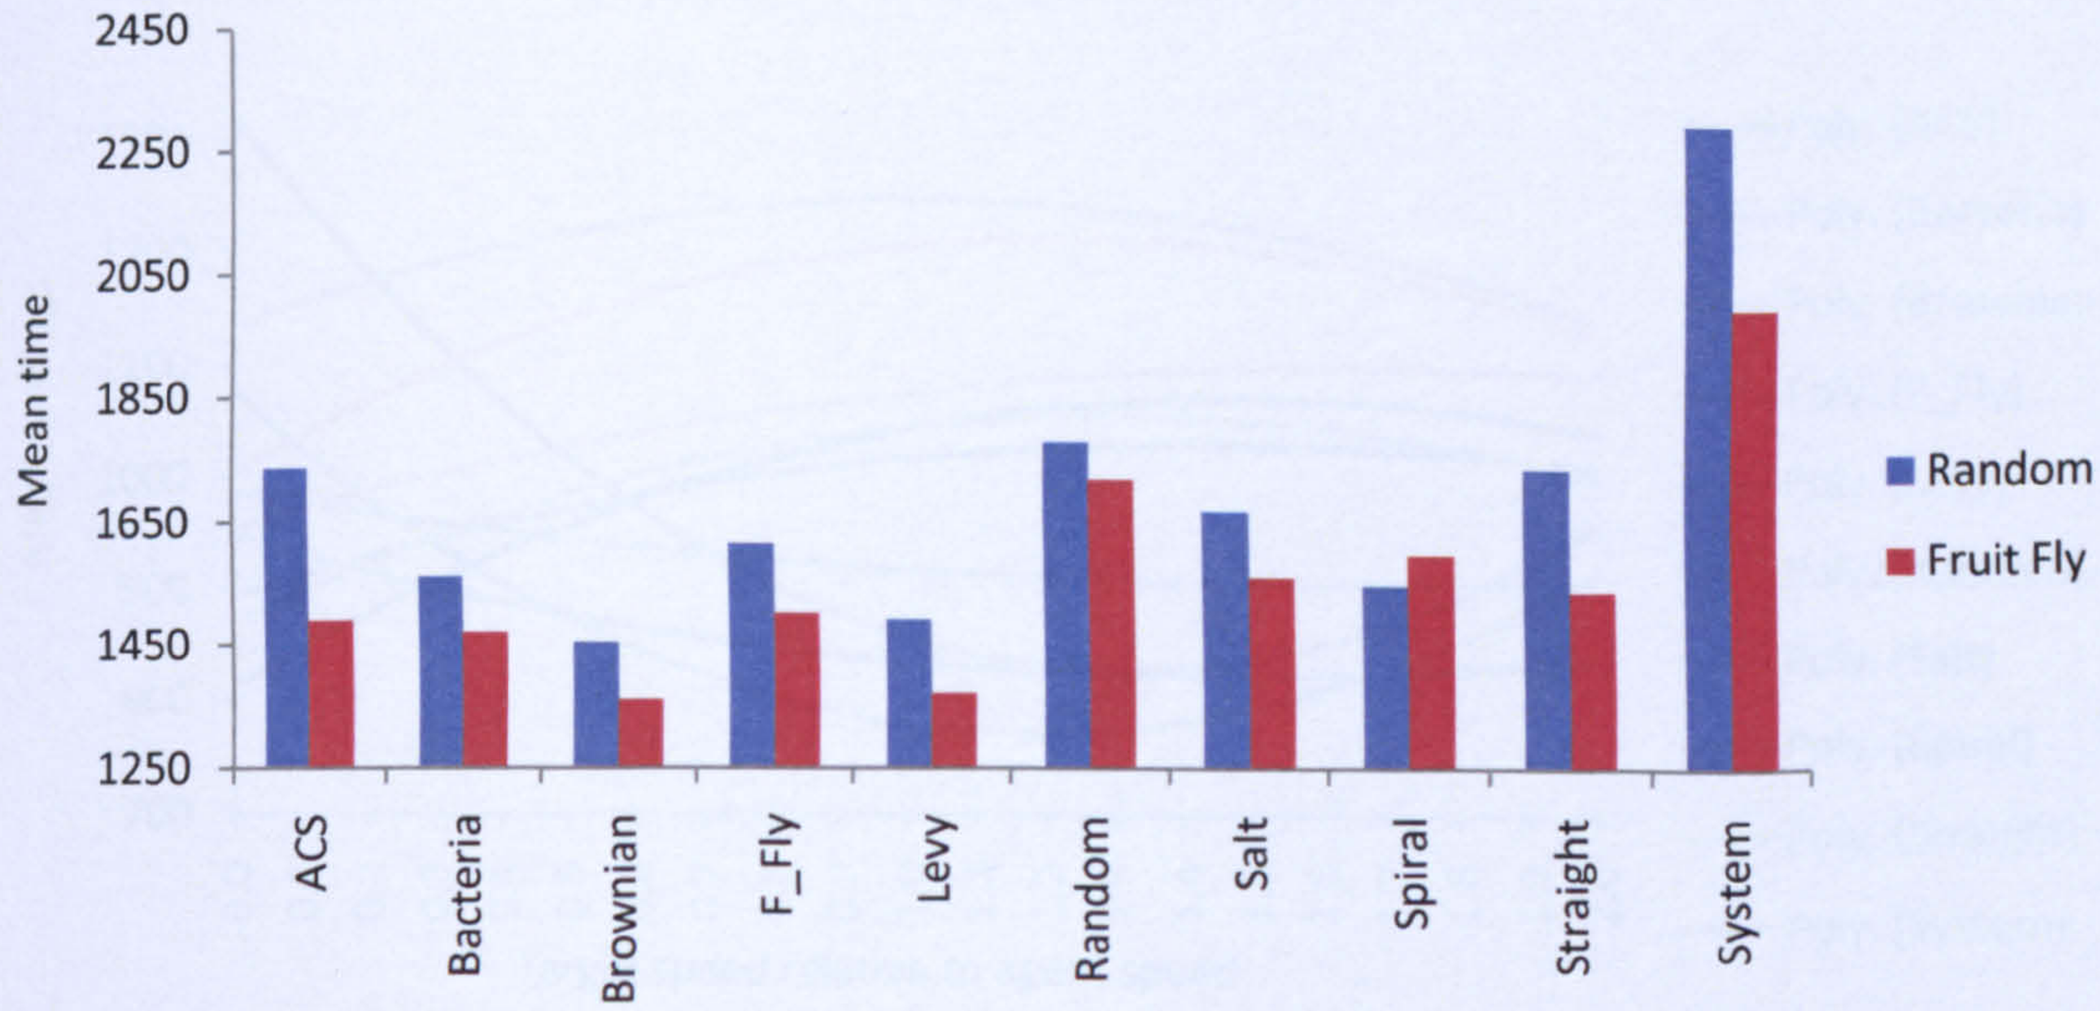

Figure 83. Summary of mean times for 10 agents to destroy 4 targets moving at maximum agent speed.

In keeping with the similarity of 10 agent performance with that of deterministic search, where much larger numbers of agents were present (Figure 84 and Figure 85) the high turn rate strategies became dominant at slower relative speed and remained so for faster speeds.

Figure 84. Summary of mean ranking results for 20 agents against 1 fruit fly target.

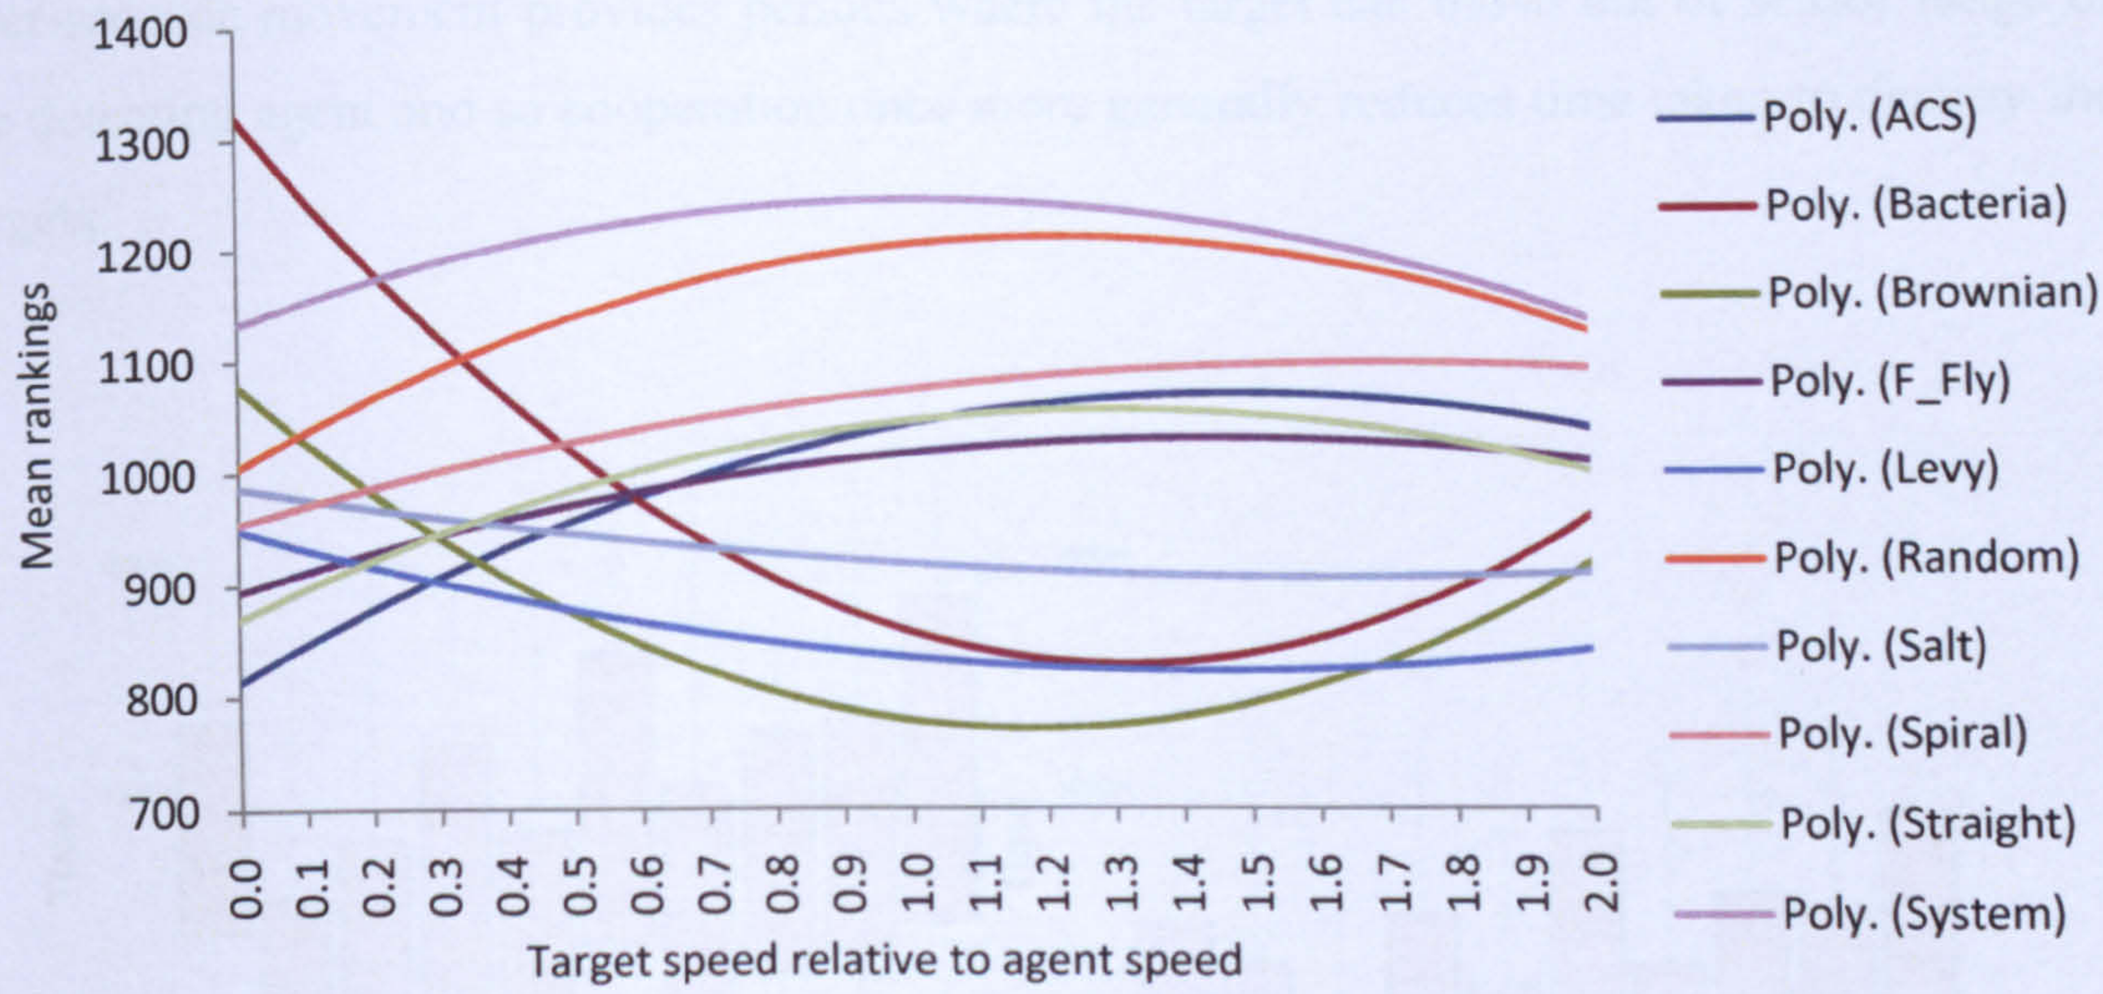

Figure 85. Summary of mean ranking results for 20 agents against 4 fruit fly targets.

5.5 Effect of Swarming

This section examines the effect of the LSZ based swarming element of the agent reaction to target detection, comparing the performance of a swarm of agents using cooperation against one where they act independently. For the comparison, 20 agents were employed to detect and destroy 4 deterministic targets moving 50% faster than agent maximum speed (20 agents were selected to emphasize the relationship). The experiment was then repeated using the remaining 3 target movement types: random, Brownian motion and fruit fly. It can be seen from Figure 86 that for deterministic targets, generally, the same relationships held between the search strategies whether or not swarming was employed: the application of the cooperative element simply improved swarm response upon target detection. This relationship also held where the agents searched for random targets (Figure 87) although the improvement was not as significant because targets turned more frequently. Targets turning during detection can reduce their chances of outrunning an agent because the latter will take a direct route, reducing the gap between itself and the target, possibly allowing itself to get close enough to attack the target in the process. The effect of the target turning is amplified still further against Brownian motion targets (Figure 88) where very little difference exists between cooperating and non-cooperating swarms. In Figure 89, however, the straight line

inter-saccade movement provides periods where the target can move out of sensor range of the detecting agent and so cooperation once more generally reduces time taken to destroy the targets.

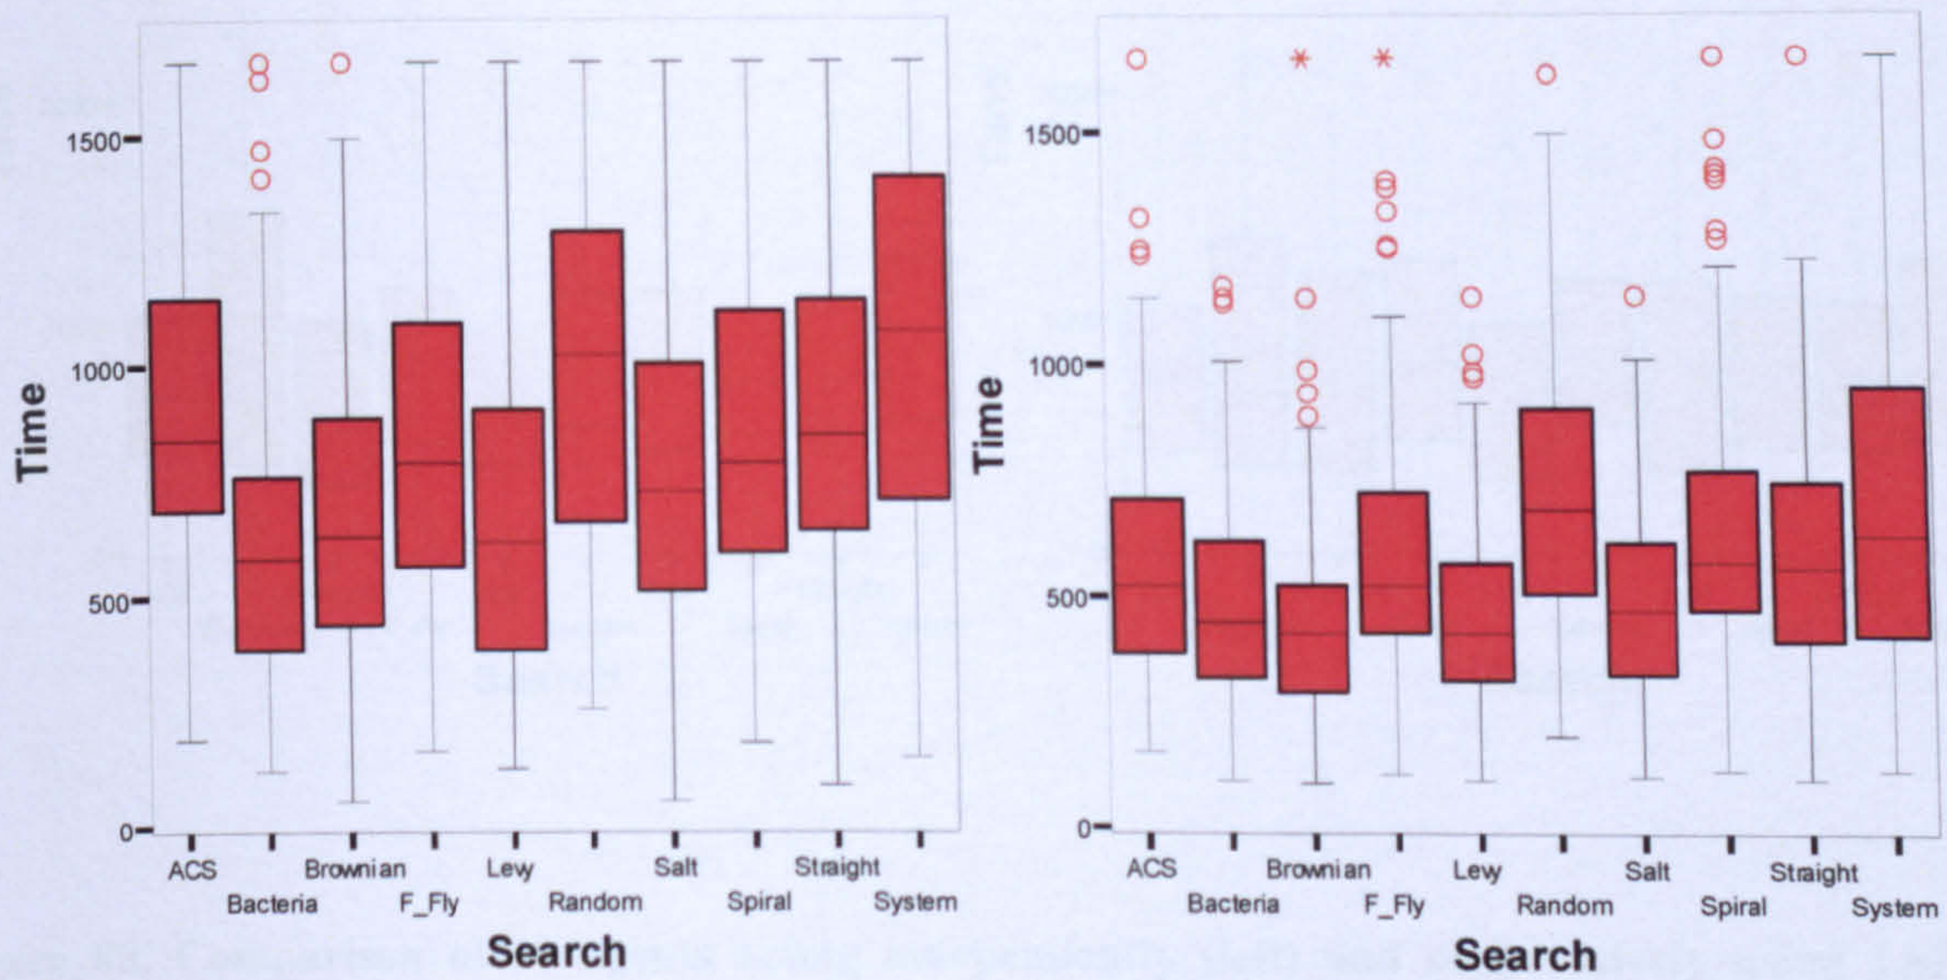

Figure 86. Comparison of 20 agents acting independently (left) and cooperatively using LSZ (right) against 4 deterministic targets travelling at 1.5 times maximum agent speed.

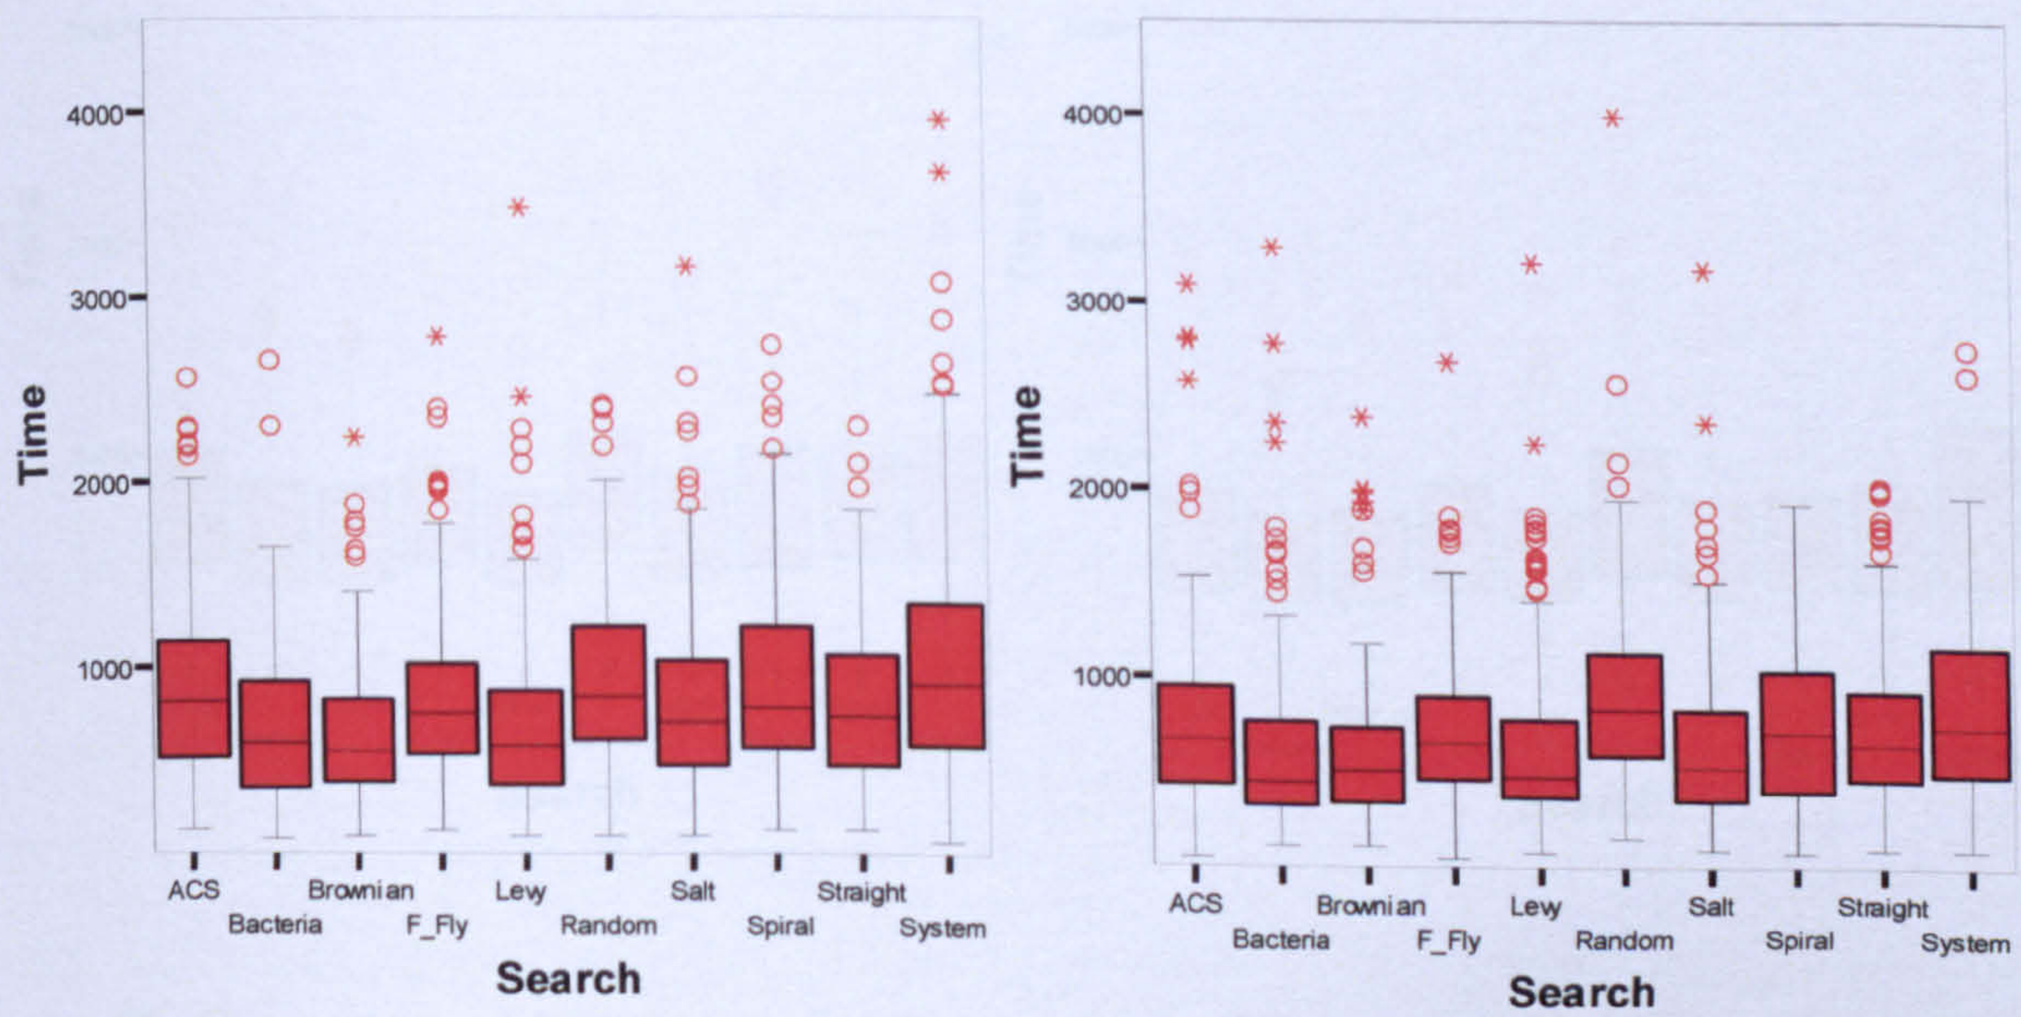

Figure 87. Comparison of 20 agents acting independently (left) and cooperatively using LSZ (right) against 4 random movement targets travelling at 1.5 times maximum agent speed.

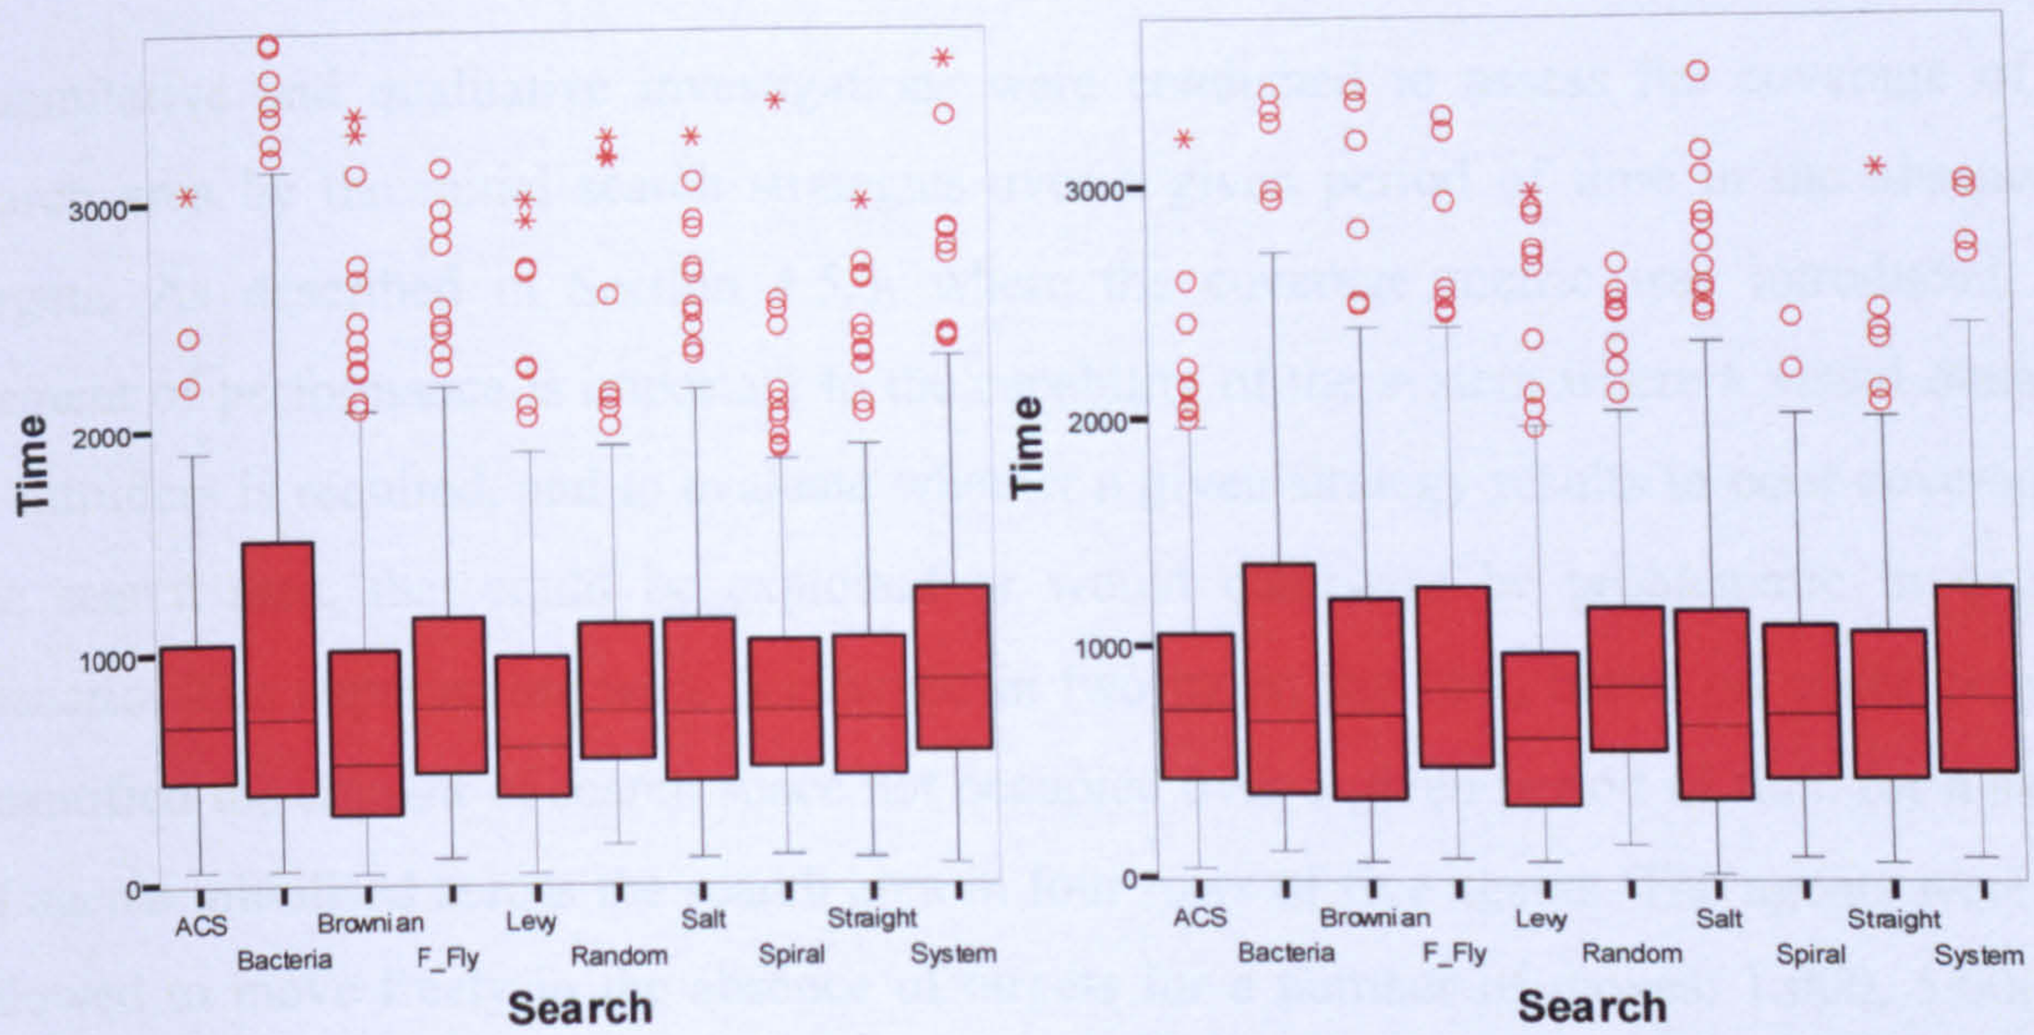

Figure 88. Comparison of 20 agents acting independently (left) and cooperatively using LSZ (right) against 4 Brownian motion movement targets travelling at 1.5 times maximum agent speed.

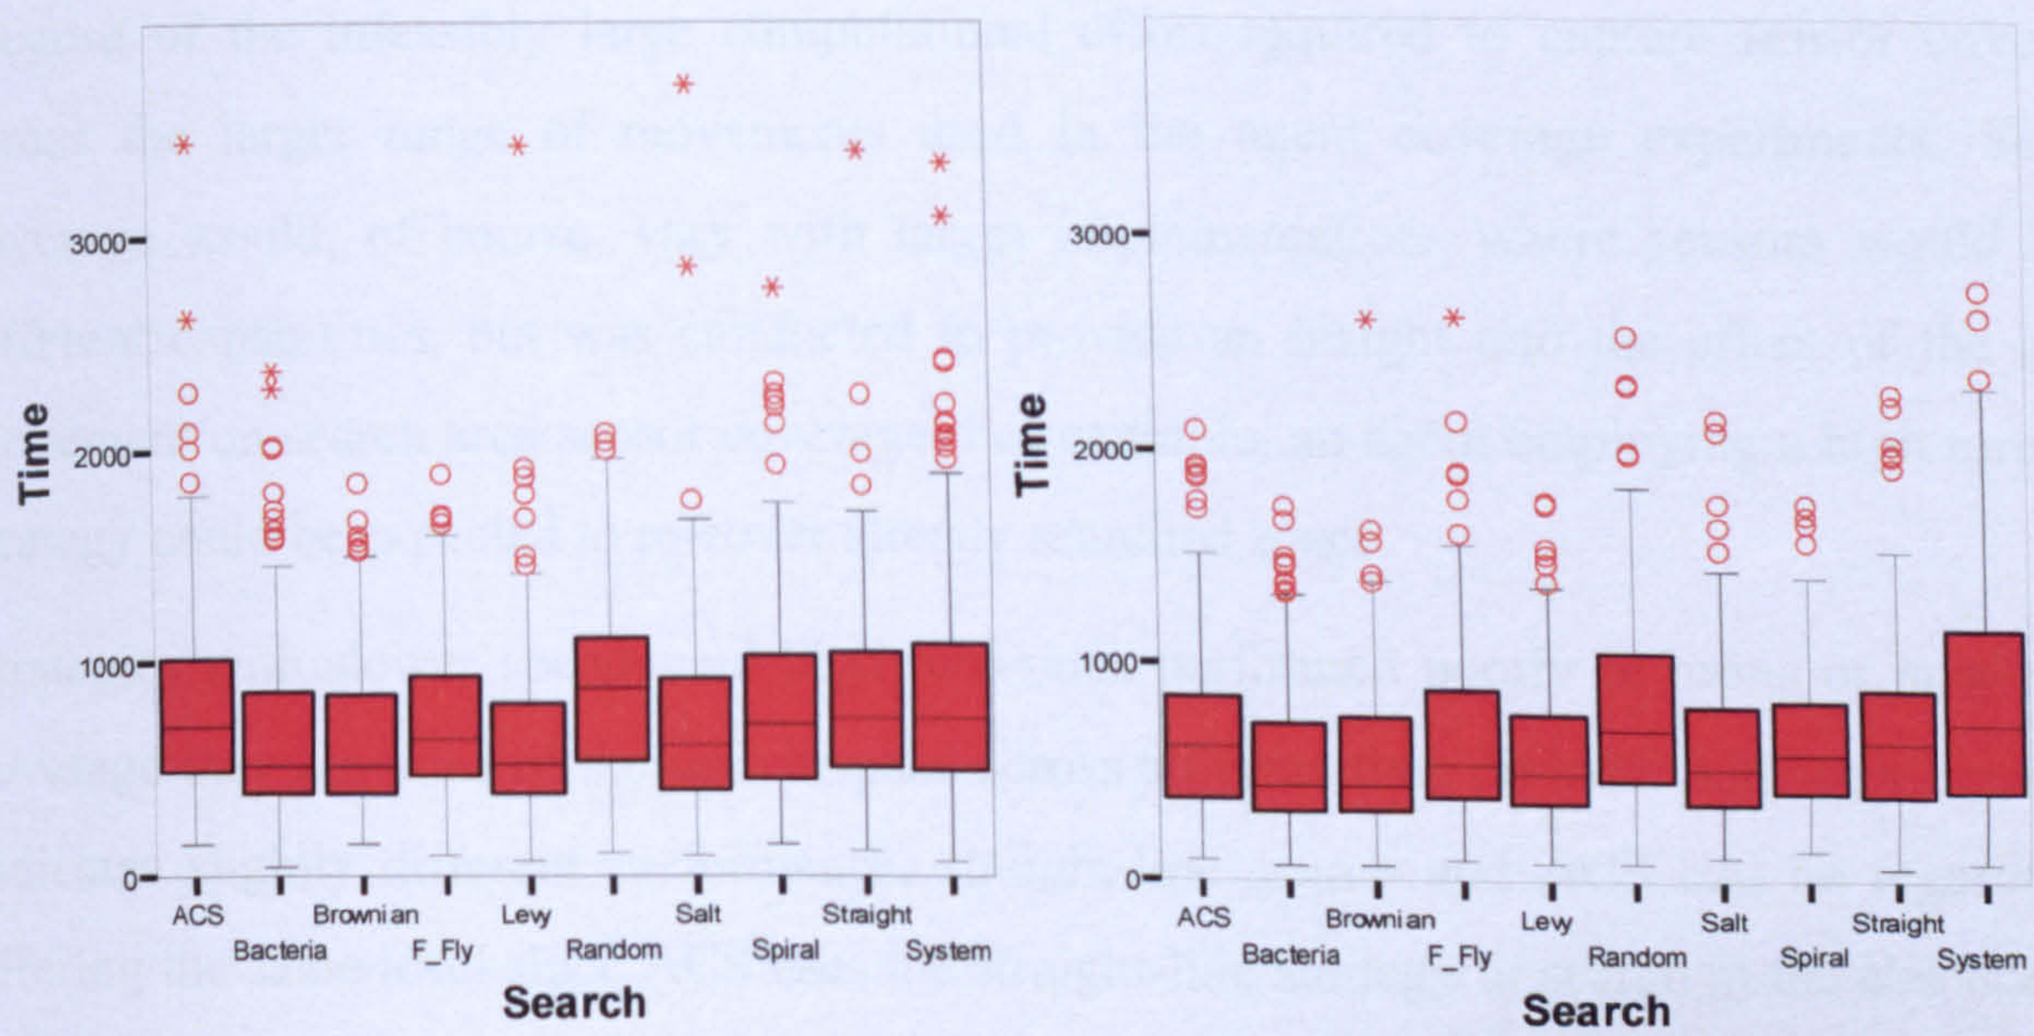

Figure 89. Comparison of 20 agents acting independently (left) and cooperatively using LSZ (right) against 4 fruit fly movement targets travelling at 1.5 times maximum agent speed.

5.6 Coverage Quality

Quantitative and qualitative investigations were conducted to assess the coverage of the search area by the initial search strategies over a given period of time in the absence of targets. As described in Section 4.5.3, where the coverage metric was introduced, this element of performance is important to the capability of the system where a visual deterrent to intruders is required, and to evaluate whether a given strategy results in poor coverage of the search area, that could be exploited or would otherwise be problematic in a given scenario. The experiments were performed in two parts; the first, based on agent position, quantified the amount of search space not occupied over a given period of time for a swarm of agents initialised across the search area in four rows of five agents. The agents were then allowed to move freely in the absence of targets for a number of moves; 1,000, 5,000 and 10,000 were selected to provide a range of time periods. The shorter periods allow the assessment of how movement may appear to the casual observer, whilst the longest period allows the deterministic patterns time to move fully across the search area. The second experiment scenario involved the production of indicative sensor coverage plots of a single agent, initialised at the centre of the search area, over 5,000 moves. This approach was taken because of the infeasibly large computational effort required to capture sensor coverage across the larger range of movements used in the agent coverage experiments. Sensor coverage would, of course, vary with target implementations where sensors would have different capabilities, but was conducted to provide an insight into the effect of the agent movement on search area sensor coverage. For example, an agent employing a high turn rate strategy could be expected to re-cover already searched space.

Strategies with slower speeds and high turn rates performed poorly in terms of how much coverage they provided of the search space across all three time periods. Although Figure 90 indicates slightly different performance, straight-line search and ACS can be regarded as offering the same level since ACS uses the straight-line strategy to search in the absence of a recently detected target. These differences were purely due to the stochastic nature of the search. Brownian search was outperformed by its near relative, Lévy search, due to Brownian motion's higher probability of direction change whereas Fruit fly search left less uncovered area than Saltatory search due to its increased speed during the intensive search phase (compare Figure 20, page 76, and Figure 25, page 80). In a reversal of the Lévy / Brownian

relationship, random search covered the area more effectively than straight-line because the long leg lengths of the latter left large areas uncovered.

It should be noted that since these results represent actual agent presence, rather than sensor coverage, the performance of the deterministic agents reduced as time increased because they began to recover space; although, as the indicative sensor coverage plots below show; sensor coverage would be more complete.

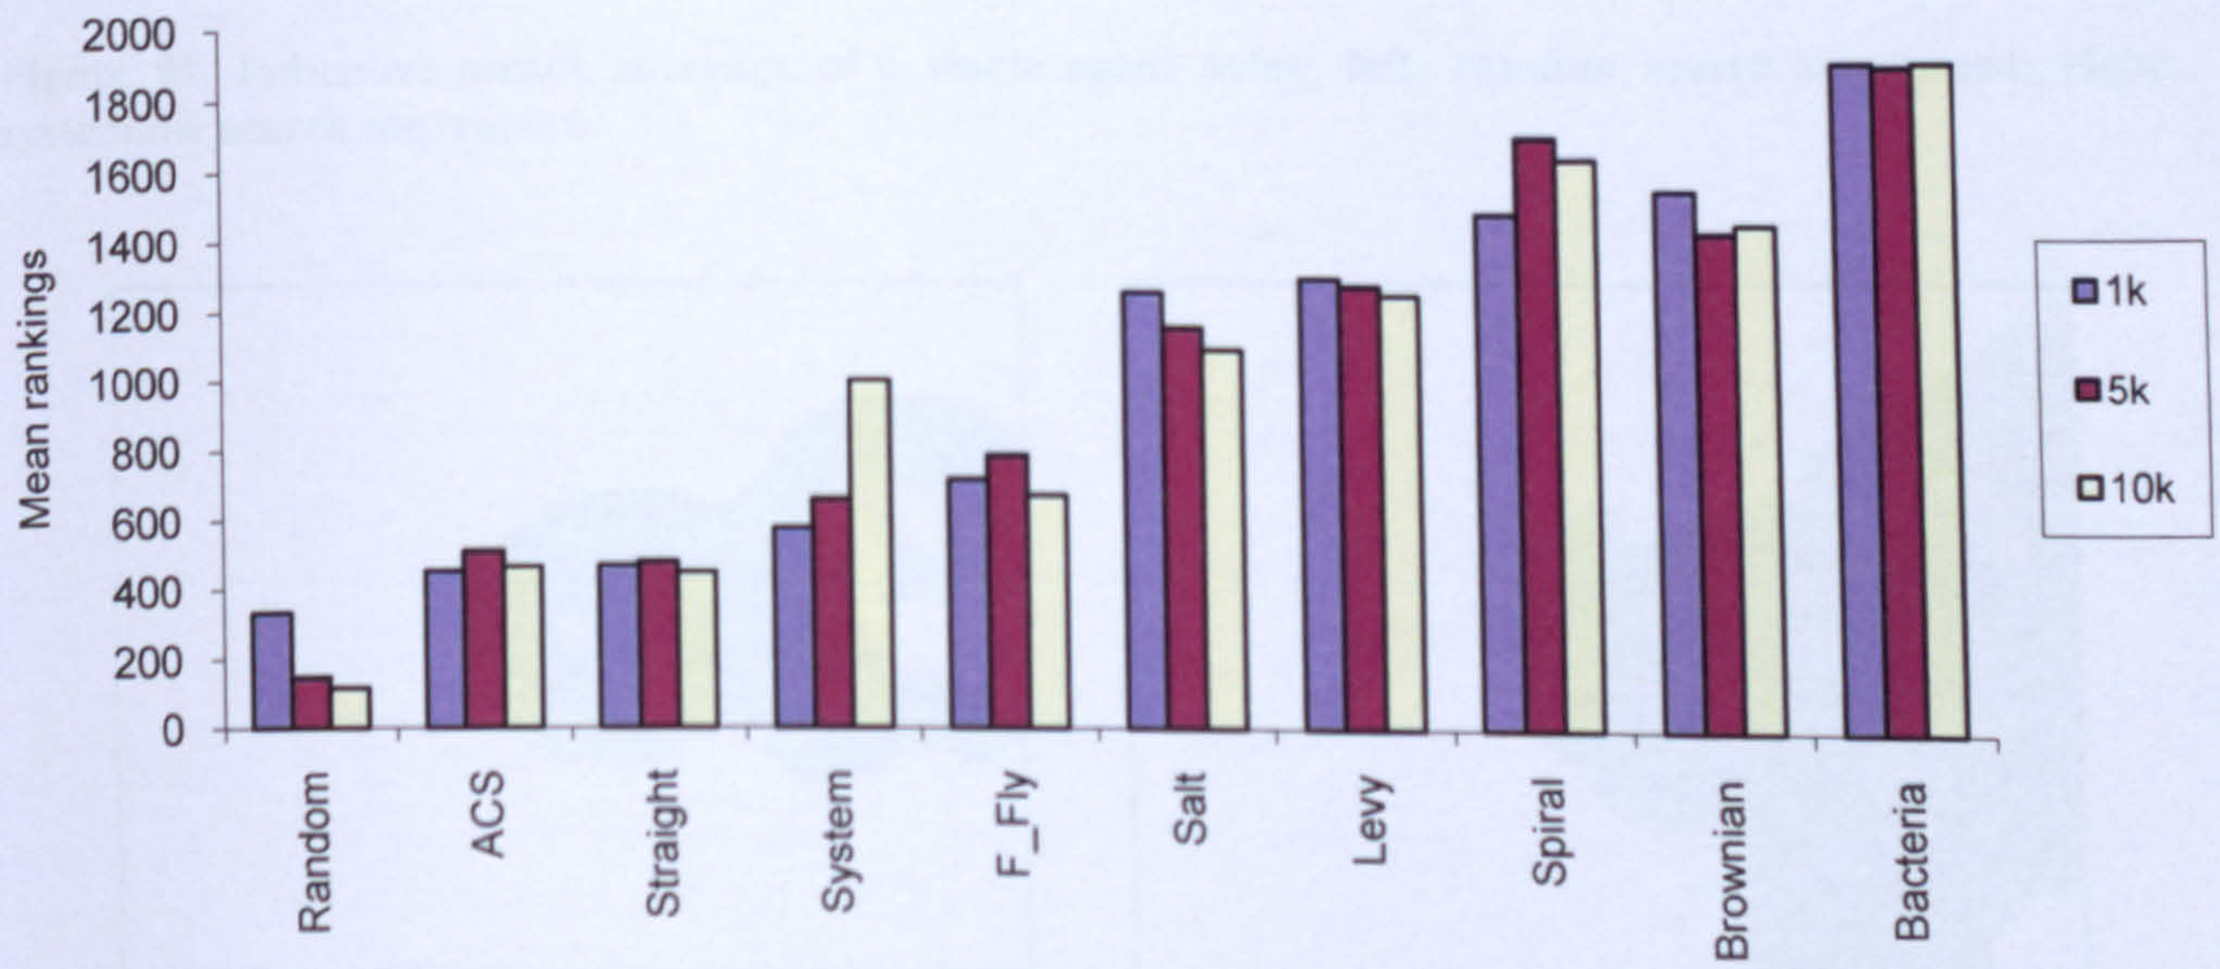

Figure 90. Summary of mean ranking results for strategy coverage quality over short (1000 moves), medium (5000 moves) and long (10000 moves) periods.

Indicative sensor coverage maps (Figure 91 to Figure 96) of a single agent over a given period were produced for all search algorithms. The plots clearly illustrate that strategies with faster speeds and reduced turn rates cover the search area with less repeat coverage than slower, high turn rate strategies; these findings generally supported the quantitative coverage quality results shown above. Worthy of note from these results is the increase in coverage shown by the bacterial search upon detection of a target (Figure 95).

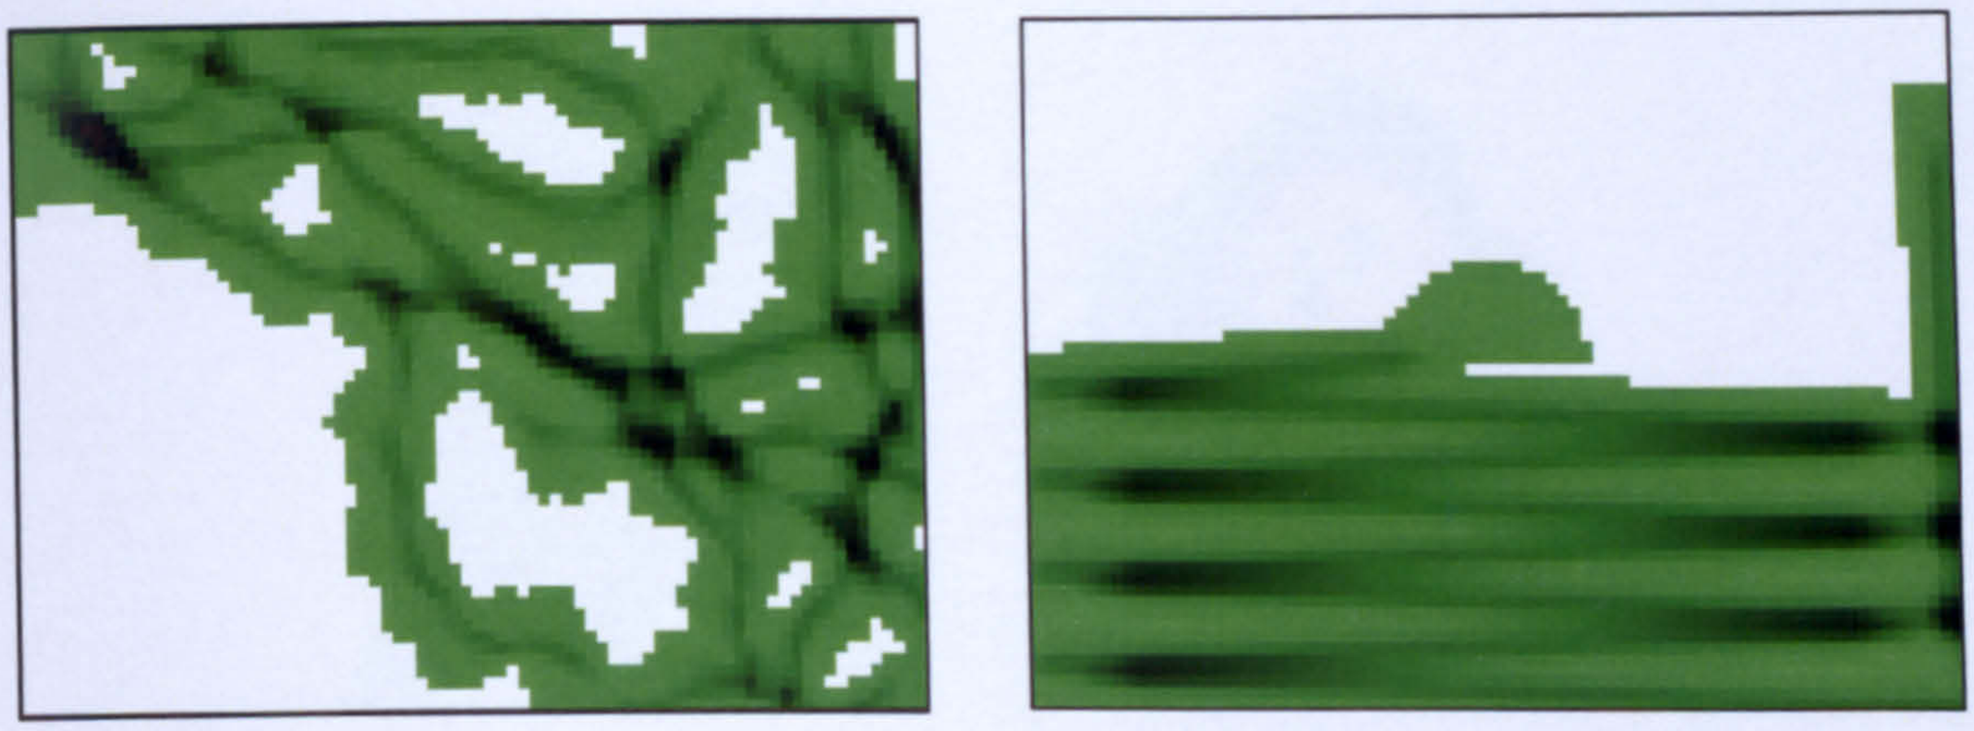

Figure 91. Indicative sensor coverage of a single agent using, left: random search movement; right: systematic search movement.

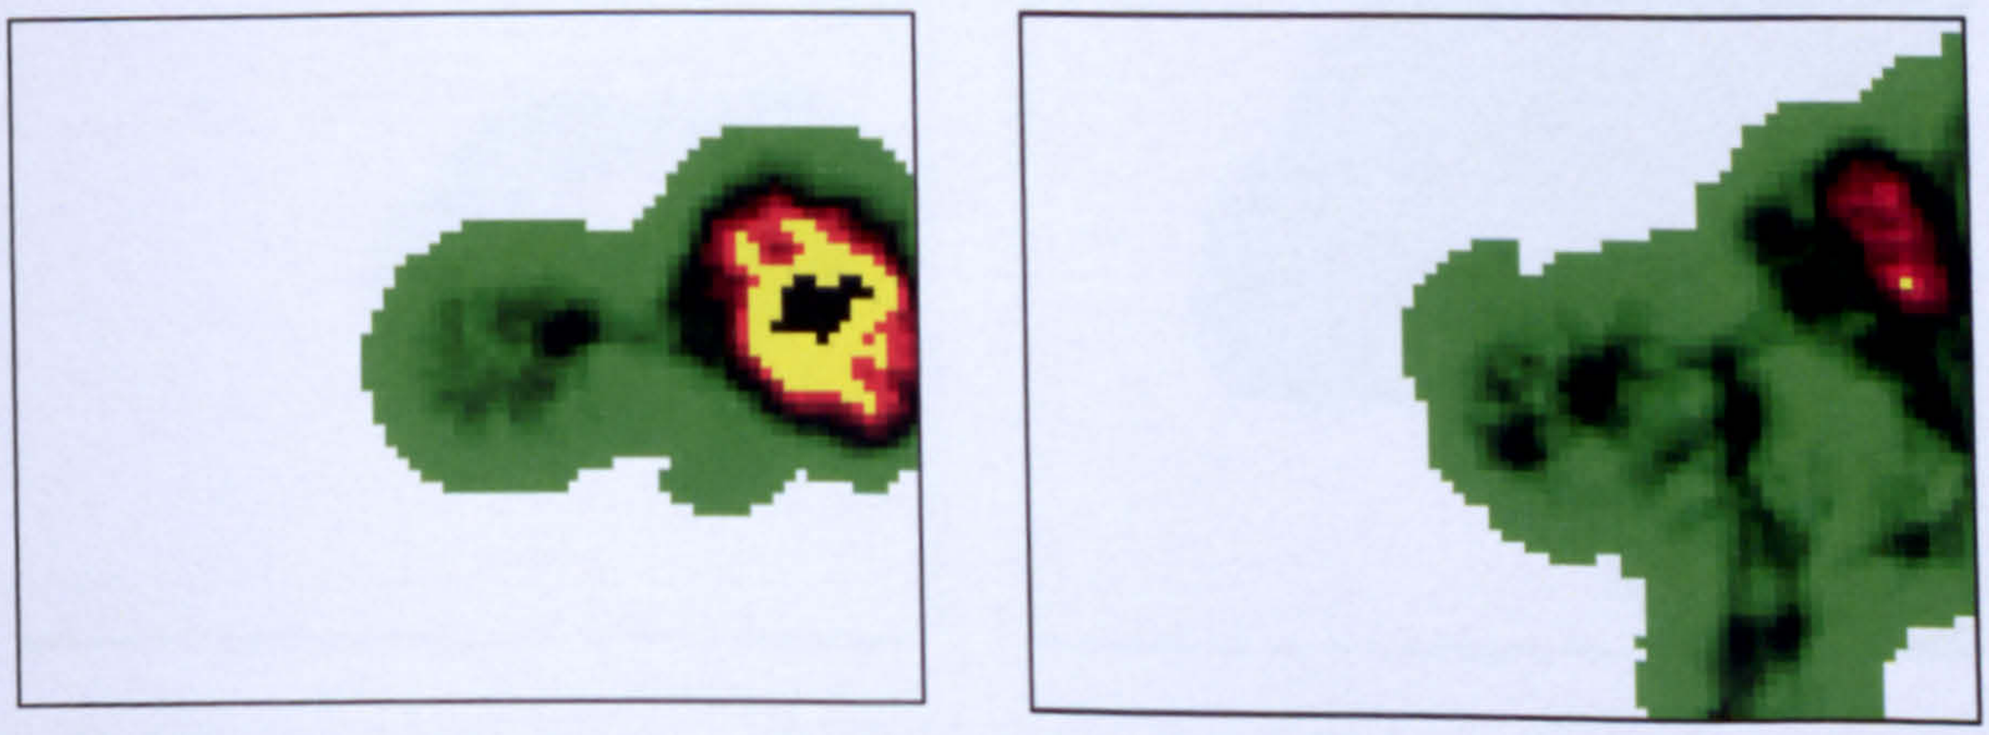

Figure 92. Indicative sensor coverage of a single agent using, left: Brownian movement search; right: Lévy path search movement.

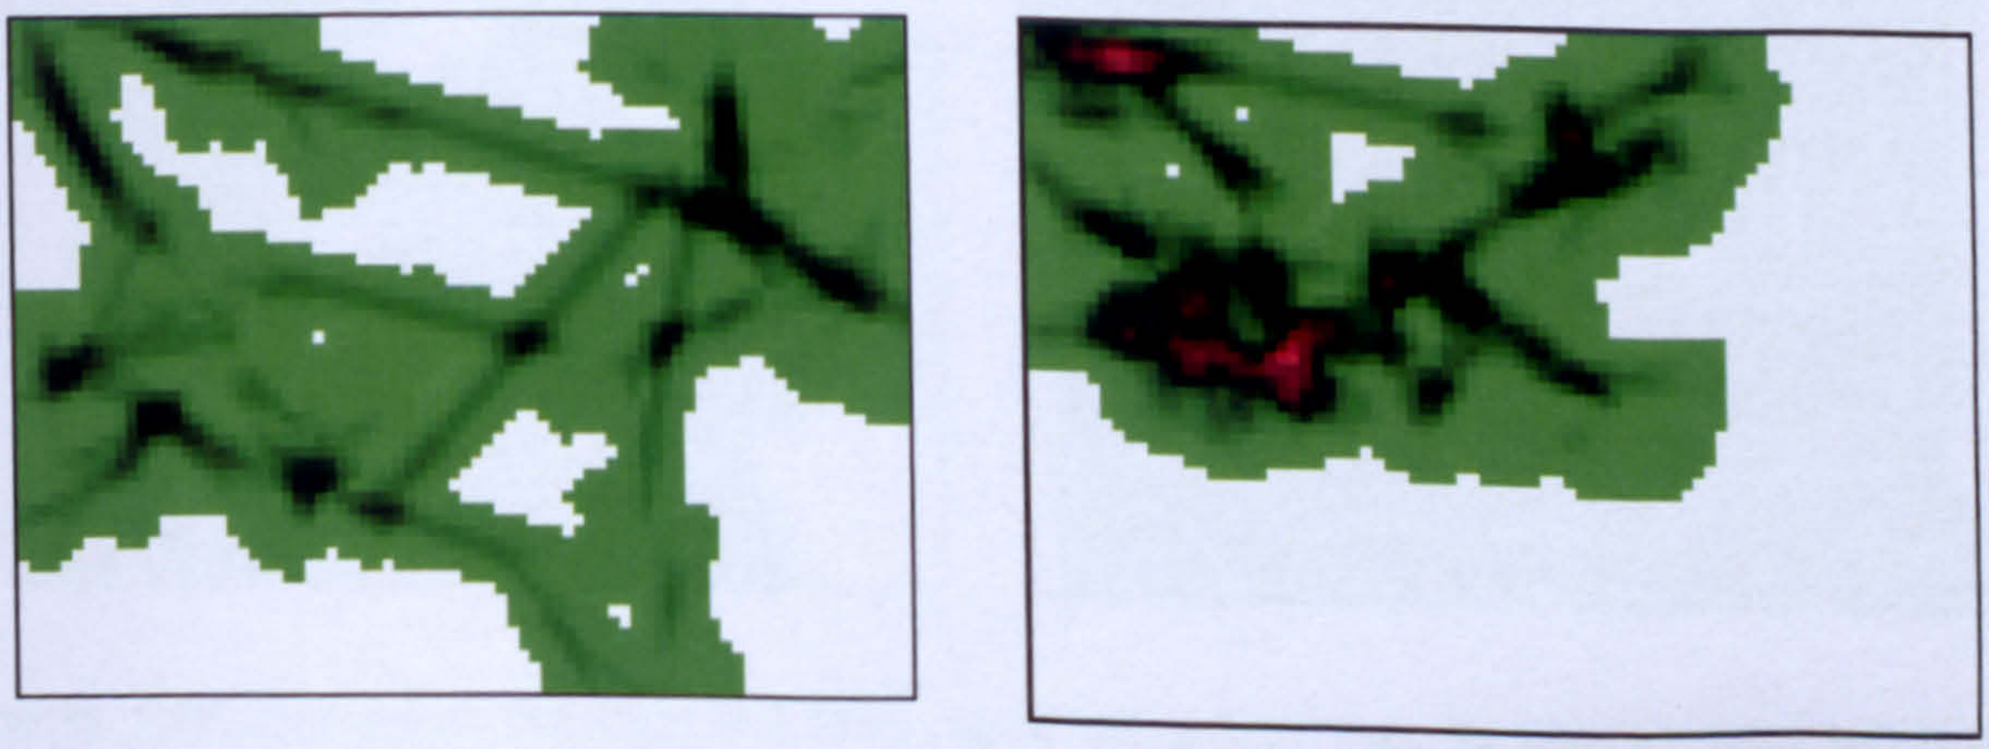

Figure 93. Indicative sensor coverage of a single agent using, left: straight-line search movement; right: saltatory search movement.

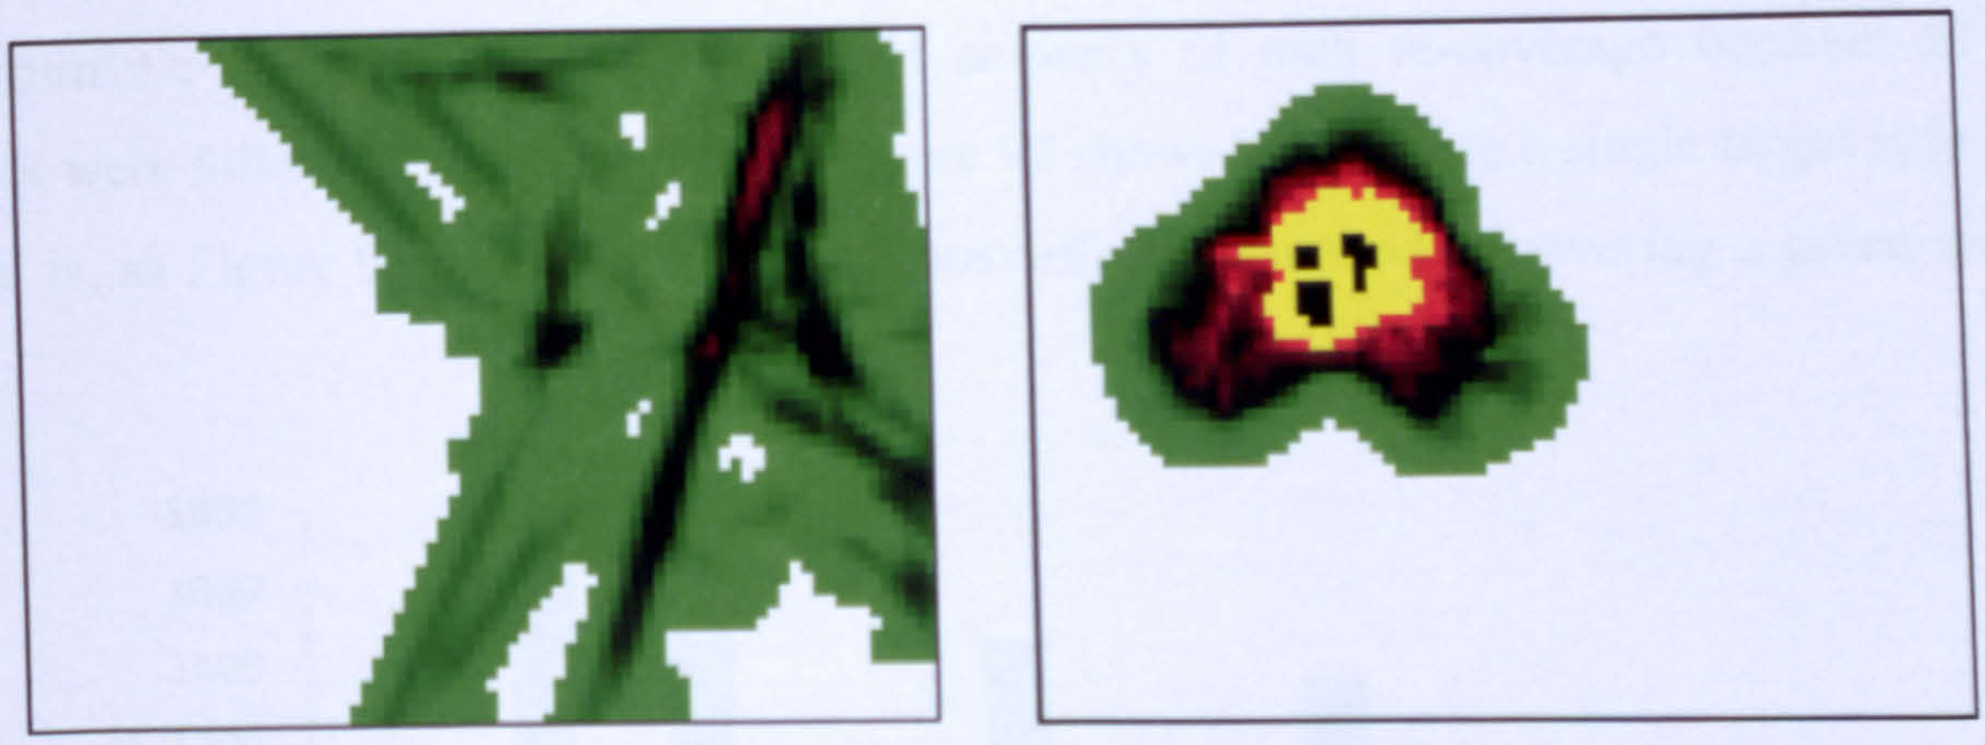

Figure 94. Indicative sensor coverage of a single agent using, left: ACS in extensive search phase (prior to target detection); right: ACS in intensive search phase (post target capture).

Figure 95. Indicative sensor coverage of a single agent using, left: bacterial search (prior to target detection; right: bacterial search (during target detection).

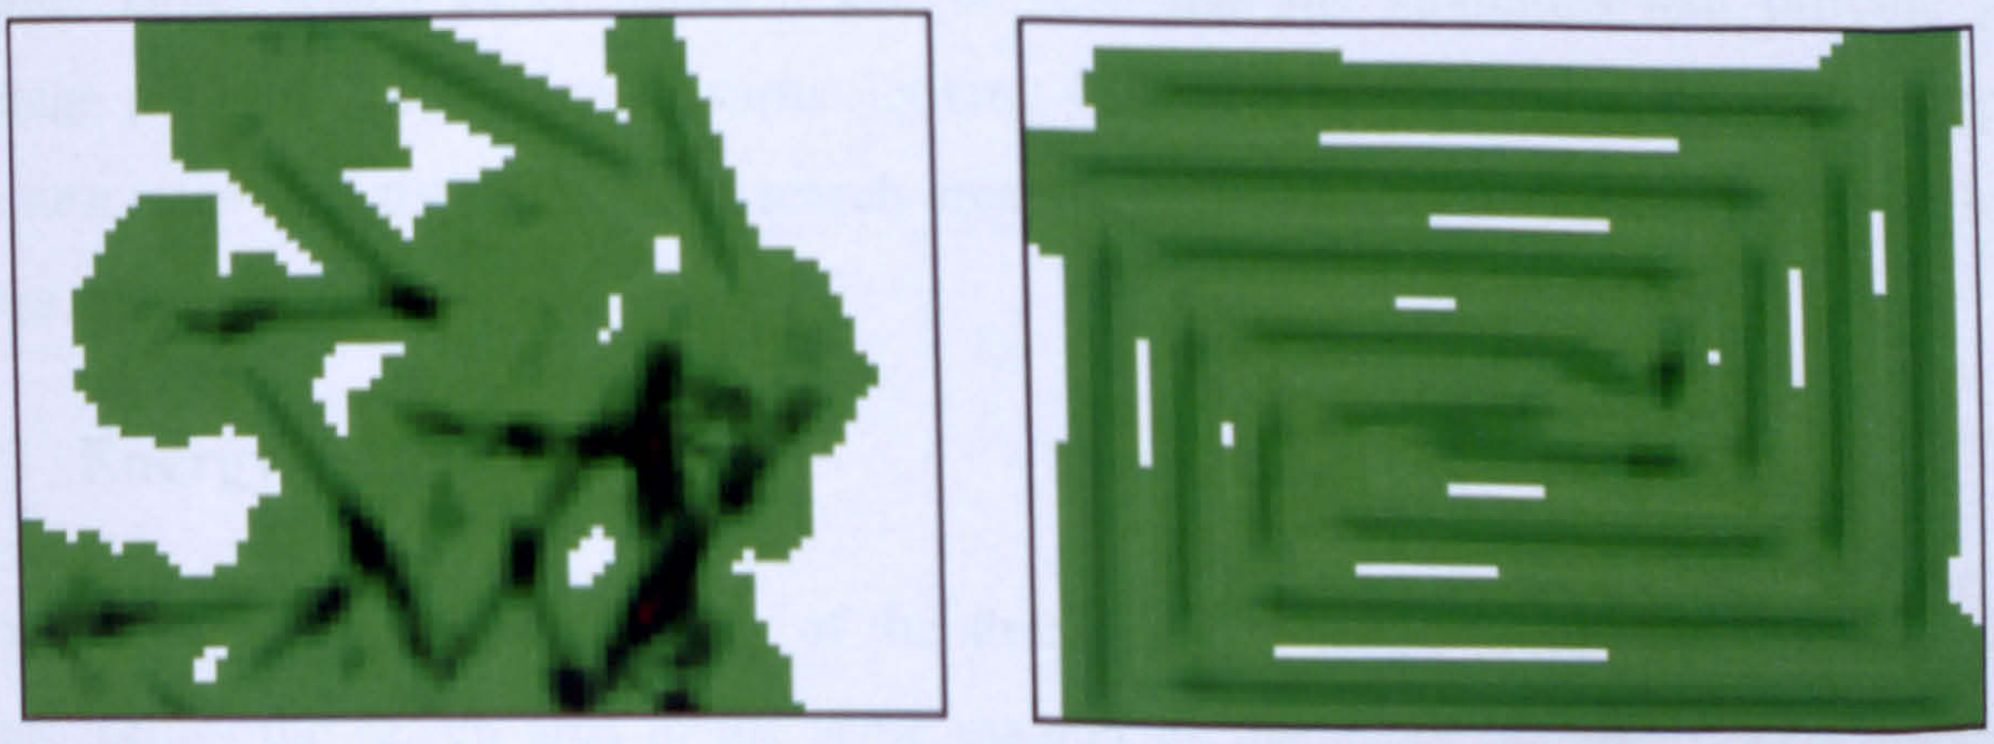

Figure 96. Indicative sensor coverage of a single agent using, left: fruit fly search; right: spiral search (note: the gaps would be covered by agents travelling on the opposite spiral).

However, notable are the deterministic search patterns which, where a single agent is employed, appear to cover the search area much better than Figure 90 would suggest. This was because where more agents were in use, as for the agent coverage experiment, the

deterministic searches included significant amounts of path re-coverage because all the agents were following the same routes. Figure 97 shows how, where a single target is in use, spiral is, as Figure 96 suggests, easily the most effective method of covering a given search area.

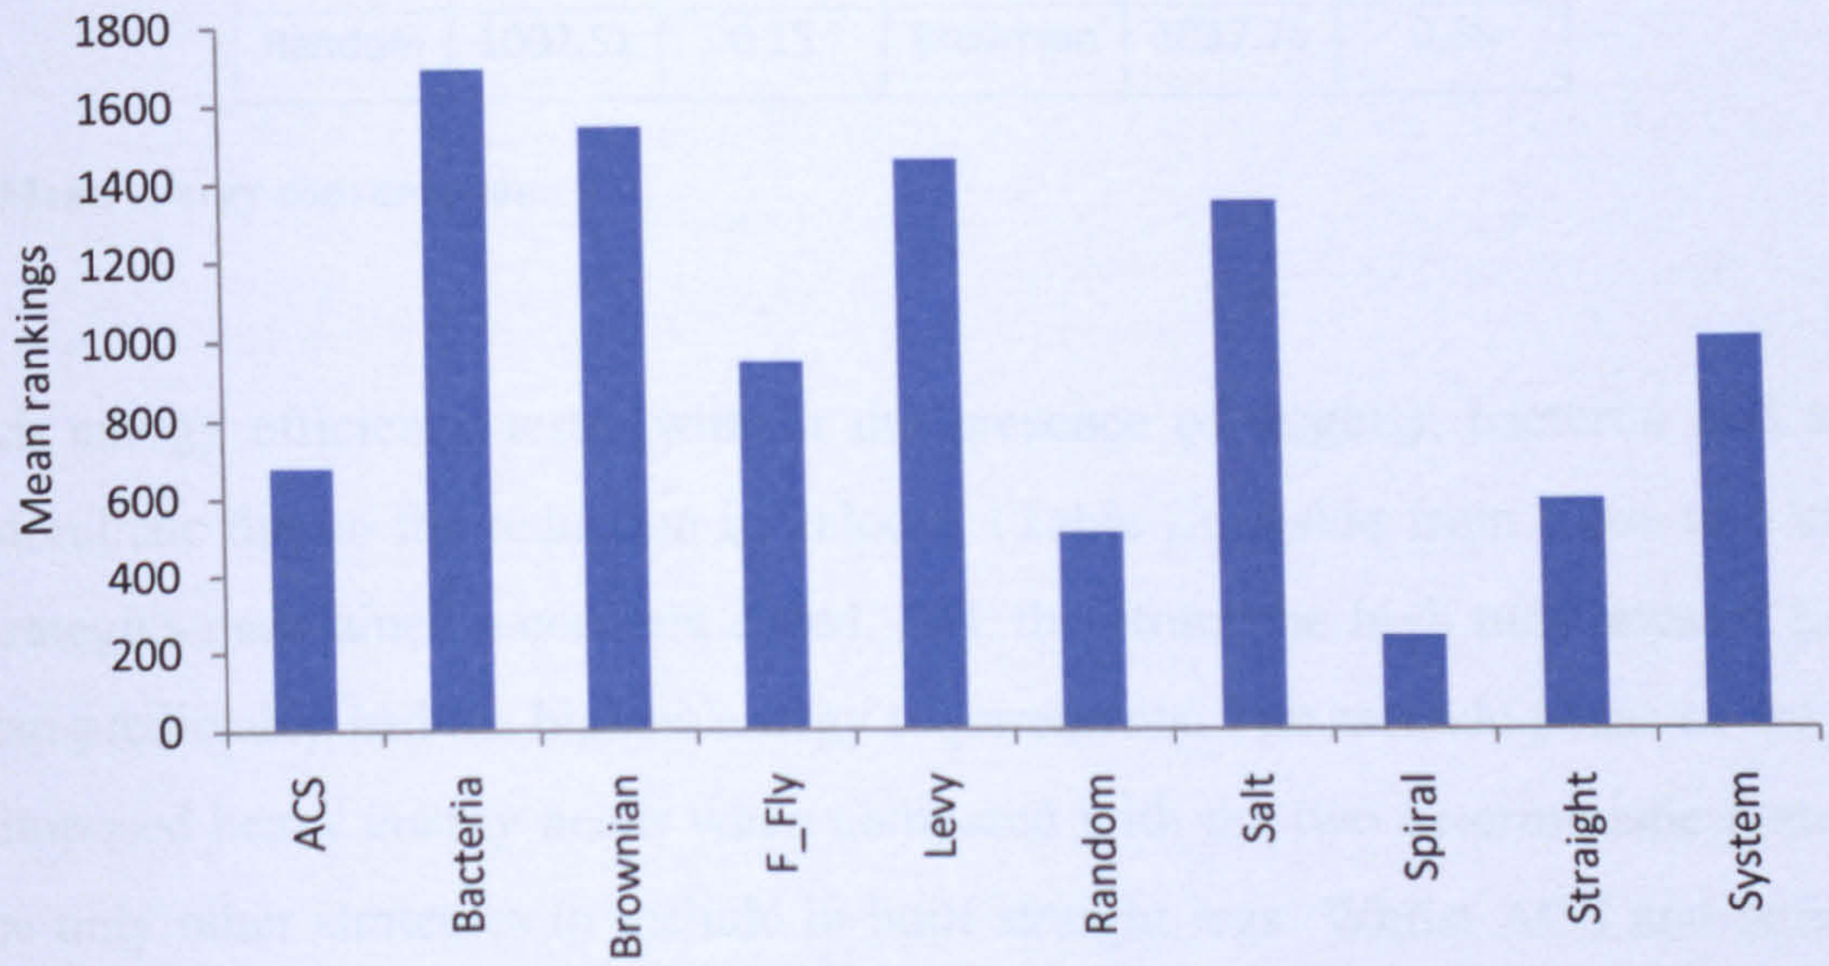

Figure 97. Summary of mean ranking results for sensor coverage of a single agent over 5,000 step wise moves.

Comparing the results of the sensor coverage experiments with those of the static and dynamic target search experiments, it can be seen that the strategies that provide good coverage perform well in small swarms looking for static targets whilst those that applied high turn rates (and therefore lower search area sensor coverage) performed better in larger groups against dynamic targets.

5.7 Energy Efficiency

To assess the relative fuel efficiency of the strategies, 20 agents were initialised in a fixed pattern across the search area in the same manner as the agent coverage experiments (as described in Section 5.6). All agents moved for 1000 iterations without the presence of a target and the energy requirements recorded, the lowest score being the most efficient (see Section 4.5.4, page 92, for a description of the metric calculation). Efficiency information is important since the agents could be required to defend the search space for long periods without the presence of targets and more efficient strategies would enable agents to remain operational for longer before returning to base for refuelling.

Search	Mean	Std. Dev.	Search	Mean	Std. Dev.
Bacteria	540.66	0.56	ACS	1009.83	0.72
Salt	691.57	7.54	Straight	1009.90	0.74
System	1002.85	0.23	F_Fly	1012.67	0.41
Spiral	1007.41	0.29	Lévy	1032.11	1.02
Random	1007.51	0.15	Brownian	1037.75	0.60

Table 2. Mean energy consumption.

In search energy efficiency tests (without the presence of targets), bacterial and saltatory search dominate due to the reduction in velocity (Table 2). Aside from those two cases, all other strategies maintained a constant speed, and, therefore, the high turn rates of Lévy and Brownian predictably had the highest energy requirements. The saccade phase of the fruit fly search imposed heavy energy needs when compared with the two deterministic patterns that were the only other strategies to include in-built straight legs. Whilst ACS and straight line searches were almost identical due to ACS utilising straight line searching during pre-detection search. Conversely, spiral and random search requirements were similar since, although random search turned more frequently, all other things being equal, it had a smaller mean turn angle. Standard deviations were all similar with the exception of saltatory search whose increased deviation was due to the stochastic, chronological switch between low and high turn rates.

It should be borne in mind that this test covers a limited period, and as search times increase, the differences would grow wider. Furthermore, in a target-laden environment ACS efficiency would increase on detection since although an agent increases its turn rate, it decreases its speed at the same time.

In real terms these results indicate that, using the energy requirements implemented for this research, there should be a direct correlation between time taken to destroy the targets and the energy requirements for the search, since no strategy has significantly larger energy consumption. To verify this, Pearson correlation coefficients were calculated for a cross section of scenarios (see Table 3) and the strength of the correlation confirms that increased performance (*i.e.* shorter times) does not come at the expense of energy.

	Single target				Four targets			
	10 at 0.5	10 at 1.5	20 at 0.5	20 at 1.5	10 at 0.5	10 at 1.5	20 at 0.5	20 at 1.5
Brownian	0.948	0.956	0.949	0.963	0.933	0.940	0.939	0.957
Deterministic	0.982	0.983	0.982	0.985	0.968	0.956	0.978	0.980
Fruit fly	0.964	0.984	0.974	0.983	0.937	0.970	0.949	0.978
Random	0.952	0.977	0.974	0.980	0.928	0.966	0.955	0.975

Table 3. Pearson Correlation Coefficients between time and energy results for swarms of 10 and 20 agents against Brownian, deterministic, fruit fly and random movement targets travelling at 0.5 and 1.5 times agent maximum speed.

5.8 Predictive Memory System Where No Occlusion or Communication Failures Occur

In previous sections, the performance of the agents has been shown to improve through the application of a variety of nature inspired search strategies without memory of targets that were lost following detection. This section reports findings of applying the PMS in conjunction with those search strategies in scenarios where no occlusion or communication failures occur; that is, detection loss only occurs due to the target moving faster than the agent can track. The five PMS strategies (predictive-C, predictive-V, direct search, crosswind casting and saccadic) were applied to examine the effect of maintaining target memory on overall performance. For this purpose, 10 agents searched for 4 dynamic targets having relative speeds greater than maximum agent speed and a range of movement patterns. Only faster than agent target speeds were considered for PMS experiments where no other factors influenced detection (*i.e.* no occlusion or communication failures) since agents are able to maintain detection of targets moving at relative speeds equal or less than the agents' capability without the use of a PMS. In summary, 10 agents applied all of the initial search strategies with LSZ neighbourhood swarming enabled and predictive-C PMS, predictive-V PMS, direct PMS, crosswind casting PMS and saccadic PMS against 4 deterministic targets, 4 Brownian motion movement targets and 4 fruit fly targets each having a relative target speed of $\{1.5, 2.0\}$.

To illustrate general findings results from all initial search strategies and relative target speeds have been combined and presented below, the full breakdown of results can be found in Appendix 3.

The first series of experiments applied the search strategies, LSZ and PMS against 4 Brownian motion targets; Figure 98 indicates that only crosswind casting PMS improved performance, whilst the remainder actually performed less well than the system with no memory. These results indicate that for higher speed targets, under circumstances where targets are only likely to 'disappear' due to their destruction, committing valuable assets to find lost targets can be counter productive since the agents would be better employed continuing to search for other targets. Crosswind casting was the exception since agents engaging in the reacquisition search continued to search space away from the predicted position and each of them could be searching a different area. That is, the search was not wasted since there was just as much chance that other targets may be found, and in cases where the target had not been destroyed, the losing agent and any other nearby agents would be likely to redetect it due to the casting movement of the search.

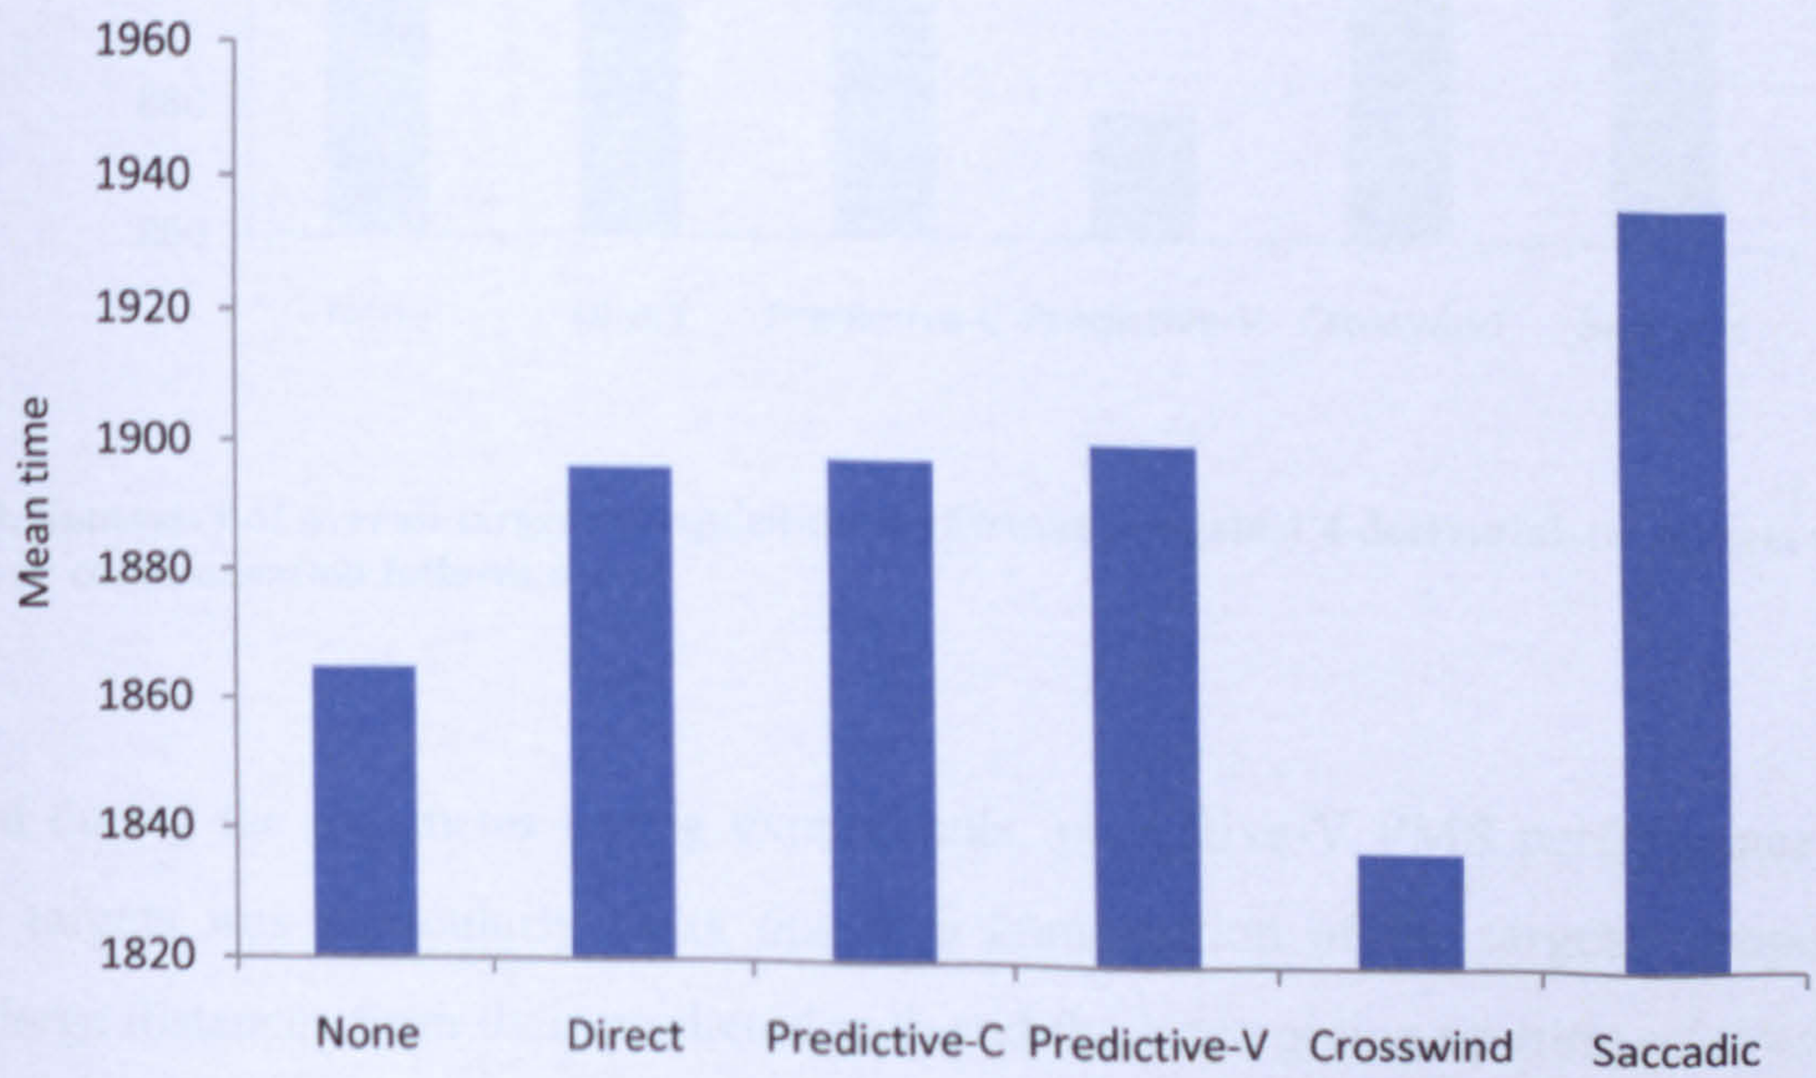

Figure 98. Summary of overall target reacquisition performance against 4 Brownian motion targets where no occlusion or communcation failures occur.

Where targets moved more predictably, the strategies that were most dependent on targets following the same path performed best (see Figure 99). It should be noted that in this

scenario, due to their more ranging movement, targets were likely to outrun initially detecting agents and as such agents engaging in reacquisition activities were more likely to looking for an existing target than where targets followed Brownian motion movement patterns.

Again the circular strategies performed worse than no PMS due to over commitment, despite shortened giving up times, whilst crosswind casting still performed better than none although by a smaller margin. This time, however, the improved performance was produced due to the occasions where the agents were moving toward the predicted path as the target passed by, allowing redetection to occur.

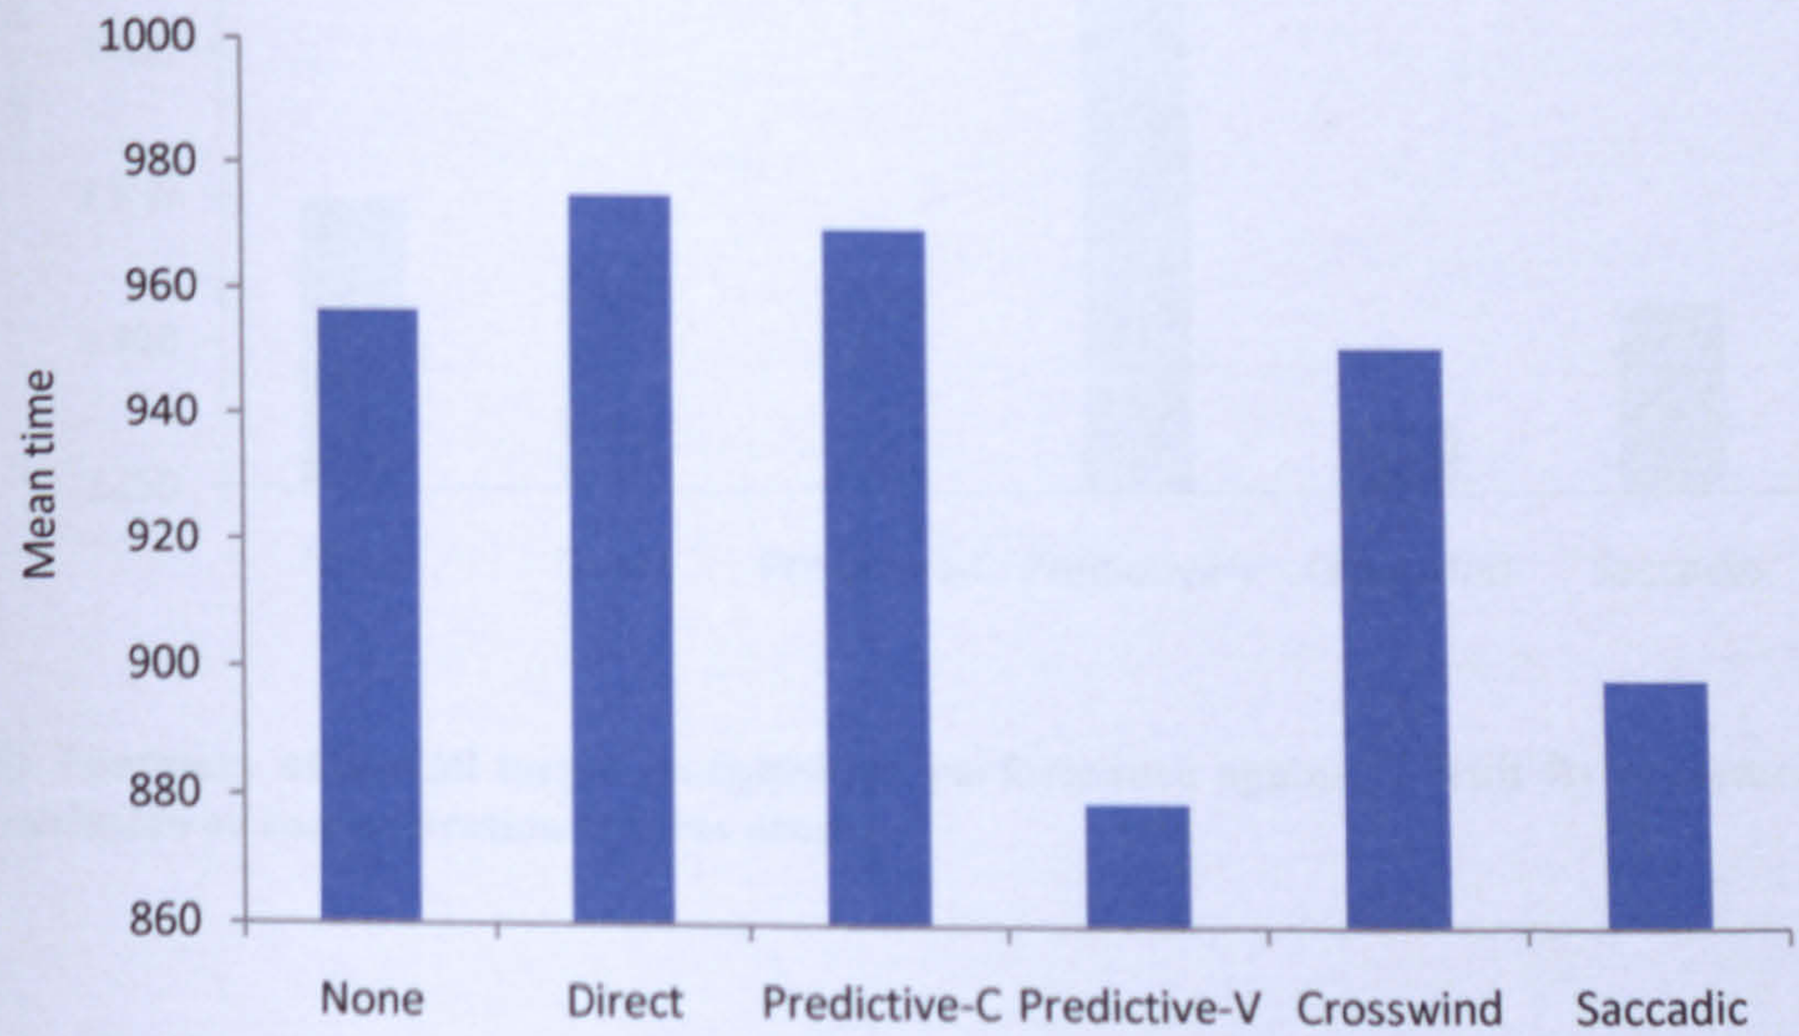

Figure 99. Summary of overall target reacquisition performance against 4 deterministic targets where no occlusion or communcation failures occur.

As noted during the parameter tuning experiments, predictive-V PMS performance against fruit fly targets was particularly weak due to a combination of the targets' propensity to deviate large distances from their predicted path and the large giving up time selected to give the advantage against predictable targets (see Figure 100). Although using the same giving up time, crosswind casting did not suffer in the same way; in fact, it was the best performer, due to the agents that were engaged in reacquisition search spending larger periods of time searching areas away from the predicted path, as noted for Brownian motion targets. The circular neighbourhood strategies (direct and predictive-C) also performed better. This was due to a combination of their short giving up time, which meant that in cases where the target

had moved out of the area following an inter-saccade move, they did not waste time searching for targets that were not there, and where targets remained in the locality but were behaving very unpredictably the circular neighbourhood meant that agents approached the area from several directions.

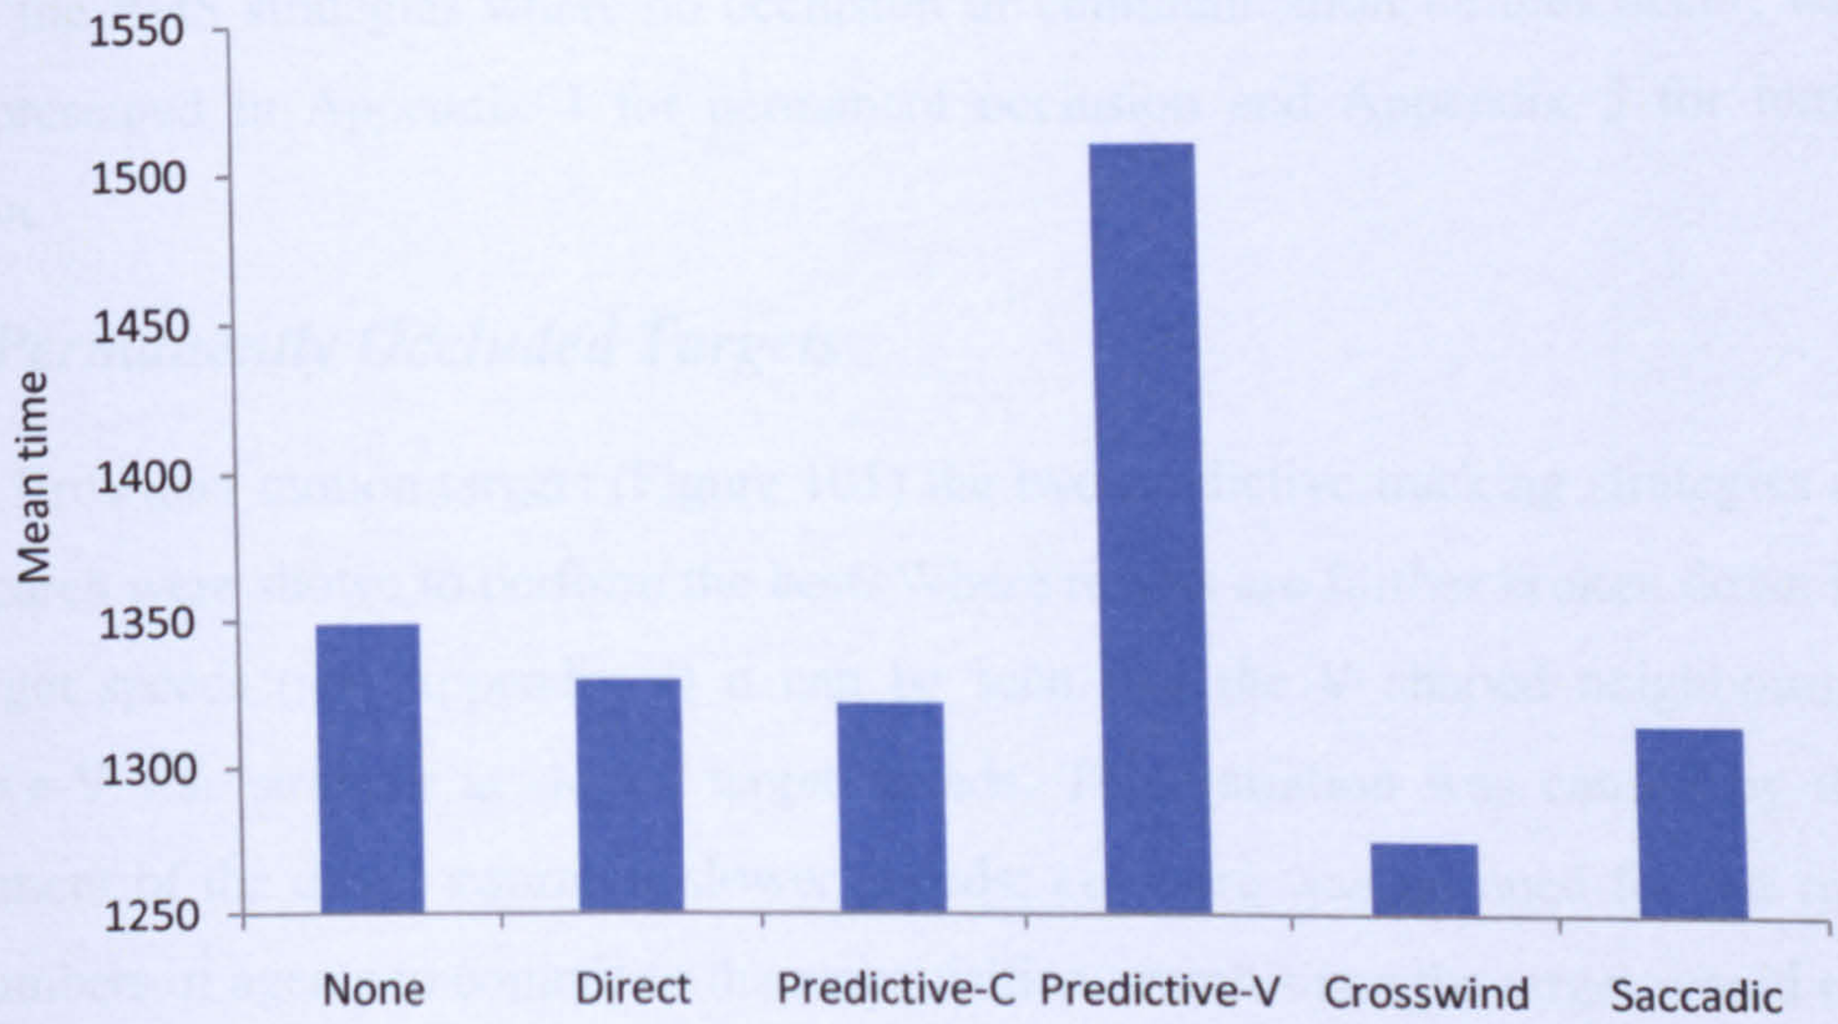

Figure 100. Summary of overall target reacquisition performance against 4 fruit fly movement targets where no occlusion or communcation failures occur.

5.9 Predictive Memory System Where Targets Become Occluded

Following experimentation without any additional factors, the five PMS strategies were applied to examine the effect of maintaining target memory on overall performance against occluded targets. Again, 10 agents were pitted against 4 dynamic targets having a range of relative speeds and movement patterns; in this part of the study, however, relative speeds of less than agent capability were included since occlusion could occur under any agent – target relationship. Upon initial detection, the targets were occluded from the detecting agent either permanently (permanent occlusion), simulating agent sensor failure during detection, or temporarily (temporary occlusion), which simulates temporary obscuration of the target to the detecting agent. In summary, 10 agents applied all of the initial search strategies with LSZ neighbourhood swarming enabled and predictive-C PMS, predictive-V PMS, direct PMS, crosswind casting PMS and saccade PMS against 4 deterministic targets, 4 Brownian motion

movement targets and 4 fruit fly targets each having a relative target speed of {0.5, 1.0, 1.5, 2.0} with permanent and temporary occlusion applied.

The following sections present the results of the application of the various PMS strategies for permanent and temporary occlusion. Findings have been combined in the same manner as the those of the PMS strategies where no occlusion or communication failures occur, with full results presented in Appendix 4 for permanent occlusion and Appendix 5 for temporary occlusion.

5.9.1 Permanently Occluded Targets

Against Brownian motion targets (Figure 101) the two predictive tracking strategies and the direct search were shown to perform the best. Where results are further broken down into the four target speeds (see Appendix 4) it can be seen that the V shaped neighbourhood of predictive-V was stronger at slower target speeds. This variation was caused by the over commitment of the direct swarm at slower speeds; *i.e.* there was no need for the relatively large numbers of agents to commit to the reacquisition search since the target would not have moved far from where it was last detected. Where speeds were increased, the target would still not have moved far from its last known location and thus the direct search was more profitable since the predictive-V would have been unreliable due to the high turn rate of the target. This same rationale explains the predictive-C strategy's competitive performance across the range of relative speeds; at lower speeds it ensured that cooperating agents were directed to the area (although moving them away from areas that may have contained as yet undetected targets), conversely, at higher speeds the prediction of the occluded target was more likely to be inaccurate and thus agents would have been directed away from the actual target location but the extended neighbourhood area increased the likelihood of the target being re-detected.

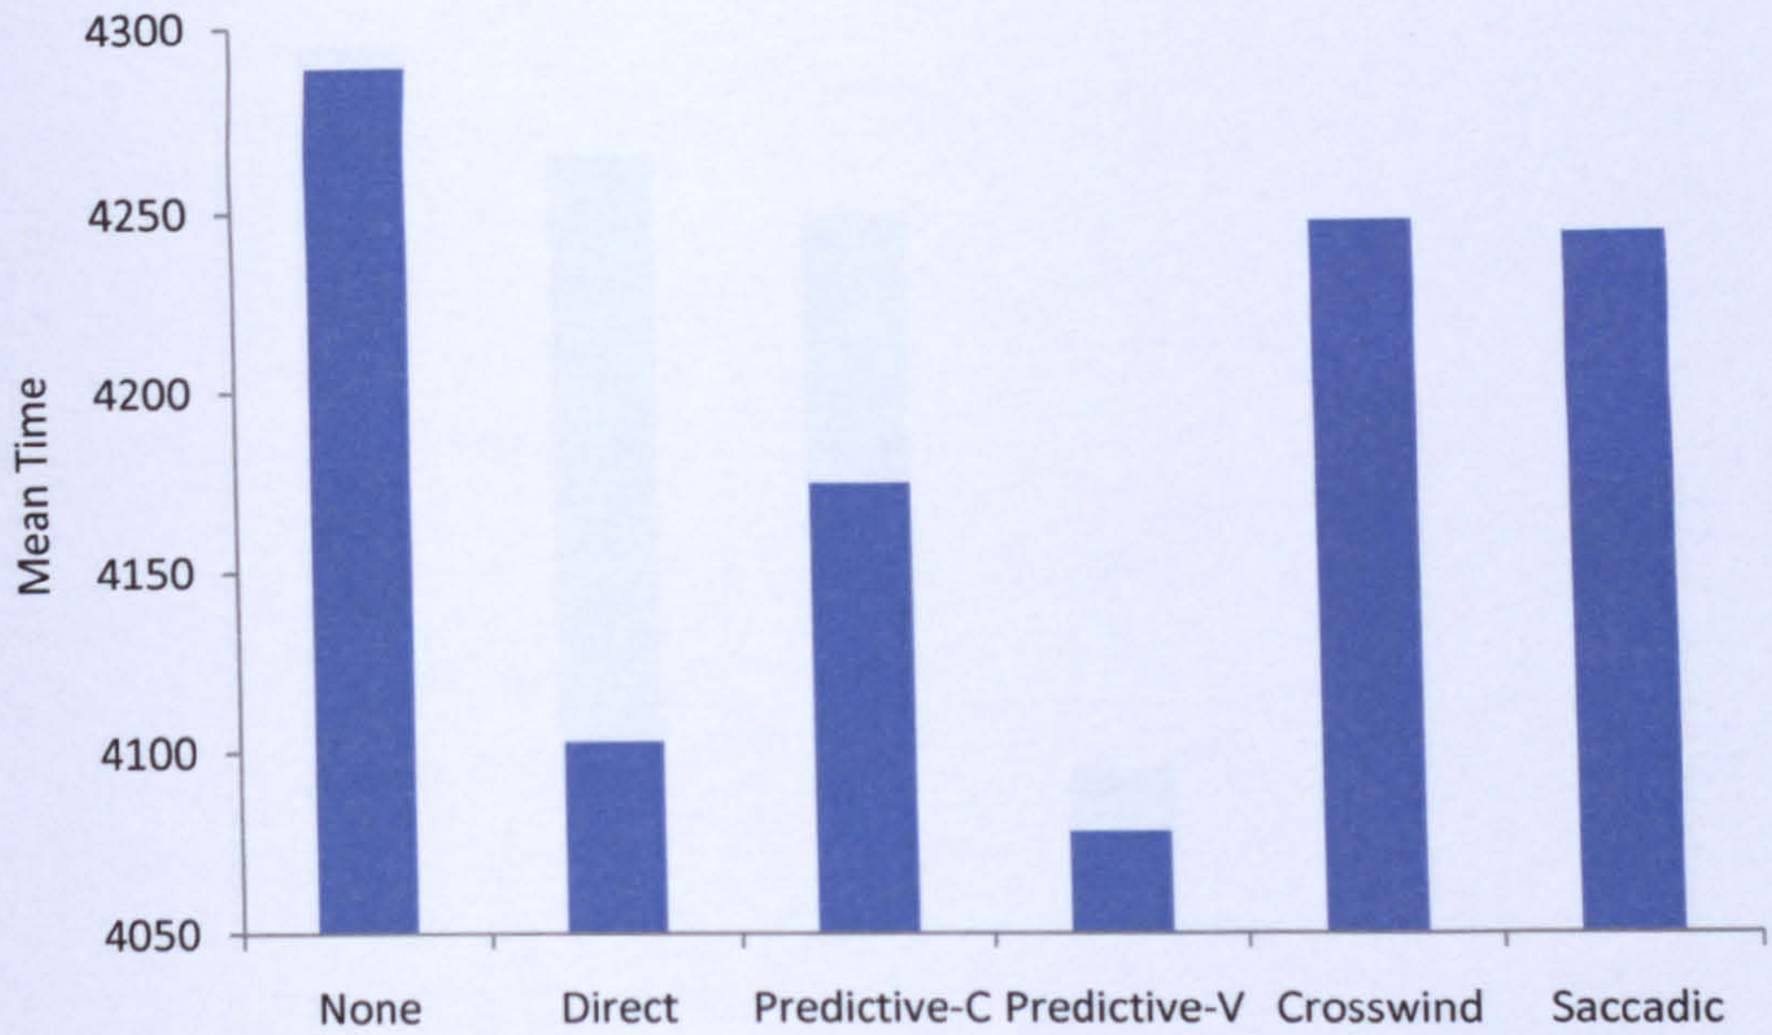

Figure 101. Summary of overall target reacquisition performance against 4 Brownian motion targets where permanent occlusion occurs.

Where target behaviour was more predictable, as for deterministic targets (see Figure 102), the strategies that only committed agents that were likely to be able to re-detect held the advantage, whilst techniques that committed agents that were unlikely to succeed. That is, the circular neighbourhoods (predictive-C and direct search), were less successful because they needlessly moved agents away from areas that may have contained other, undetected, targets.

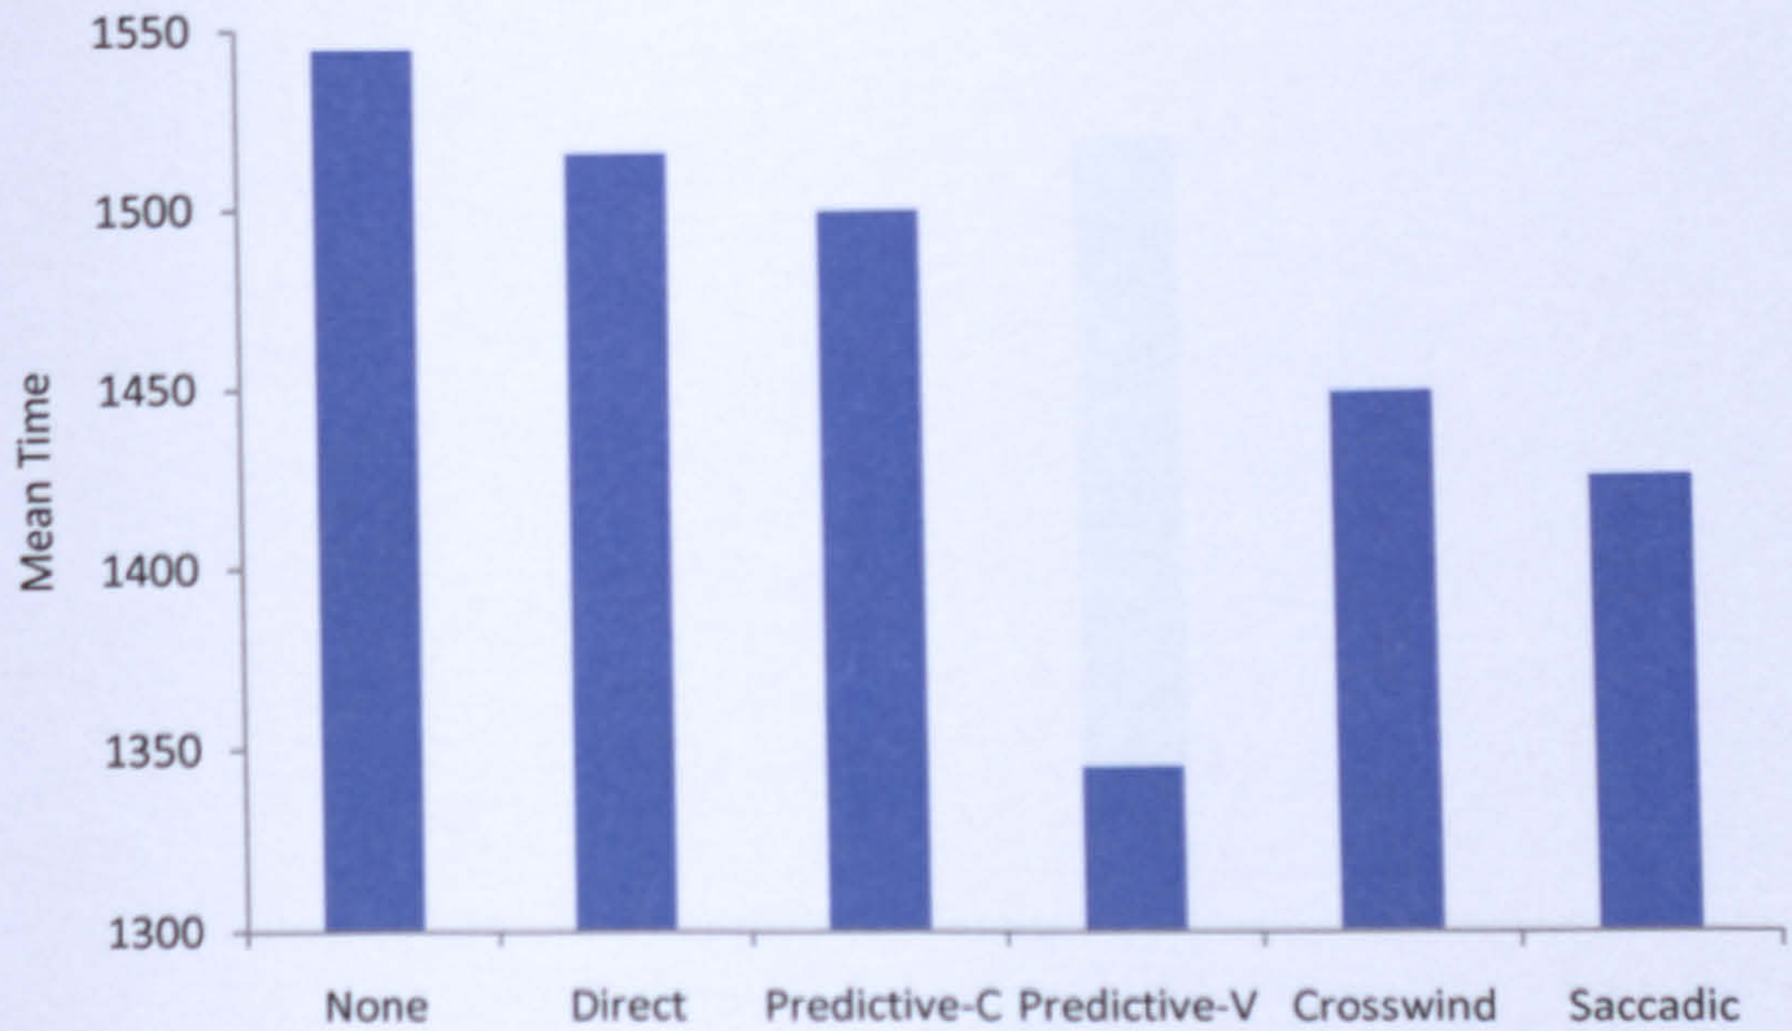

Figure 102. Summary of overall target reacquisition performance against 4 deterministic targets where permanent occlusion occurs.

Unsurprisingly, following the tuning experiments reported in Section 5.2.4, the predictive-V strategy did not perform well against fruit fly targets (Figure 103). Such performance occurred because, despite the movement type's long inter-saccade leg lengths, the unpredictable heading changes and switches to the short saccades mean the target is unlikely to be in the vicinity of its predicted location. This was in contrast to the other predictive neighbourhood, predictive-C, which, although not the best, remained competitive due to the target's propensity to switch into its intensive search mode (and therefore agents being drawn in from behind the target's last known position would be likely to detect it on their way to the predicted position, as was the case against Brownian motion targets). Agents applying saccadic reacquisition search were successful for similar reasons, except where targets remained on a long leg the agents that were able to intercept them changed their route to do so; where the target switched to intensive search – or was already in that phase when they became occluded – the agents would have moved back along the predicted track to the intensive search area, thus increasing the likelihood of re-detecting the target.

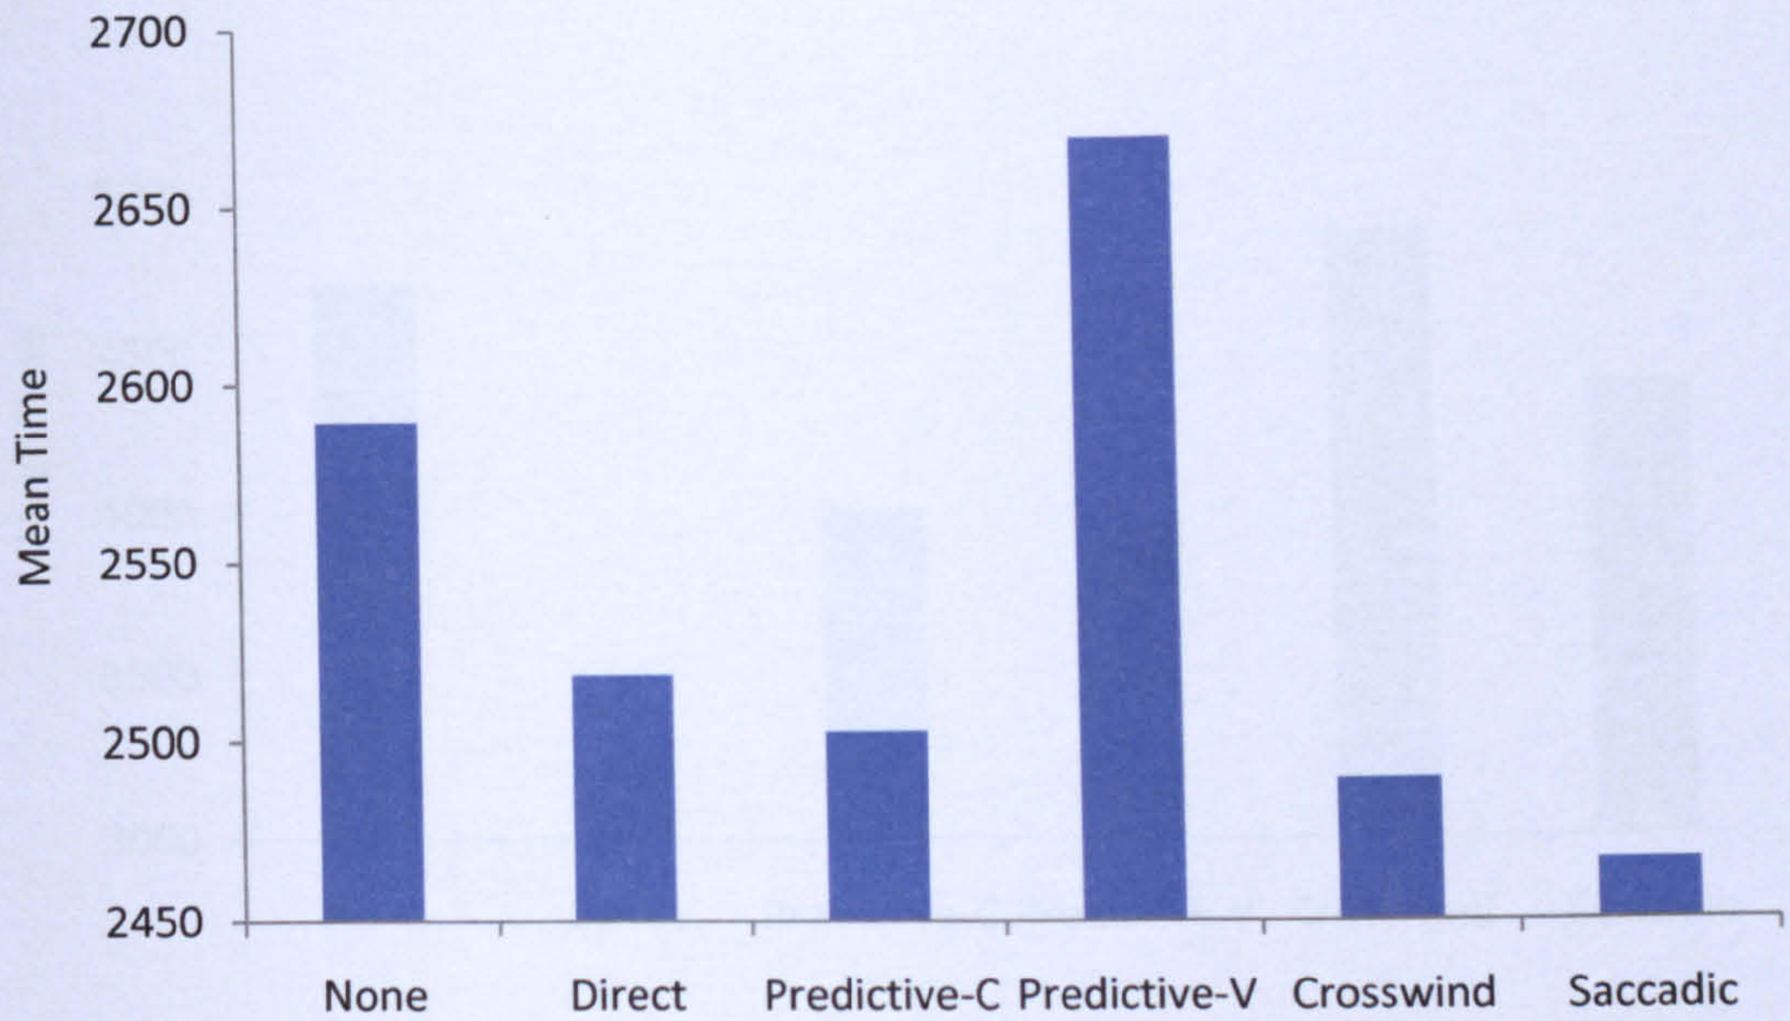

Figure 103. Summary of overall target reacquisition performance against 4 fruit fly targets where permanent occlusion occurs.

5.9.2 Temporarily Occluded Targets

When occlusion was less permanent, as simulated in temporary occlusion, the direct search became more prominent against Brownian targets (Figure 104, full results can be found in Appendix 5) because the agent that last detected the target maintained its focus on the area where the target was lost; in Brownian motion targets this was likely to remain a good area to search because the target would not move far away from the area.

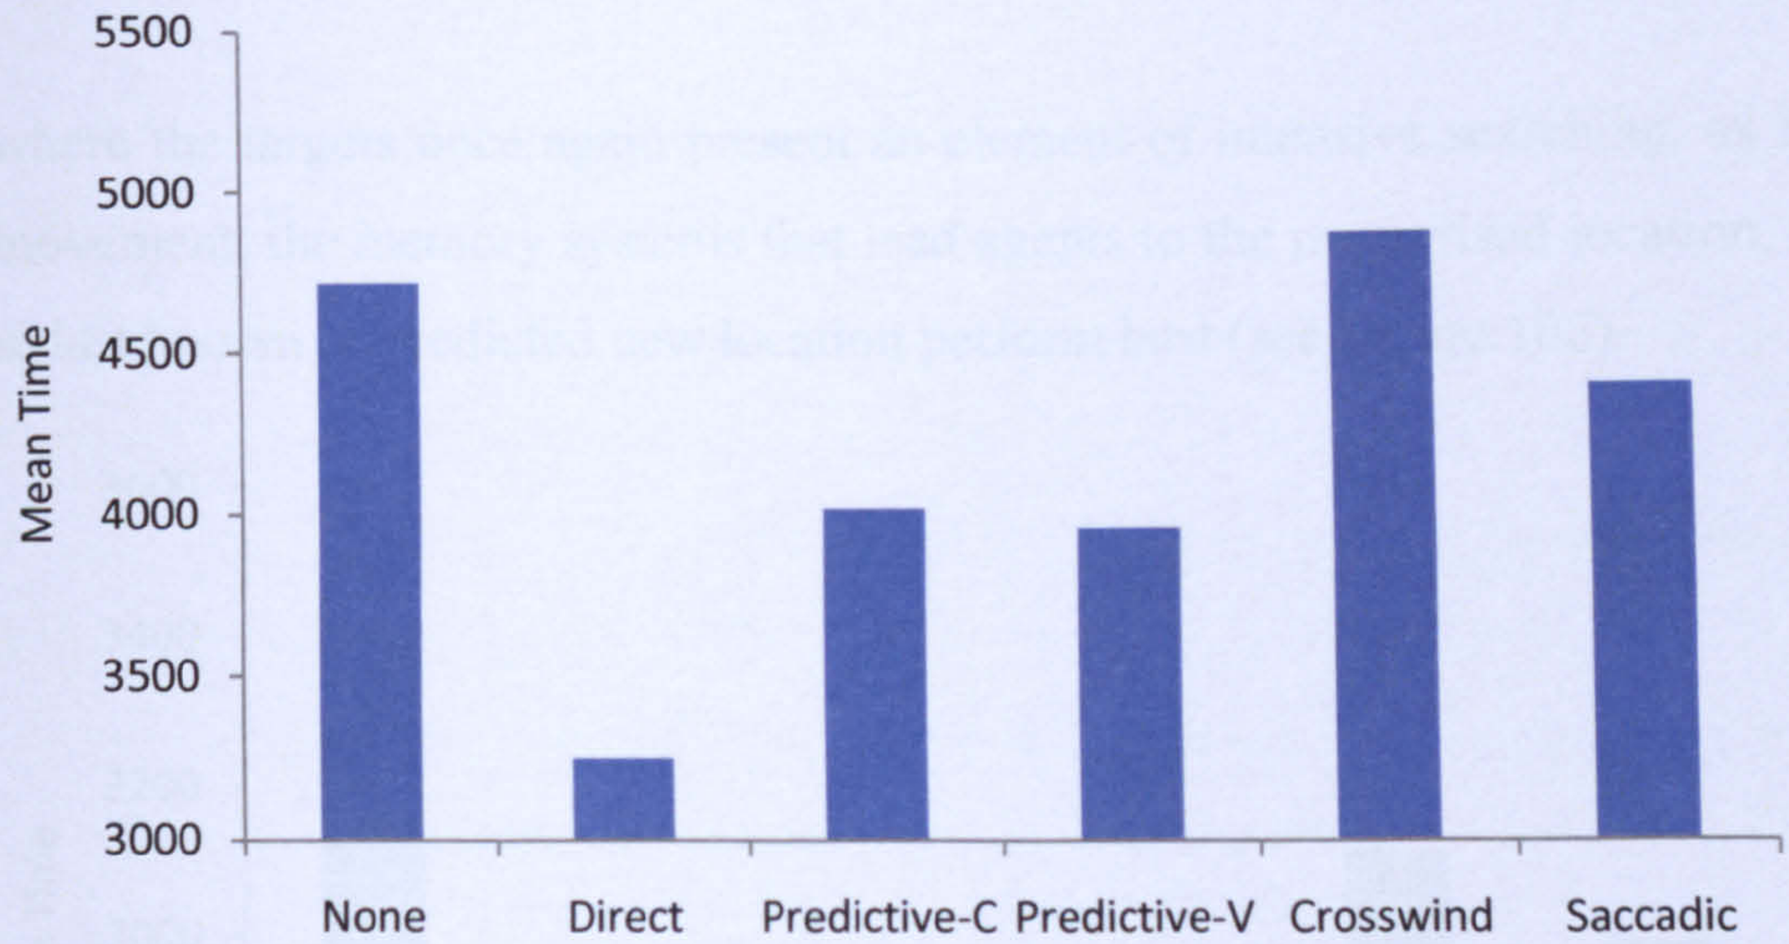

Figure 104. Summary of overall target reacquisition performance against 4 Brownian motion targets where temporary occlusion occurs.

Against deterministic targets (Figure 105), the results remained similar to permanent occlusion due to the predictability of the targets. Although, the direct search strategy did not perform as strongly for the same reason that it improved against the Brownian targets: the losing agent remained in the area, this time unproductively, where it could have moved on to detect other targets.

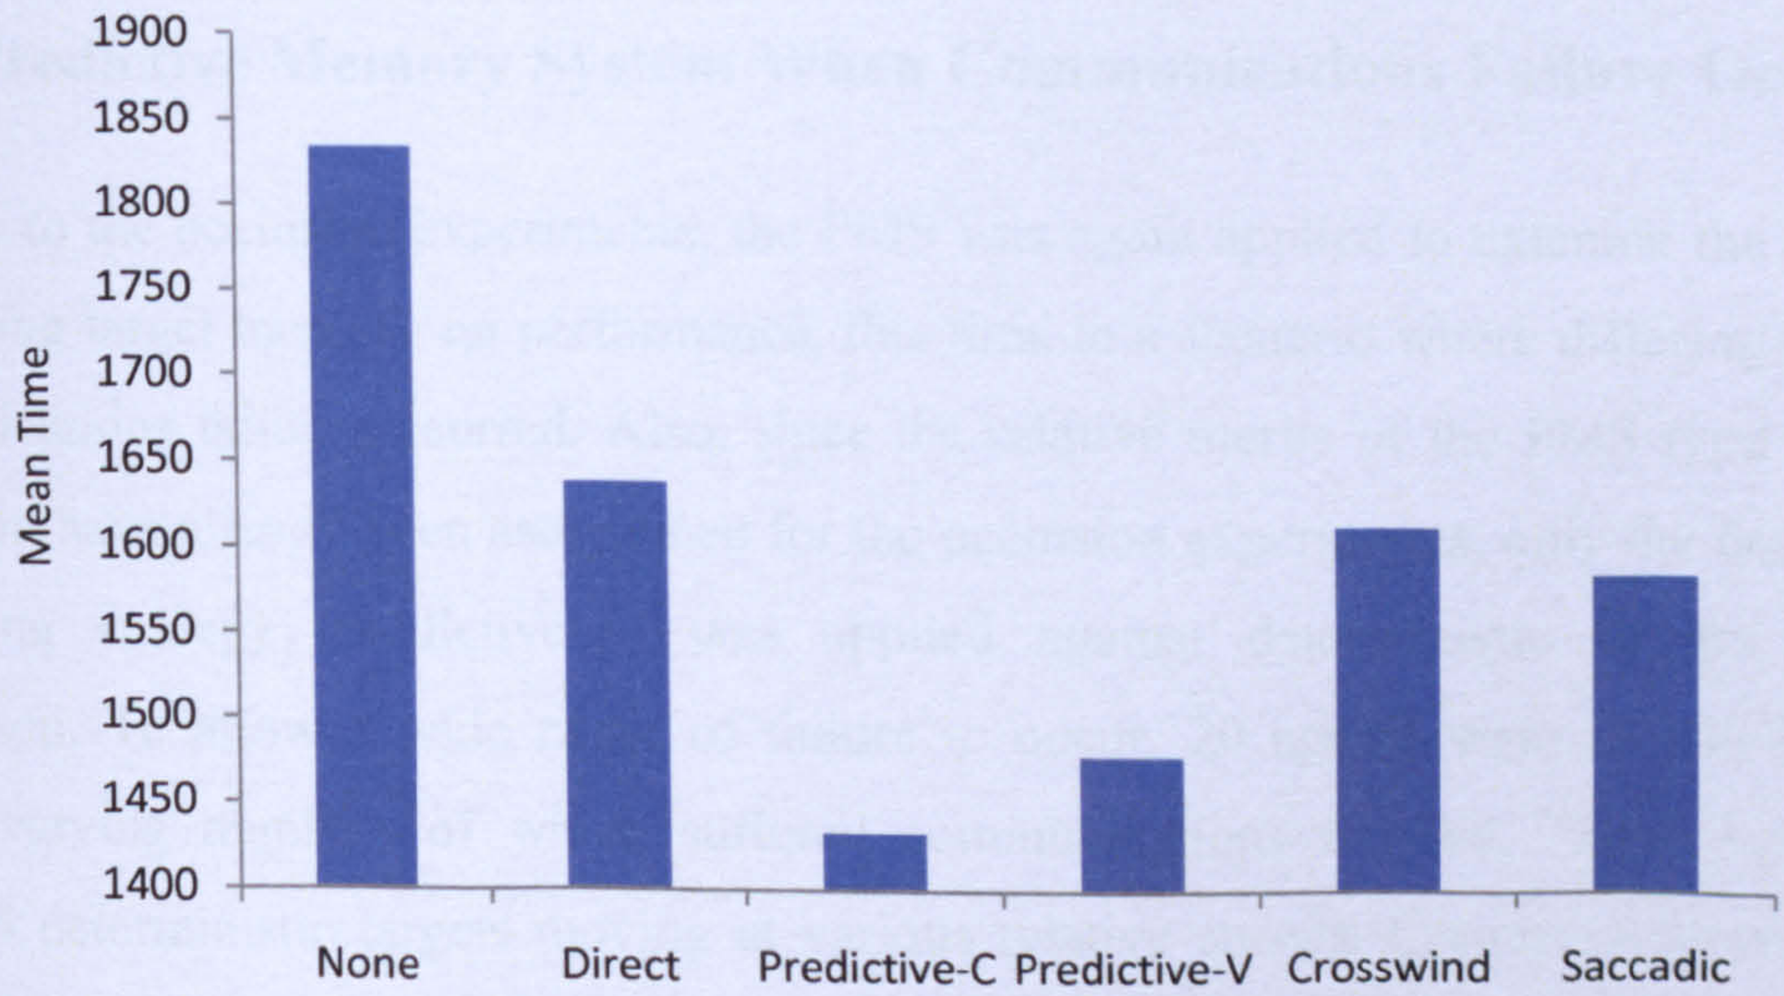

Figure 105. Summary of overall target reacquisition performance against 4 deterministic targets where temporary occlusion occurs.

Finally, where the targets once again present an element of intensive searching, as found in fruit fly movement, the memory systems that lead agents to the memorised location, whether that be the last known or predicted new location perform best (see Figure 106).

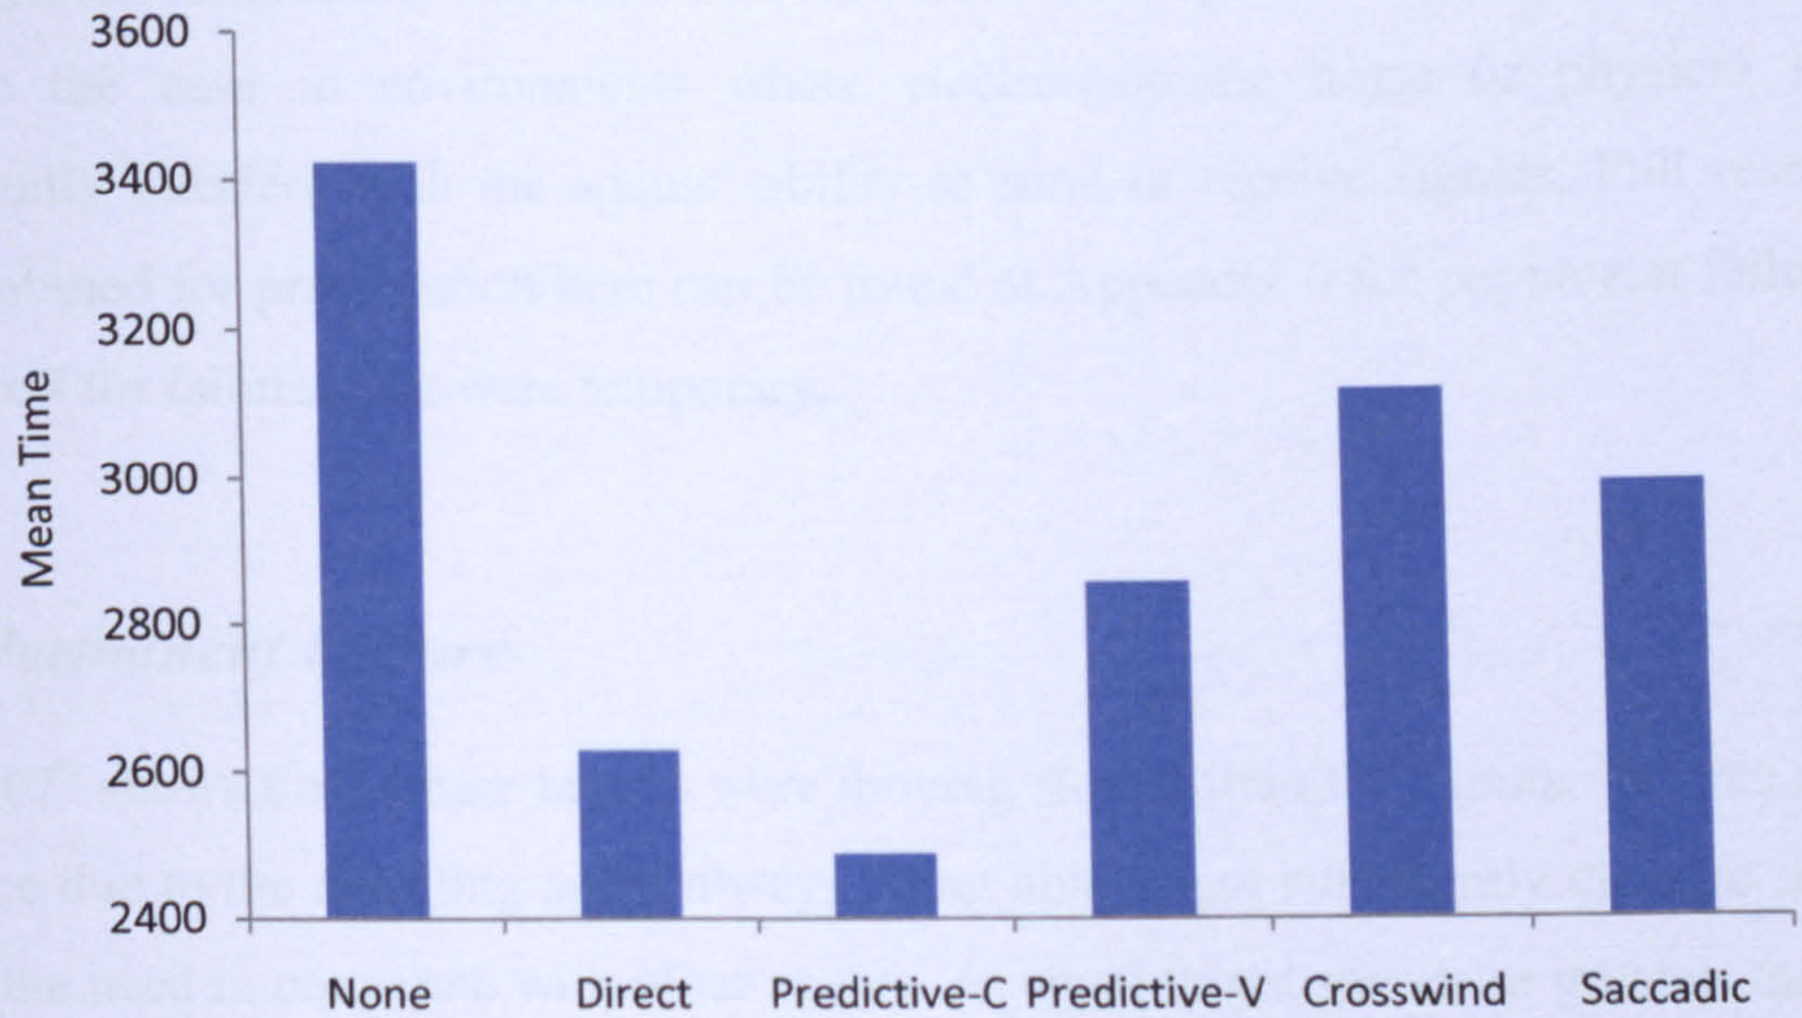

Figure 106. Summary of overall target reacquisition performance against 4 fruit fly targets where temporary occlusion occurs.

5.10 Predictive Memory System When Communications Failure Occurs

Similarly to the occlusion experiments, the PMS was again applied to examine the effect of maintaining target memory on performance, this time in a scenario where differing levels of communications failure occurred. Also, since the relative merits of the PMS type to target movement have already been ascertained for the occlusion experiments, only the best overall performing strategy (predictive-V) was applied against deterministic targets for this comparison. To allow a wide range of failure to occur, 20 agents were employed in the swarm, varying numbers of which suffered communications failures. These were pitted against 4 deterministic targets moving at various relative speeds. Communications failures were either intermittent or permanent; where they were intermittent each random outage period (in the range [10, 20]) was followed by an equal period of serviceable communication. In summary: 20 agents applied all of the initial search strategies with LSZ neighbourhood

swarming and predictive-V PMS enabled; 4 deterministic targets moved with a relative target speed of {0.5, 1.0, 1.5, 2.0}; and temporary and permanent communication failures occurred where 4, 8 and 12 agents were affected.

Subsection 5.10.1 provides results where differing numbers of agents have permanent communication failures and Subsection 5.10.2 those where agents temporarily lose contact as might be the case in environments where electromagnetic noise or physical features intermittently interfere with the agents' ability to send or receive signals. Full results that were combined for presentation here can be found at Appendix 6 for permanent failures and Appendix 7 for failures that were temporary.

5.10.1 Permanent Failure

Figure 107[6] shows that, where targets were moving slower than the agents, failures made no difference due to the detecting agent always being able to get sufficiently close to the target without the need to cooperate with other agents. At equal target speeds or greater, the system always took longer to detect and destroy targets with fewer cooperating agents (those with serviceable communications systems). At greater speeds, the cooperation requirement is obvious since the target could simply outpace the detecting agent; at equal speed, cooperation was still required since the only time an agent could reduce the distance between it and the target sufficiently for the agent to get close enough to destroy it was when the target turned and the agent could 'cut the corner'. Furthermore, and possibly counter-intuitively, at greater speeds, the PMS had a greater advantage with fewer failing agents. This was because the PMS can only operate when potential cooperators have received a signal informing them of the target presence prior to it escaping detection, thus with more communicating agents the likelihood of this occurring was greater.

[6] It should be noted that the significance analysis, reported in Appendix 11, indicates that the differences between the PMS and the non-PMS results at equal and lower speeds were not significant.

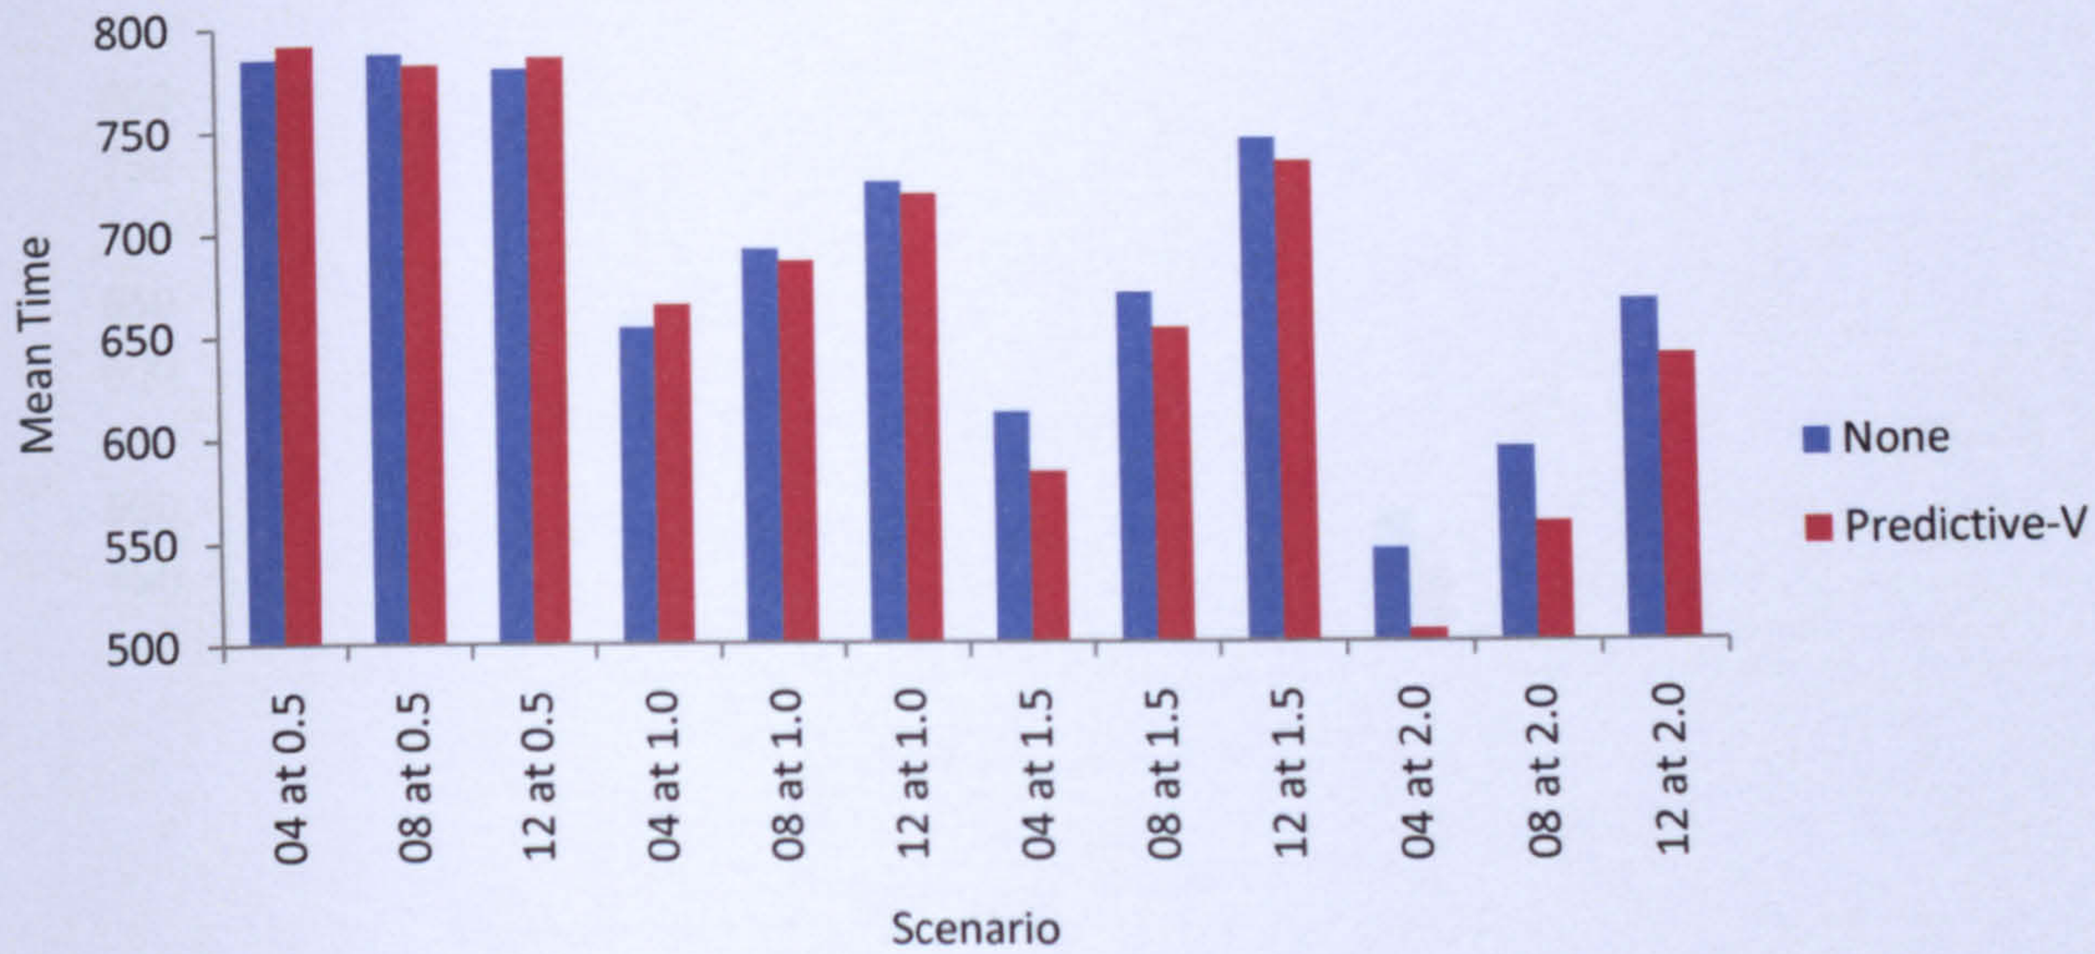

Figure 107. Summary of performance where permanent communication failure occurs in 4, 8 and 12 agents against targets moving at 0.5, 1.0, 1.5 and 2.0 times agent speed.

5.10.2 Temporary Failure

Where temporary failure occurs (see Figure 108), performance is not improved at target speeds of less than or equal to that of the agent. At slower speeds, cooperation was not required and at equal speeds the target could not out pace the agent and therefore intermittent communication was sufficient to enable cooperation. Once targets could out pace agents the memory system held a clear advantage since, in scenarios where target speed was faster, reacquisition can assist agent cooperation because their continued active pursuit during the intermittant communications loss places them in a stronger position to intercept, or at least redetect, the target. Notable in Figure 108 is the PMS strategy's consistent performance regardless of the numbers of failing agents. This was because of the intermittent nature of signal loss, when communication was restored the agents could update their knowledge and behaviour based on known data rather than prediction.

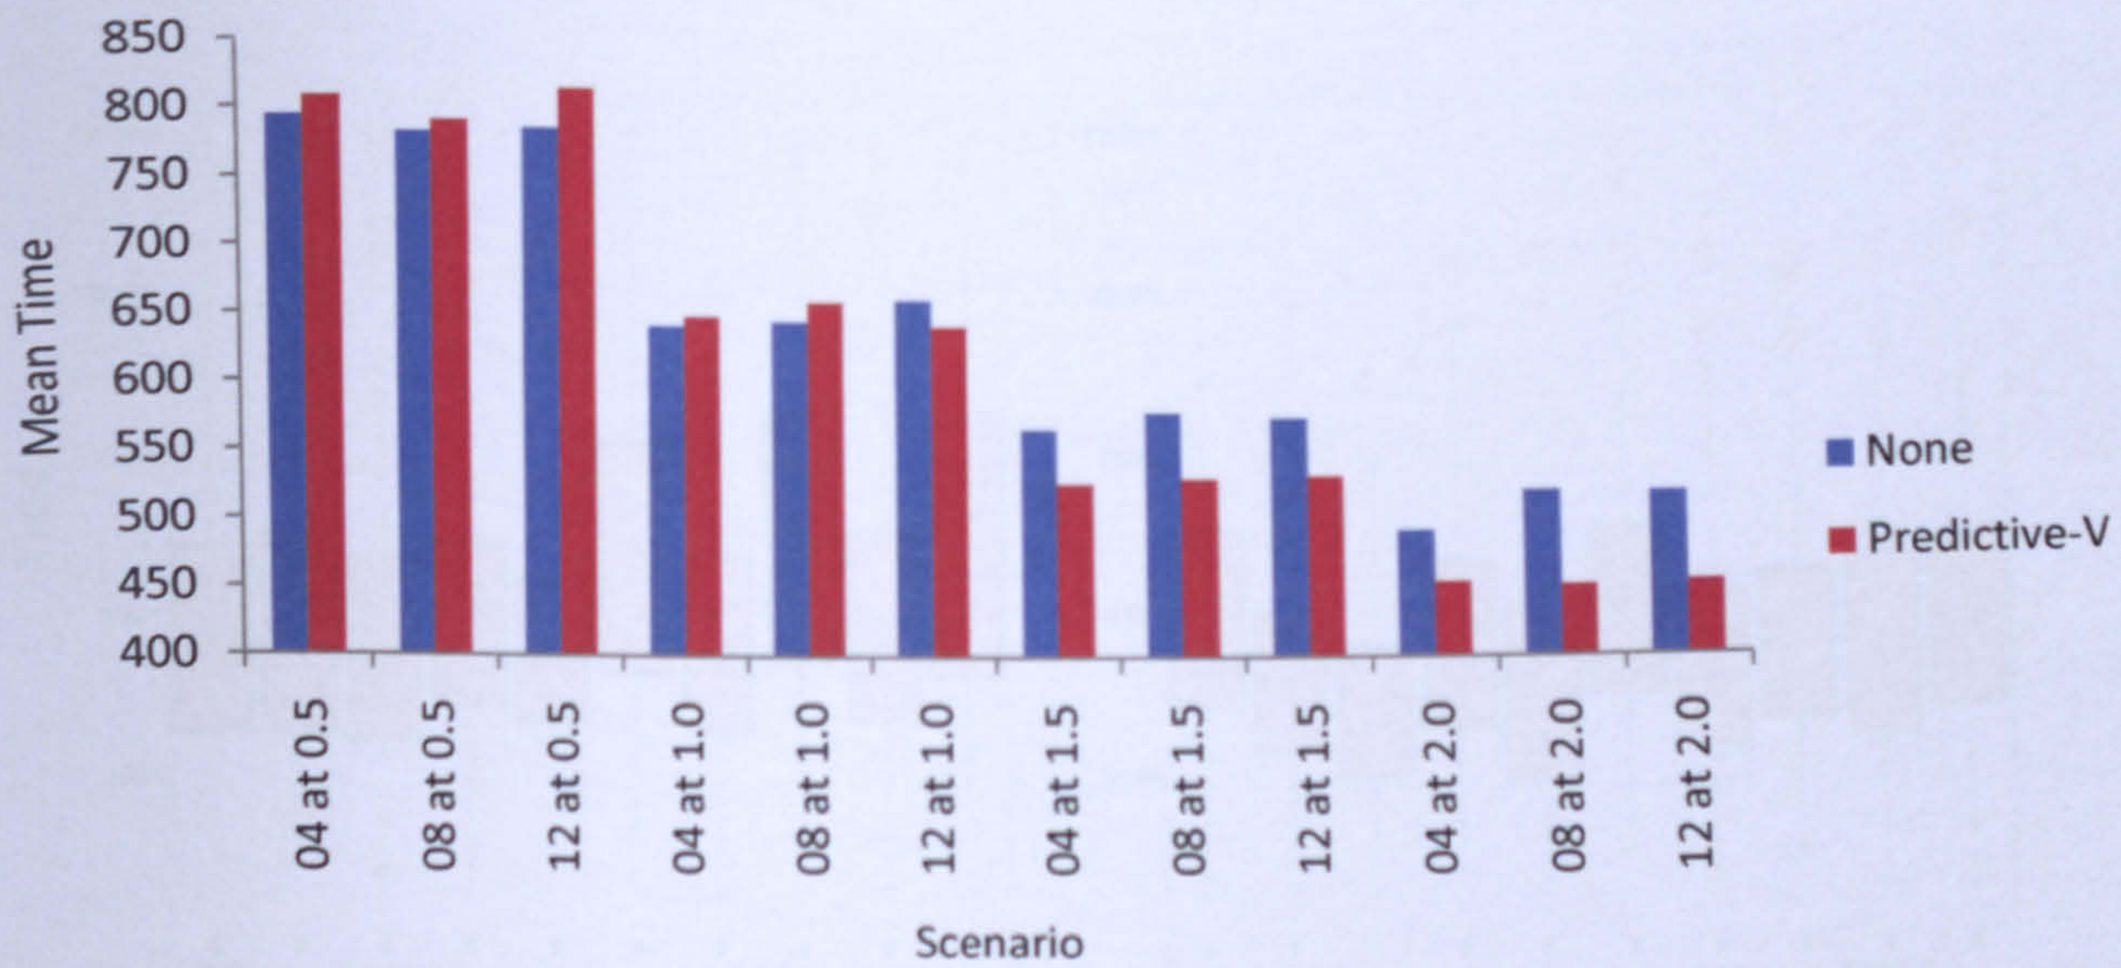

Figure 108. Summary of performance where temporary communication failure occurs in 4, 8 and 12 agents against targets moving at 0.5, 1.0, 1.5 and 2.0 times agent speed.

If the actual results for an indicative scenario are examined (Figure 109), it can be seen that some initial search strategies benefit more than others from the application of the PMS. This is because strategies such as straight-line (and therefore also ACS, because it initially applies straight-line search) and fruit fly search would continue on the same path for a short period of time following loss of detection. Hence, cooperating agents that behave in this manner, *i.e.* continuing on the same path during intermittant detection loss, have a secondary benefit of increased likelihood of intercepting targets where information regarding their location is lost intermittantly without having to apply a PMS.

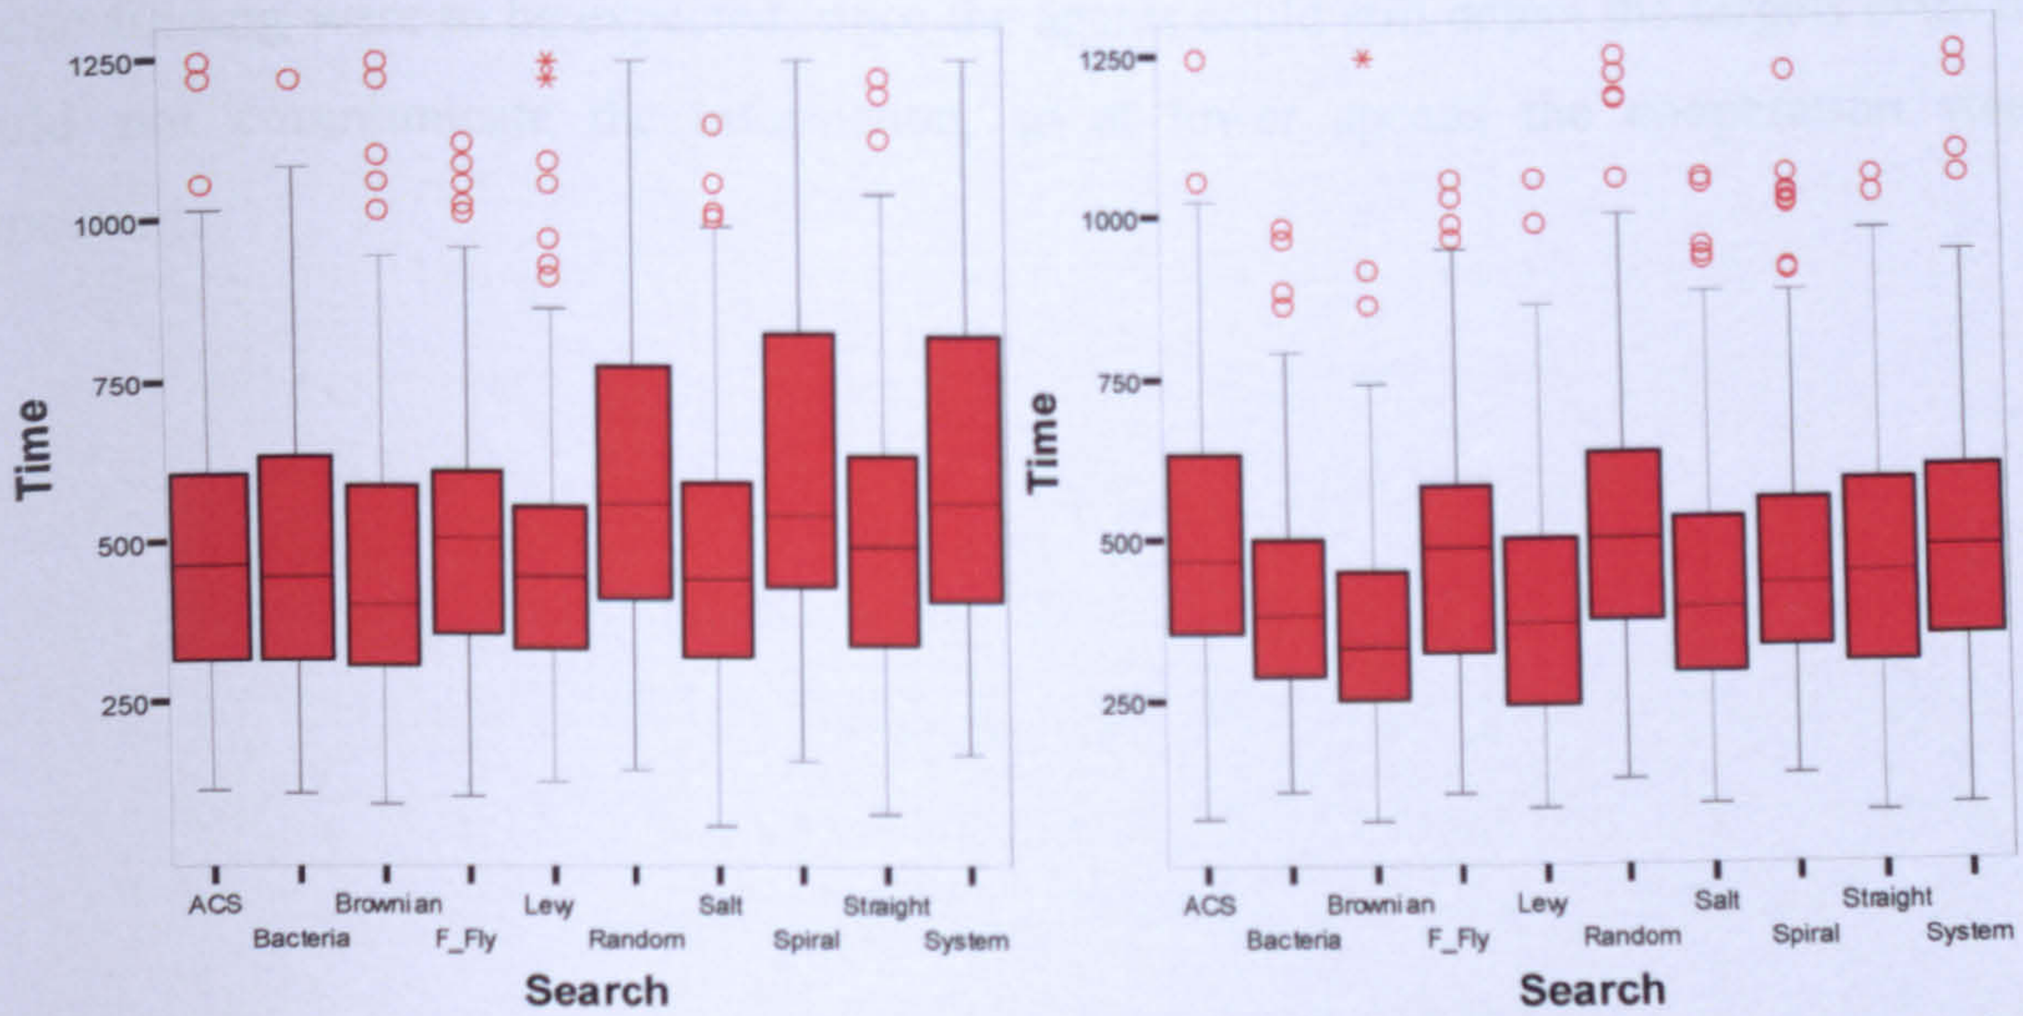

Figure 109. Results for 20 agents against 4 determinstic targets travelling at 2.0 agent speed where temporary comunication failure occurs in 8 agents.

5.11 Statistical Significance

Statistical significance analysis was performed on the test results as described in Section 4.6 (page 92) and tables of significance test outcomes are included in Appendix 8 to Appendix 11. Findings are summarised here.

Results for all static target experiments were significant, as were dynamic targets with the exception of the results where the difference between better and worst performers were compressed, for example, in cases where targets applied Brownian motion movement. In such instances, the differences between best and worst performers were always significant; however, comparisons between strategies that were less differentiated became less significant, sometimes to the point of being insignificant.

All efficiency results were significant, as were the PMS results where no occlusion or communication failures occurred, with the one exception of none against predictive-V where targets were moving using Brownian motion. Similarly, results were significant for permanent and temporary occlusion. Communication failures were found to provide the least significant results of all the experiments, with relative speeds of 1.0 or less not being

significant, and only the largest relative speeds always producing significant differences. These finding were to be expected, since the agents could still detect the targets even if they could not communicate the information, so at lower speeds the cooperation was less important.

6 Discussion

This section reviews the findings of the research. It should be borne in mind that these findings can only be regarded as general trends that should hold across a broad spectrum of autonomous vehicle search. In real-world applications a variety of practical considerations, including environmental conditions, vehicle characteristics and sensor capabilities would need to be made.

Before pursuing the main body of the investigation, a variety of tuning experiments were carried out. This work in itself yielded some surprising results, particularly the sensitivity of the V shaped PMS neighbourhood compared to that of the LSZ neighbourhood. When viewed as if the predictive tracking of the PMS were an agent tracking the target and communicating the target location to cooperating agents, one would expect the optimal settings for both mechanisms to be similar, given the identical environments, with the PMS possibly requiring a wider arc of inclusion to allow for target heading changes. This did not prove to be the case; both wider and narrower angles producing only small gains that were very different from the optimal, whereas in the LSZ mechanism the settings produced much smaller differences from the optimum. The difference between the two can be explained by the fact that in the LSZ case, the target is definitely present. Agents finding themselves in the LSZ have a good chance of detecting the target, wider or narrower angles simply reduce the availability of agents for further target detection (less agents in initial search mode, although once the target is destroyed they immediately return to initial search mode) and reduce the likelihood of an agent receiving cooperative assistance respectively. When the PMS was in operation, the target was not guaranteed to be at the predicted location and, therefore, whilst a wider angle may have assisted when the target did not remain on the predicted path but only moved a short distance away, *e.g.* for a switch to intensive search, where the target changed direction completely, the additional resources were being committed needlessly. To compound matters further, the agents were committed to the search for the entire giving-up period. When the angle was narrowed the effectiveness was reduced due to the unpredictability of the target; common sense would dictate that where targets can be guaranteed to follow a given path then narrower angles would be suitable.

During GUT tuning it became necessary to switch from fruit fly target behaviour to deterministic movement. This was again due to unpredictability; where an occluded target is

unlikely to continue on its path it becomes more effective to give up quickly and re-allocate resources to initial search behaviour than to continue with prediction. Following the switch of target behaviour to a more predictable movement pattern the optimal GUT increased substantially. These findings held for V shaped neighbourhoods (predictive-V and crosswind casting); where circular neighbourhoods were employed, both direct search and predictive-C, shorter GUT times remained optimal. For direct search, this was because the search focuses on where the target was last detected; the only requirement is that agents conduct a passing search of that area. These findings are in keeping with observed natural behaviours such as observed with the short-tailed hawk (Ogden, 1974). In circular predictive tracking, predictive-C, the GUT time had been expected to increase with target predictability inline with that of the predictive-V neighbourhood. However, this did not occur because the additional resource commitment required by the neighbourhood in situations where targets were no longer present, or unassailable because they were moving away from the agents, incurred a penalty for the extended GUT that made it less attractive.

Moving on to the findings from the static target search, results indicate that where items are clustered, the switch to an intensive search strategy increased performance markedly. However, the difference in performance between the locate-only and locate-and-destroy results for the ACS indicate that it is not effective to switch strategy until any other potential targets are likely to fall within sensor range. This was the factor that separated the ACS performances in the two experimental settings; during locate-only the agents switched to intensive search and unless the additional targets happened to be close by, the intensive search did not range sufficiently for the additional clustered targets to be detected. Where the agents were required to get closer to the target to destroy them the switch occurred nearer to the centre of the cluster and the short ranging nature of the intensive search was sufficient to detect targets most of the time. Comparison of actual results (as opposed to mean rankings) between locate-only and locate-and-destroy indicated that where search strategies did not move agents across the search space quickly, for example in Brownian motion search, if targets were clustered the strategies benefited from being drawn into the cluster. This was because once the target was destroyed, the agent returned to the high turn rate movement and generally located other nearby targets.

In ACS, the switch to a more appropriate search technique was based on the knowledge that there was an increased likelihood of target presence; however, the switch does not have to be

restricted to target detection. Saltatory search in the simulation was time based, and whilst this does occur in nature, as seen in polecats (Lode, 2000), other factors could be used. For example, Zamon (2001) describes seals switching to a more localised intensive search based on tidal patterns (to exploit salmon runs), whilst some animals switch to stationary intensive search due to sensory deficiencies whilst moving (O'Brien, 1989) and others due to environmental factors such as vegetation densities, as occurs in marmosets (Souto, 2007). As noted in the Section 4.2, saltatory search was denied its ideal environment where targets are hidden. Considering this, the strategy performed consistently across the range of agent – target relationships investigated. This consistent 'mid-table' performance was due to the two-phase nature of the behaviour, with the extensive search using straight-line search and the intensive phase using Brownian motion. It would have been possible to improve the overall performance using a Lévy path intensive search, which would have maintained some plausibility but would have moved away from the literature and arguably could also have been misleading since the marginal difference in performance gain in this environment could be lost in a target environment where the more intensive nature of Brownian motion could be more effective in finding hidden targets in, for example, dense vegetation.

When searching for dynamic targets, the results could be considered counter-intuitive in that where a target ranges about the search area at faster speeds it becomes more effective for the agent to range less and expend more energy in ensuring that when a target does enter within sensor range that it is detected, *i.e.* the sensor is covering the maximum new area in each move. This should not be confused with the overall search coverage as described in Section 5.5; in the case under discussion it refers to the amount of coverage within a small number of time steps. For a given time step there will be an amount of repeated sensor coverage; where an agent is travelling in a straight line the repeated coverage will be higher than if it turns (assuming sensor range is longer than its field of view). Where a target is moving quickly through the search area it is more likely to be detected the more sensor coverage that is achieved during the short time it could be detected. This relationship holds across the entire range of dynamic experiments: it is the amount of target ranging across the search area rather than the actual target speed that dictates the most appropriate strategy to detect it. A comparison of the strategies' performance against deterministic targets and Brownian motion targets highlights this.

The ratio of swarm size to target numbers was also found to influence relative strategy effectiveness, which was particularly evident against random moving targets. Compare Figure 71 and Figure 72 (page 128), where high turn rate strategies became more effective at much lower speeds when only one target was present than against four. Observations of the searches in progress revealed that this was because there was a higher chance of at least one of the targets operating in a sparsely populated area of the search space. This was supported where greater numbers of agents were in use; the crossover point where high-turns become more profitable occurred at slower speeds because there was less chance of targets being remote from all agents. This was particularly true when the bacterial movement search strategy was in use because of its reduced speed prior to target detection.

When agents applied a Predictive Memory System to attempt target reacquisition, with no occlusion or communications failures occuring, it was found that oftentimes the PMS retarded search performance. This was usually where the target was moving unpredictably and fast such that where detection was lost the agents prediction of target location was innaccurate. Where targets behaved more predictably, strategies that had been tuned to take advantage of predictability increased agent performance dramatically. Conversely, where that predictability was not forthcoming strategies that relied less heavily on prediction accuracy, particularly crosswind casting, held the advantage.

During permanent occlusion, and where the target was following Brownian motion, it was seen that all PMS types improved performance over using no memory. Intuitively for Brownian motion targets one might have expected a direct strategy to be the best performer since the target would not move far from the last detection location. However, the agent to target ratio was low and therefore the neighbourhood shape could commit a relatively large number of the agents to the reacquisition search. In the predictive-V, although the target may have changed direction, it would not have moved far, therefore, the co-operating agents would move into the right area, and as such, even where the reacquisition search was unsuccessful the switch back to the initial search could re-detect the target. These arguments are supported by the performance of the predictive-C system which, having the drawbacks of both related systems, performed worse than both. The two other systems, crosswind and saccadic suffered because of the unpredictable nature of the target, the former would not move the agents toward the target's location and the latter altered the agents' route down the route least probable to encounter the target for the target movements adopted.

Against deterministic targets the predictive-V continued to dominate since the agent was usually predicting the target's path. Crosswind and saccadic systems performed relatively better because the target was now moving toward the agents. Unfortunately, for the crosswind system the agents were often not facing the target as it passed by. In nature, crosswind casting is used to re-detect targets that have a broad signature (*e.g.* a plume of odour) by changing to a perpendicular movement the animal increases its chance of regaining detection. Where the target is relatively small, as in the simulation, such behaviour actually reduces the possibility as the target passes by. A further noteworthy point to be made regarding the positioning of the agent during crosswind casting (in a real implementation) is that it could place the vehicle in a position vulnerable to counter-attack by the target since the search places the vehicle perpendicular to the target track. The direct and predictive-C systems once again suffered due to the over commitment of system resources, with the latter performing better due to the generally correct prediction of the target location.

The surprisingly poor performance of the V shaped system against fruit fly targets where permanent occlusion occurred could be regarded as a warning about over-tuning a system to a particular type of target. During the tuning experiments the GUT was found to be very short against fruit fly targets but very long against deterministic targets. Since the agents did not have a learning mechanism that allowed them to select an appropriate GUT based on the target movement type they applied an inappropriately long amount against targets that were often not following the path they were last detected on. This was not so problematic for the saccadic system since those agents gave up after finding the target was not following the predicted path from its last known location. However, oftentimes the fruit fly targets did apply long leg lengths which meant that the saccadic system was worthwhile. The other three systems, whilst better than nothing, still suffered from the same issues as discussed for other target types (over commitment, agents facing wrong direction, *etc.*).

When an agent could re-detect the target it had lost, results changed significantly. Against Brownian motion targets the direct approach worked best because the agent remained searching in the area that it last detected the target and when the occlusion ceased it was able to attack; of course, the same was true for permanent occlusion except that re-detection never occurred and the opportunity was lost.

Continuing with temporary occlusion scenarios, where deterministic targets were operating, the predictive-C system held the advantage over the predictive-V system because when an

gent detected a target that was moving away from it, remembering that results include ırgets moving slower than agents, it remained in the reacquisition neighbourhood; in the ›redictive-V system that was not the case and even when an agent was capable of 'chasing lown' a target it had lost it would not attempt to because it was not in the qualifying ıeighbourhood. The remainder of the systems held their relative positions.

[n the final occlusion based PMS experiments, against fruit fly targets where temporary occlusion occurred, the circular neighbourhoods were best performers. Again, this was because of the inter-saccade phase where targets did not range far (in the case of direct search) and the ability of the predictive-C scheme to chase down targets travelling away from them at the time of occlusion. V shaped neighbourhood systems did not suffer as much where targets could be reacquired by the losing agent because quite often the agent would be in the neighbourhood and be able to reacquire the target before it moved out of sensor range.

Looking at situations where communications failure occurs, the benefits of applying a PMS are diminished where target speed is less than that of the agent. At these lower target speeds, the initally detecting agent has the capability to continue with the interception without the assistance of other agents; it is at faster relative speeds where cooperation becomes necessary. Where cooperation is required, the use of the PMS allows cooperating agents to continue with the interception, with the same caveats as placed on reacquisition with occluded targets. That is, the success of the strategy is dependent on target predictability and speed, and, where a PMS is applied unnecessarily, the commitment of cooperating agents can reduce the overall performance of the system because those agents would have been better used searching for other possible targets. Where the information supplied to cooperating agents is intermittant, it was found that some strategies did not benefit from applying the PMS, since their intial search behaviour ensures their movement continues toward a target and the intermittant information is sufficient to keep them on course to intercept the target.

Efficiency of operation is vital for all vehicles working in hostile environments, even more so for autonomous vehicles. Indeed, in an interview with Avionics magazine (Carey 2007) Dr Darryl Pine, Professor and Chair of the Department of Aerospace Engineering at Maryland University, argued that the biggest challenge in developing NAVs was the "efficiency in power conversion, whatever the power source...". In this work, a generic vehicle model was implemented based loosely on existing aircraft heuristics. These would not hold in many target applications and could vary largely; *e.g.* a tracked ground based vehicle would have

larger turn penalties than Honeywell's, ducted fan driven, T-Hawk. Such considerations could assist potential natural search implementers in their strategy selection. For example, let us consider a ground based search problem where the target is likely to be following straight trajectories, could be moving faster than the search vehicle and the terrain is such that a tracked vehicle would be most appropriate. Reviewing the search results for deterministic targets (Figure 64 *et al.*, page 124), Brownian motion would seem to be the most appropriate selection. However, in terrain that requires tracked vehicles, it is likely that refuelling could be a non-trivial task and therefore a Lévy path based search may be attractive since it offers similar performance with lower turn rates and is thus more efficient. Where turning is less expensive, a bacterial inspired search could be a more favourable option despite being slightly sub-optimal in terms of the search scenario, since it requires almost 50% less energy due to its speed reduction.

7 Summary and Conclusions

This section provides a general summary of the research conducted, key conclusions, and a review of the contribution to knowledge.

7.1 Summary

This work has reviewed some of the natural search strategies that have evolved to assist animal predatory behaviour; it then implemented the more pertinent behaviours in a simple 2-dimensional simulation environment where their performance was compared against a variety of target movement types (agents had no *a priori* knowledge of target behaviour).

Performance against static targets was shown to be dependent on target distribution, swarm size and, to a lesser extent, target density. Furthermore, there are conditions upon which the type of search becomes irrelevant, such as where defensive groups are oversized and target speeds are too fast; it could be argued that under these circumstances it would be most profitable to use the most fuel efficient strategy.

Against uniformly distributed static targets, agents that apply strategies to cover unchecked space performed best (*i.e.* spiral for small groups and straight line for large groups). Somewhat surprisingly, the systematic grid pattern did not perform well; this was considered to be due to the amount of path crossing, and, hence, overlapping search effort.

Where items were clustered, ACS dominated due to its switch to intensive search following target destruction, supporting previous results from research in ecological modelling (Baum and Grant, 2001). In keeping with its natural inspiration, the initial ACS search phase was implemented as the ranging straight line search (hence, the two strategies had similar performance against uniformly distributed targets); however, where the goal is to produce the most effective overall strategy, initial target detection for small groups would be more expedient using the spiral search, as discussed above.

Giving-up-time (GUT) is important when implementing a bi-phase intensive–extensive strategy, since fruitless intensive search is temporally inefficient and, conversely, insufficient intensive search in a patchy environment can also be detrimental to performance. GUT is not the only factor in a bi-phase strategy's success and, although generally competitive, saltatory search was not as successful as may have been expected, since it is designed to work where

targets are hidden. It could have been improved by including a switch to intensive search based on experience as well as time. Nature does provide examples of this (*e.g.* harbour seals utilising a combination of tidal pattern and ocean topography to switch foraging strategy (Zamon, 2001)). *A priori* information is useful in such situations, although long term searches in stable environments could implement a learning strategy based on experience. Krakauer and Rodríguez-Gironés (1995) indicated that such an approach was possible using a Bayesian learning rule.

Consistently high turn rates are not good in search of uniformly distributed static targets because they tend to re-cover search space. Bacterial search particularly suffers due to its low effective speed. This effect is not so problematic against clustered targets, since the low turn rate strategies tend to miss targets, whereas higher turn rates will usually result in the destruction of the whole cluster.

In dynamic environments, the increase in sensor coverage caused by the increased turn rates of some strategies made them more effective. As target speed increased with respect to the agents' maximum speed this became more important since the targets were more likely to move into sensor FOV. Where targets were slower, *i.e.* less than the agents' speed, the agents' speed was also influential, since it increased the likelihood of the agent moving to an area containing a target. Across all ranging target movement strategies it can be seen that high turn rate search strategies improved their relative performance rapidly, and dominance was maintained though to target speeds that were twice that of the agents. Saltatory search, with its time based bi-phase search, maintained good performance across all target types and speeds. In addition, it was noted that this strategy conserved energy through slower search speed, and, therefore, could increase efficiency with minimal loss of effectiveness.

Where detection was lost, the application of predictive memory systems was shown to improve systemic capability, although once again the agent – target relationship was seen to be important in how much improvement was gained. When targets behaved predictably, then predictive tracking only applied by agents finding themselves in a V shaped neighbourhood performed best. This was also true for targets that were less predictable but did not range far from their last known position. Targets that did move away from the area quickly were best dealt with by predictive tracking with circular neighbourhoods or saccadic behaviour with shorter giving up times. However, the application of a PMS was seen to be inefficient when applied unnecessarily since, once committed to the pursuit of a memorised target, the

likelihood of an agent detecting other targets is diminished due to the reduction in the turn rate of the agent; *i.e.* the agent starts to move in a straight line to intercept the target which, it has been found, is generally less effective in scenarios with moving targets.

It should also be noted that this study was not intended to be exhaustive in its exploration of natural search strategies. For example, it could be beneficial to model local enhancement (Hinde, 1956), in which animals monitor other animals (not necessarily conspecific) for successful foraging in environments where prey is dense but patchy and those patches are widely dispersed, to improve the reaction to the location of clustered targets.

The strategies explored can be divided broadly into high turn rate searches that do not explore the search area widely and low turn rate strategies that range the area with less repeated coverage. The results from the sensor based agent simulations concur with the findings from ethological research: agents employing high turn rate strategies were more effective against faster targets that had low turn rates, and agents using low turn rate strategies were more effective against targets that were slower or had high turn rates (ranged less). Furthermore, agents that applied multiple strategies, switching between them based on environment knowledge could harness the benefits of each. For example, against static targets the agent may search for an initial target using a low turn rate strategy before switching to a high turn rate where target clustering was suspected.

The variety of behaviours found in natural systems is reflected in the vast amount of research that has only just begun to reveal the complexities that make some species so successful, even in challenging environments. As the examples have shown, the success of such species can be exploited to produce effective artificial solutions to large and/or complex computing problems, such as solving abstract mathematical problems or robot control. In keeping with Wolpert and Macready's "No Free Lunch Theorems" (1997), there is no solution that is better than any other across all problems, but there may be a given class of problem for which an optimal solution can be identified. For example, in an environment where prey items are clustered, predators have phylogenetically developed ACS strategies that maximise their exploitation of the prey behaviour. Such predatory behaviour, however, is only useful because it matches the likely distribution of the prey. If the environment becomes less predictable, the predator must be able to adapt or it may well die out. The same could be argued when looking to solve the problem of autonomous vehicles search for potential targets: where *a priori* knowledge is available, the most suitable strategy could be employed,

but where there is a paucity of such information, either more general strategies have to be applied or the vehicles need the ability to learn from the environment in order to be able to select or adapt a suitable strategy; where this is not possible the approach may well be sub-optimal or not succeed at all.

7.2 Conclusion

The aim of this work was to look to naturally occurring search strategies, which have evolved over millenia to assist predators in locating prey, apply them in a simulated environment and assess whether such strategies could be used to guide autonomous vehicle search. From the outset, it was not supposed that there would be a single strategy that would outperform all others across the range of target types but that where analogous predator – prey relationships existed for a given vehicle – target relationship then, where *a priori* knowledge exists of the target type, that the vehicle (generally termed 'agent' during this work) could apply the behaviour seen in nature to improve its search performance. The overall conclusions for the most appropriate search types are shown in Table 4, for static targets and Table 5, for dynamic targets.

Target Distribution	Agent / Target Ratio	Search Type	Example Search
Uniform	Low	Ranging	Spiral
	High	High turn rate	Lévy flight
Clustered	Low	Ranging	Spiral
	High	Ranging switching to high turn rate following target destruction	ACS

Table 4. Summary of the most appropriate search types for static target scenarios.

Target Type	Example Target	Target To Agent Speed Ratio	Initial Search Type	Example Search	Reacquisition Search
Predictable ranging	Deterministic	Low	High turn rate	Lévy flight	Predictive-V
		High	High turn rate	Brownian motion	Predictive-V
Unpredictable ranging	Fruit fly movement	Low	Ranging	ACS	Saccadic
		High	High turn rate	Brownian motion	Predictive-C
High turn rate	Brownian motion	Low	Ranging	Straight line	Predictive-V
		High	High turn rate	Lévy flight	Direct

Table 5. Summary of the most appropriate search types and reacquisition strategies for dynamic targets.

In general, it can be concluded that nature inspired search behaviours are potentially valuable for autonomous vehicle control, across a wide variety of agent – target relationships. Although the experiments reported here are based on scenarios drawn from a particular (military) application domain, it is felt that the use of a minimal simulation environment would render the findings more generally applicable to a wide range of domains, military or otherwise.

7.3 Contribution To Knowledge

This work has provided the first exploratory investigation into the use of natural search strategies for improving the initial search performance of autonomous vehicles operating in a search role. It has contributed the first review of ethological research into natural search strategies from the perspective of potential utilisation in natural computing paradigms. Building on that review, and previous work in autonomous vehicle search, advances in knowledge have been made in understanding how natural search strategies can be used to improve the initial search of autonomous vehicles for a range of target types. Moreover, where target detection is lost, either to the detecting vehicle due to sensor failure or occlusion, or to cooperating vehicles due to communication failure, novel contributions have also been made in understanding how natural reacquisition strategies can assist in re-detection.

8 Further Work

Considering the richness of the underlying natural metaphors and complexities of the target application one can easily imagine many directions for further work. These can be loosely grouped into: further development of the strategies directly researched in this thesis; extension of the ideas into improved modelling; application of the research into physical vehicles; and further mining of ethological research for additional metaphors that could improve performance still further. More specifically, further work could include:

- Further development of the simulation environment to model the 3-dimensional world, including environmental features, and 6 degree of freedom vehicle models.
- The difference in performance between initial search and reacquisition strategies across the various target movements and relative speeds indicate that an intelligent target could exploit knowledge of agent behaviour to employ movement patterns that would decrease the likelihood of detection (the development of which could also be an area for further work). Development of agent based learning strategies could enable swarm agents to switch their own strategy accordingly, making it more difficult to exploit the search area.
- The giving up time of the ACS was fixed and, whilst that was sufficient to show the strategy's strength in the simulated environment, in target applications target patchiness may be less predictable. Under such circumstances, it would be beneficial for the agents to dynamically adjust their GUT based on their experience of encountered target distribution or other relevant factors as observed in the natural world; for example, *inter alia*, likelihood of predation upon the predator, as in the juvenile plaice (Hill *et al.*, 2000), or seasonal prey abundance, as in European polecats (Lode, 2000).
- Bi-phase strategies were seen to benefit from the advantages of each component behaviour and could be employed by all the strategies where *a priori* information indicted that target clustering (or other features that could be exploited by phase switching) was likely.
- The search strategies had to be adapted to suit the scale of the model; the adoption of scale-free dynamics (see Reynolds 2005) could produce robust algorithms for use

across a variety of platforms operating at different scales, from nanobots to large vehicles.

- Notably, none of the search strategies discussed make use of pre-detection cooperation (cooperation is only used post target detection). Searching cooperatively could improve the effectiveness of the search but relies on adequate communication. If cooperative modifications were applied the results from this work could be used to identify effective reversionary strategies for situations in which inter-agent communication was lost through environments that were subject to heavy jamming, other intense levels of electro-magnetic noise, or other causes of practical communications difficulties.
- In the simulation environment there is nowhere for the target to hide after escaping detection, and as such the direct search is much simplified. In more complex environments, where the target may be able to hide, it could be beneficial to search for longer, possibly including a switch to an intensive search.
- In this work, GUT and inter-search movements for saltatory and ACS search were adapted for the simulation environment. More complex relationships can be observed in nature, and these may be of further benefit to autonomous vehicle search. For example, in an ACS such as that of hawk owls (Sonerud, 1992) the search variables such as move distance between searches and GUT were assigned dynamically depending on the bird's current search environment, *e.g.* a higher perch position led to longer GUT and larger move distances between searches. This correlation between height and time meant that the available area was searched for a sufficient time but not excessively. Moreover, the positive correlation between perch height and inter-search move yielded very little overlap occurred between search areas. This strategy was believed to have developed because of the bird's poor ability to search whilst moving.
- Finally, the linear prediction model used by the PMS is simplified, and several more realistic models are available such as those applied by Erwin (2004) to model sonar tracking in bats. Erwin's biomemetic model was unsuccessful in its predictive tracking due to the real-time response requirements, which were too short for the unscented Kalman filter, and the implementation model, which again was too slow to react. However, these issues would be unlikely to present themselves in the typical

target applications due to larger timeframes and as such the utilisation of more elaborate predictive models could improve system performance still further.

References

AdaCore, 2005. Ada academic initiative. Available from: http://www.adacore.com/academic_overview.php [accessed 26 January 2005].

Aerospace International, 2008. 100 years of British Aviation, The next 100 years. Aerospace International, July 2008. Royal Aeronautical Society.

Anderson, C., and Franks, N.R., 2001. Teams in animal societies. Behavioral Ecology, Vol. 12, No 5, pp 534-540.

Anderson, C., Theraulaz, G. and Deneubourg, J.L., 2002. Self-assemblages in insect societies. Insectes Sociaux. 49, pp. 99-110

Andrews, B.W., Passino, K.M. and Waite, T.A., 2007. Foraging theory for autonomous vehicle decision-making system design. Journal of Intell. Robot Syst. Vol. 49, pp 39-65. Springer.

Atkinson, R. P. D., Rhodes, C. J., Macdonald, D. W. & Anderson, R. M., 2002. Scale-free dynamics in the movement patterns of jackals. Oikos, 98, 134–140.

Banks, A., 2005. A particle swarm guidance system for autonomous, unmanned airborne vehicles in an air defence role. Masters Dissertation, Bournemouth University, December, 2005.

Banks, A., Vincent, J. and Anyakoha, C., 2007. A review of particle swarm optimization. Part I: background and development. Natural Computing, Vol. 6, No.4, pp 467-484. Springer.

Banks, A., Vincent, J. and Anyakoha, C., 2008a. A review of particle swarm optimisation. Part II: hybridisation, combinatorial, multicriteria and constrained optimization, and indicative applications. Natural Computing, Vol. 7, No. 1, pp. 109-124. Springer.

Badler, J.B. and Heinen, S.J., 2006. Anticipatory movement timing using prediction and external cues. The Journal of Neuroscience, 26(17), pp. 4519-4525.

Banks, A., Vincent, J. and Phalp, K., 2008b. Particle swarm guidance system for autonomous, unmanned aerial vehicles in an air defence role. Journal of the Royal Institute of Navigation, Vol. 61, No.1, pp. 9-29.

Banks, A., Vincent, J. and Phalp, K., 2008c. Natural strategies for search. Natural Computing, Springer.

Banks, A. and Vincent, J., 2009. Exploring the performance of natural search strategies for the control of unmanned autonomous vehicles. Journal of the Royal Institute of Navigation, Vol. 62, No.2. pp. 283-301

Bartumeus, F., Catalan, J., Fulco, U.L., Lyra, M.L. and Viswanathan, G.M., 2002. Optimizing the encounter rate in biological interactions: Lévy versus Brownian strategies. Physical Review Letters, Vol. 88, No. 9. American Physical Society.

Baum, K.A. and Grant W.E., 2001. Hummingbird foraging behavior in different patch types: simulation of alternative strategies. Ecological Modelling 137 (2001), pp 201-209.

Beekman, M., Sumpter, D.J.T. and Ratnieks, F.L., 2001. Phase transition between disordered and ordered foraging in Pharoah's ants. PNAS, Vol. 98, No. 17, pp 9703-9706.

Bell, W.J., 1991. Searching behaviour: the behavioural ecology of finding resources. Chapman and Hall Animal Behaviour Series, Chapman and Hall, London.

Bénichou, O., Coppey, M., Moreau, M., Suet, P-H. and Voituriez, R., 2005. Optimal search strategies for hidden targets. Phys. Rev. Lett. 94, 198101, (2005).

Bergbreiter, S. and Pister K.S.J., 2007. Design of an autonomous jumping microrobot. In Proceedings of IEEE International Conference on Robotics and Automation, 2007, pp 447 – 453.

Borrelli, F., Falcone, P., Keviczky, T., Asgari, J. and Hrovat, D., 2005. MPC based approach to active steering for autonomous vehicle systems. International Journal of Vehicle Autonomous Systems, Vol. 3, No. 2/3/4, pp. 265-291, 2005.

Boyer, D., Miramontes, O., Ramos-Fernandez, G., Mateos, J.L. and Cocho, G., 2004. Modeling the searching behavior of social monkeys. Physica A-Statistical Mechanics and its Applications 342 (1-2): 329-335 Oct 15 2004.

Budick, S.A. and Dickinson, M.H., 2006. Free-flight responses of *Drosophila melanogaster* to attractive odors. The Journal of Experimental Biology, 209, pp. 3001 – 3017.

Burns, J.G, 2005. Impulsive bees forage better: the advantage of quick, sometimes inaccurate foraging decisions. Animal Behaviour, 2005, 70, e1-e5.

Butler, M.A., 2005. Foraging mode of the chameleon, *Bradypodian pumilum*: a challenge to the sit-and-wait versus active forager paradigm? Biological Journal of the Linnean Society, 2005, 84, pp. 797-808.

Carey, B., 2007. Miniature air vehicles. Avionics Magazine, March 2007. Access Intelligence, Rockville, MD, USA.

Charnov, E.L., 1976. Optimal foraging, the Marginal Value Theorem. Theoretical Population Biology Vol.9, No.2, April 1976.

Chien, S., Sherwood, R., Tran, D., Cichy, B., Rabideau, G., Castaño, R., Davies, A., Frye, S., Trout, B., D'Agostino, J., Shulman, S., Mandl, D., Boyer, D., Hayden, S., Sweet, A. and Christa, S., 2005. The autonomous sciencecraft embedded systems architecture. Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC 2005), Waikoloa, HI, October 2005.

Chitre, M., 2008. DSAAV - A Distributed Software Architecture for Autonomous Vehicles. IEEE/MTS OCEANS'08, Quebec City, Canada, 2008

Colorni, A., Dorigo, M. and Maniezzo, V., 1991. Distributed optimization by ant colonies. Proc. European Conference on Artificial Life, Paris, France, pp. 134-142.

Conte, G., Rudol, P., Wzorek, M., Petitti, P., Iocchi, L. and Doherty, P., 2008. Use of geo-referenced images with unmanned aerial systems. Workshop Proceedings of SIMPAR 2008, International Conference on Simulation, Modeling and Programming for Autonomous Robots, Venice Italy 2008, pp. 444-454.

Cope, J.M. and Fox, C.W., 2003. Oviposition decisions in the seed beetle, *Callosobruchus maculatus* (*Coleoptera*: *Bruchidae*): effects of seed size on superparatism. Journal of Stored Products Research 39 (2003) pp. 355-365.

Cory, R. and Tedrake, R., 2008. Experiments in fixed-wing UAV perching. In Proceedings of the AIAA Guidance, Navigation, and Control Conference, 2008.

Courbon, J., Mezeouar, Y., Lequievre, L., and Eck, L., 2008. Navigation of urban vehicle: An efficient visual memory management for large scale environments. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Nice, France, September 2008.

De Jong, K., 1975. An analysis of the behaviour of a class of genetic adaptive systems. PhD thesis, University of Michigan.

De Meyer, K., Bishop, J.M. and Nasuto, S.J., 2003. Stochastic diffusion: using recruitment for search. Symposium on Evolvability & Interaction. 8 – 10th October 2003. Queen Mary. The University of London, UK.

De Nardi, R. and Holland, O., 2007. Ultraswarm: a further step towards a flock of miniature helicopters. Swarm Robotics, pp. 116-128. Springer, Berlin Heidelberg.

De Wagter, C., Bijnens, B. and Mulder, J.A., 2007. Vision-only control of a flapping MAV on Mars. AIAA Guidance, Navigation and Control Conference and Exhibit, 2007.

Dechaume-Moncharmont, FX., Dornhaus, A., Houston, A.I., McNamara, J.M., Collins, E.J. and Franks, N.R., 2006. The hidden cost of information in collective foraging. Proceedings B of The Royal Society Biological Sciences 1(1):1689 (2006)

Dejean, A., Le Breton, J., Suzzoni, J.P., Orival, J. and Saux-Moreau, C., 2005. Influence of interspecific competition on the recruitment behavior and liquid food transport in the tramp ant species *pheidole megacephala*. Naturwissenschaften, 2005, 92, pp. 324- 327.

Delestrade, A., 1999. Foraging strategy in a social bird, the alpine chough: effect of variation in quantity and distribution of food. Animal Behaviour, 1999, 57, pp. 299-305.

Dobbs, R.C. and Martin, T.E., 1998. Variation in foraging behavior among nesting stages of female red-faced warblers. The Condor, 100, pp. 741-745. The Cooper Ornithological Society.

Dornhaus, A. and Chittka, L., 2004. Why do honey bees dance? Behav Ecol Sociobiol (2004) 55:395-401

Dugatkin, L.A., 1997. The evolution of cooperation. Bioscience, June 1997, Vol. 47, Issue 6, p. 355.

Eberhart, R.C. and Shi, Y., 2000. Comparing inertia weights and constriction factors in particle swarm optimization. Proc. 2000 Congr. Evolutionary Computing, pp. 84-88, 2000.

Edwards, S., Rutkowski, A.J., Quinn, R.D., Willis, M.A., 2005. Moth-inspired plume tracking strategies in three-dimensions. Proceedings of the 2005 IEEE International Conference on Robotics and Automation, Vol. 2, pp 1669-1674.

English, J., 2000. About JEWL. Available from: http://www.it.bton.ac.uk/staff/je/jewl/ [Accessed 22 June 2005].

Erwin, H.R., 2004. Algorithms for sonar tracking in biomimetic robots. In Proceedings of the 5th International Conference on Recent Advances in Soft Computing, pp. 68-73, Nottingham, United Kingdom, December 2004.

Etiene, A.S. and Jeffery, K.J., 2004. Path integration in mammals. Hippocampus 14 pp. 180-192. Wiley-Liss Inc.

Fabiani, P., Fuertes, V., Piquereau, A., Mampey, R. and Teichteil-Königsbuch, F., 2007. Autonomous flight and navigation of VTOL UAVs: from autonomy demonstrations to out-of-sight flights. Aerospace Science and Technology 11 (2007), pp. 183-193.

Farina, W.M., Wainselboim, A.J., 2005. Trophallaxis within the dancing context: a behavioural and thermographic analysis in honeybees (*Apis mellifera*). Apidologie 36 (2005) pp. 43-47.

Fitzpatrick, J.W., 1980. Foraging behavior of neotropical tyrant flycatchers. Condor, 82, pp.43-57. The Cooper Ornithological Society.

Fogel, L.J., Owens, A.J. and Walsh, M.J., 1964. On the evolution of artificial intelligence. Proc. 5th National Symp. On Human Factors in Engineering, IEEE, San Diego, CA, pp. 63-76.

Freake, M.J., 2001. Homing behavior in the sleepy lizard (*Tiliqua rugosa*): the role of visual cues and the parietal eye. Behav. Ecol. Sociobiol (2001) 50:563-569.

Freeman, M.C. and Grossman, G.D, 1992. Group foraging by a stream minnow: shoals or aggregations? Animal Behaviour, 1992, 44, pp 393-403.

Fresneau, D., 1985. Individual foraging and path fidelity in a ponerine ant. Insectes Sociaux, Paris, 1985, Volume 32, n 2, pp 109-116.

Friedberg, R.M., 1958. A learning machine: Part I. IBM Journal, pp. 2-13.

Fruergaard-Pedersen, R., 2006. Optimal demining using a swarm of low-cost robotic units. Phd Thesis, Department of Computer Science – Daimi, Faculty of Science, University of Aarhus, Denmark.

Frye, M.A., Tarsitano, M. and Dickinson, M.H., 2003. Odor localization requires visual feedback during free flight in *Drosophila melanogaster*. The Journal of Experimental Biology, vol. 206, pp. 843 – 855.

Gage, D.W., 1990. Security considerations for autonomous robots. Proceedings of Symposium on Security and Privacy, Oakland CA, 22-24 April 1985, pp 224-228. Reprinted in Computer Security Journal, vol 6, 1990, pp. 95-99.

Ghumman, U., Kunwar, F. and Benhabib, B., 2008. Guidance-based on-line motion planning for autonomous highway overtaking. International Journal on Smart Sensing and Intelligent Systems, Vol.1, No. 2, June 2008.

Gil, M. and De Marco, R.J., 2005. Olfactory learning by means of trophallaxis in *Apis mellifera*. The Journal of Experimental Biology, 208, pp. 271-280.

Goodrich, M.A., Morse, B.S., Gerhardt, D., Cooper, J.L., Quigley, M., Adams, J.A. and Humphrey, C., 2008. Supporting wilderness search and rescue using a camera-equipped mini UAV. Journal of Field Robotics, Vol. 25, No.1, pp. 89-110.

Gordon, D.M., 2003. The organization of work in social insect colonies. Complexity, Vol. 8, No.1, pp. 43- 46.

Grassé, P.P., 1959. La reconstruction du nid et les coordinations interindividuelles chez bellicositermes natalensis et cubitermes sp. La theorie de la stigmergie: essai d'interpretation du comportament des termites constructeurs. Insects Sociaux, 6, pp. 41-81.

Green, W.E. and Oh, P.Y., 2006. A fixed-wing aircraft for hovering in caves, tunnels, and buildings. IEEE American Control Conference (ACC), Minneapolis, MN, June 2006.

Guinn, J.R., Williams, B.G., Wolff, P.J., Fennessey, R. and Lt. West, T, 1995. TAOS orbit determination results using global positioning satellites. Advances in the Astronautical Sciences, Spaceflight Mechanics, Vol. 89.

Gurevich, G., Bell, R. and Wertz, J.R., 2000. Autonomous on-board orbit control: flight results and applications. AIAA 2000 Conference and Exposition, Long Beach, California, USA, 19th – 21st September 2000.

Hauert, S., Zufferey, J-C. and Florean, D., 2009. Evolved swarming without positioning information: an application in aerial communication relay. Autonomous Robots, Vol. 26, Issue 1. Kluwer Academic Publishers, USA.

Hauser, M.D., 1992. Costs of deception: cheaters are punished in rhesus monkeys (*Macaca mulatta*). Proc. Natl. Acad. Sci. USA Vol. 89, December 1992, pp. 12137-12139.

Hill, S., Burrows, M.T. and Hughes, R.N., 2000. Increased turning per unit distance as an area-restricted search mechanism in a pause-travel predator, juvenile plaice, foraging for buried bivalves. Journal of Fish Biology (2000) 56, pp 1497–1508.

Hinde, R.A., 1956. The biological significance of the territories of birds. Ibis 98:340-369.

Holland, J.H., 1962. Outline for a logical theory of adaptive systems. Journal of the Association for Computing Machinery, 3, pp. 297-314.

Holland, J.H., 1975. Adaptation in natural and artificial systems. University of Michigan Press, Ann Arbour.

Holland, O.E. and Melhuish, C., 1999. Stigmergy, self-organization, and sorting in collective robots. Artificial Life, 5:2, 173-202, 1999.

Honeywell, 2008. Honeywell wins first production contract for its THawk™ micro air vehicle. Honeywell press release, 04 November 2008.

Hsiao, F-B., Lai, Y-C., Tenn, H-K., Hsieh, S-Y., Chen, C-C., Chan, W-L. and Hirst, R., 2006. The development of an unmanned aerial vehicle system with surveillance, watch, autonomous flight and navigation capability. The 21st Bristol UAV Systems Conference – April 2006.

Iwasa, Y., Higashi, M. and Yamamura, N., 1981. Prey distribution as a factor determining the choice of optimal foraging strategy. Am. Nat. 117, pp. 710-723.

Jameson, S., Satterfield, B., Bolden, C., Allen, N. and Youngren, H., 2007. SAMARAI nano air vehicle – A revolution in flight. In the Proceedings of the Association for Unmanned Vehicle Systems International (AUSVI) Unmanned Systems North America 2007, Washington, DC, August 6-9, 2007.

Jin, Y., Liao, Y., Minai, A.A. and Polycarpou, M.M., 2006. Balancing search and target response in cooperative Unmanned Aerial Vehicle (UAV) teams. IEEE Transactions on Systems, Man, and Cybernetics, Vol. 36, Issue 3, June 2006 pp 571-587.

Johnson, N.K. and Peeters, H.J., 1963. The systematic position of certain hawks in the genus *Buteo*. The Auk, Vol.80, No.4, pp 417-446.

Jung, D. and Tsiotras, P., 2007. Inertial attitude and position reference system development for a small UAV. AIAA Infotech@Aerospace Conference 2007, Rohnert Park, California, USA.

Kennedy, J. and Eberhart, R.C., 1995. Particle swarm optimization. Proc. IEEE Int. Conf. Neural Networks, Piscataway, NJ, pp. 1942–1948.

Kennedy, J., 1997. The particle swarm: social adaptation of knowledge. Proc. Int. Conf. Evolutionary Computation, IEEE, pp. 303–308, Piscataway, NJ.

Keviczky, T., Falcone, P., Borrelli, F., Asgari, J. and Hrovat, D., 2006. Predictive control approach to autonomous vehicle steering. American Control Conference, Minneapolis, Minnesota, June 14-16, 2006. IEEE.

Klam, F., Petit, J., Grantyn, A., Berthoz, A., 2001. Predictive elements in ocular interception and tracking of moving targets by untrained cats. Experimental Brain Research, 139, pp. 233-247.

Kolmogorov, A.N., 1933. On the empirical determination of a distribution function. Giornale dell'Instituto Italiano degli Attuari, 4, pp. 83-91.

Koza, J., 1992. Genetic programming: on the programming of computers by means of natural selection. MIT Press, Cambridge, MA.

Krakauer, D.C. and Rodríguez-Gironés, M.A., 1995. Searching and learning in a random environment. Journal of Theoretical Biology, 177, pp. 417-419.

Kruskal, W.H. and Wallis, W.A., 1952. Use of ranks in one-criterion variance analysis. Journal of the American Statistical Association, 47, pp. 583-621.

Kurt, A. and Özgüner, Ü., 2008. Hybrid state system development for autonomous vehicle control in urban scenarios. Proceedings of the 17th World Congress, The International Federation of Automatic Control, Seoul, Korea, 2008.

Lange, S., Sünderhauf, N. and Protzel, P., 2008. Autonomous landing for a multirotor UAV using vision. Workshop Proceedings of SIMPAR 2008, International Conference on Simulation, Modeling and Programming for Autonomous Robots, Venice, Italy 2008, pp. 482-491.

Lee, U., Magistretti, E., Gerla, M., Bellavista, P., Lió, P. and Lee, K-W., 2009. Bio-inspired multi-agent data harvesting in a proactive urban monitoring environment. Ad Hoc Networks, 7 (2009) pp. 725-741.

Liang, S.X., Reibling, L.A. and Betts, J., 2008. re-Ada: Reliable Ada-based descriptive architecture for C4ISR via a quantitative interoperating model. Proceedings of the ACM SIGAda Annual International Conference (SIGAda 2008), pp. 39-56.

Lindauer, M. and Kerr, W.E., 1960. Communication between the workers of stingless bees. BeeWorld 41, pp. 29-41and pp. 65-71.

Lochmatter, T. and Martinoli, A., 2008. Tracking odor plumes in a laminar wind field with bio-inspired algorithms. 11th International Symposium on Experimental Robotics 2008 (ISER 2008); Athens, Greece, July 14-17, 2008.

Lockheed Martin, 2007. The nano air vehicle: nature's ultimate spiral development. Insights Magazine, Ed. Madden TS Jr., Vol.4 No.2, 2nd quarter, 2007.

Lode, T., 2000. Functional response and area-restricted search in a predator: seasonal exploitation of anurans by the European polecat, *Mustela putorius*. Austral Ecology (2000) 25, pp. 223-231.

Mann, H.B. and Whitney, D.R., 1947. On a test of whether one of two random variables is stochastically larger than the other. Annals of Mathematical Statistics, 18, pp. 50-60.

Marchand, D. and McNeil, J.N., 2004. The importance of behavioural plasticity for maximizing foraging efficiency in Frugivorous *Lepidopteran* Larvae. Journal of Insect Behavior, Vol.17, No.5, September 2004.

Marthaler, D., Bertozzi, A., Schwartz, I., 2004. Lévy searches based on a priori information: The biased Lévy walk. UCLA CAM Report, (04-50), 2004.

Maschwitz, U., Steghaus-Kovac, S., Gaube, R. and Hänel, H., 1989, A South East Asian ponerine ant of the genus Leptogenys (Hym. Form.) with army ant life habits, Behavioral Ecology and Sociobiology, 24, pp 305–316.

Maynard Smith, J., 1982. Evolution and the theory of games. Cambridge University Press, Cambridge

Matos, A., Cruz, N., Martins, A. and Pereira, F.L., 1999. Development and implementation of a low-cost lbl navigation system for an AUV. In Proceedings of the MTS/IEEE Conf. Oceans'99, Seatle, WA, USA. September 1999.

Menzel, R., Brandt, R., Gumbert, A., Komischke, B., Kunze, J., 2000. Two spatial memories for honeybee navigation. Proc. R. Soc. Lond. B. Biol. Sci. 267 (2000), pp. 961-968

Merkle, T., Rost, M. and Alt, W., 2006. Egocentric path integration models and their application to desert arthropods. Journal of Theoretical Biology 240 (2006), pp. 385-399.

Miasnikov, E., 2005. Threat of terrorism using unmanned aerial vehicles: technical aspects. Center for Arms Control, Energy and Environmental Studies Moscow Institute of Physics and Technology 2005.

Minsky, M., 1986. The society of mind. Simon and Schuster, New York, 1986.

Mittelstaedt, H., 2000. Triple-loop model of path control by head direction and place cells. Biol. Cybern. 83, pp. 261-270.

Monmarche, N., Venturini, G. and Slimane, M., 2000. On how *Pachycondyla apicalis* ants suggest a new search algorithm. Future Generation Computer Systems, 16, pp. 937 – 946.

Morse, T. M., Ferrée, T. C., and Lockery, S. R., 1998. Robust spatial navigation inspired by chemotaxis in *Caenorhabditis elegans*. Adaptive Behavior. 6:393-410. 1998.

Morrison, M.L., With, K.A., Timossi, I.C. and Milne, K.A. 1987. Composition and temporal variation of flocks in the Sierra Nevada. The Condor 89:739-745. The Cooper Ornithological Society.

Motro, U. and Shmida, A., 1995. Near-far search – an evolutionary stable search strategy. J. Theor. Biol. (1995) 173, pp 15-22.

Müller, M. and Wehner, R., 1994. The hidden spiral: systematic search and path integration in desert ants, *Cataglyphis fortis*. Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology. Volume 175, Number 5 / November, 1994

NATO, 2003. Urban operations in the year 2020. RTO Technical Report 71. RTO Studies, Analysis and Simulation Panel Study Group SAS-030, NATO.

Nicholls, K.W., Abrahamsen, P.E., Heywood, K.J., Stansfield, K. and Østerhus, S., 2008. High-latitude oceanography using the Autosub autonomous underwater vehicle. Limnol. Oceangr., 53(5, part 2), pp 2309-2320.

Nishimura, K., 1999. Exploration of optimal giving-up time in uncertain environment: a sit-and-wait forager. Journal of Theoretical Biology (1999) 199, pp. 321-327.

O'Brien, J.W., Evans, B.I. and Browman, H.I., 1989. Flexible search tactics and efficient foraging in saltatory searching animals. Oecologia, 80, pp. 100 – 110.

O'Brien, J.W., Browman, H.I. and Evans, B.I., 1990. Search strategies of foraging animals. American Scientist Vol.78, pp. 152-160.

Ogden, J.C., 1974. The short-tailed hawk in Florida I. Migration, habitat, hunting techniques, and food habits. The Auk, 91, pp. 95-110. January 1974.

Orban de Xivry, J.J. and Lefèvre, P.L., 2007. Saccades and pursuit: two outcomes of a single sensorimotor process. Journal of Physiology, 584, pp. 11-23.

Papaj, D.R. and Rausher, M.D., 1983. Individual variation in host location by phytophagous insects. Herbivorous Insects, Host-Seeking Behavior and Mechanisms. S. Ahmad (ed), pp. 77-124.

Passino, K.M., 2002. Biomimicry of bacterial foraging for distributed optimisation and control. IEEE Control Systems Magazine, June 2002, pp. 52-67.

Pearson, K., 1896. Mathematical contributions to the theory of evolution.-III. Regression, heredity and panmixia. Philosophical Transactions of the Royal Society of London, Series A, 187, pp. 253-318.

Pianka, E.R., 1966. Convexity, desert lizards, and spatial heterogeneity. Ecology, 47, pp. 1055-1059.

Pugh, J. and Martinoli, A., 2007. Inspiring and modeling multi-robot search with particle swarm optimization. Swarm Intelligence Symposium, 2007. SIS 2007. IEEE.

Pyk, P., Bermúdez, i., Badia. S., Bernardet, U., Knüsel, P., Carlsson, M., Gu, J., Chanie, E., Hansson, B.S., Pearce, T.C. and Verschure P.F.M.J., 2006. An artificial moth: chemical source localization using a robot based neuronal model of moth optomotor anemotactic search. Auton Robot, vol. 20, pp. 197-213, 2006.

Pyke, G.H., 1978. Optimal foraging: movement patterns of bumblebees between inflorescences. Theoret. Pop. Biol., 13, pp. 72-98.

Quigley, M., Barber, B., Griffiths, S., and Goodrich, M. A., 2005. Towards real-world searching with fixed-wing mini-uavs. In Proceedings of the IEEE International Conference on Robotics and Automation, 2005, pp 3028-3033.

Ramos-Fernández, G., Mateos, J.L., Miramontes, O., Cocho, G., Larralde, H., Ayala-Orozco, B., 2004. Lévy walk patterns in the foraging movements of spider monkeys (*Ateles geoffroyi*). Behavioral Ecology and Sociobiology 55, 223, 2004.

Rankin, A.L. and Matthies, L.H., 2008. Daytime mud detection for unmanned ground vehicle autonomous navigation. Proceedings of the 26th Army Science Conference, December 1-4, 2008.

Rayman, M.D., Varghese, P., Lehman, D.H. and Livesey, L.L., 2000. Results from the Deep Space 1 technology validation mission. 50th International Astronautical Congress, Amsterdam, The Netherlands, 4th – 8th October 1999.

Rechenberg, I., 1965. Cybernetic solution path of an experimental problem. Royal Aircraft Establishment, Library Translation No. 1122.

Reed, H.C. and Landolt, P.J., 2000. Application of alarm pheromone to targets by southern yellowjackets (*Hymenoptera: Vespidae*). *Florida Entomologist* 83(2), pp. 193-196.

Reed, S., Cormack, A., Hamilton, K., Tena Ruiz, I. and Lane, D., 2006. Automatic ship hull inspection using unmanned underwater vehicles. Proceedings from the Seventh International Symposium on Technology and the Mine Problem, 2006.

Reilly, J., 2008. Missing Chad drone may have tried to fly home. Sunday Independent newspaper article, Vol. 103, No. 49, dated 07 December, 2008. Published in Republic of Ireland by Independent News & Media PLC, Dublin, Republic of Ireland.

Reynolds, A.M., 2005. Scale-free movement patterns arising from olfactory driven foraging, Physical Review E 72 (4) (2005).

Robinson, T., 2006. Out of the black. Aerospace International, April 2006. Royal Aeronautical Society.

Royal Aeronautical Society, 2008. The House of Commons Select Committee UAS and ISTAR – A submission from the Royal Aeronautical Society, 17 April 2008. The Aerospace Professional, December 2008. The Royal Aeronautical Society, UK.

Sandia National Laboratories, 2001. Developing technologies for asymmetrical warfare and homeland defense. Sandia Technology Vol.3, No.2. Sandia National Laboratories. Albuquerque, USA.

Santos, N., Matos, A. and Cruz, N., 2008. Navigation of an autonomous underwater vehicle in a mobile network. Proceedings of the MTS-IEEE Conference Oceans'2008, Quebec, Canada, September 2008.

Saripelli, S., Montgomery, J.F. and Sukhatme, G.S., 2002. Vision-based autonomous landing of an unmanned aerial vehicle. IEEE International Conference on Robotics and Automation (IRCA(2002)), pp. 2799-2804.

Schmickl, T. and Crailsheim, K., 2006. Trophallaxis among swarm-robots: A biologically inspired strategy for swarm robotics. BioRob 2006, The First IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics, 2006. February 20-22, 2006, pp. 377 - 382

Secor, S.M., 1994. Ecological significance of movements and activity range for the sidewinder, *Crotalus cerastes*. Copeia, Vol. 1994, No. 3. (Aug. 17, 1994), pp. 631-645.

Shahabudeen, S., Chitre, M. and Motani, M., 2007. A multi-channel MAC protocol for AUV networks. Proceedings of IEEE Oceans'07 Europe, Aberdeen, 18-22 June 2007.

Shi, Y, and Eberhart, R.C., 1998. A modified particle swarm optimizer. Proceedings of the IEEE international conference on evolutionary computation, pp 69–73. IEEE Press, Piscataway, NJ.

Skutelsky, O., 1995. Flexibility in foraging tactics of *Buthus Occitanus* scorpions as a response to above-ground activity of termites. The Journal of Arachnology 23, pp. 46-47.Smirnov, N.V., 1939. On the estimation of the discrepency between empirical curves of distribution for two independent samples. Bulletin of Moscow University, Vol. 2, pp. 3-16.

Sonerud, G.A., 1992. Search tactics of a pause-travel predator: adaptive adjustments of perching times and move distances by hawk owls (*Surnia ulula*). Behavioral Ecology and Sociobiology 30, pp. 207-217.

Souto, A., Bezerra, B.M., Schiel, N. and Huber, L., 2007. Saltatory search in a free-living *Callithrix jacchus*: environmental and age influences. International Journal of Primatology 28, pp. 881-893.

Stahl, J.C. and Sagar P.M., 2000. Foraging strategies if southern Buller's albatrosses *Diomedea b. bulleri* breeding on The Snares, New Zealand. Journal of The Royal Society of New Zealand, Vol. 30, No. 3, pp. 299-318.

Steinmetz, I., Schmolz, E., Ruther, J., 2003. Cuticular lipids as trail pheromone in a social wasp. Proceedings of the Royal Society of London B 270: 385-291.

Stephens, D.W., Brown, J.S. and Ydenberg, R.C., 2007. Foraging, behaviour and ecology. The University of Chicago Press, Chicago, USA.

Stevens, J.R., Cushman, F.A., Hauser, M.D., 2005. Evolving the psychological mechanisms for cooperation. Annu. Rev. Ecol. Evol. Syst. 2005. 36, pp. 499-518

Tammero, L.F. and Dickinson, M.H., 2002. The influence of visual landscape on the free flight behavior of the fruit fly *Drosophila melanogaster*. The Journal of Experimental Biology, Vol. 205, pp. 327-343.

Thakoor, S., Chahl, J., Bouffant, N., Stange, G., Srinivasan, M., Hine, B. and Zornetzer, S., 2003. Bio-inspired engineering of exploration systems: A horizon sensor/attitude reference system based on the dragonfly ocelli for Mars exploration applications. Journal of Robotic Systems, vol. 20(1), pp. 35-42.

Vacek, S., Schimmel, C. and Dillman, R., 2007. Road-marking analysis for autonomous vehicle guidance. Proceeding of 3rd European Conference on Mobile Robots (ECMR'07), pp. 72-77.

Vacek, S., Gindele, T., Zöllner, J.M. and Dillman, R., 2007. Using case-based reasoning for autonomous vehicle guidance. Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems, 2007, IROS 2007, San Diego, USA, pp. 4271-4276.

Vincent, I. and Trentini, M., 2007. Shape-shifting tracked robotic vehicle for complex terrain navigation. Technical Memorandum, DRDC Suffield TM 2007 – 190, December 2007. Defence Research and Development Canada.

Viswanathan, G.M., Buldyrev, S.V., Havlin, S., Da Luz, M.G.E., Raposo, E.P. and Stanley, H.E., 1999. Optimizing the success of random searches. Nature 401, pp. 911-914.

Viswanathan, G.M., Afanasyev, V., Buldyrev, S.V., Havlin, S., da Luz, M.G.E., Raposo, E.P. and Stanley, H.E., 2001. Lévy flights search patterns of biological organisms. Physica A 295 (2001), pp. 85-88.

Viswanathan, G.M., Bartumeus, F., Buldyrev, S.V., Catalan, J., Fulco, U.L., Havlin, S., da Luz, M.G.E, Lyra, M.L., Raposo, E.P. and Stanley, H.E., 2002. Lévy flight random searches in biological phenomena. Physica A 314 (2002) pp. 208-213.

Vladusich, T., Hemmi, J.M., Zeil, J., 2006. Honeybee odometry and scent guidance. The Journal of Experimental Biology, 209, pp. 1367-1375.

Von Frisch, K., 1967. The dance language and orientation of bees. Harvard University Press, Cambridge, Massachusetts, USA.

Waddington, K.D., 1980. Flight patterns of foraging bees relative to density of artificial flowers and distribution of nectar. Oecologia (Berl.) 44, pp. 199-204.

Waibel, M., Floreano, D., Magnenat, S. and Keller, L., 2006. Division of labour and colony efficiency in social insects: effects of interactions between genetic architecture, colony kin structure and rate of perturbations. Proceedings of the Royal Society B (2006) 273, pp. 1815-1823.

Walmsley, N.S. and Syms, P.R., 2004. A novel rapid ISTAR assessment method. 2004 Command and Control Research and Technology Symposium - The Power of Information Age Concepts and Technologies.

Ward, P. and Zahavi, A., 1973. The importance of certain assemblages of birds as 'information centres' for food finding. Ibis, 117, pp. 517-534.

Wertz, J.R., 2003. Autonomous navigation and autonomous orbit control in planetary orbits as a means of reducing operations cost. 5th International Symposium on Reducing Cost of Spacecraft Ground Systems and Operations, Pasadena, USA 8th –11th July, 2003.

Wiess, M.R. and Papaj, D.R., 2003. Colour learning in two behavioural contexts: how much can a butterfly keep in mind? Animal Behaviour, 65, pp. 425-434.

Willis, M.A. and Arbas, E.A., 1998. Variability in odor-modulated flight by moths. Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology, vol. 182, pp. 191-202.

Wolpert, D.H. and Macready, W.G., 1997. No free lunch theorems for optimization. IEEE Trans. Evolutionary Computation, 1(1), pp. 67-82.

Wright, J., Stone, R.E. and Brown, N., 2003. Communal roosts as structured information centres in the raven, *Corvus corax*. Journal of Animal Ecology 2003 72, pp 1003-1014.

Young, L.A., Aiken, E., Lee, P. and Briggs, G., 2005. Mars rotorcraft: possibilities, limitations, and implications for human/robotic exploration. 2005 IEEE Aerospace Conference, 5th – 12th March 2005, pp. 300-318.

Zamon, J.E., 2001. Seal predation on salmon and forage fish schools as a function of tidal currents in the San Juan Islands, Washington, USA. Fisheries Oceanography, 10:4, pp. 353-366, 2001.

Appendices

Appendix 1 - Analysis of Bacterial Search Behaviour

As described in the background literature, in nature, each bacterium in a group of bacteria responds individually on detecting a positive nutrient gradient. However, because the nutrient usually extends quickly across the environment containing the bacteria the overall response appears to be cooperative with all of the bacteria changing behaviour, to longer and faster periods of swimming, almost simultaneously. In this work, the agent behaviour during bacterial search was cooperative and as such whilst the group response was arguably the same, in that the whole group responds in the presence of a target (where a target represents a positive nutrient gradient), the individual behaviour is not the same, since the agents do not have to be experiencing the targets to respond to their presence in the search area.

To assess the impact of including all agents in the switch from a non-target mode of slow speed and high tumbling rates to the faster more ranging movement when any of the agents detect a target, representative experiments were carried out that applied circular neighbourhoods ranging in size around the detected target simulating nutrient patches of differing sizes. Neighbourhood sizes began at zero, *i.e.* as per the natural behaviour where only individual bacterium responded to positive nutrient gradients, then increased in size as multiples of the sensor range, finally the full agent swarm was included as used in the main body of experiments. In the static target scenario a medium sized swarm of 10 agents searched for 20 uniformly distributed targets. This setup was used to highlight differences in the ranging across the search space during target detection, which could affect system performance. For dynamic targets, 20 agents searched for 4 fruit fly targets travelling at maximum agent speed (*i.e.* relative target speed of 1.0). These targets were used because of their two phase movement that both ranged across the search area and included periods of high turn rates. To summarise: 10 agents were used against 20 uniformly distributed static targets and 20 agents against 4 fruit fly targets with a relative target speed of 1.0; all initial search strategies were applied in each scenario (for comparison) along with LSZ

neighbourhood swarming; and bacterial search used either a circular neighbourhood with radius of {0, 1, 2, 3, 4, 5} times sensor range or a whole search area response.

Results shown in Figure 1 - 1 indicate that, in this simulation environment, the changes from the natural behaviours made no difference to the performance of the group in both static and dynamic target configurations.

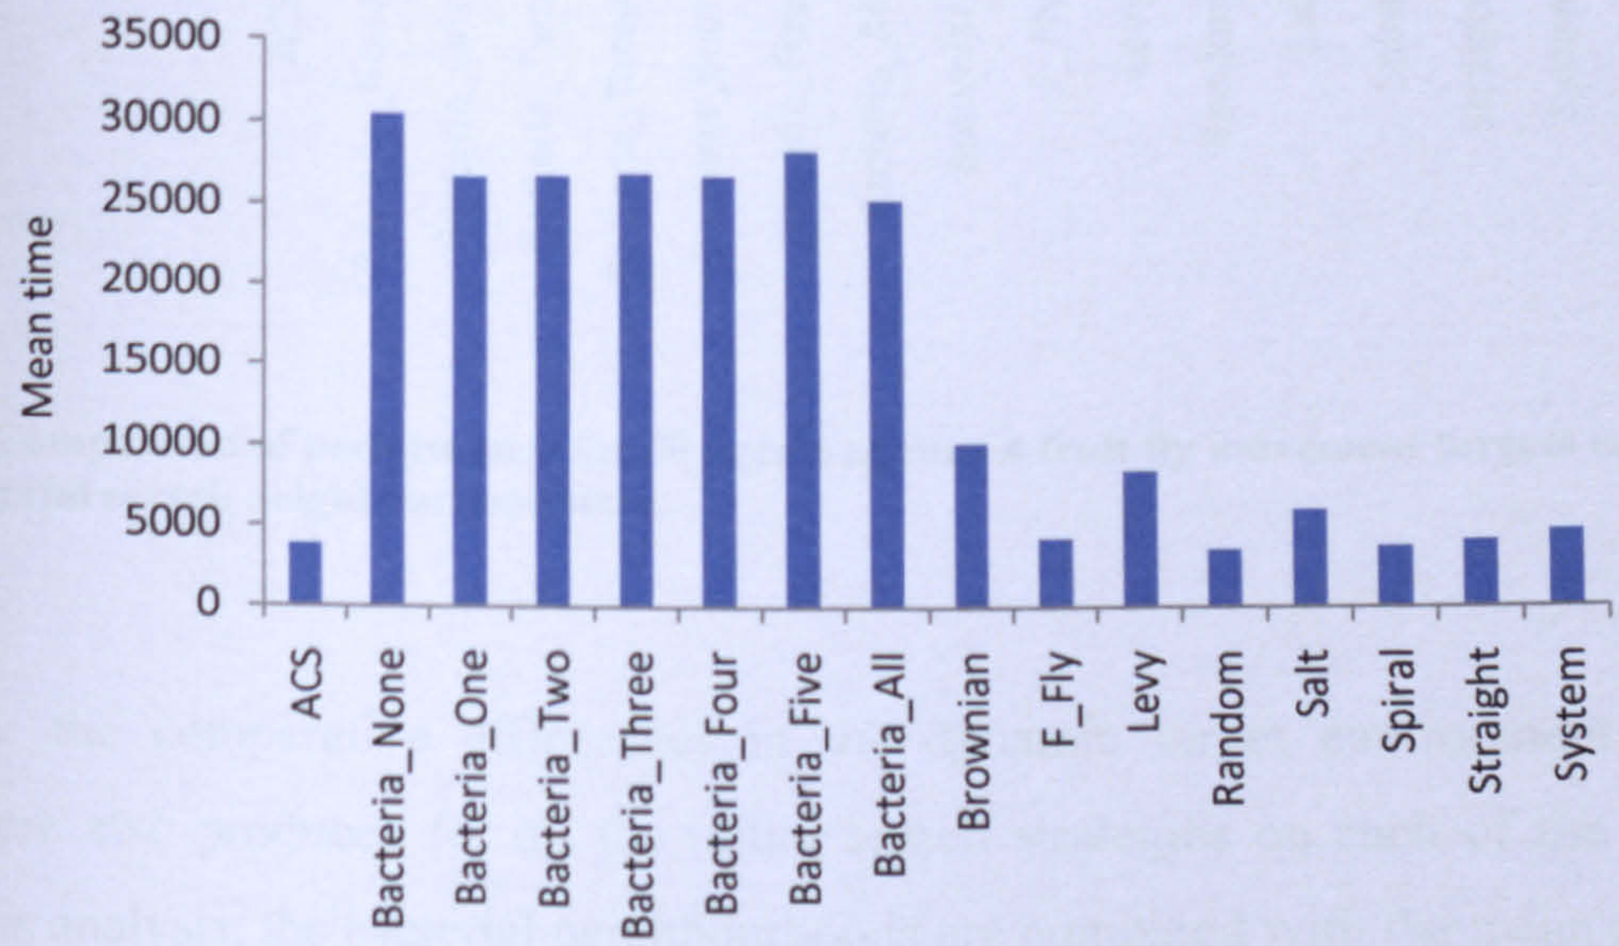

Figure 1 - 1. Comparison of performance for 10 agents against 20 uniformly distributed static targets including a range of bacterial search neighbourhood sizes.

In the dynamic target scenario (see Figure 1 - 2), differences in the mean time were less obvious although the difference between the ranging search behaviours and the bacterial search were still apparent.

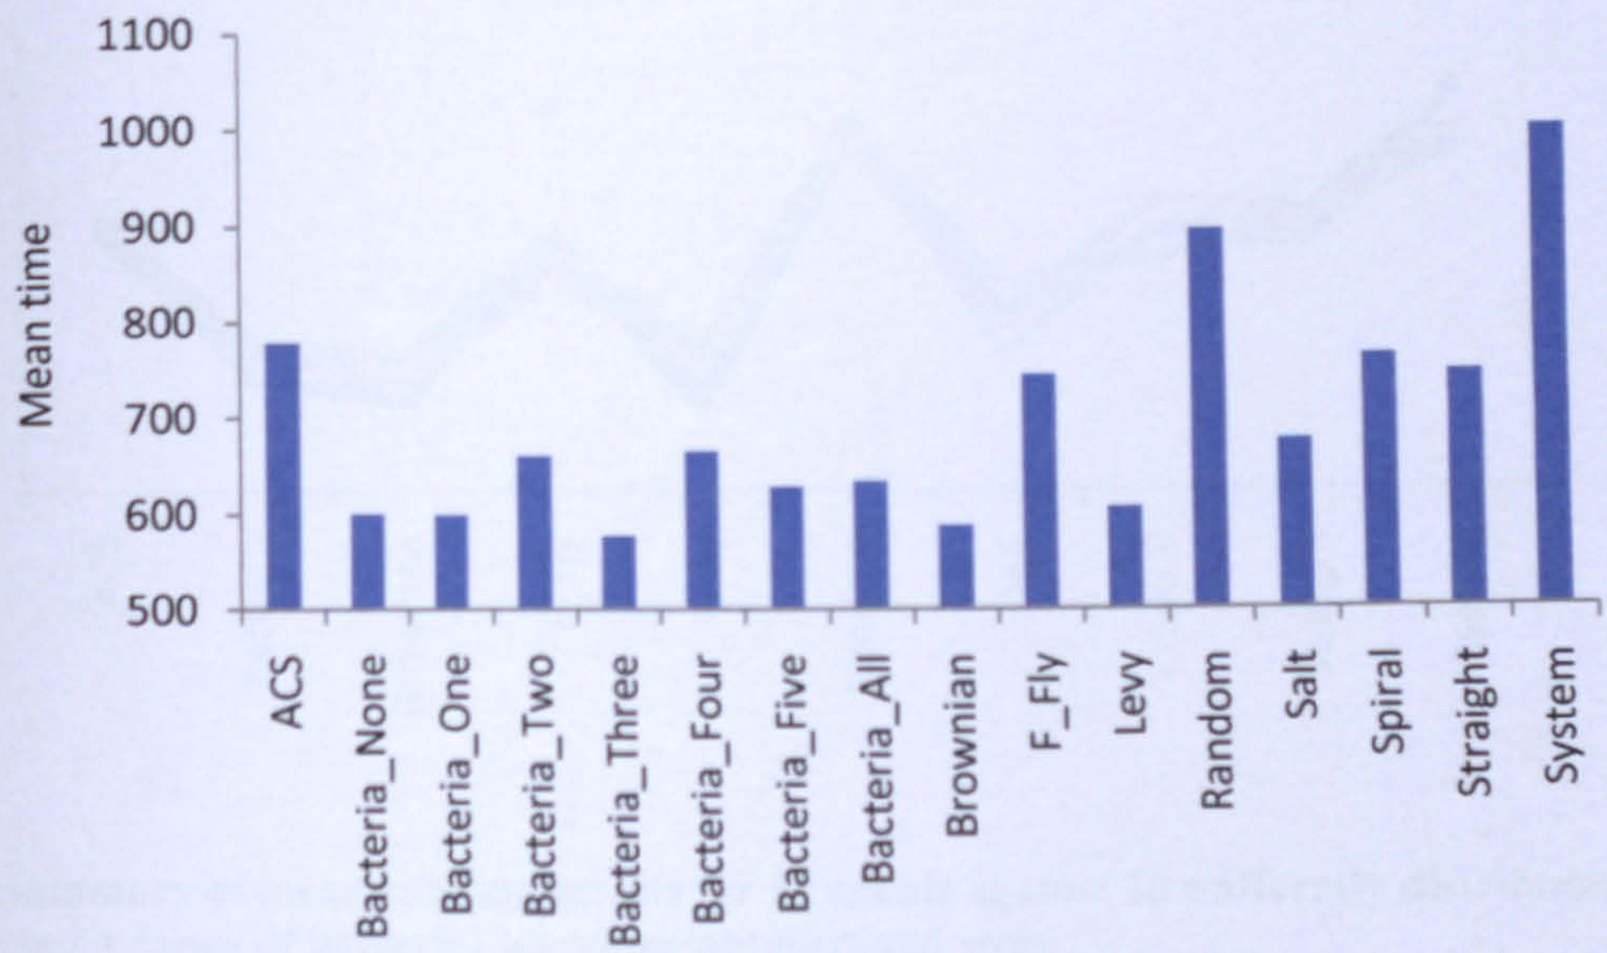

Figure 1 - 2. Comparison of performance for 20 agents against 4 fruit fly movement targets including a range of bacterial search neighbourhood sizes.

To examine the comparative differences in the dynamic target environment, the mean rankings were also produced for all the initial search strategies on each of the experiment runs. For this analysis, the bacterial neighbourhoods are compared with the mean rankings for each of the other strategies in that experimental run, the rankings for bacterial search being for the neighboourhood size shown in the key. It can be seen that the results (Figure 1 - 3) were as consistent as could be expected from a stochastic environment.

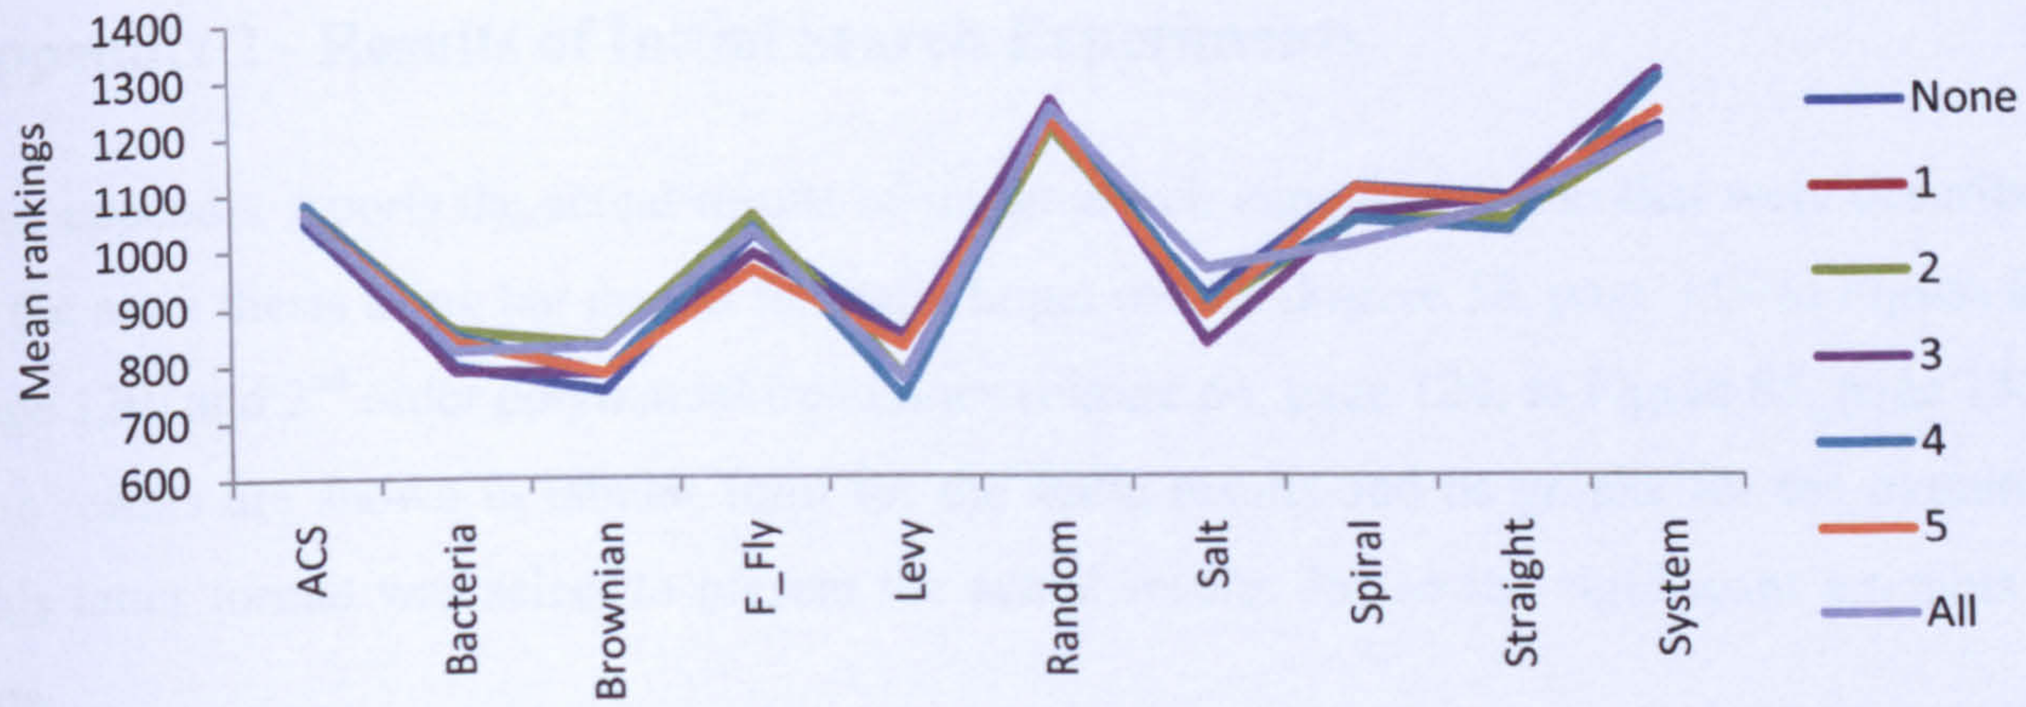

Figure 1 - 3. Summary of mean ranking results for 10 agents against 20 uniformly distributed static targets including a range of bacterial search neighbourhood sizes.

These consistent results may seem surprising, it may be expected that a switch in behaviour would affect the overall performance in some manner, especially against static targets where the additional ranging provided during detection could move the agents closer to remote targets. However, the detection period is relatively short before the target is destroyed and, furthermore, the switch to a ranging movement does not occur. The overall effect of the switch, therefore, is insufficient to impact on the results in this research and as such any of the behaviours could have been used without affecting the results.

The individual response as seen in nature was not used because the detecting agent, upon detection, moves directly to destroy the target. This would effectively remove the 'nutrient present' bacterial behaviour from the simulation which was undesirable in terms of modelling the natural behaviour. Neighbourhoods were also not applied since, although it could be argued that they could simulate the spread of a nutrient gradient across the search space, it would be difficult to reconcile the size of the neighbourhood with the natural metaphor, a nutrient patch could be any size. Hence, a decision was made in light of these findings to maintain the overall group response to a nutrient gradient and for the experimentation all agents respond to the presence of detected targets.

Appendix 2 - Results of Initial Search Experiments.

This appendix reports the actual results of initial search experimentation that were described in the main thesis using bar graphs for static target results (Figure 58, page 117 to Figure 61, page 120) and 2nd order polynomial trend lines (Figure 64, page 124, to Figure 85, page 137). The results are shown in tabular form for the static results and as graphs for the dynamic. This latter format was select to present the actual results due to the significant amounts of data.

	5 Against 4	5 Against 20	10 Against 4	10 Against 20
ACS	774.77	818.66	822.00	732.31
Bacteria	1743.38	1864.10	1562.39	1846.15
Brownian	1342.58	1493.15	1161.29	1429.67
F Fly	851.94	823.04	897.07	785.24
Levy	1221.56	1402.13	1129.86	1274.10
Random	773.96	580.95	772.58	596.06
Salt	1051.16	1153.13	1020.22	1144.72
Spiral	658.04	427.50	870.57	603.87
Straight	655.23	729.05	718.36	723.36
System	932.39	713.31	1050.69	869.53

Table 2 - 1. Actual results of locate-only strategies against static uniformly distributed targets.

	5 Against 4	5 Against 20	10 Against 4	10 Against 20
ACS	760.46	742.32	819.07	698.91
Bacteria	1393.33	1788.06	1266.65	1731.31
Brownian	1026.80	1329.88	962.32	1265.58
F Fly	902.87	787.35	952.93	793.55
Levy	984.03	1181.44	861.97	1040.29
Random	966.90	798.04	1076.51	799.87
Salt	1021.53	1093.74	960.38	1031.83
Spiral	952.72	638.37	1080.42	847.20
Straight	918.38	767.14	856.22	716.38
System	1077.96	878.66	1168.55	1080.10

Table 2 - 2. Actual results of locate-only strategies against static clustered targets.

	5 Against 4	5 Against 20	10 Against 4	10 Against 20
ACS	807.82	764.36	767.30	736.86
Bacteria	1670.01	1832.01	1627.64	1791.89
Brownian	1363.19	1481.44	1162.15	1392.01
F Fly	813.95	778.82	864.08	783.89
Levy	1189.45	1355.55	1023.65	1195.86
Random	764.42	651.51	835.71	623.79
Salt	1033.79	1077.12	1045.92	1115.75
Spiral	672.17	536.83	835.41	729.62
Straight	734.19	808.73	760.22	679.54
System	956.03	718.66	1082.94	955.79

Table 2 - 3. Actual results of locate-and-destroy strategies against static uniformly distributed targets.

	5 Against 4	5 Against 20	10 Against 4	10 Against 20
ACS	648.80	519.64	678.95	543.31
Bacteria	1076.83	1586.94	904.96	1353.09
Brownian	926.29	1122.19	788.10	1019.21
F Fly	979.54	839.95	1037.77	883.19
Levy	869.88	1034.26	893.20	968.48
Random	1237.95	984.89	1286.76	1064.17
Salt	899.05	922.07	950.79	891.34
Spiral	1069.52	811.70	1079.69	972.53
Straight	1090.84	1036.52	1142.33	1031.00
System	1206.33	1146.85	1242.46	1278.68

Table 2 - 4. Actual results of locate-and-destroy strategies against static clustered targets.

Figure 2 - 1. Summary of mean ranking results for 10 agents against 1 target moving deterministically.

Figure 2 - 2. Summary of mean ranking results for 10 agents against 4 targets moving deterministically.

Figure 2 - 3. Summary of mean ranking results for 20 agents against 1 target moving deterministically.

Figure 2 - 4. Summary of mean ranking results for 20 agents against 4 targets moving deterministically.

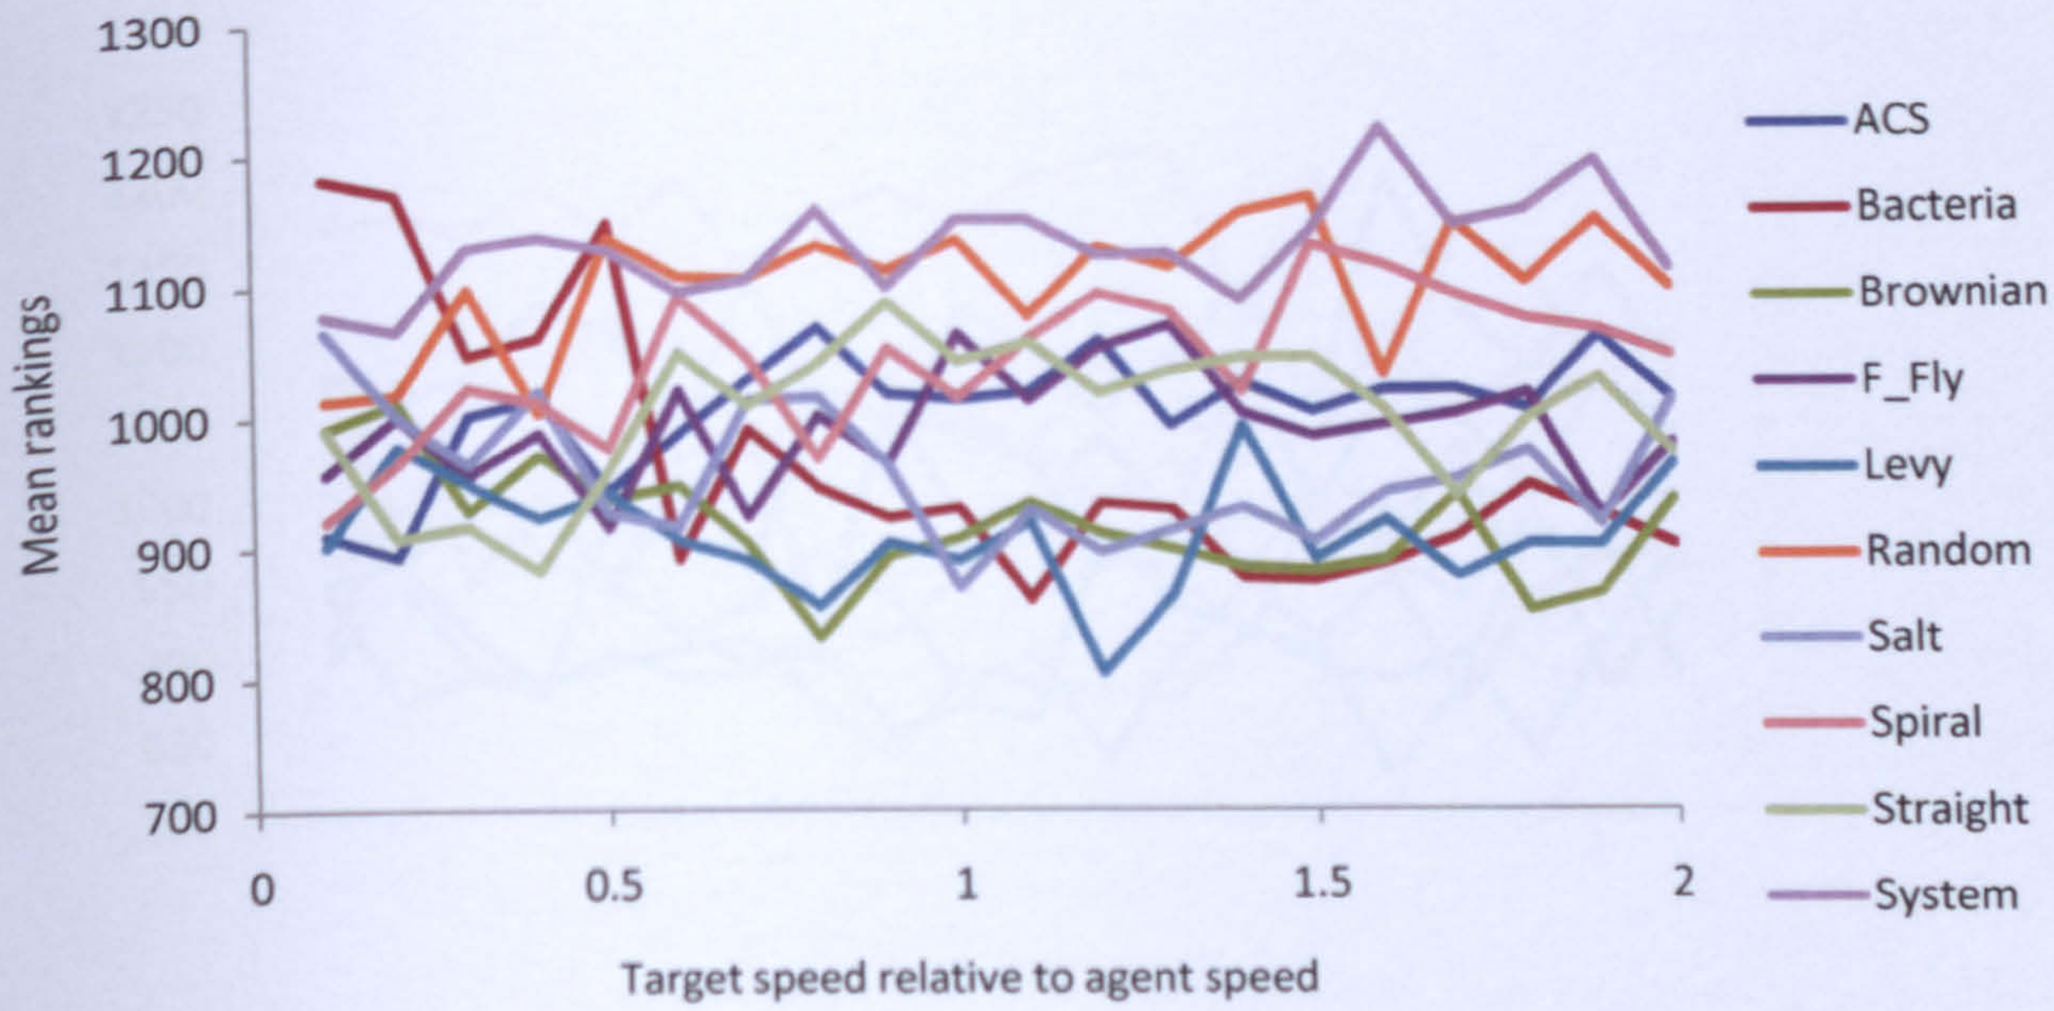

Figure 2 - 5. Summary of mean ranking results for 10 agents against 1 target moving randomly.

Figure 2 - 6. Summary of mean ranking results for 10 agents against 4 targets moving randomly.

Figure 2 - 7. Summary of mean ranking results for 20 agents against 1 target moving randomly.

Figure 2 - 8. Summary of mean ranking results for 20 agents against 4 targets moving randomly.

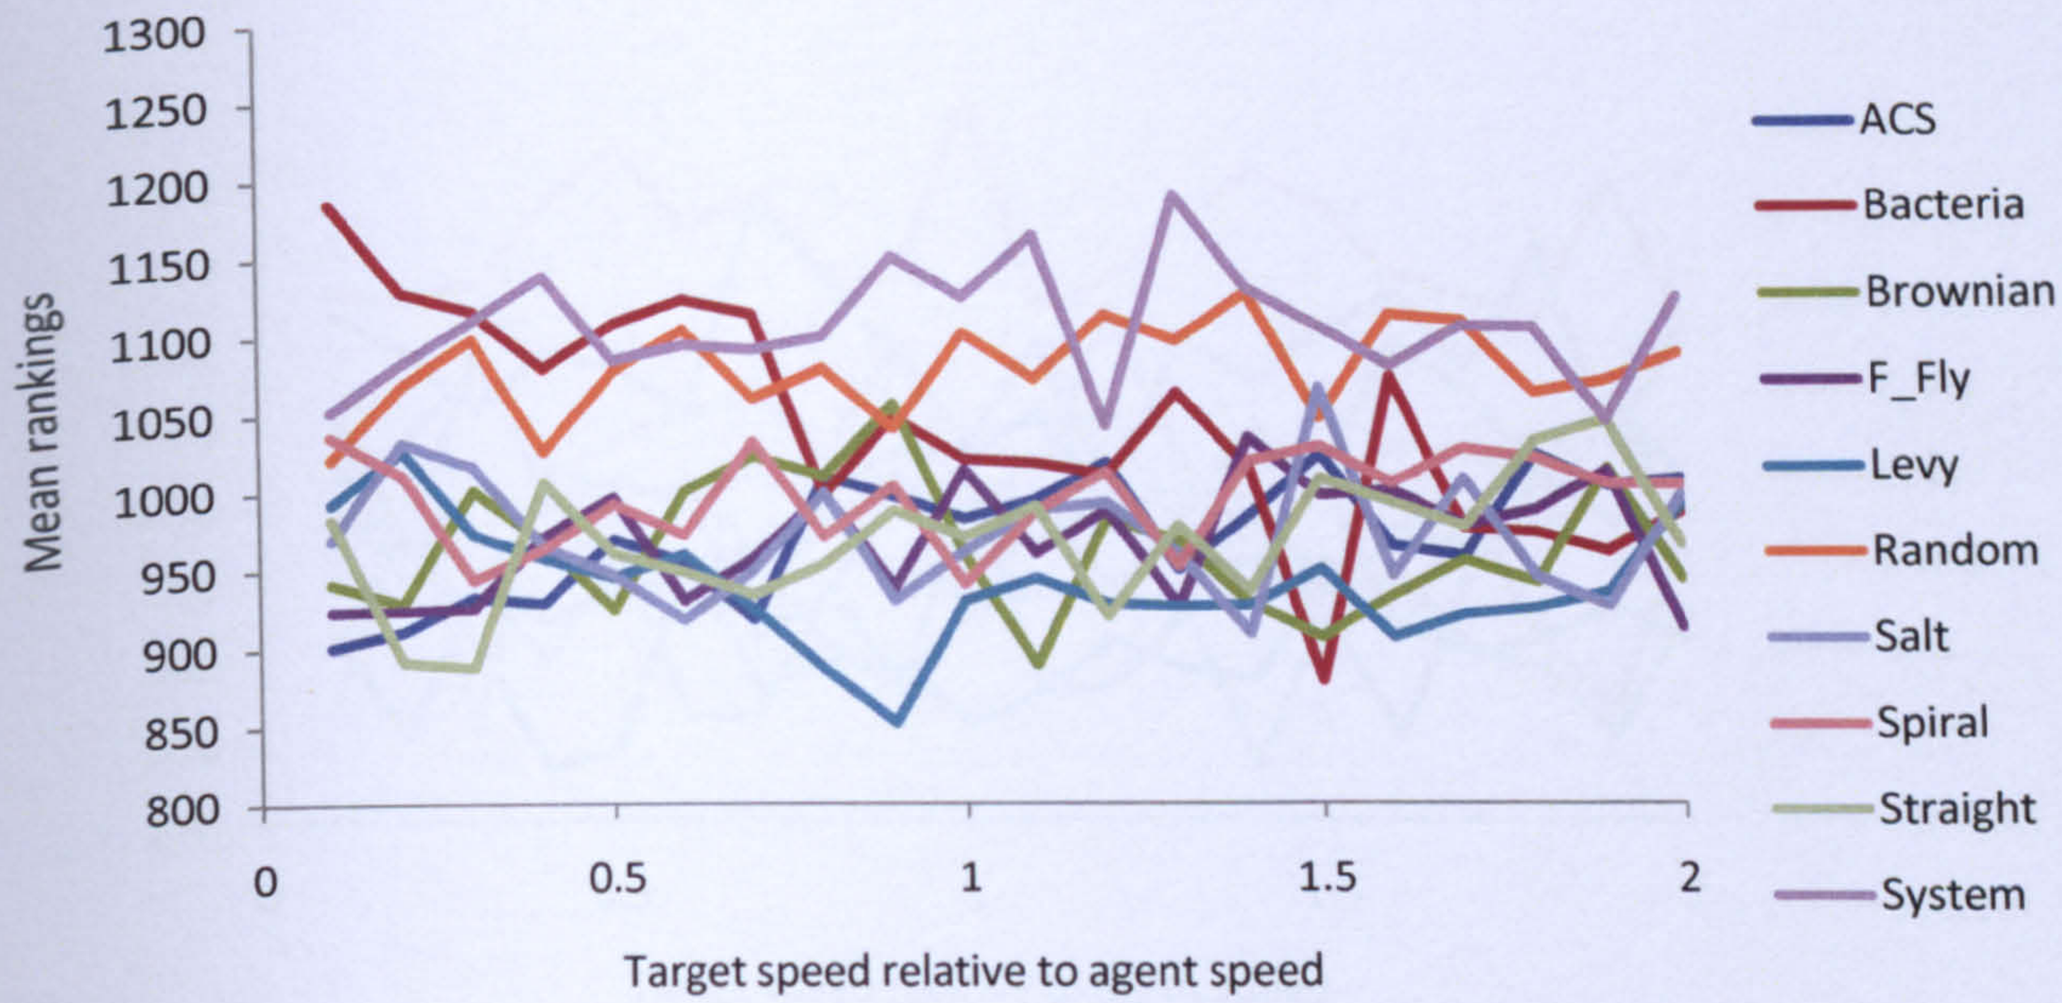

Figure 2 - 9. Summary of mean ranking results for 10 agents against 1 Brownian motion target.

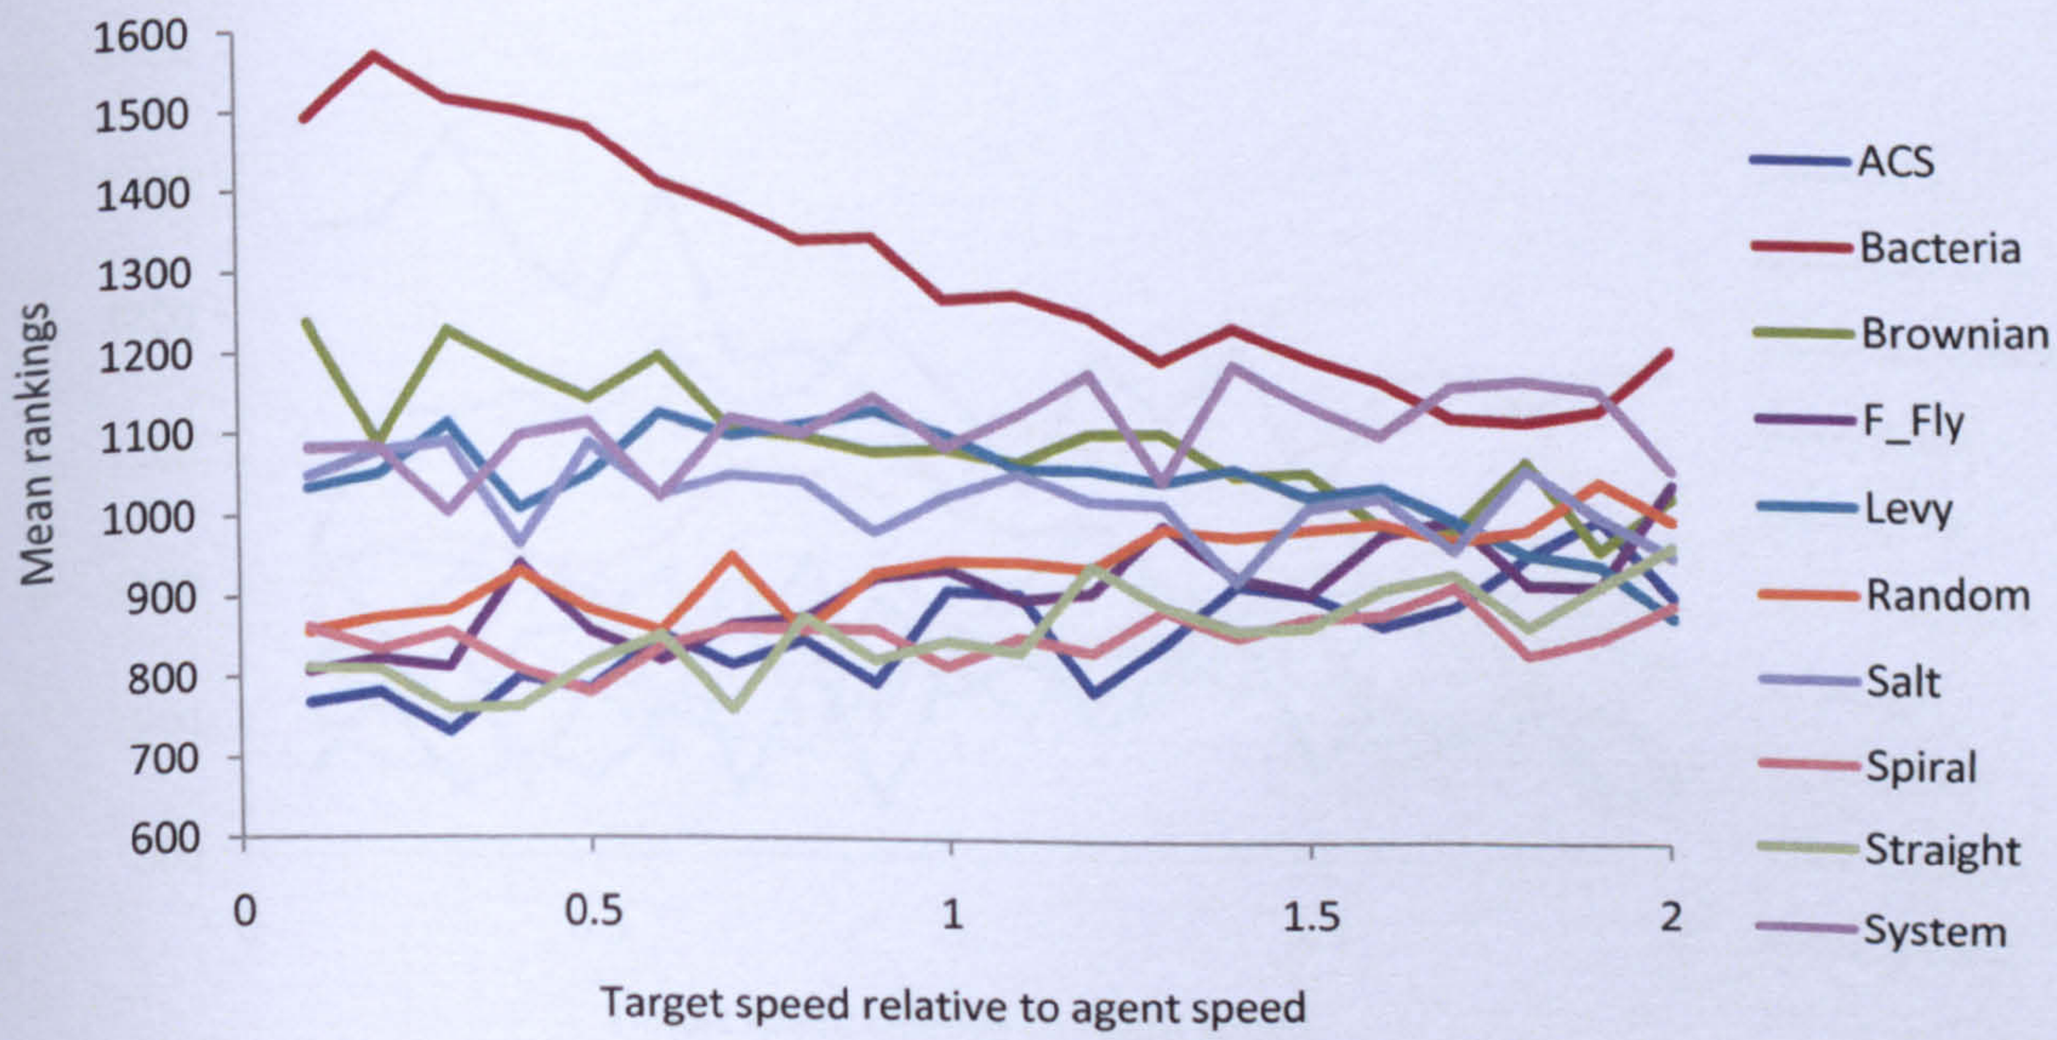

Figure 2 - 10. Summary of mean ranking results for 10 agents against 4 Brownian motion targets.

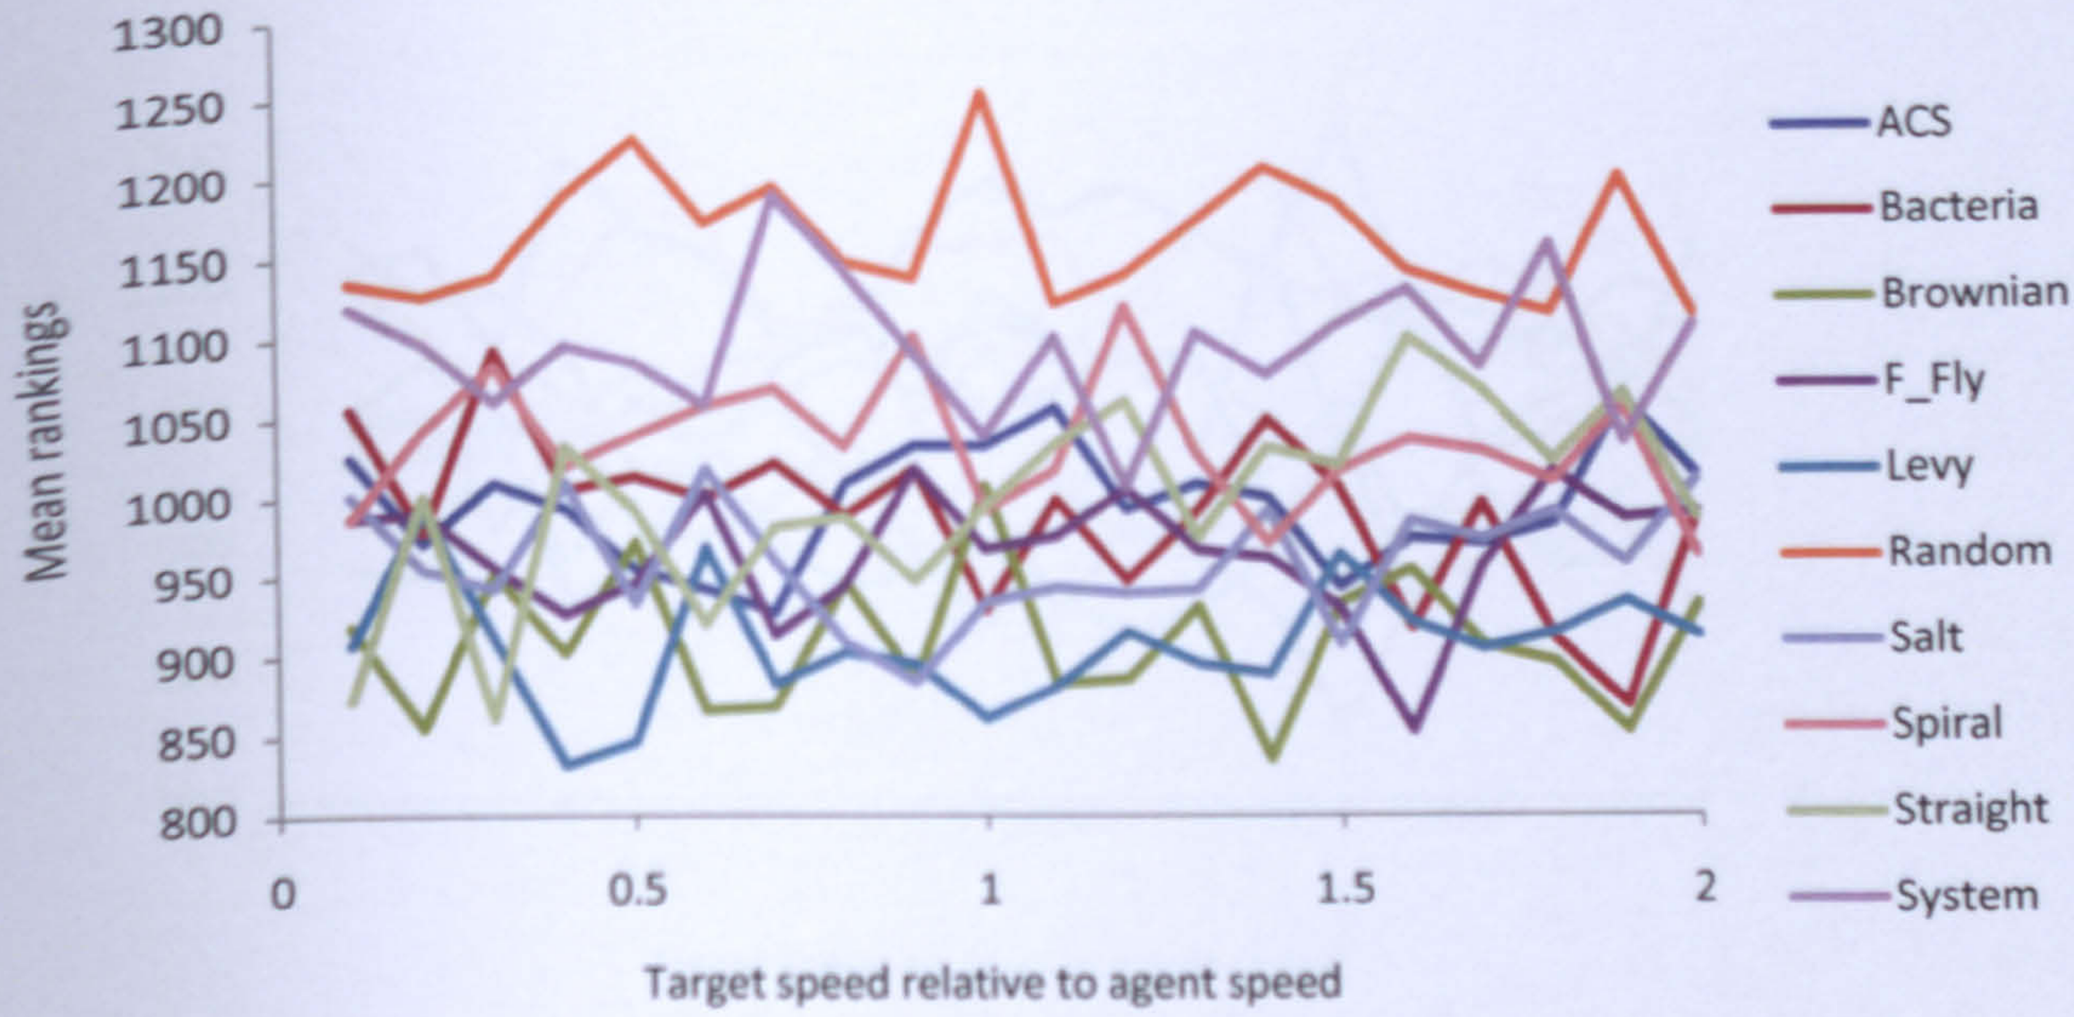

Figure 2 - 11. Summary of mean ranking results for 20 agents against 1 Brownian motion target.

Figure 2 - 12. Summary of mean ranking results for 20 agents against 4 Brownian motion targets.

Figure 2 - 13. Summary of mean ranking results for 10 agents against 1 fruit fly target.

Figure 2 - 14. Summary of mean ranking results for 10 agents against 4 fruit fly targets.

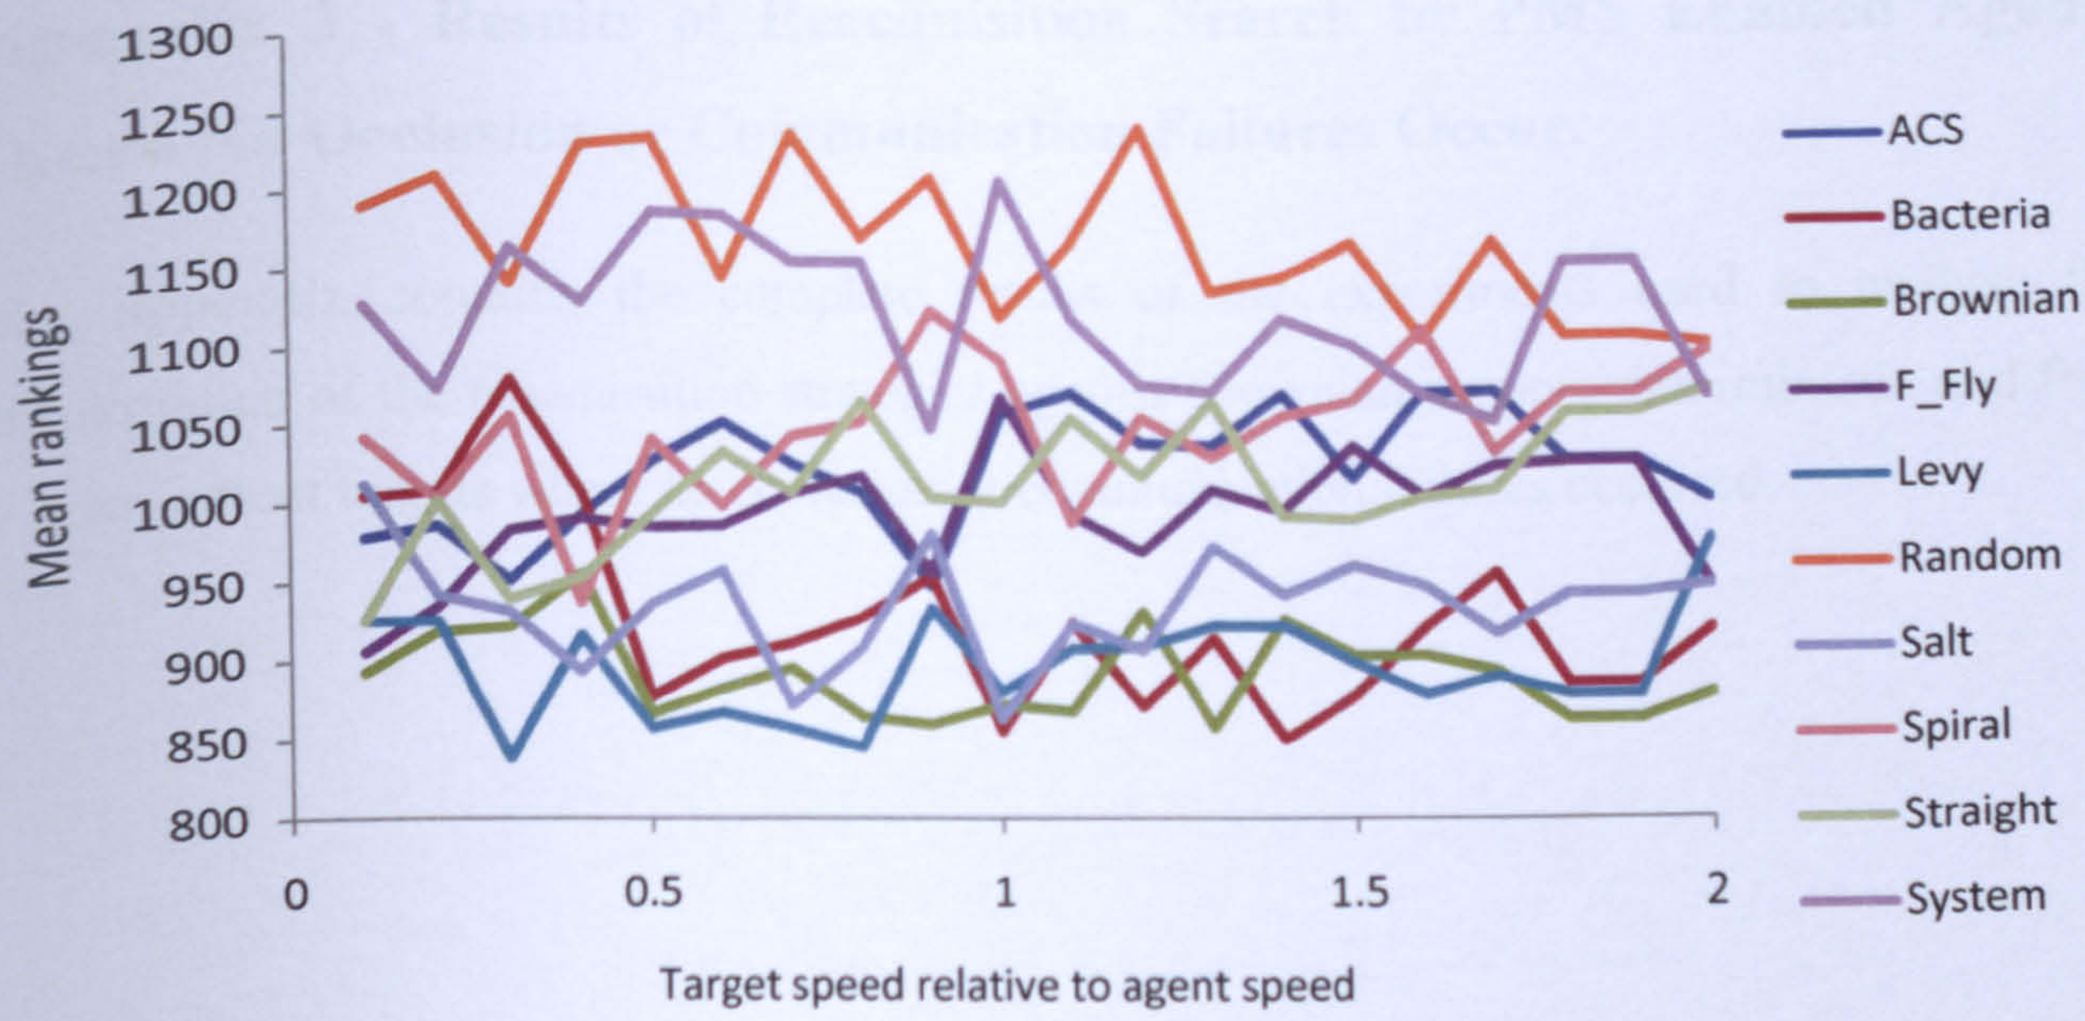

Figure 2 - 15. Summary of mean ranking results for 20 agents against 1 fruit fly target.

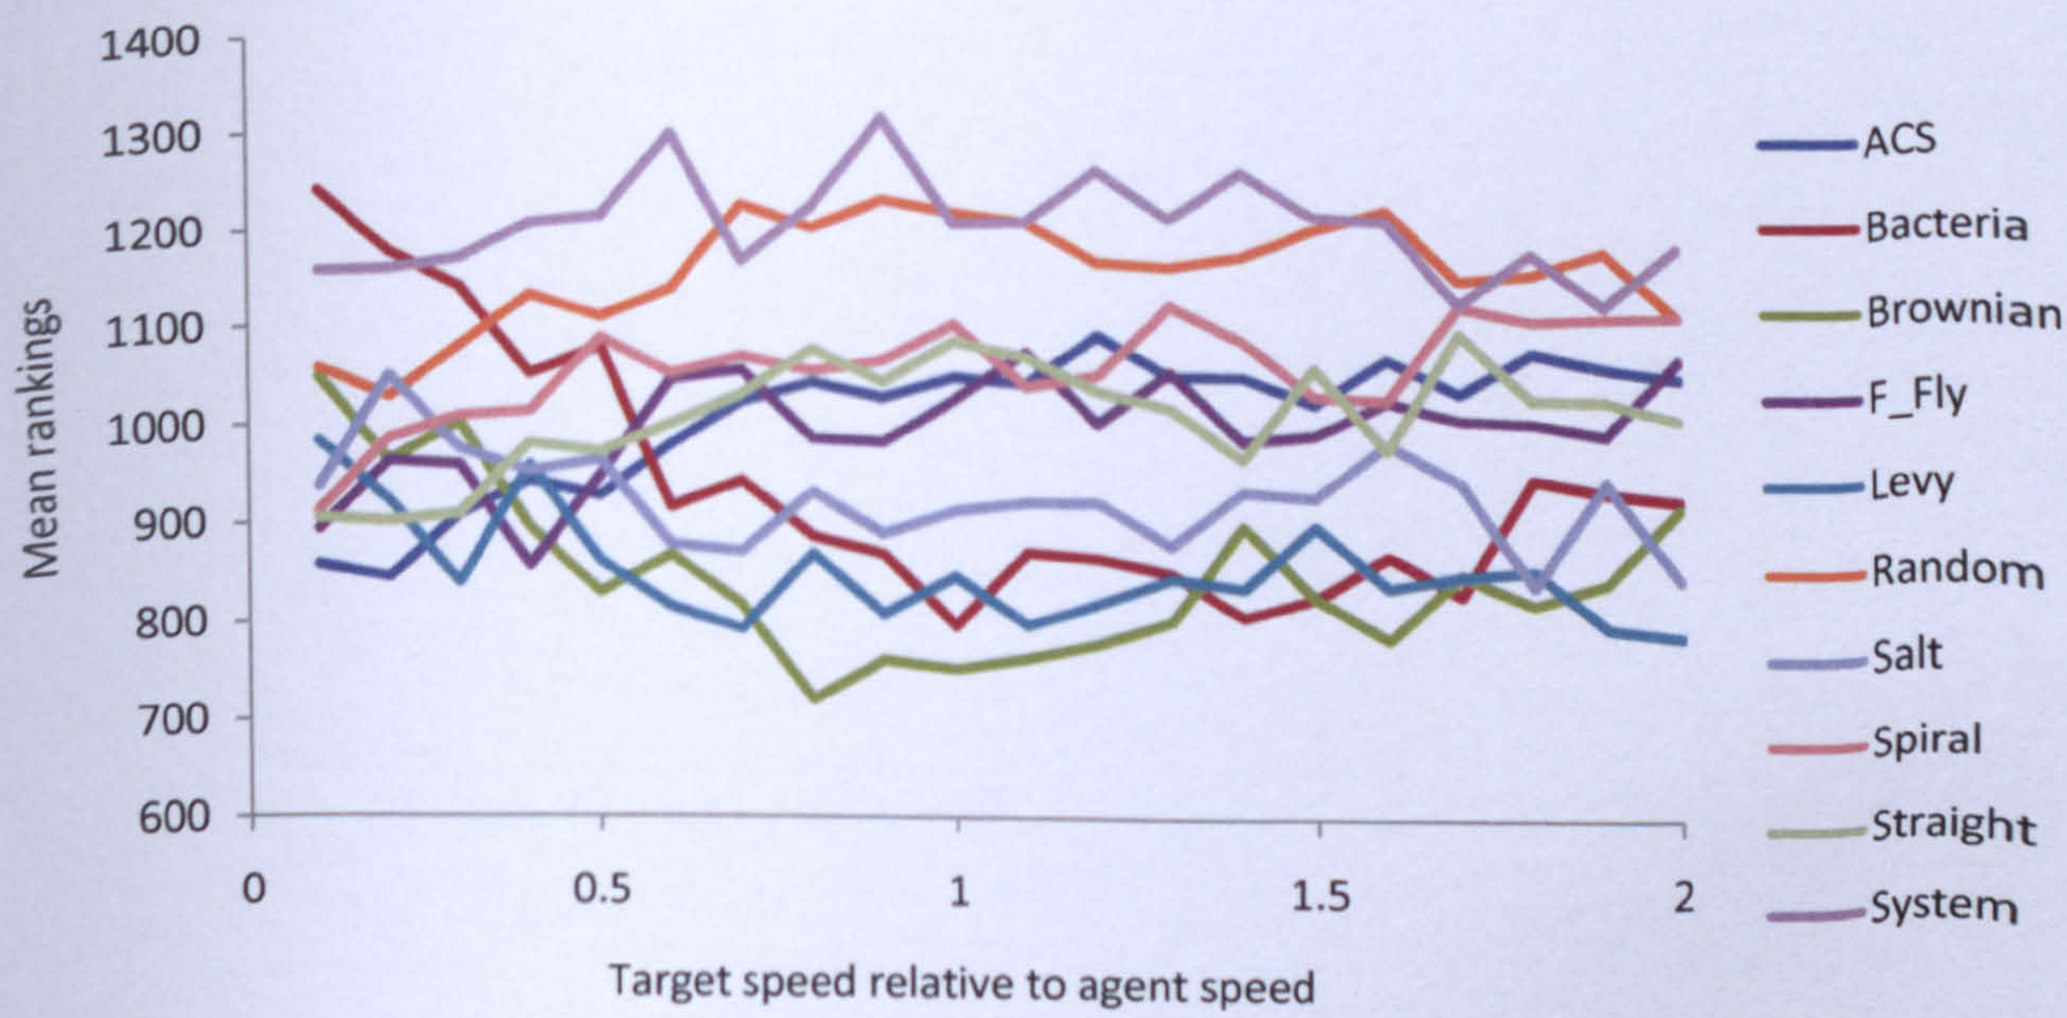

Figure 2 - 16. Summary of mean ranking results for 20 agents against 4 fruit fly targets.

Appendix 3 - Results of Reacquisition Search by PMS Enabled Agents Where No Occlusion or Communication Failures Occur.

This appendix contains the complete results of the experiments used to analyse the performance of the reacquisition strategies against Brownian motion, deterministic and fruit fly movement targets where no occlusion or communication failures occurred.

Figure 3 - 1. Mean time for 10 agents against 4 Brownian motion targets travelling at 1.5 agent speed (left) and 2.0 agent speed (right).

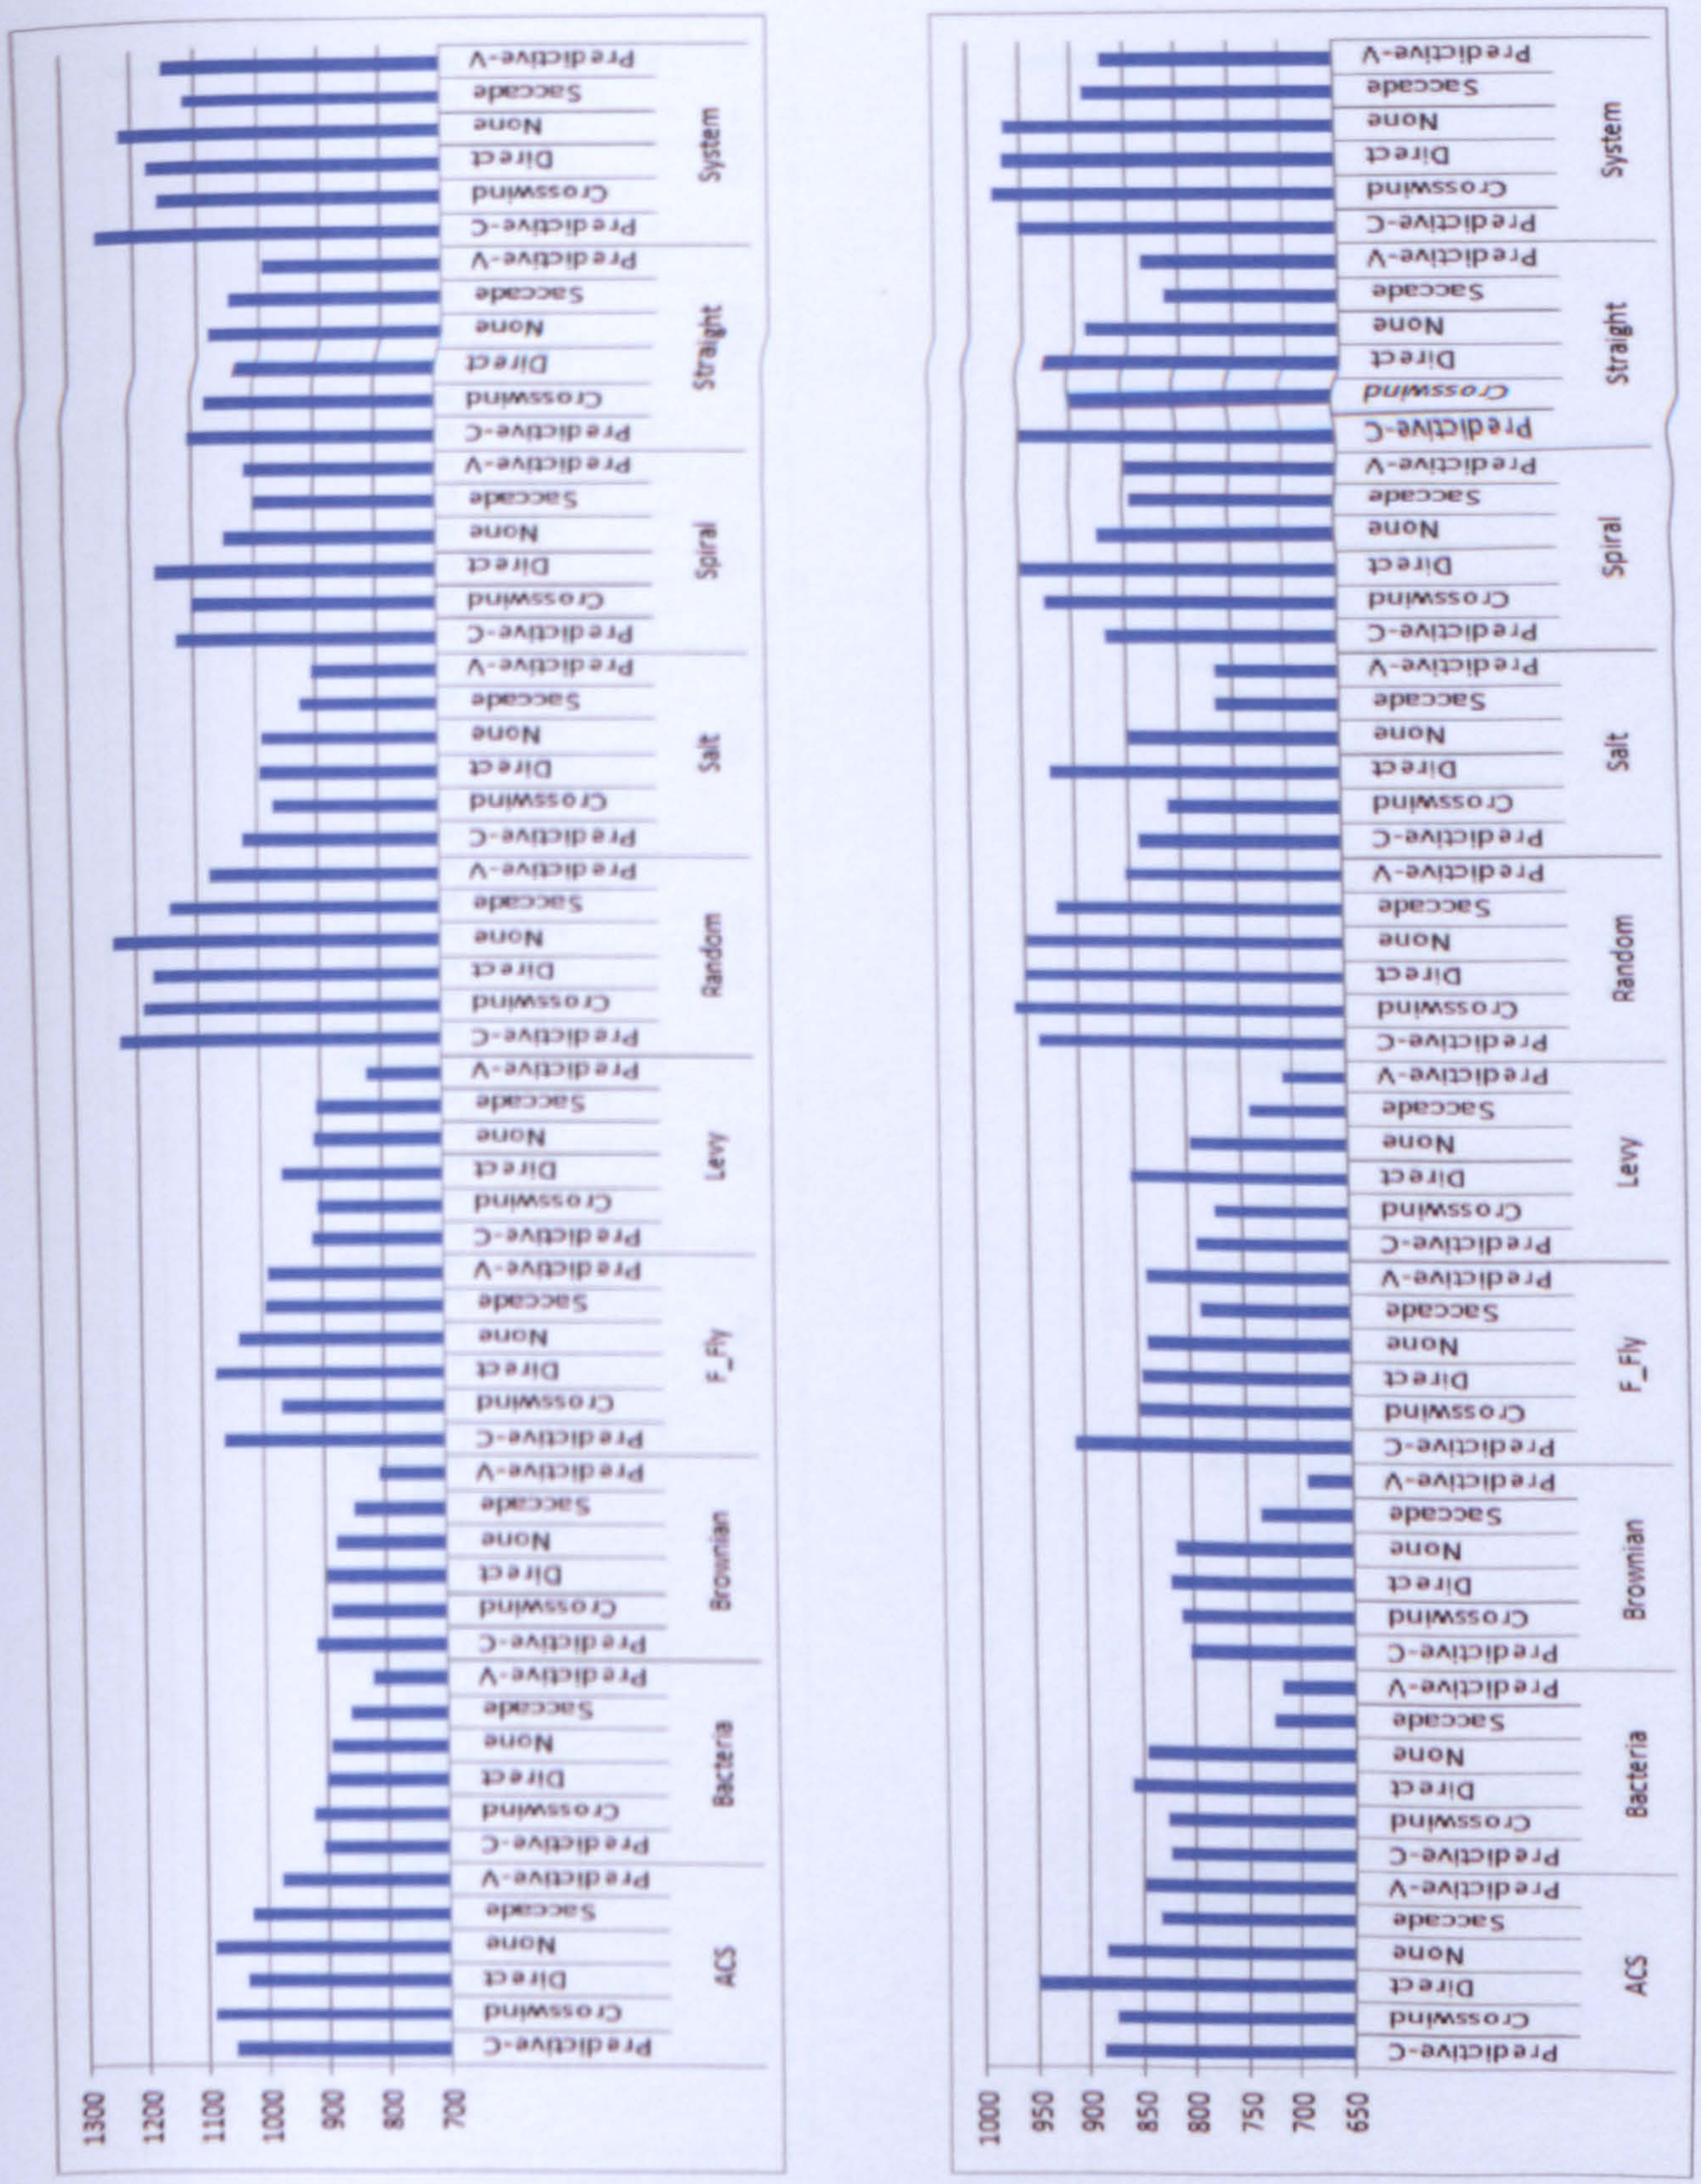

Figure 3 - 2. Mean time for 10 agents against 4 Determinstic targets travelling at 1.5 agent speed (left) and 2.0 agent speed (right).

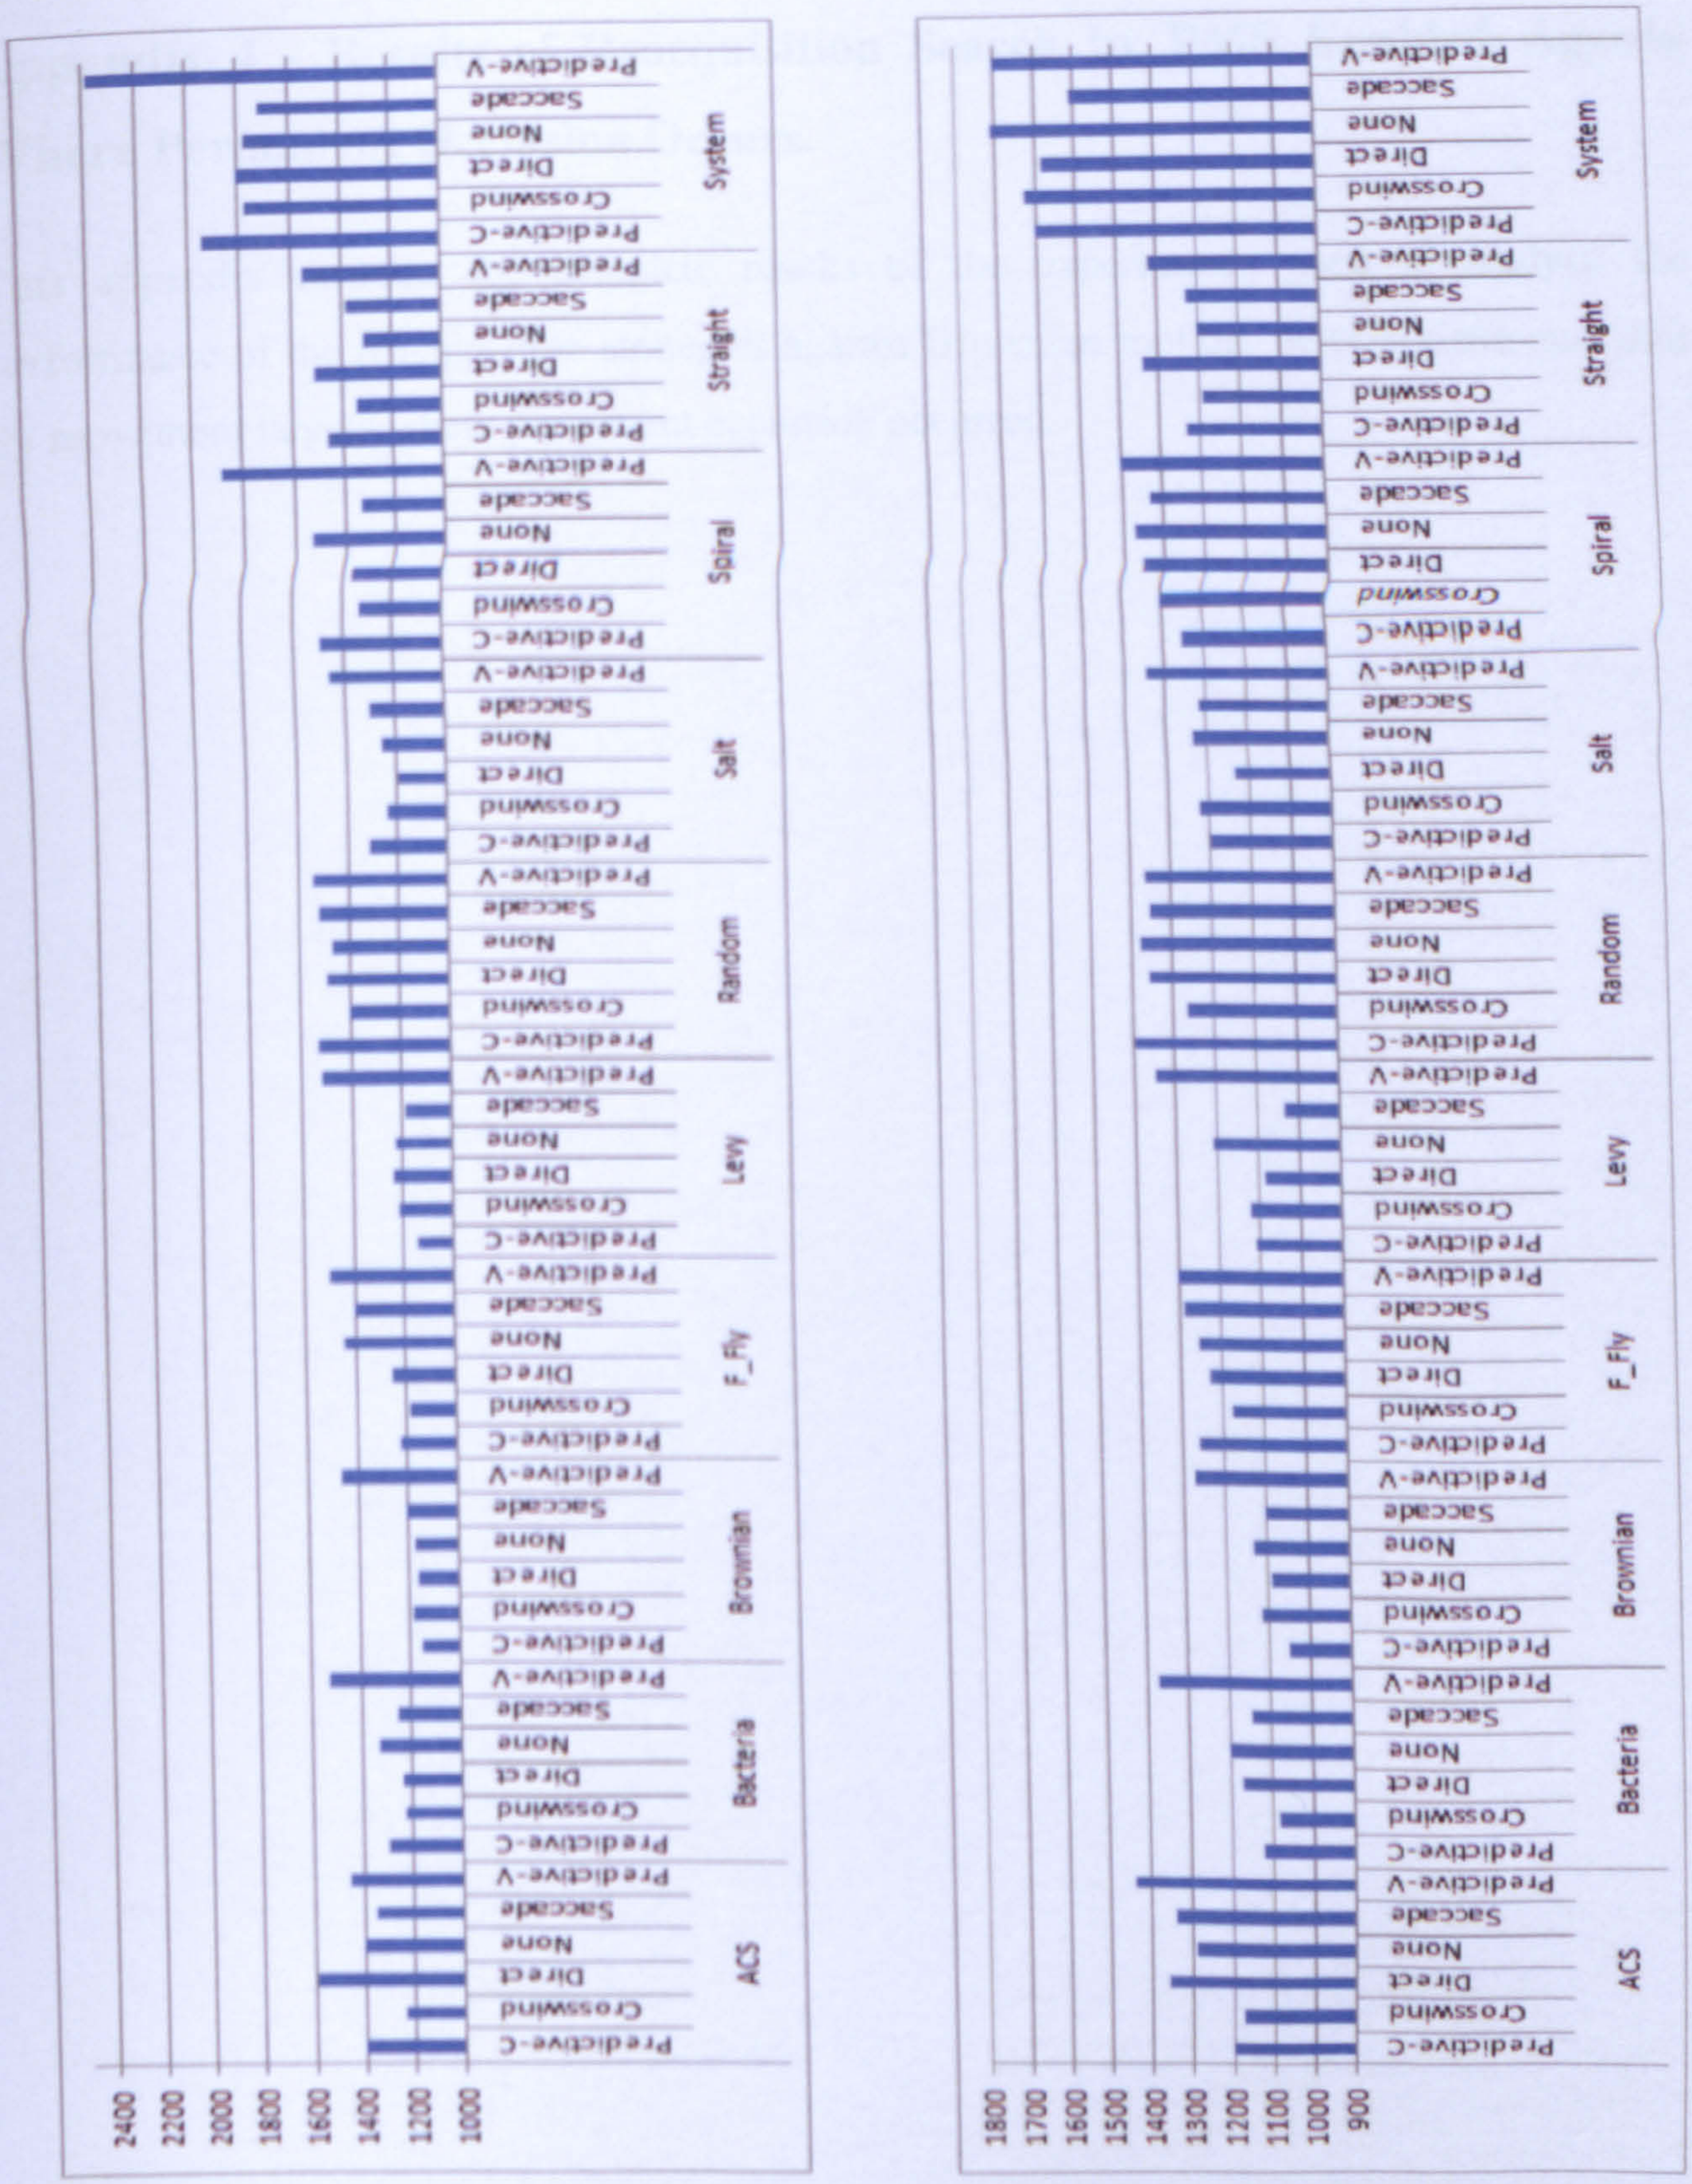

Figure 3 - 3. Mean time for 10 agents against 4 fruit fly movement targets travelling at 1.5 agent speed (left) and 2.0 agent speed (right).

Appendix 4 - *Results of Reacquisition Search by PMS Enabled Agents Where Permanent Occlusion Occurs.*

This appendix contains the complete results of the experiments used to analyse the performance of the reacquisition strategies against Brownian motion, deterministic and fruit fly movement targets where permanent occlusion occurred.

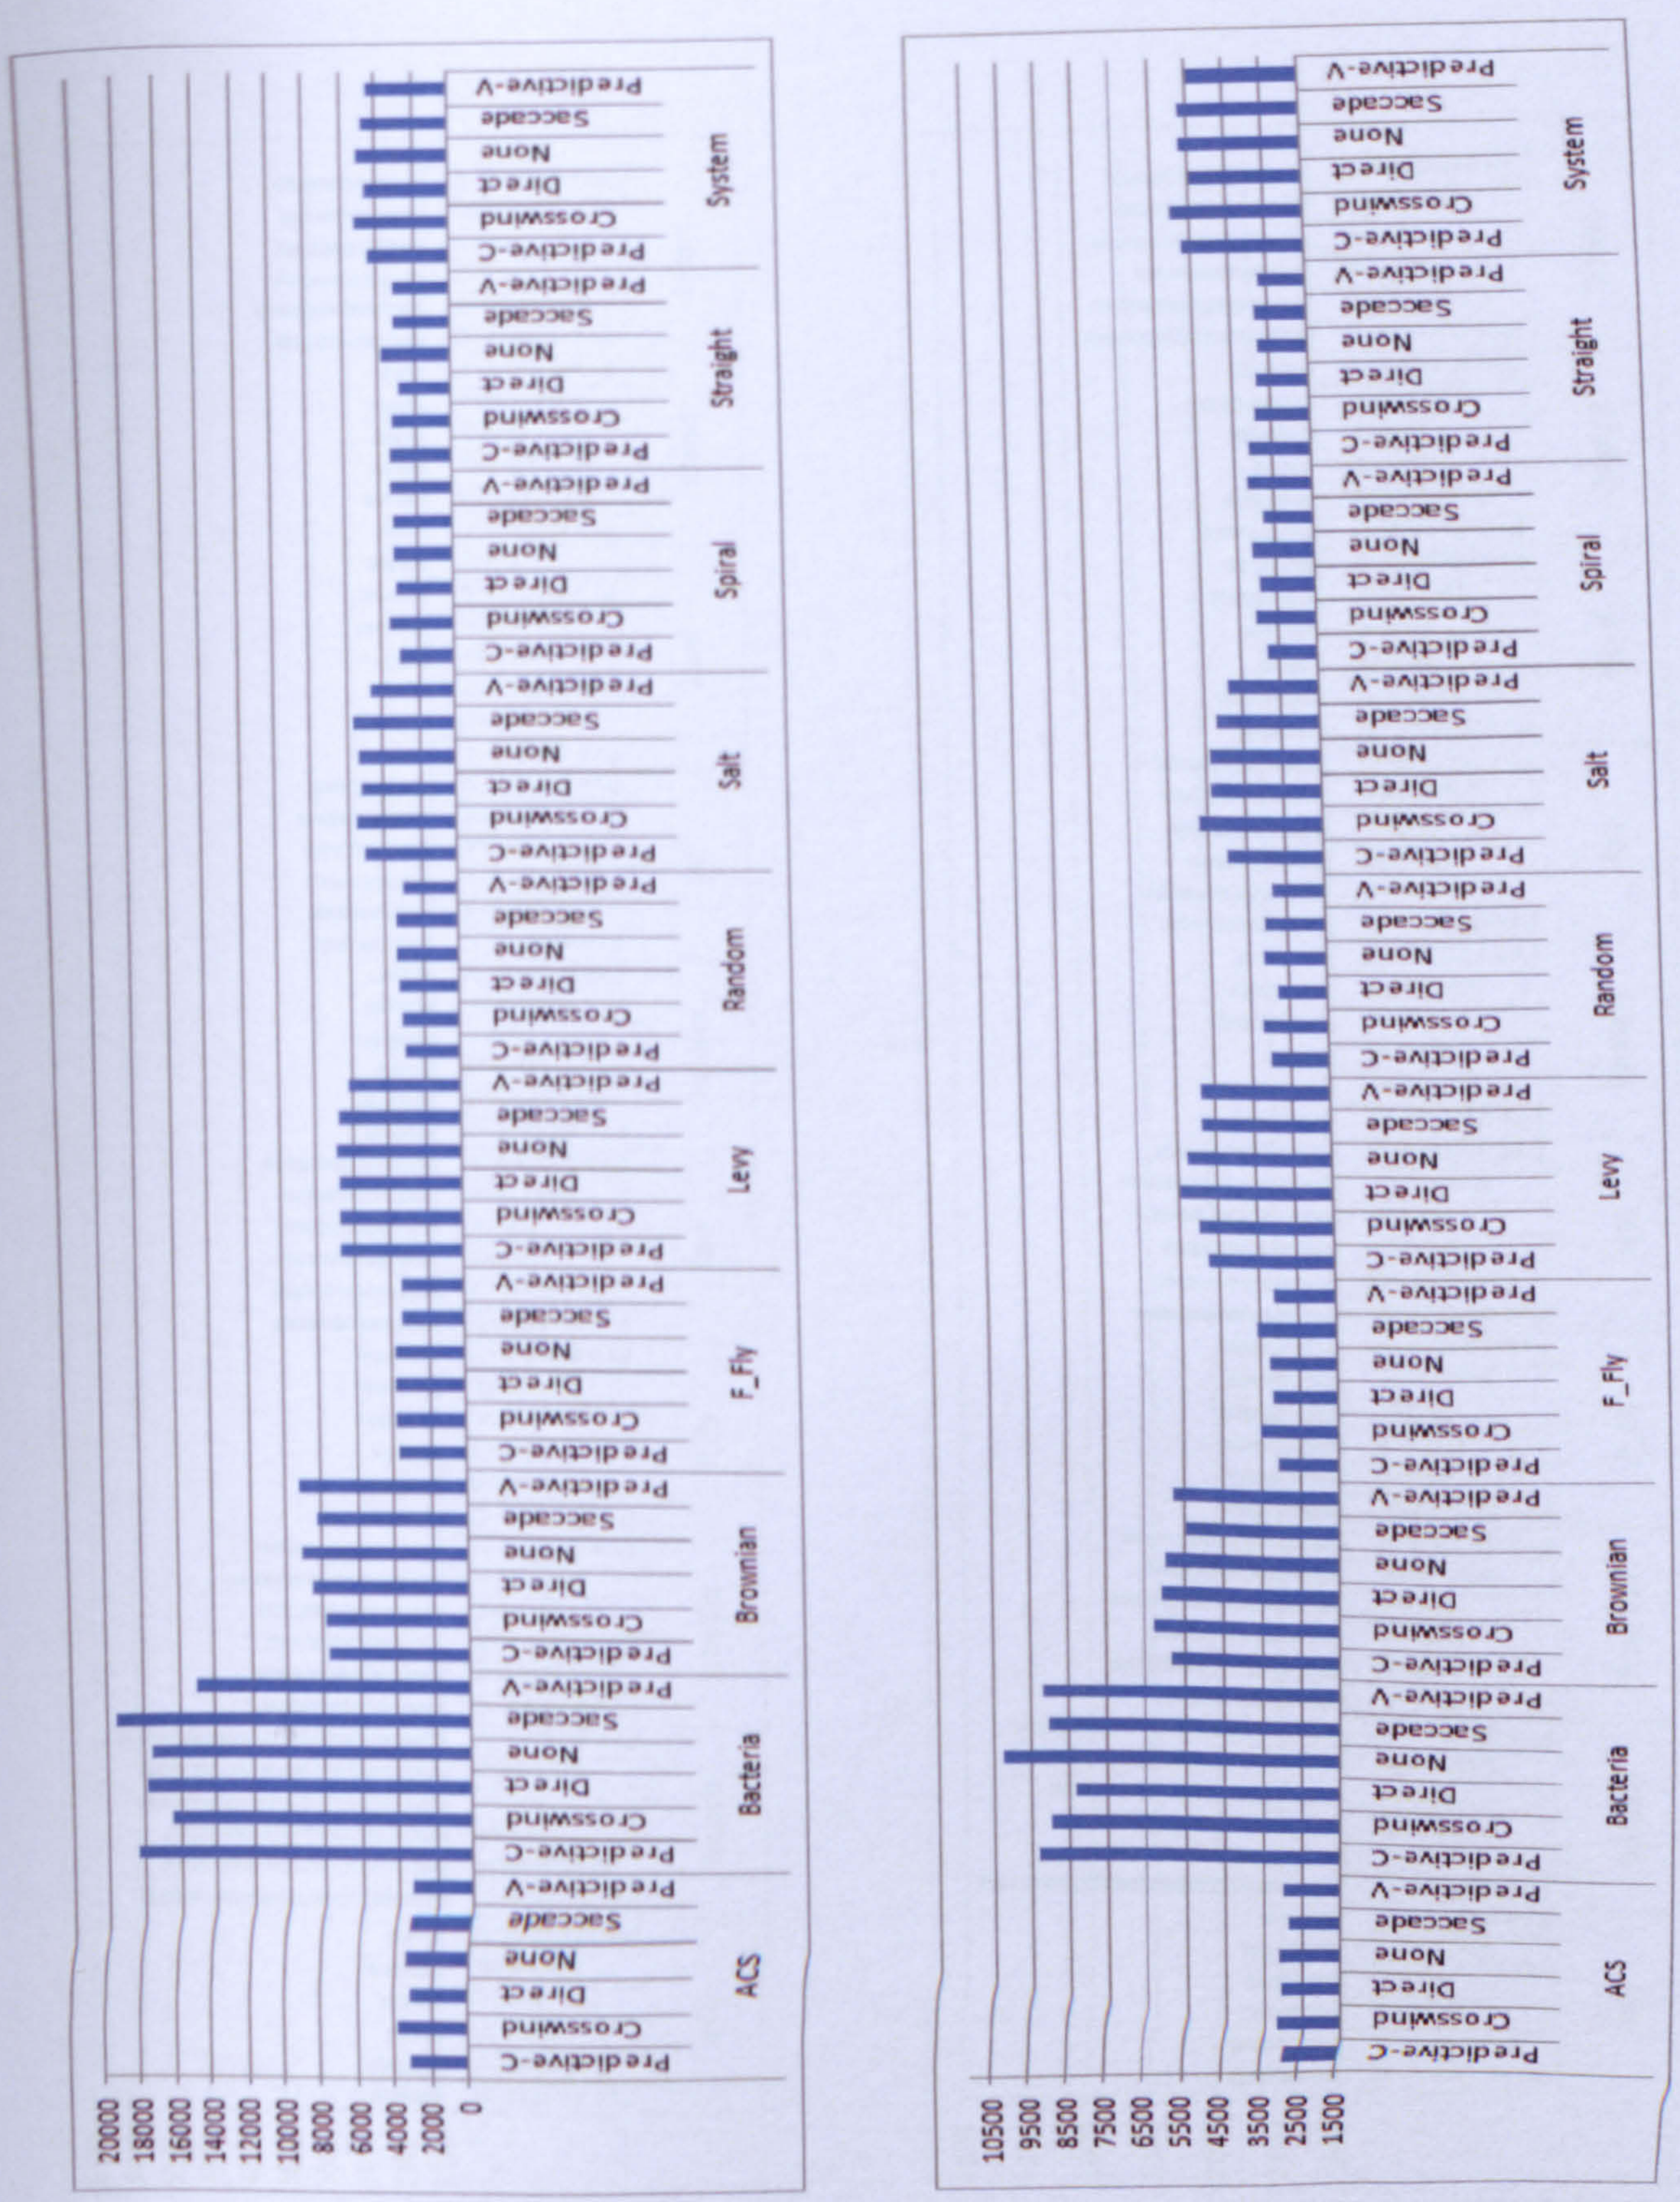

Figure 4 - 1. Mean time for 10 agents against 4 Brownian motion targets travelling at 0.5 agent speed (left) and 1.0 agent speed (right).

Figure 4 - 2. Mean time for 10 agents against 4 Brownian motion targets travelling at 1.5 agent speed (left) and 2.0 agent speed (right).

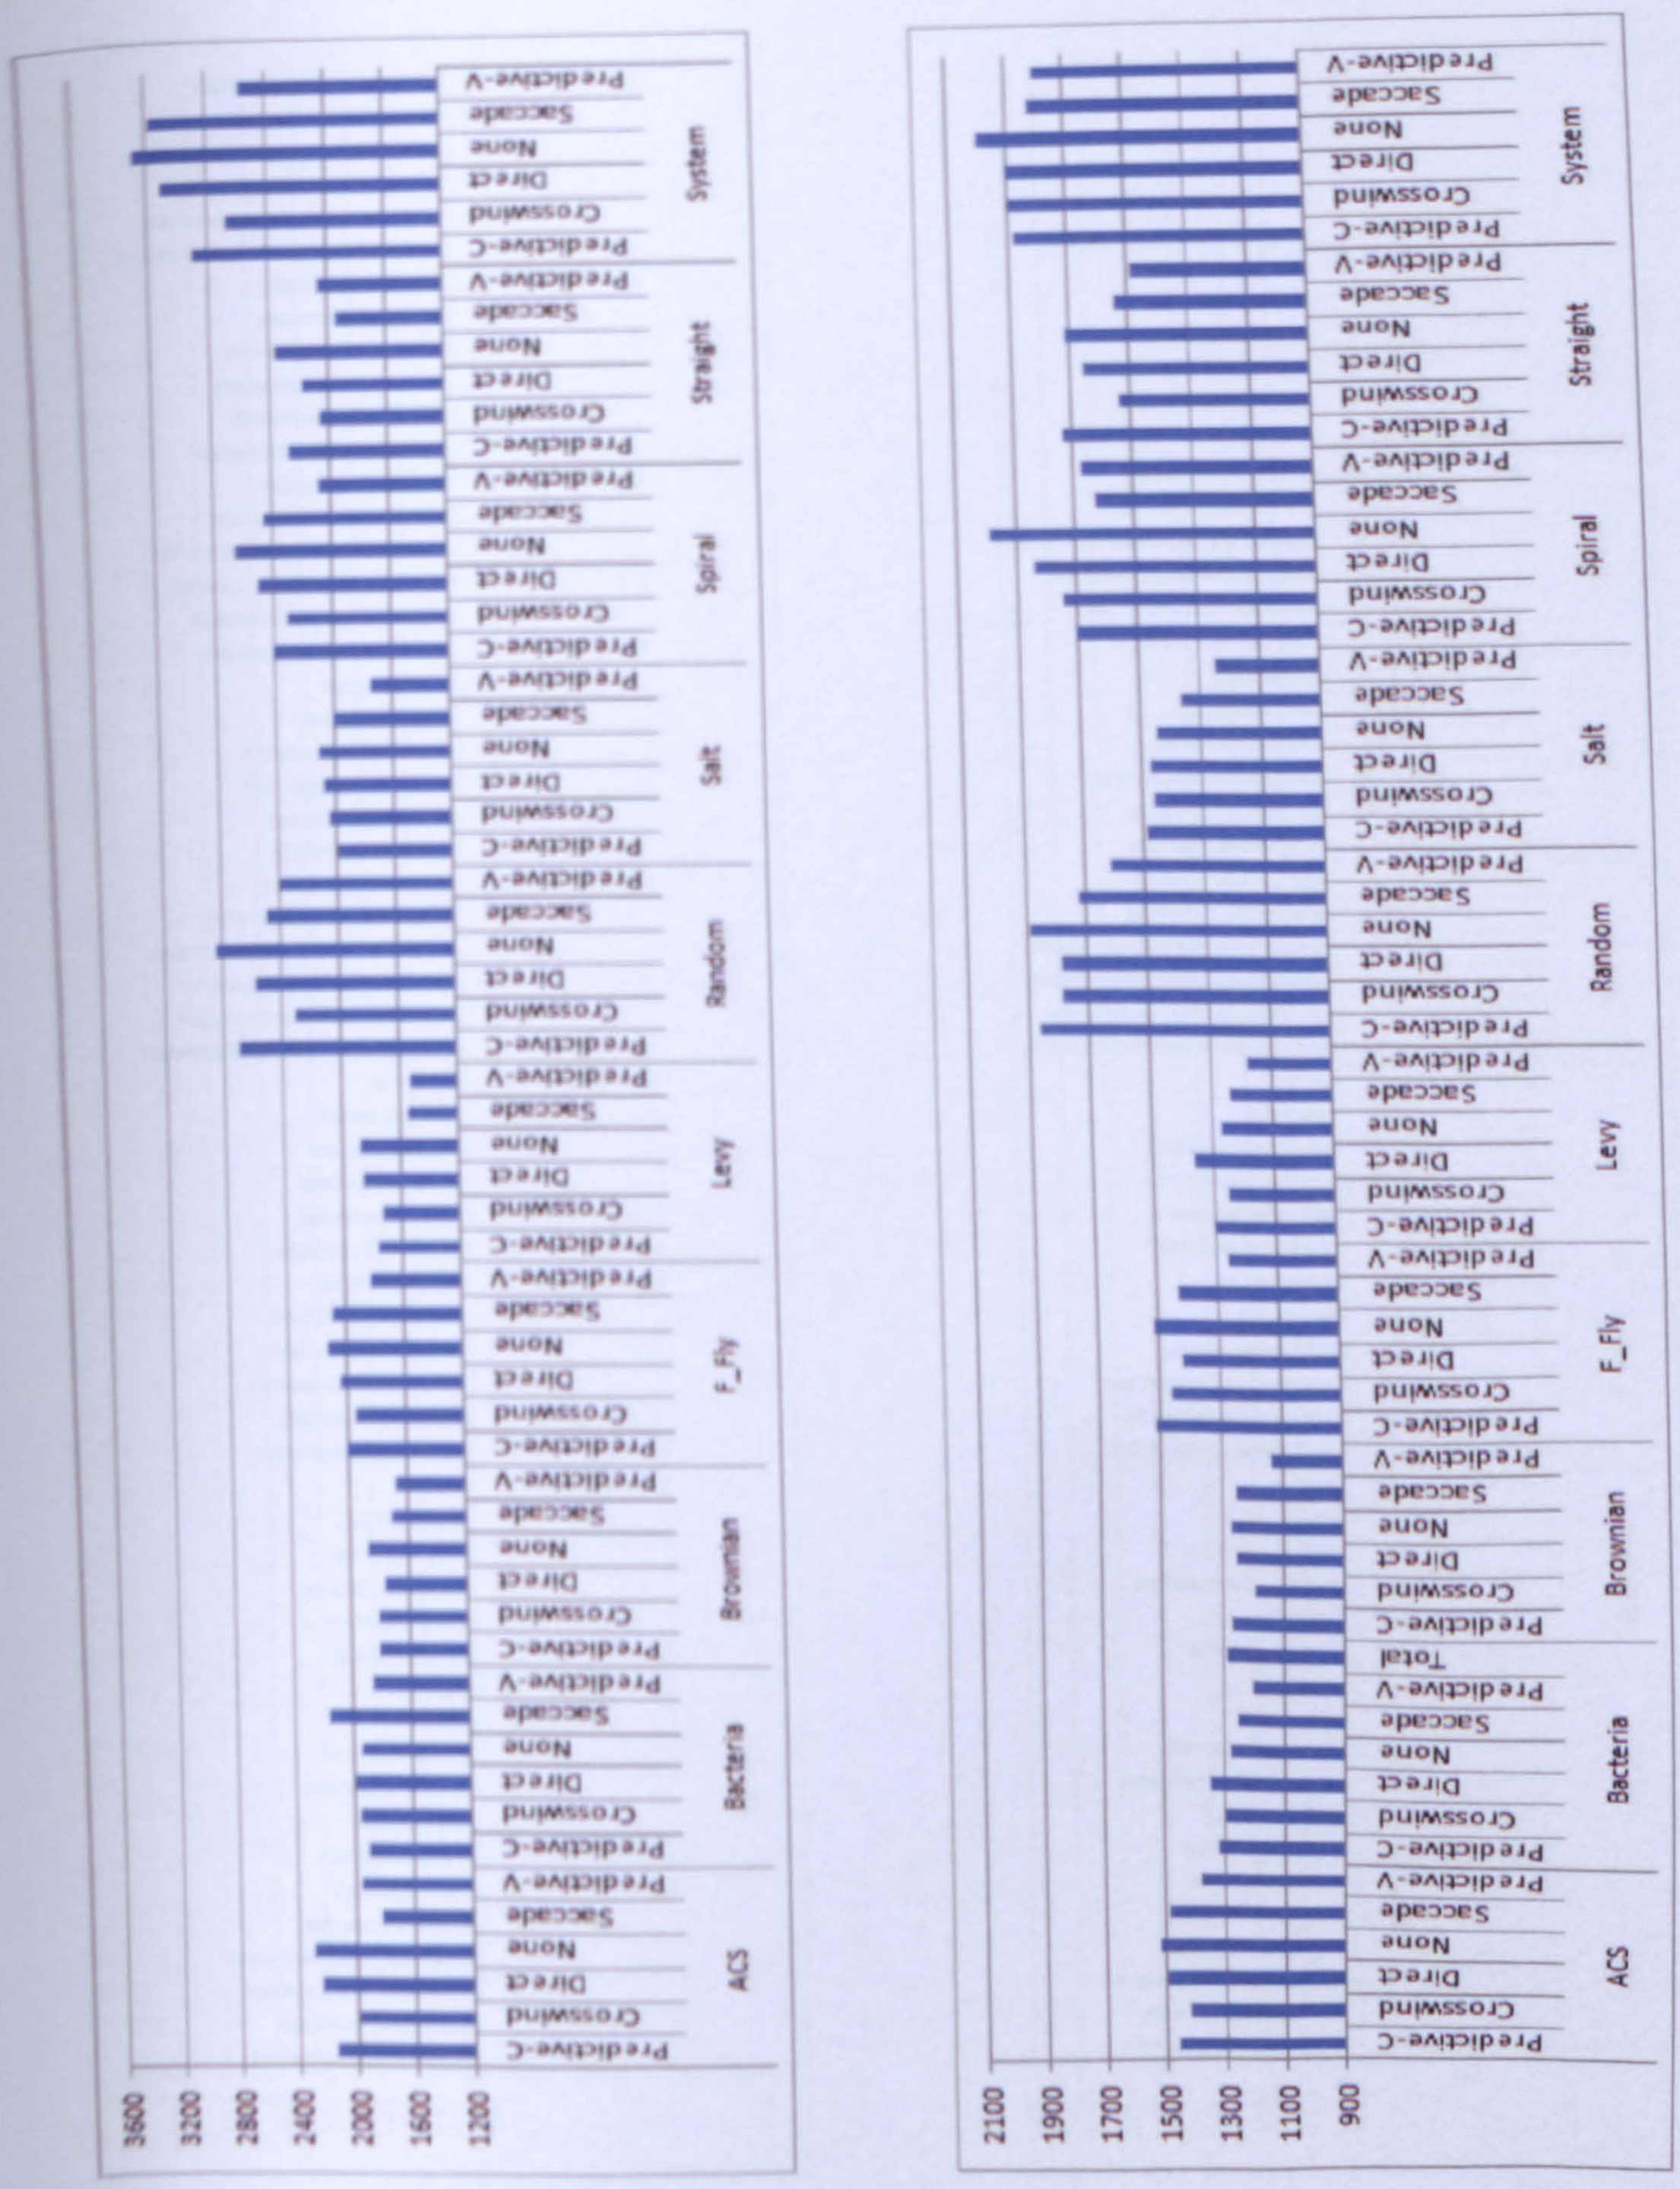

Figure 4 - 3. Mean time for 10 agents against 4 Deterministic targets travelling at 0.5 agent speed (left) and 1.0 agent speed (right).

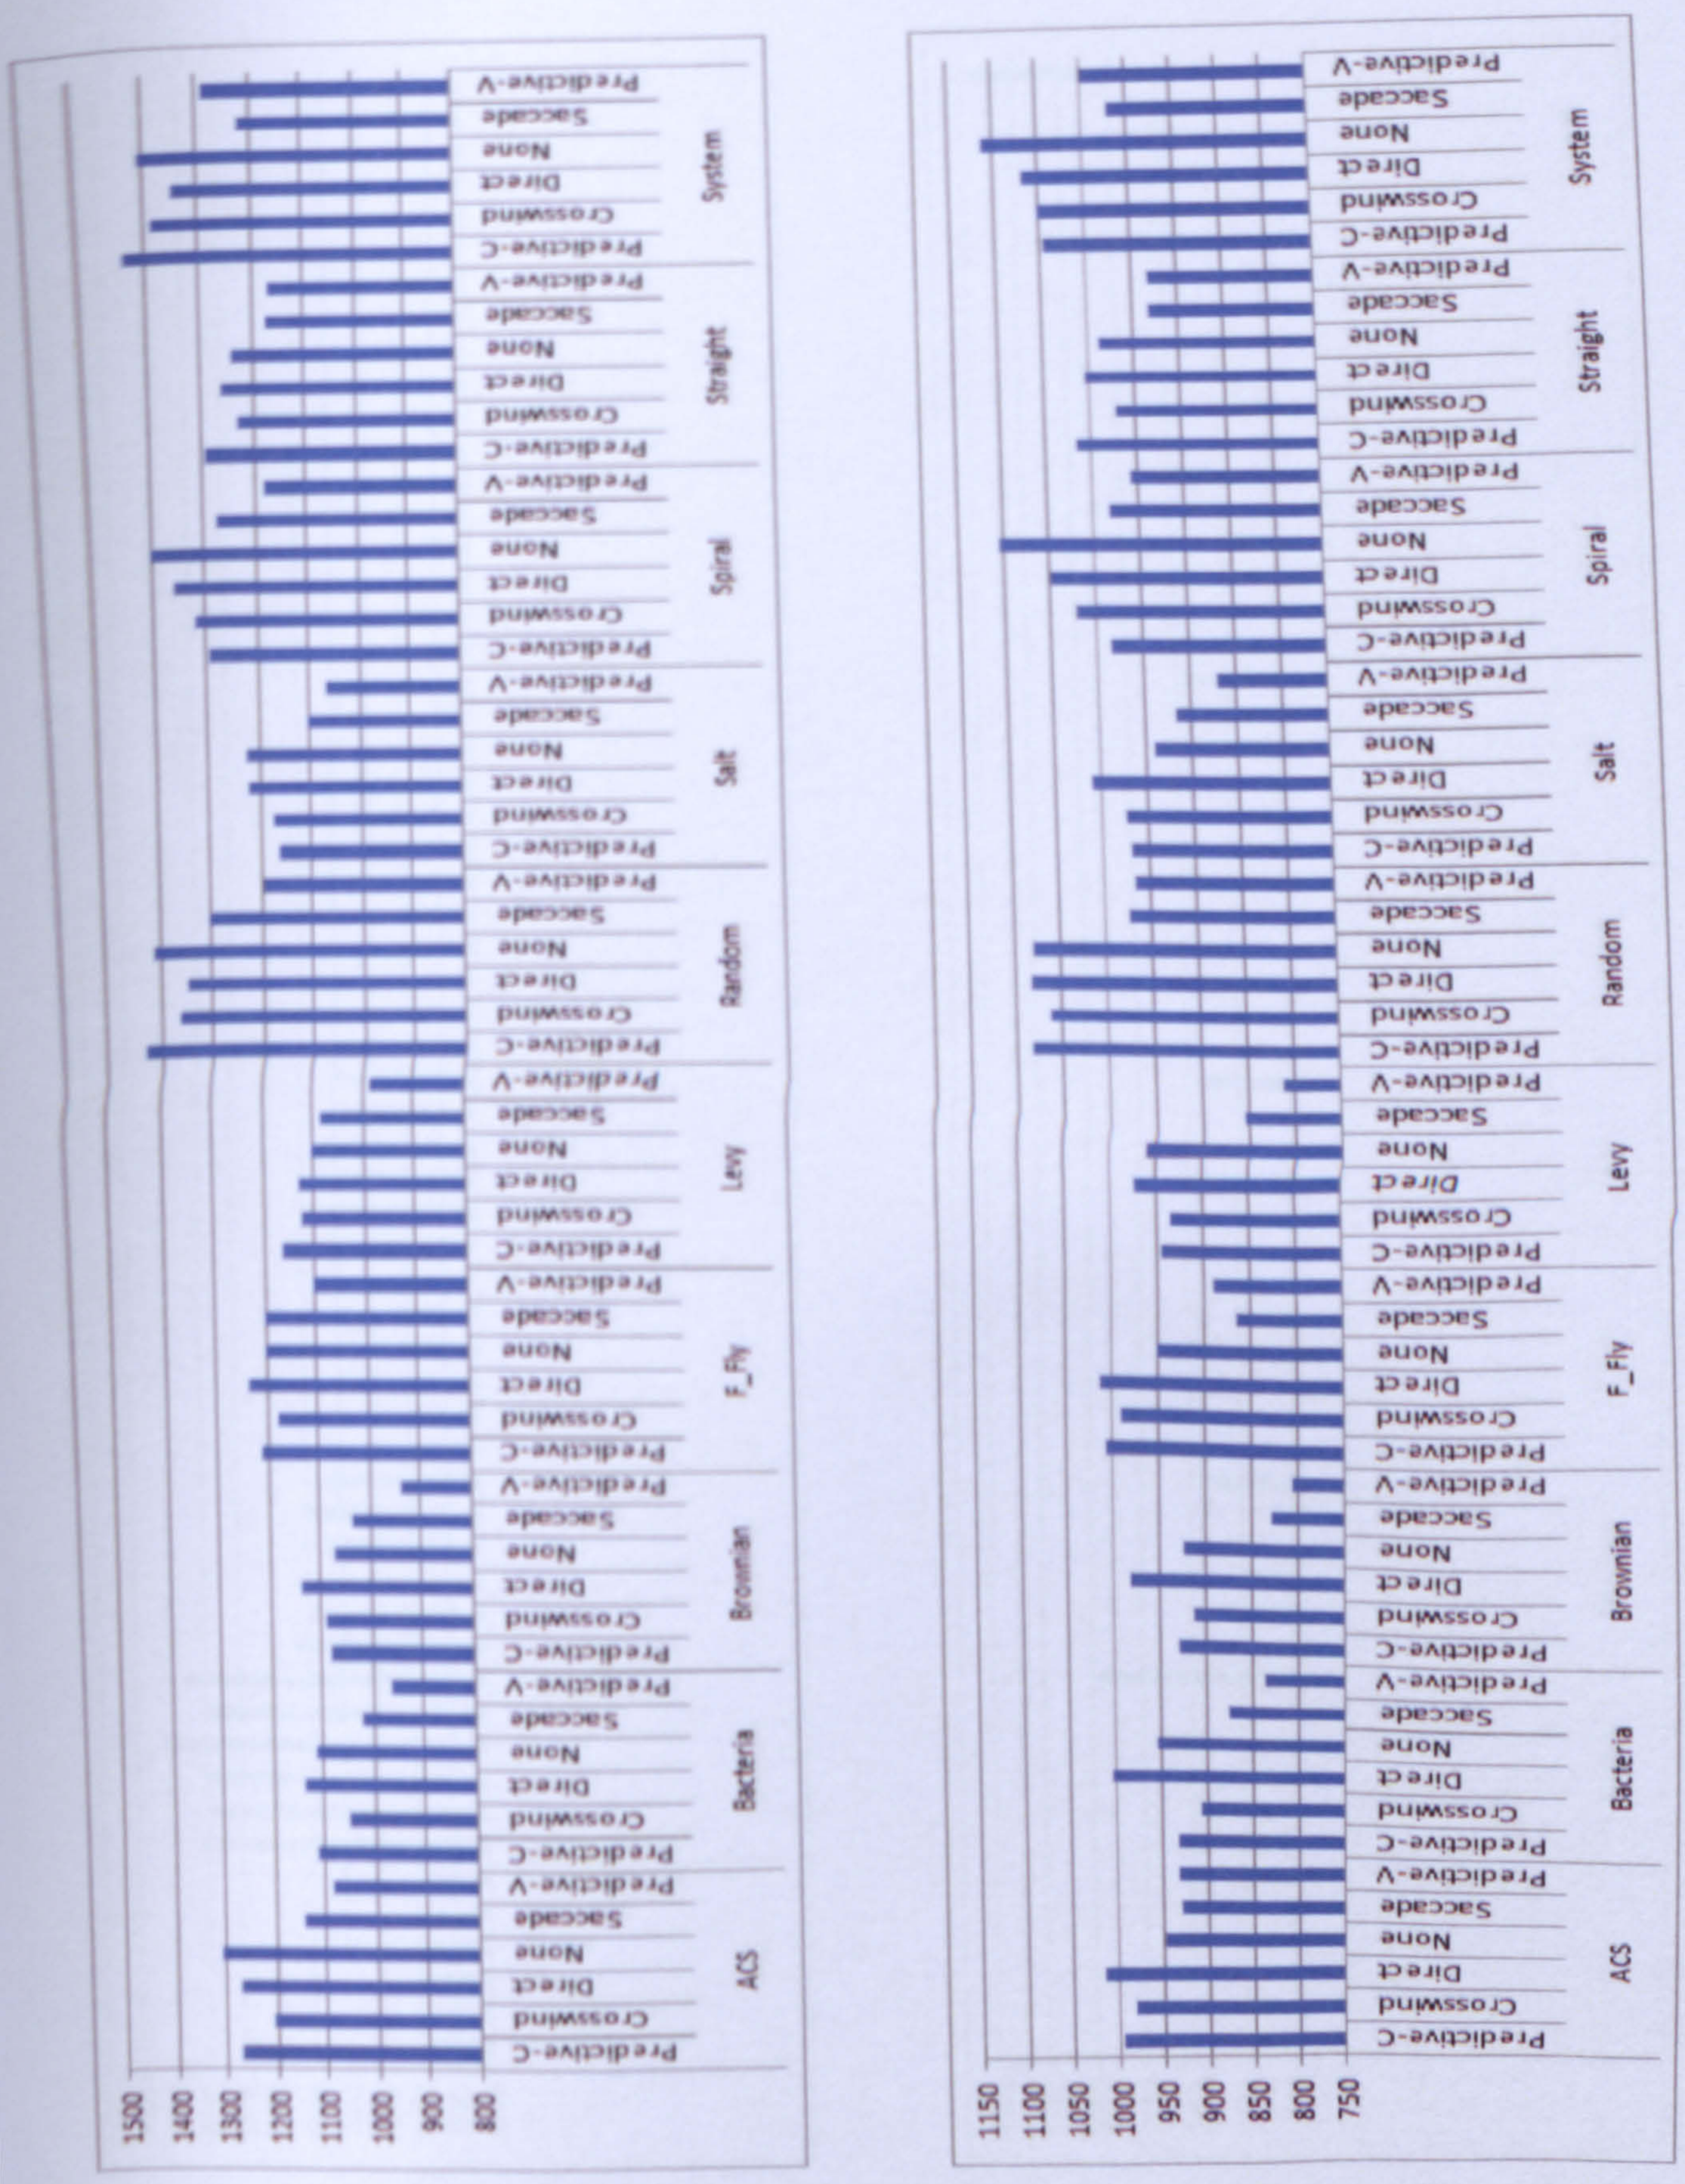

Figure 4 - 4. Mean time for 10 agents against 4 deterministic targets travelling at 1.5 agent speed (left) and 2.0 agent speed (right).

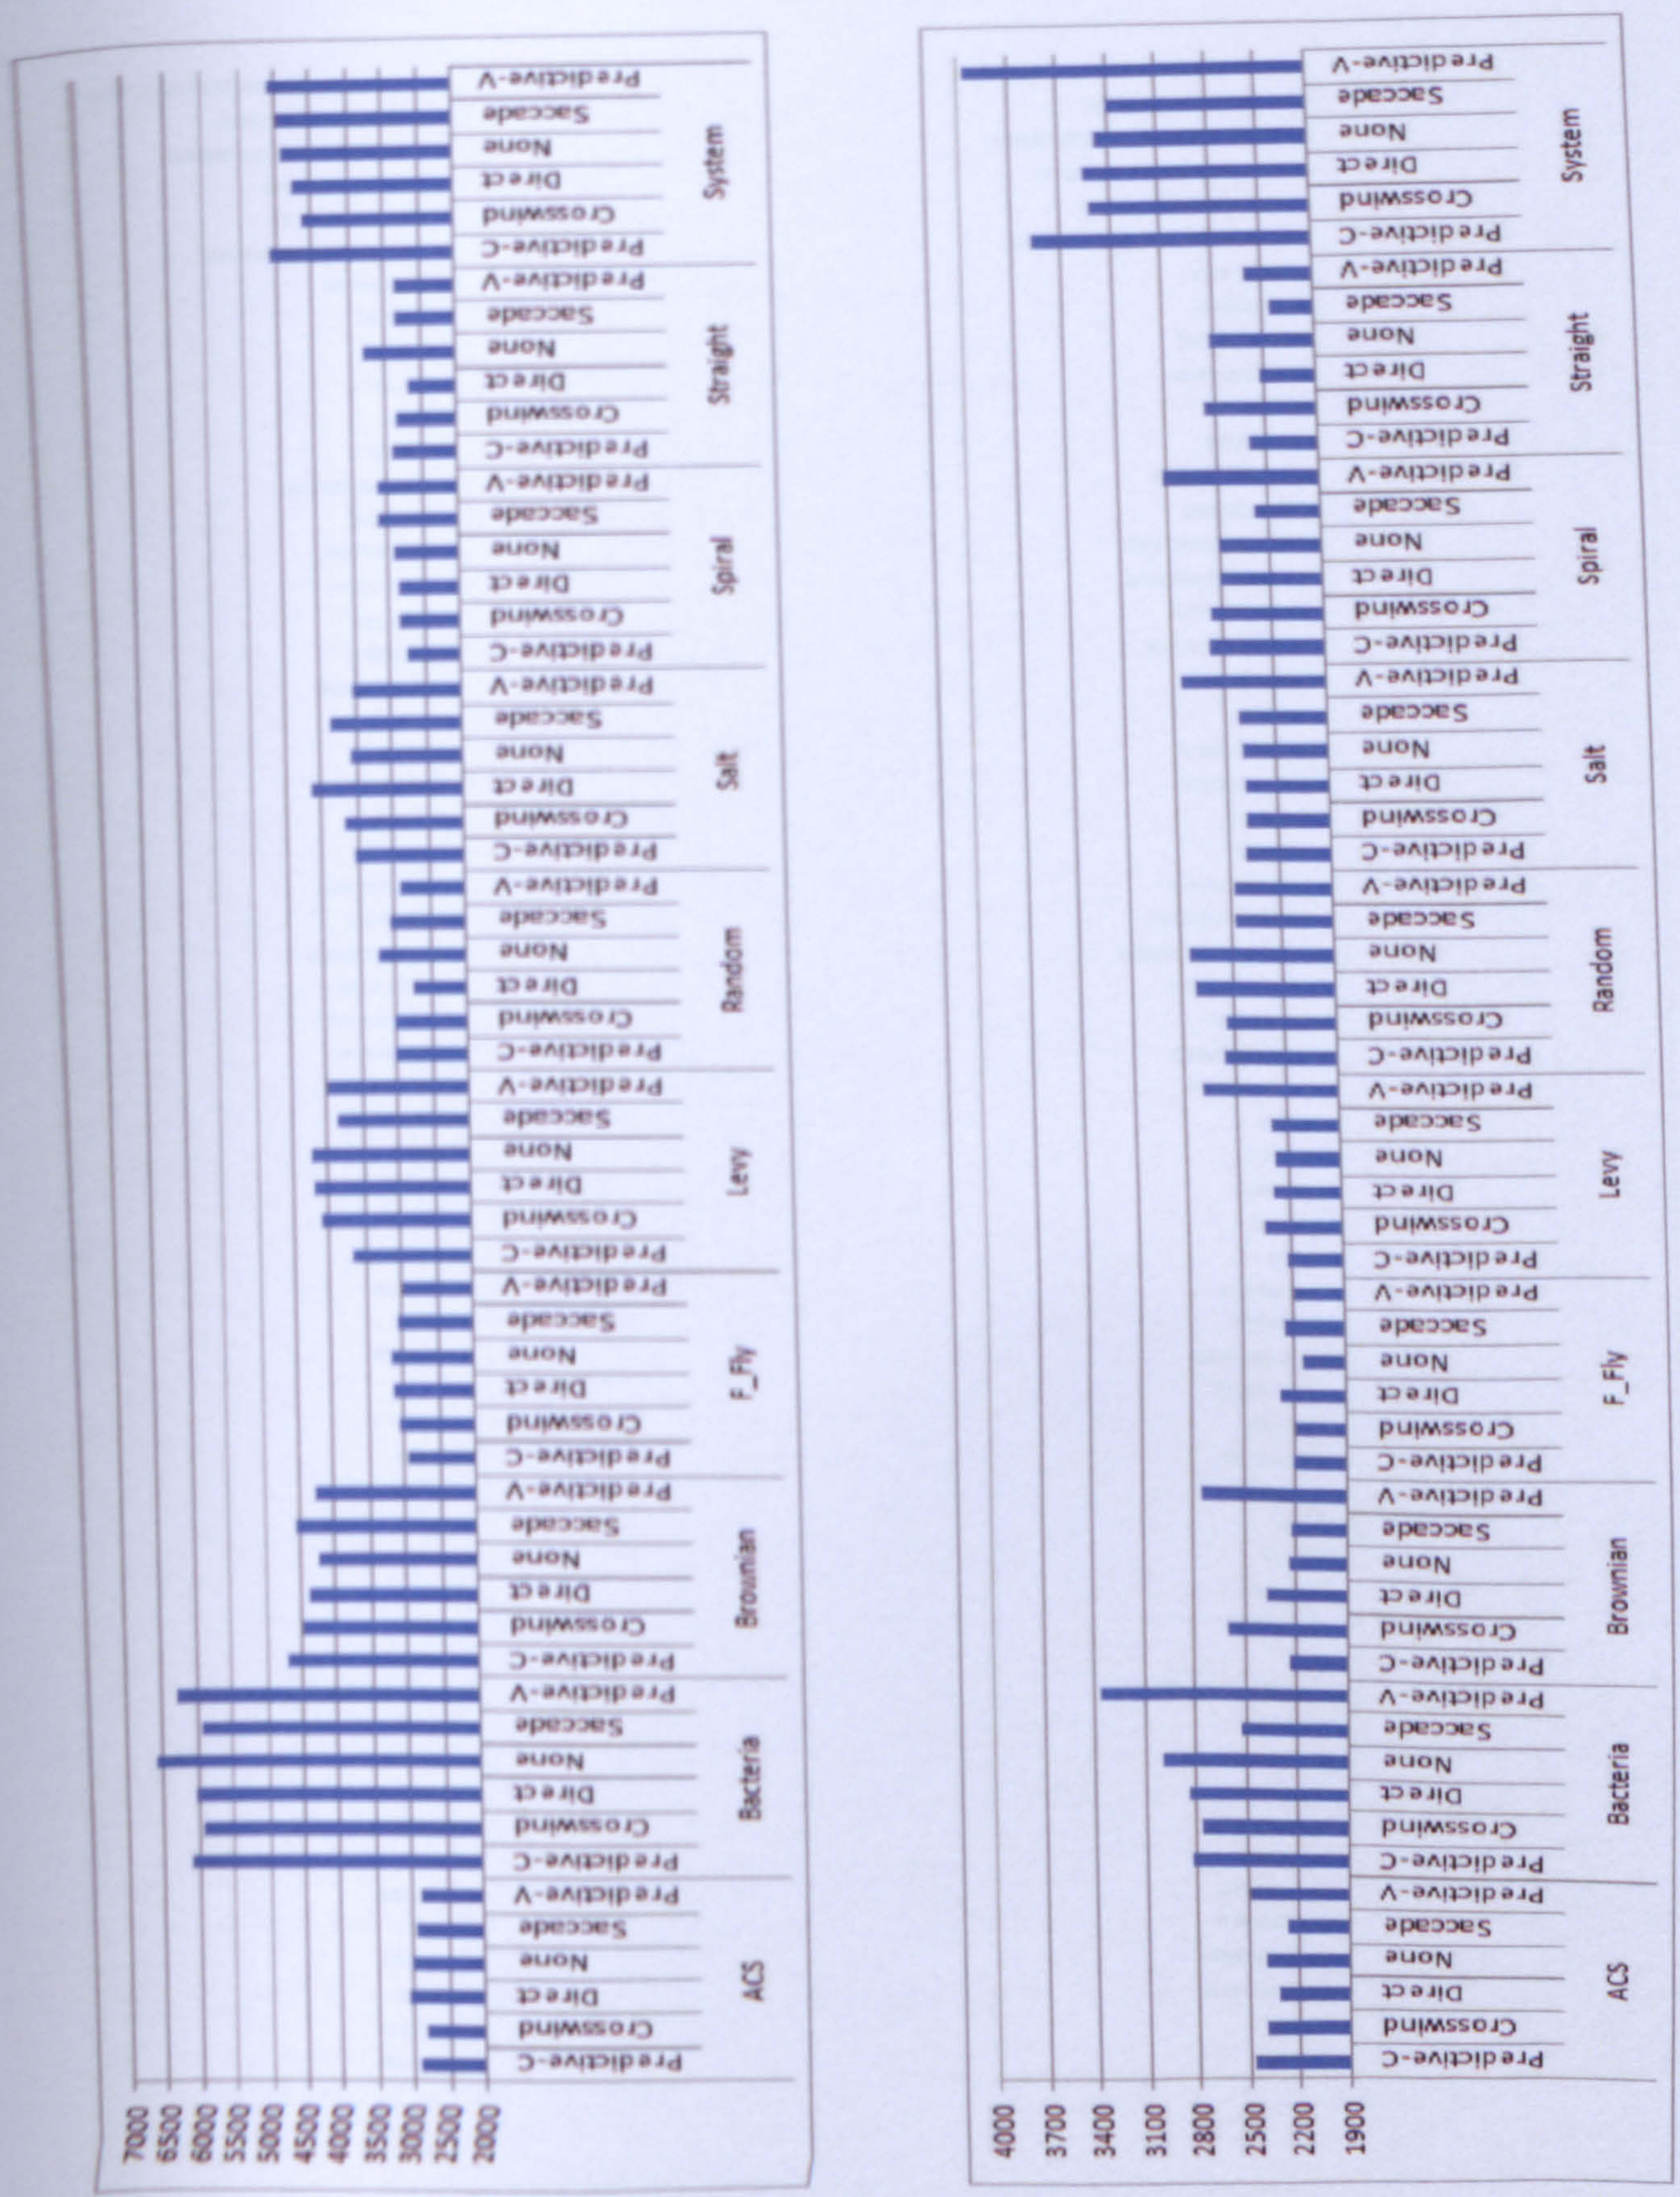

Figure 4 - 5. Mean time for 10 agents against 4 fruit fly targets travelling at *0.5 agent* speed (left) and 1.0 agent speed (right).

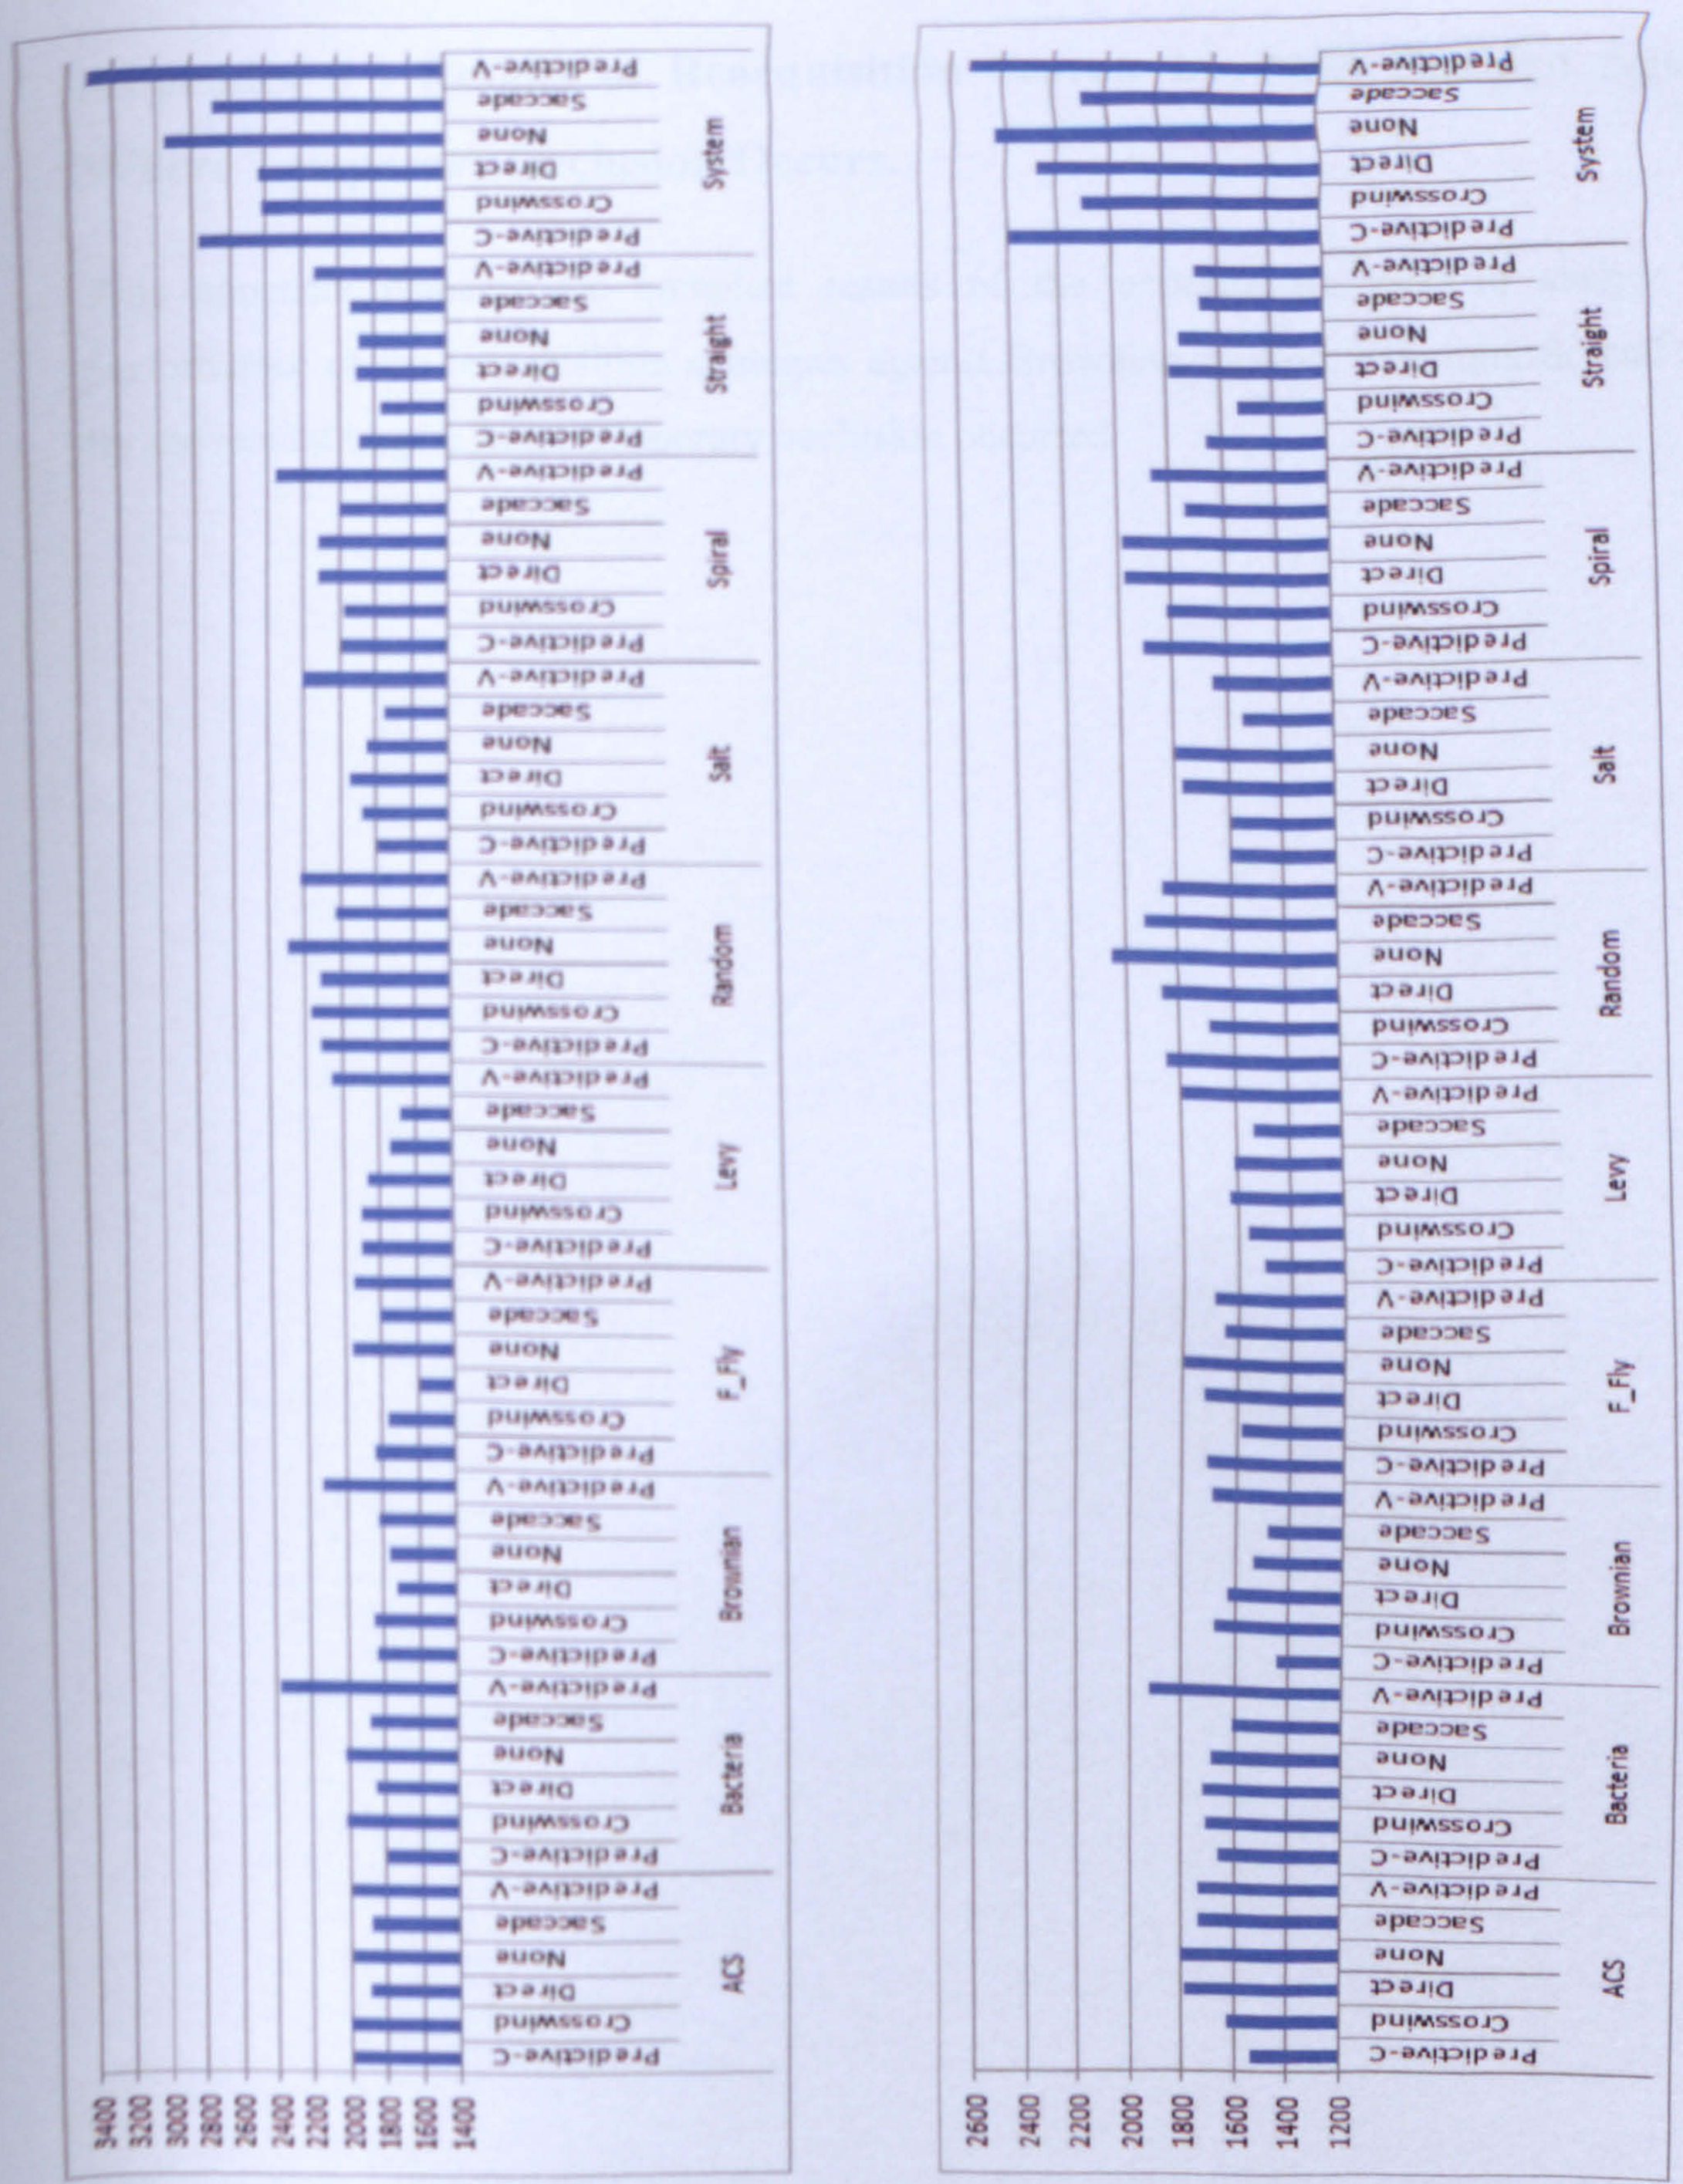

Figure 4 - 6. Mean time for 10 agents against 4 fruit fly targets travelling at 1.5 agent speed (left) and 2.0 agent speed (right).

Appendix 5 - Results of Reacquisition Search by PMS Enabled Agents Where Temporary Occlusion Occurs.

This appendix contains the complete results of the experiments used to analyse the performance of the reacquisition strategies against Brownian motion, deterministic and fruit fly movement targets where temporary occlusion occurred.

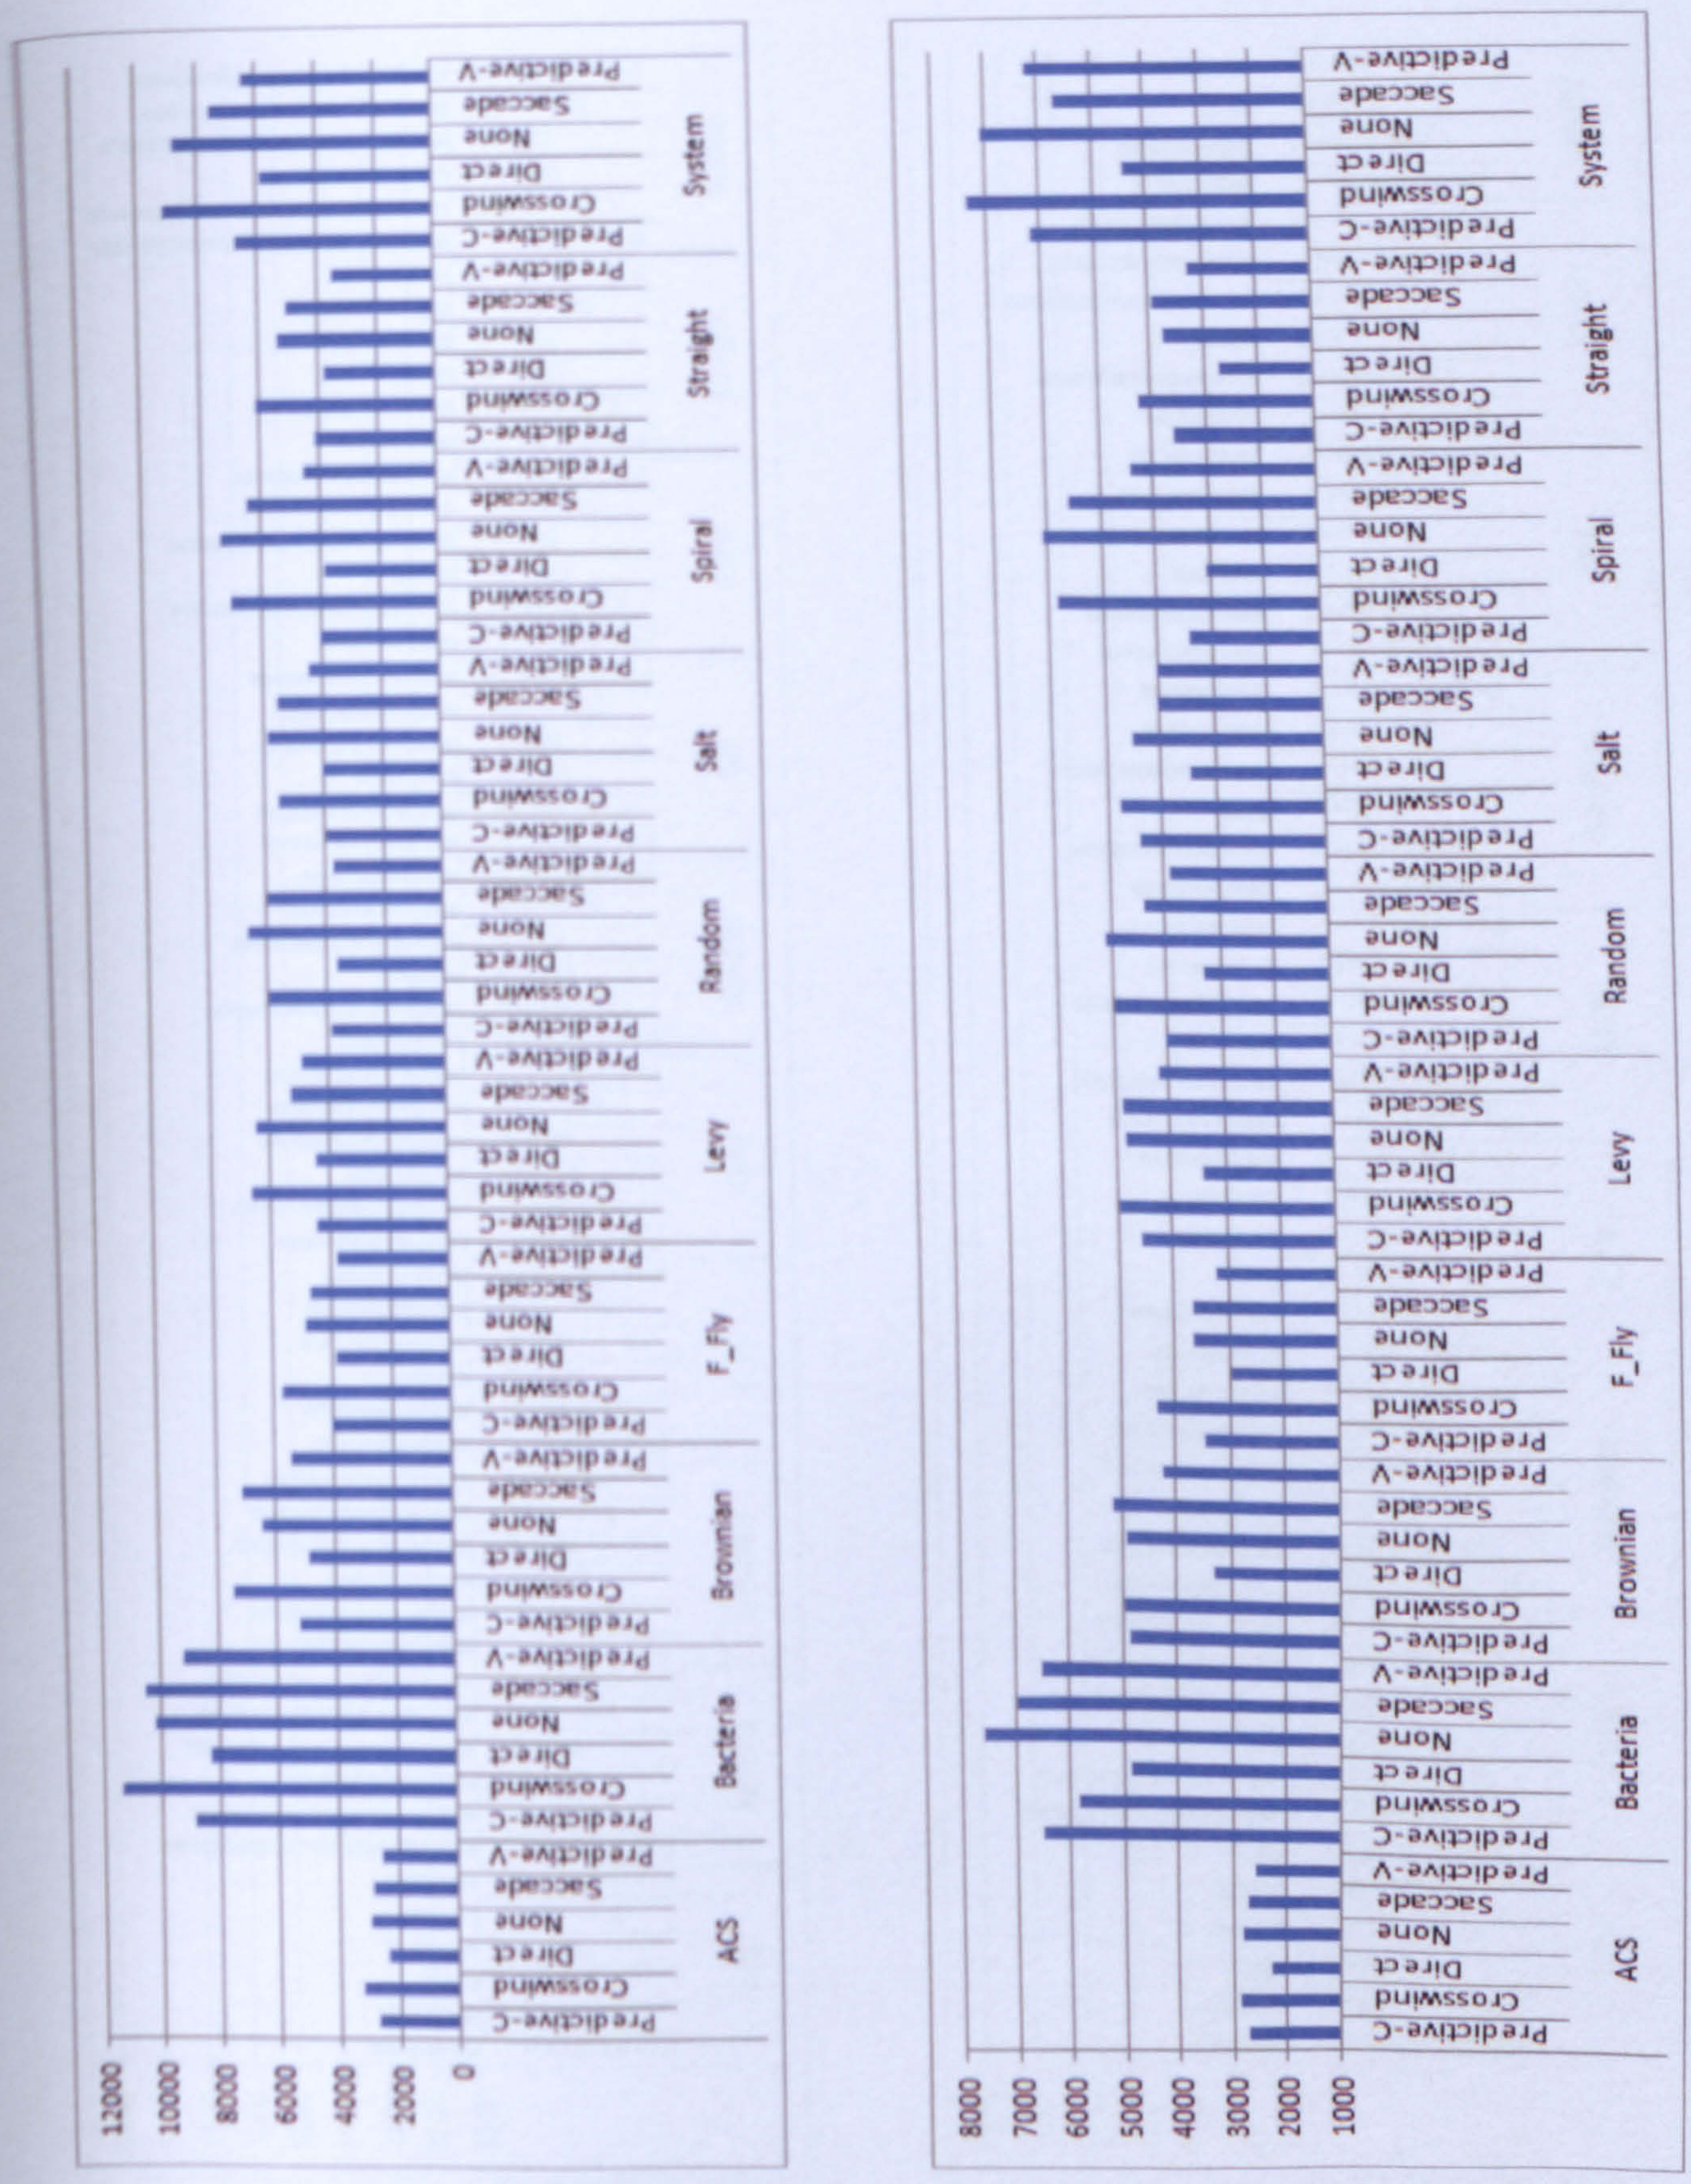

Figure 5 - 1. Mean time for 10 agents against 4 Brownian motion targets travelling at 0.5 agent speed (left) and 1.0 agent speed (right).

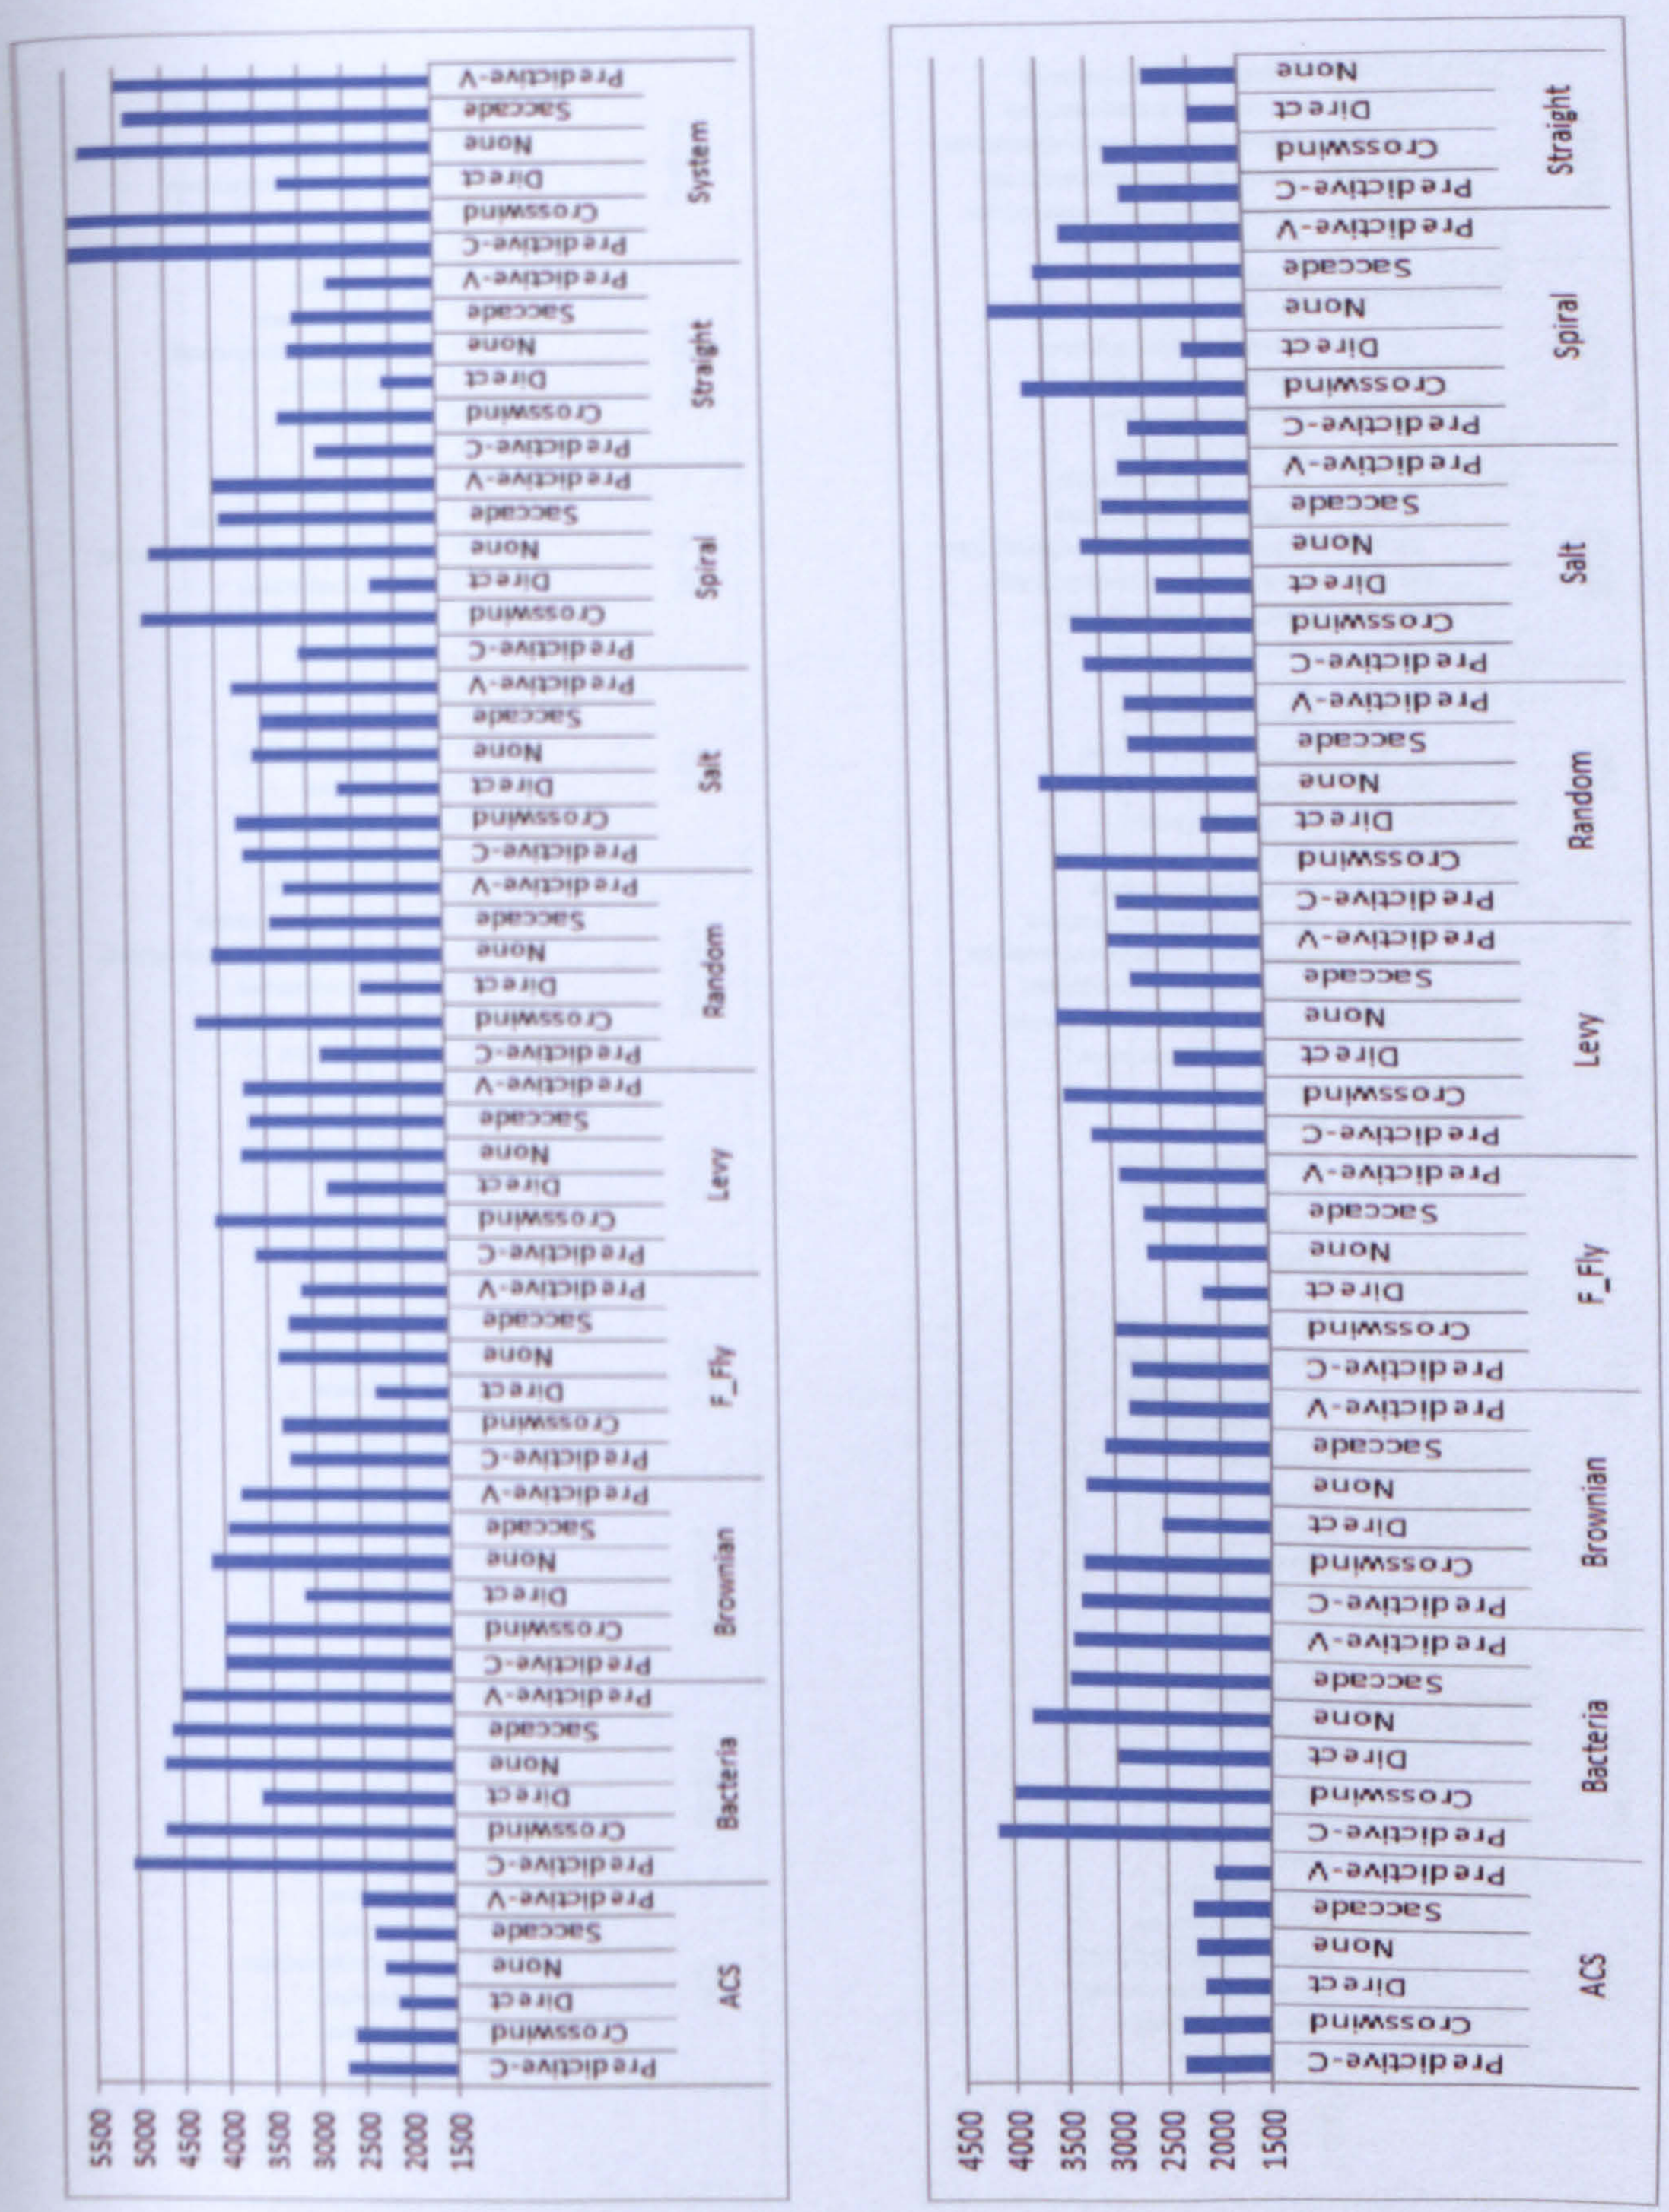

Figure 5 - 2. Mean time for 10 agents against 4 Brownian motion targets travelling at 1.5 agent speed (left) and 2.0 agent speed (right).

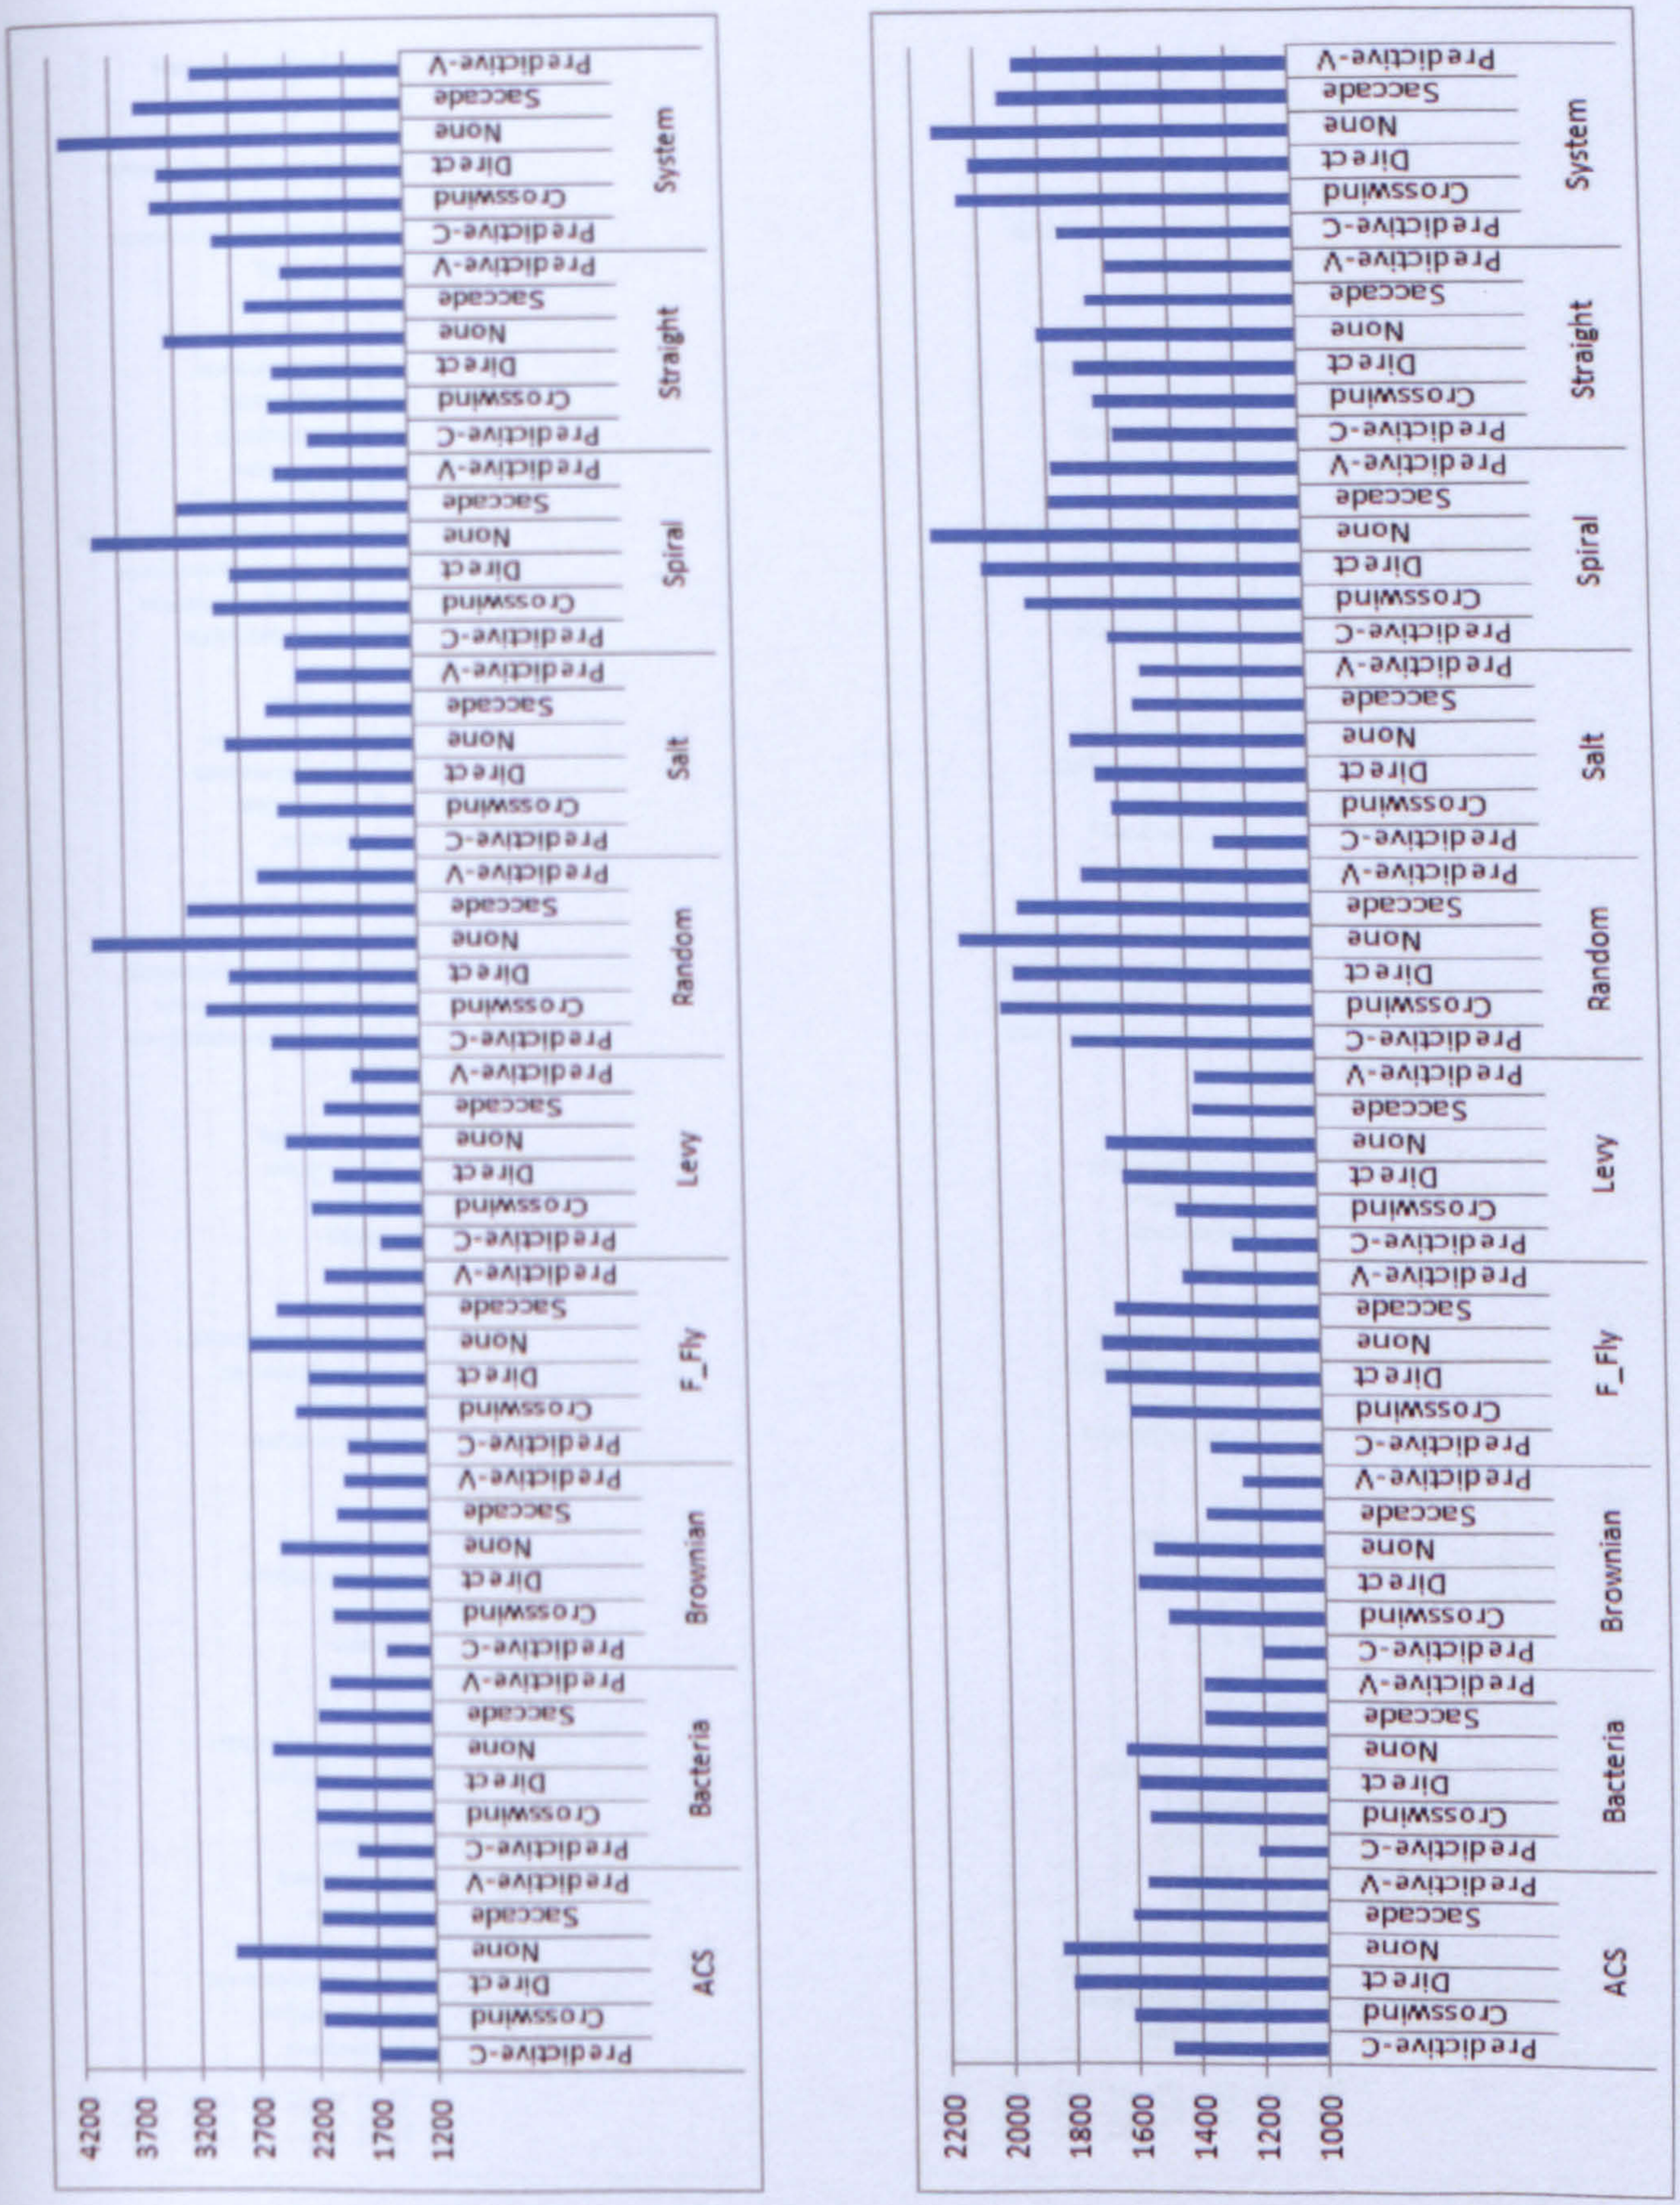

Figure 5 - 3. Mean time for 10 agents against 4 deterministic targets travelling at 0.5 agent speed (left) and 1.0 agent speed (right).

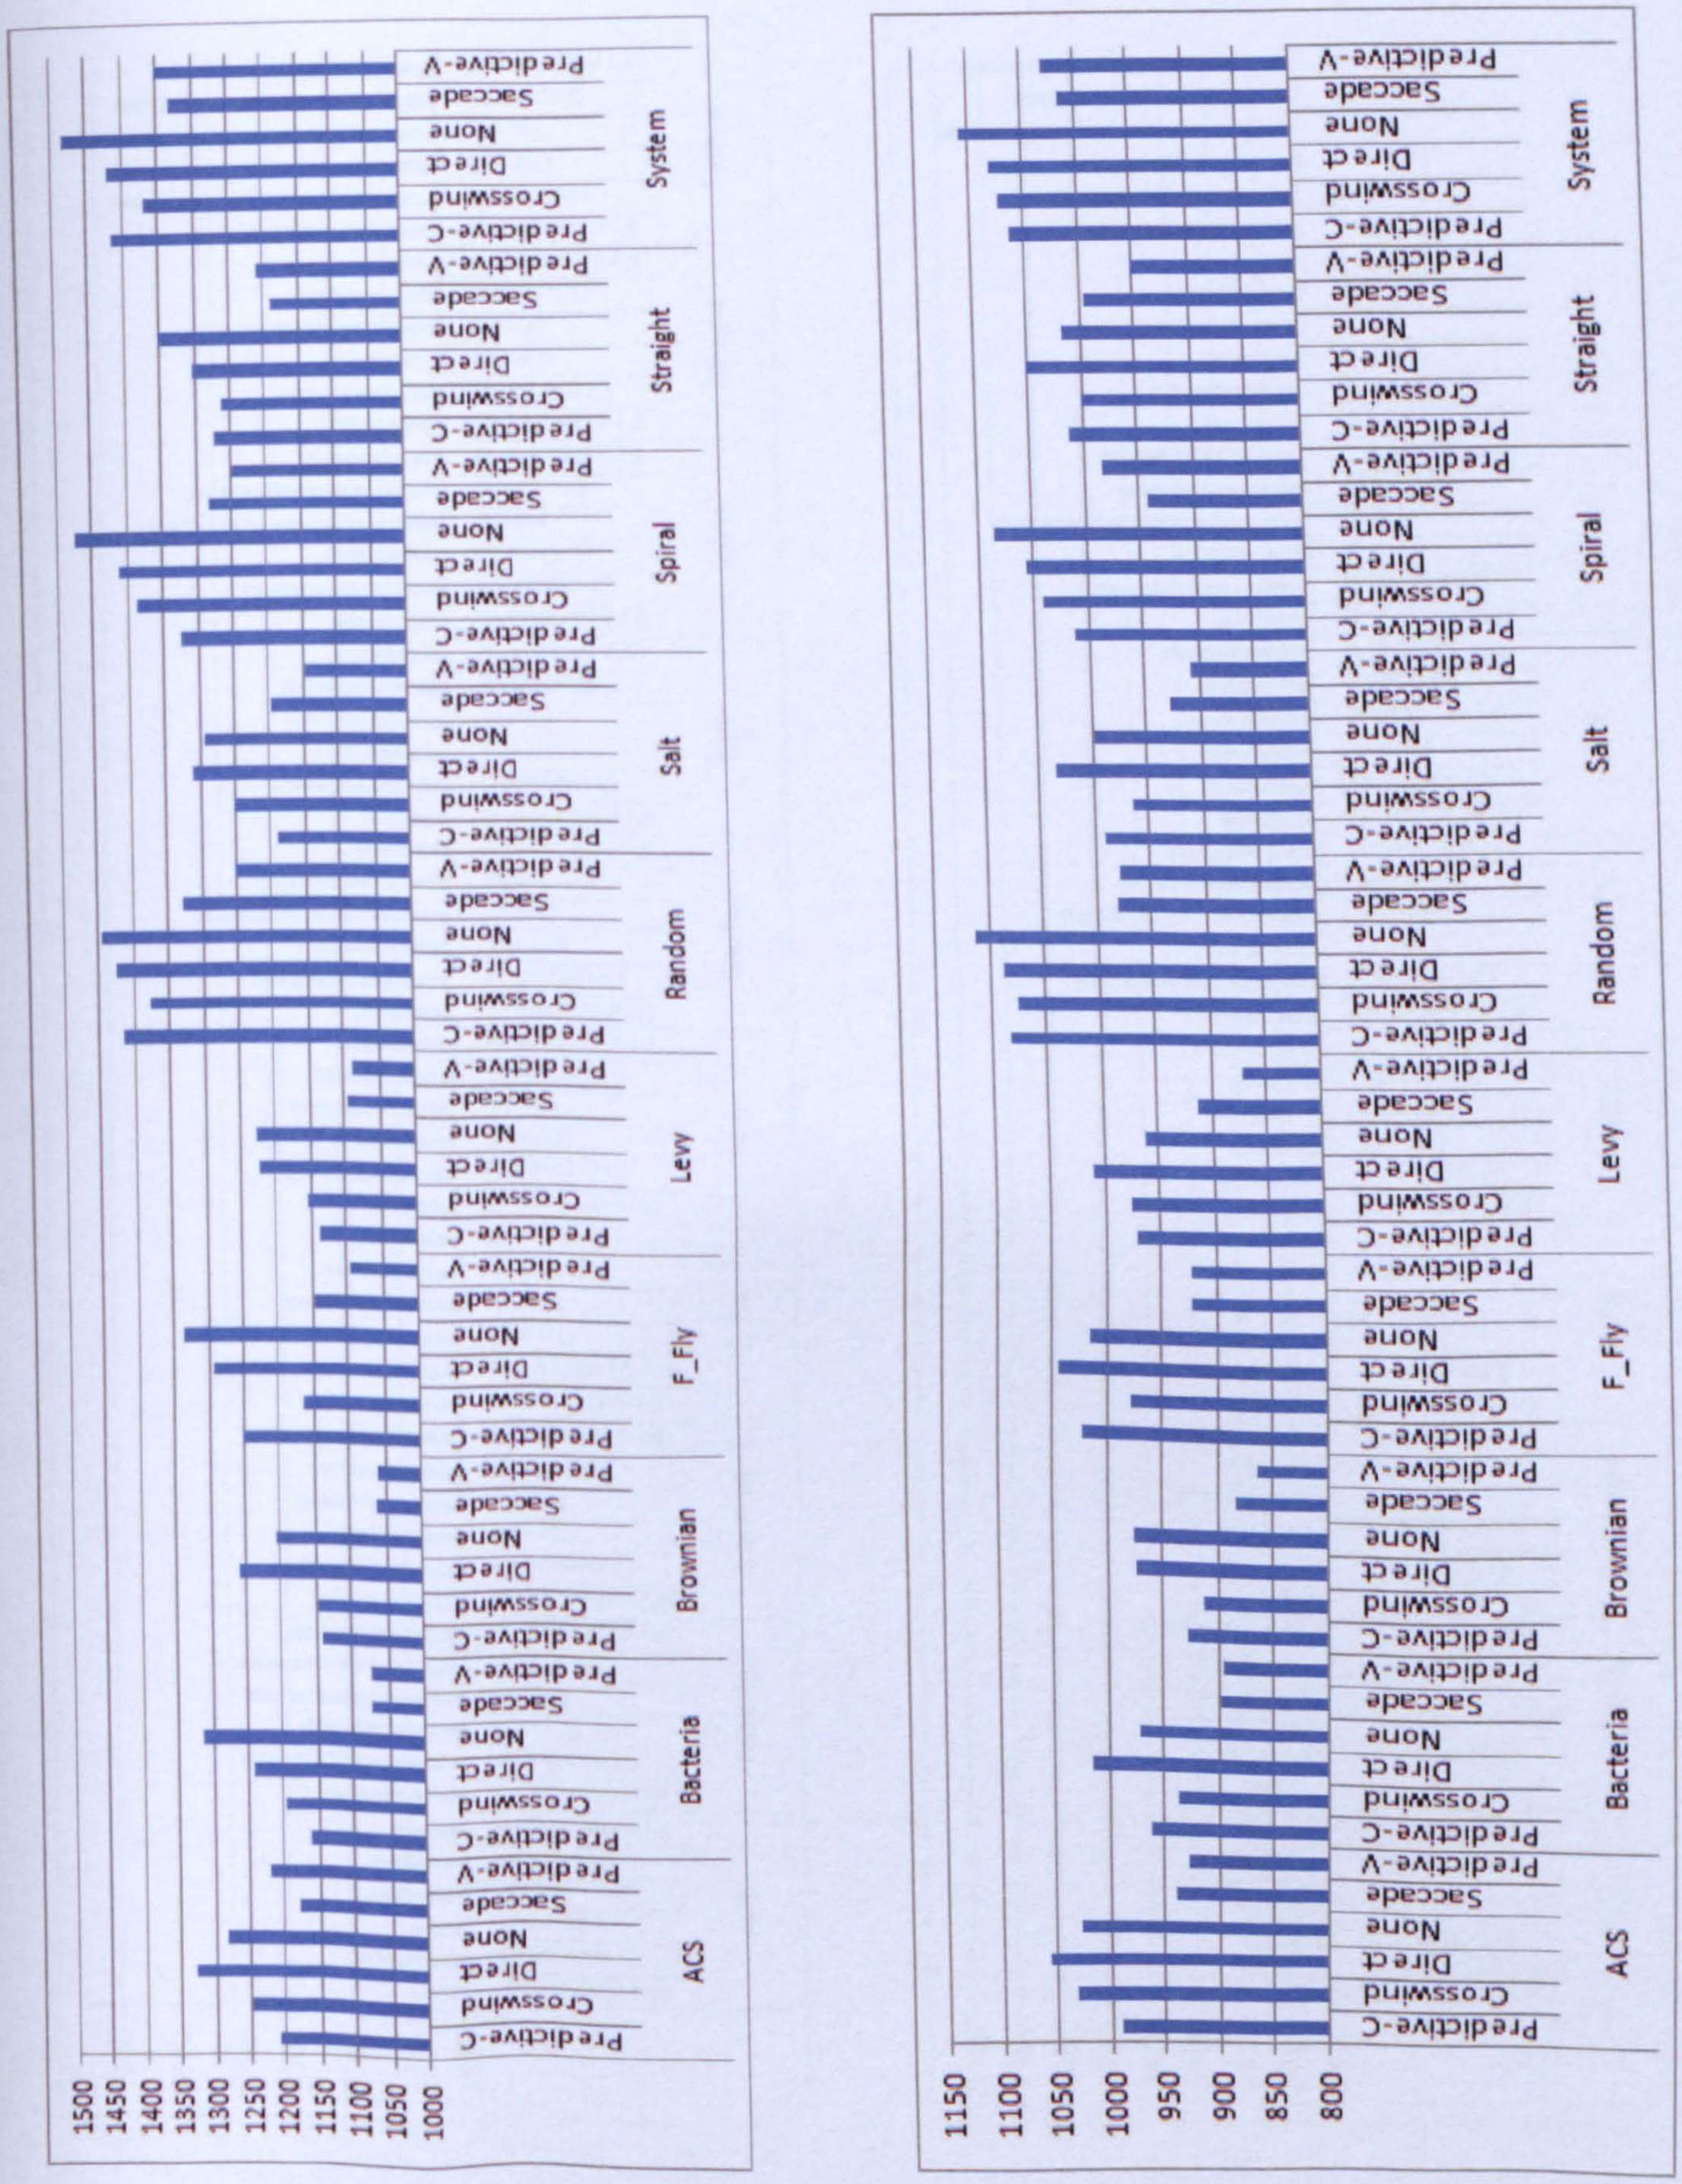

Figure 5 - 4. Mean time for 10 agents against 4 deterministic targets travelling at 1.5 agent speed (left) and 2.0 agent speed (right).

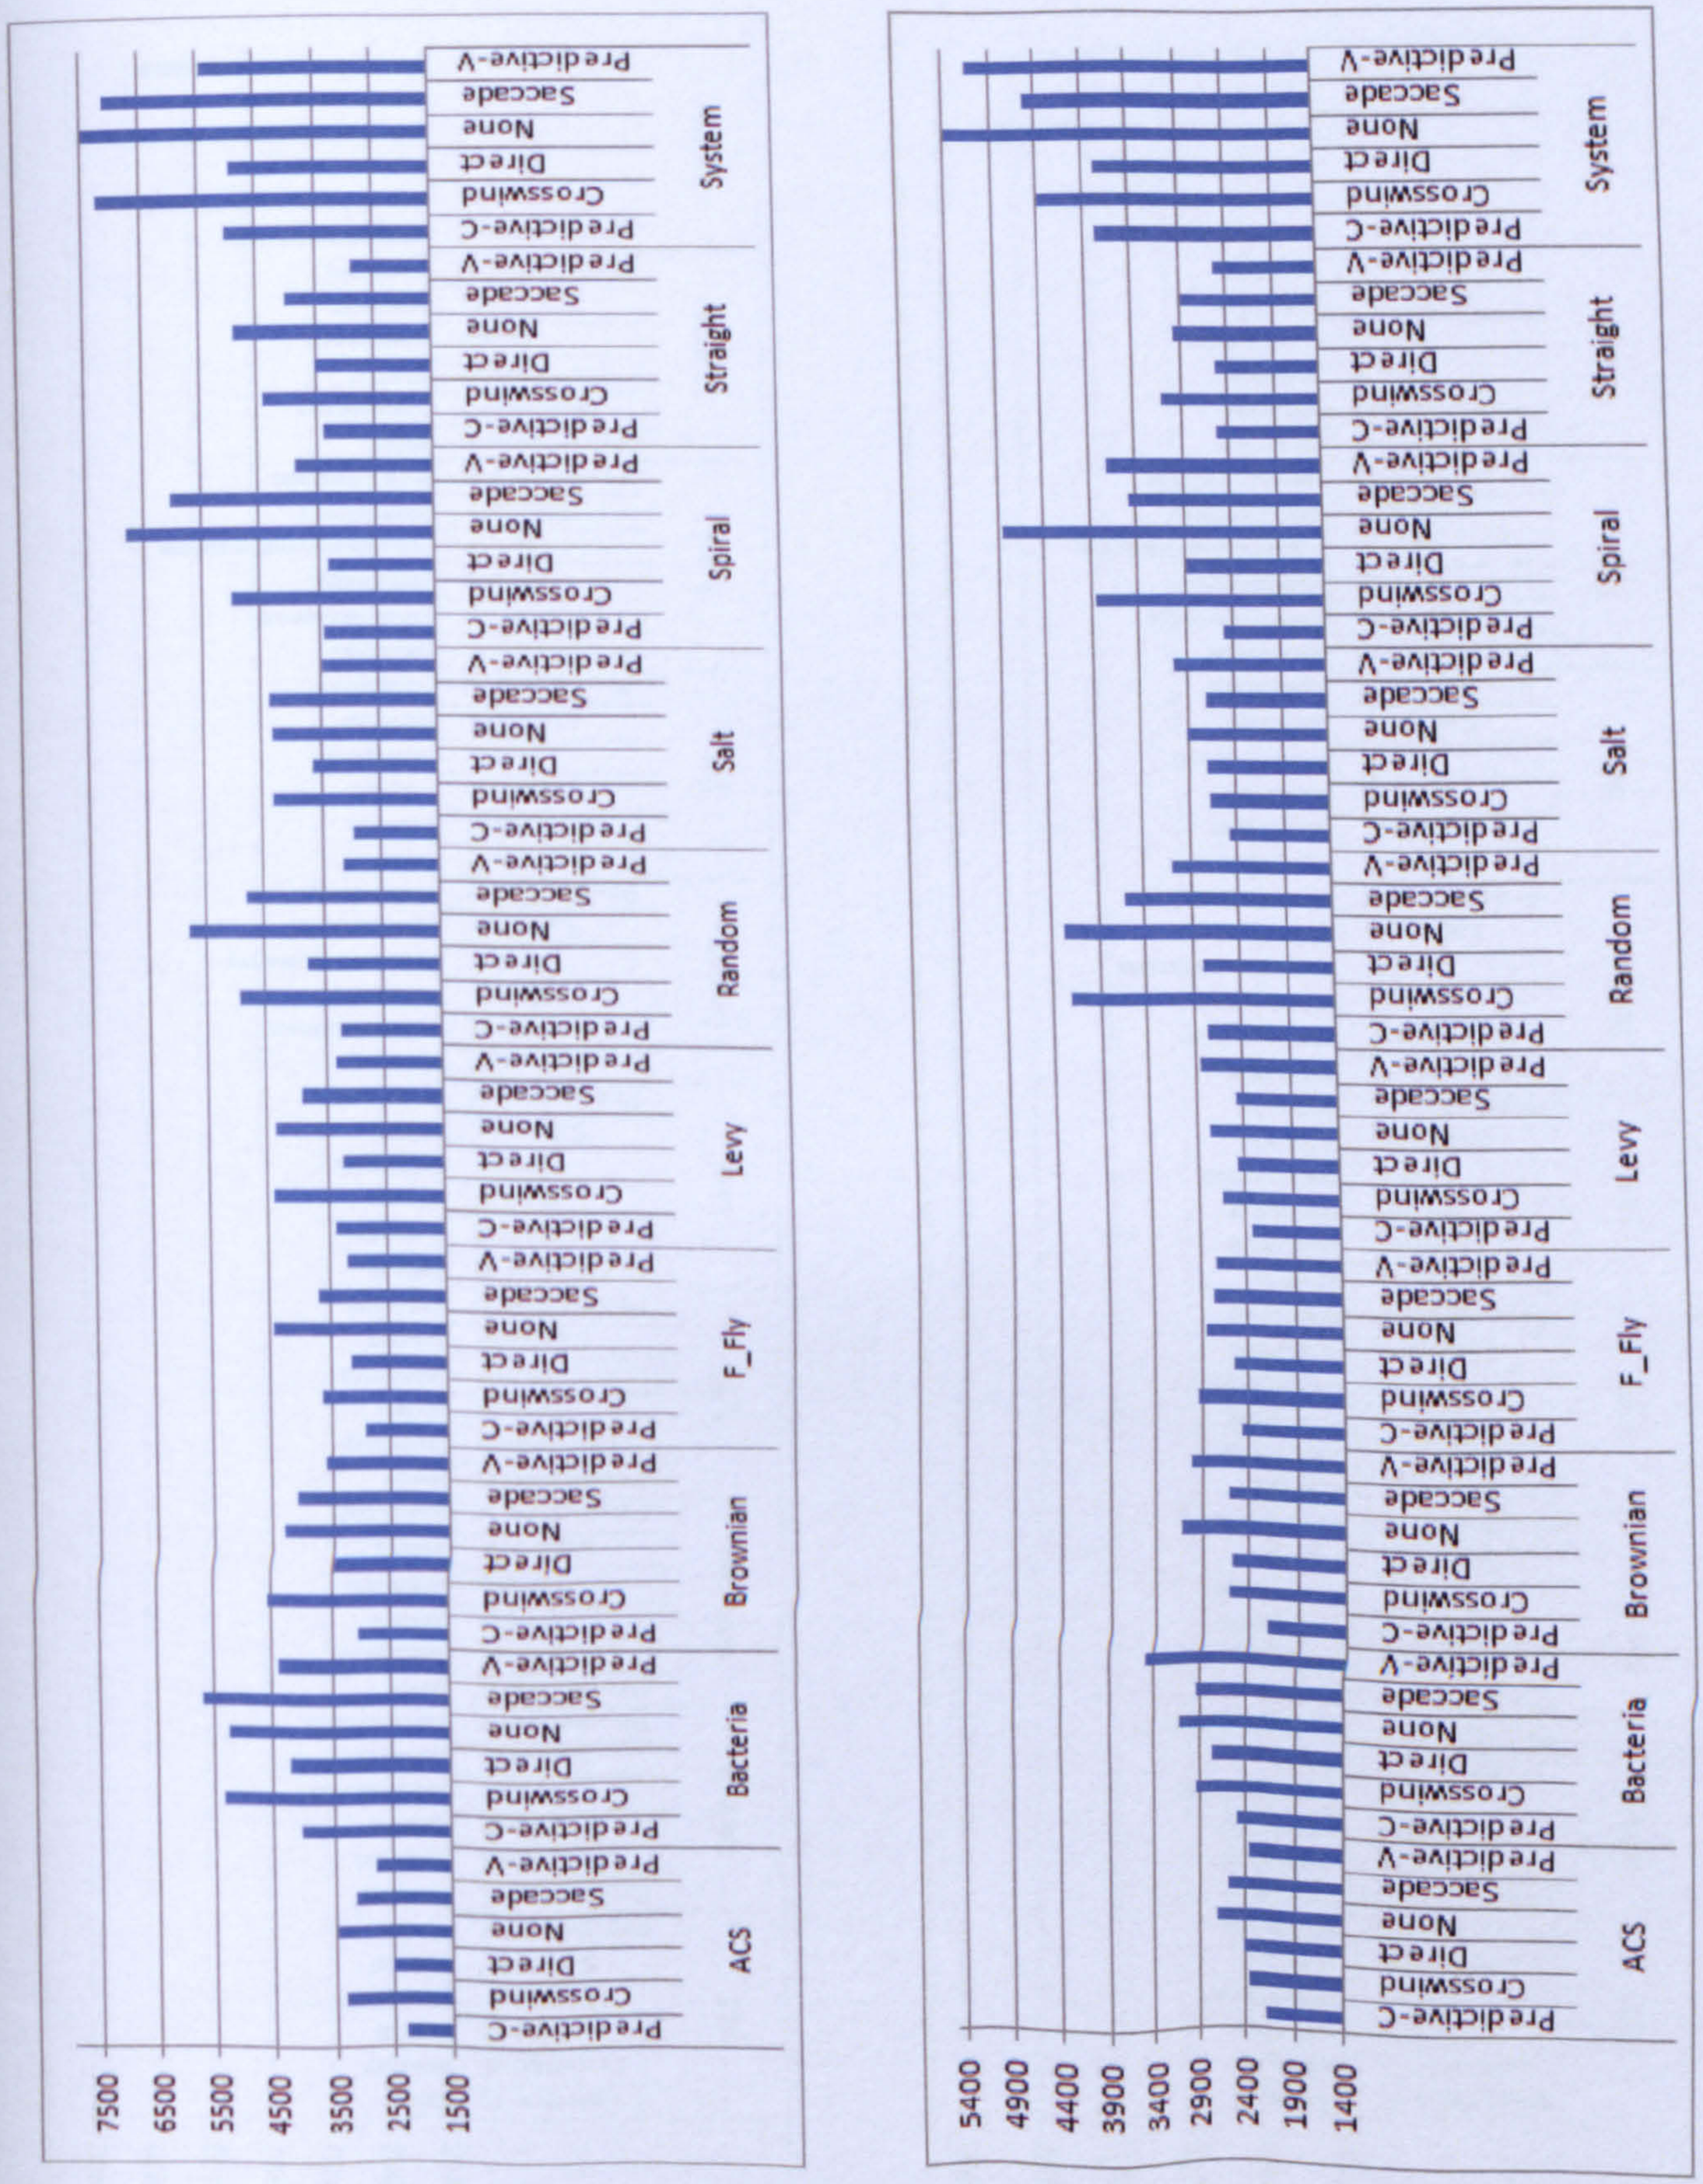

Figure 5 - 5. Mean time for 10 agents against 4 fruit fly targets travelling at 0.5 agent speed (left) and 1.0 agent speed (right).

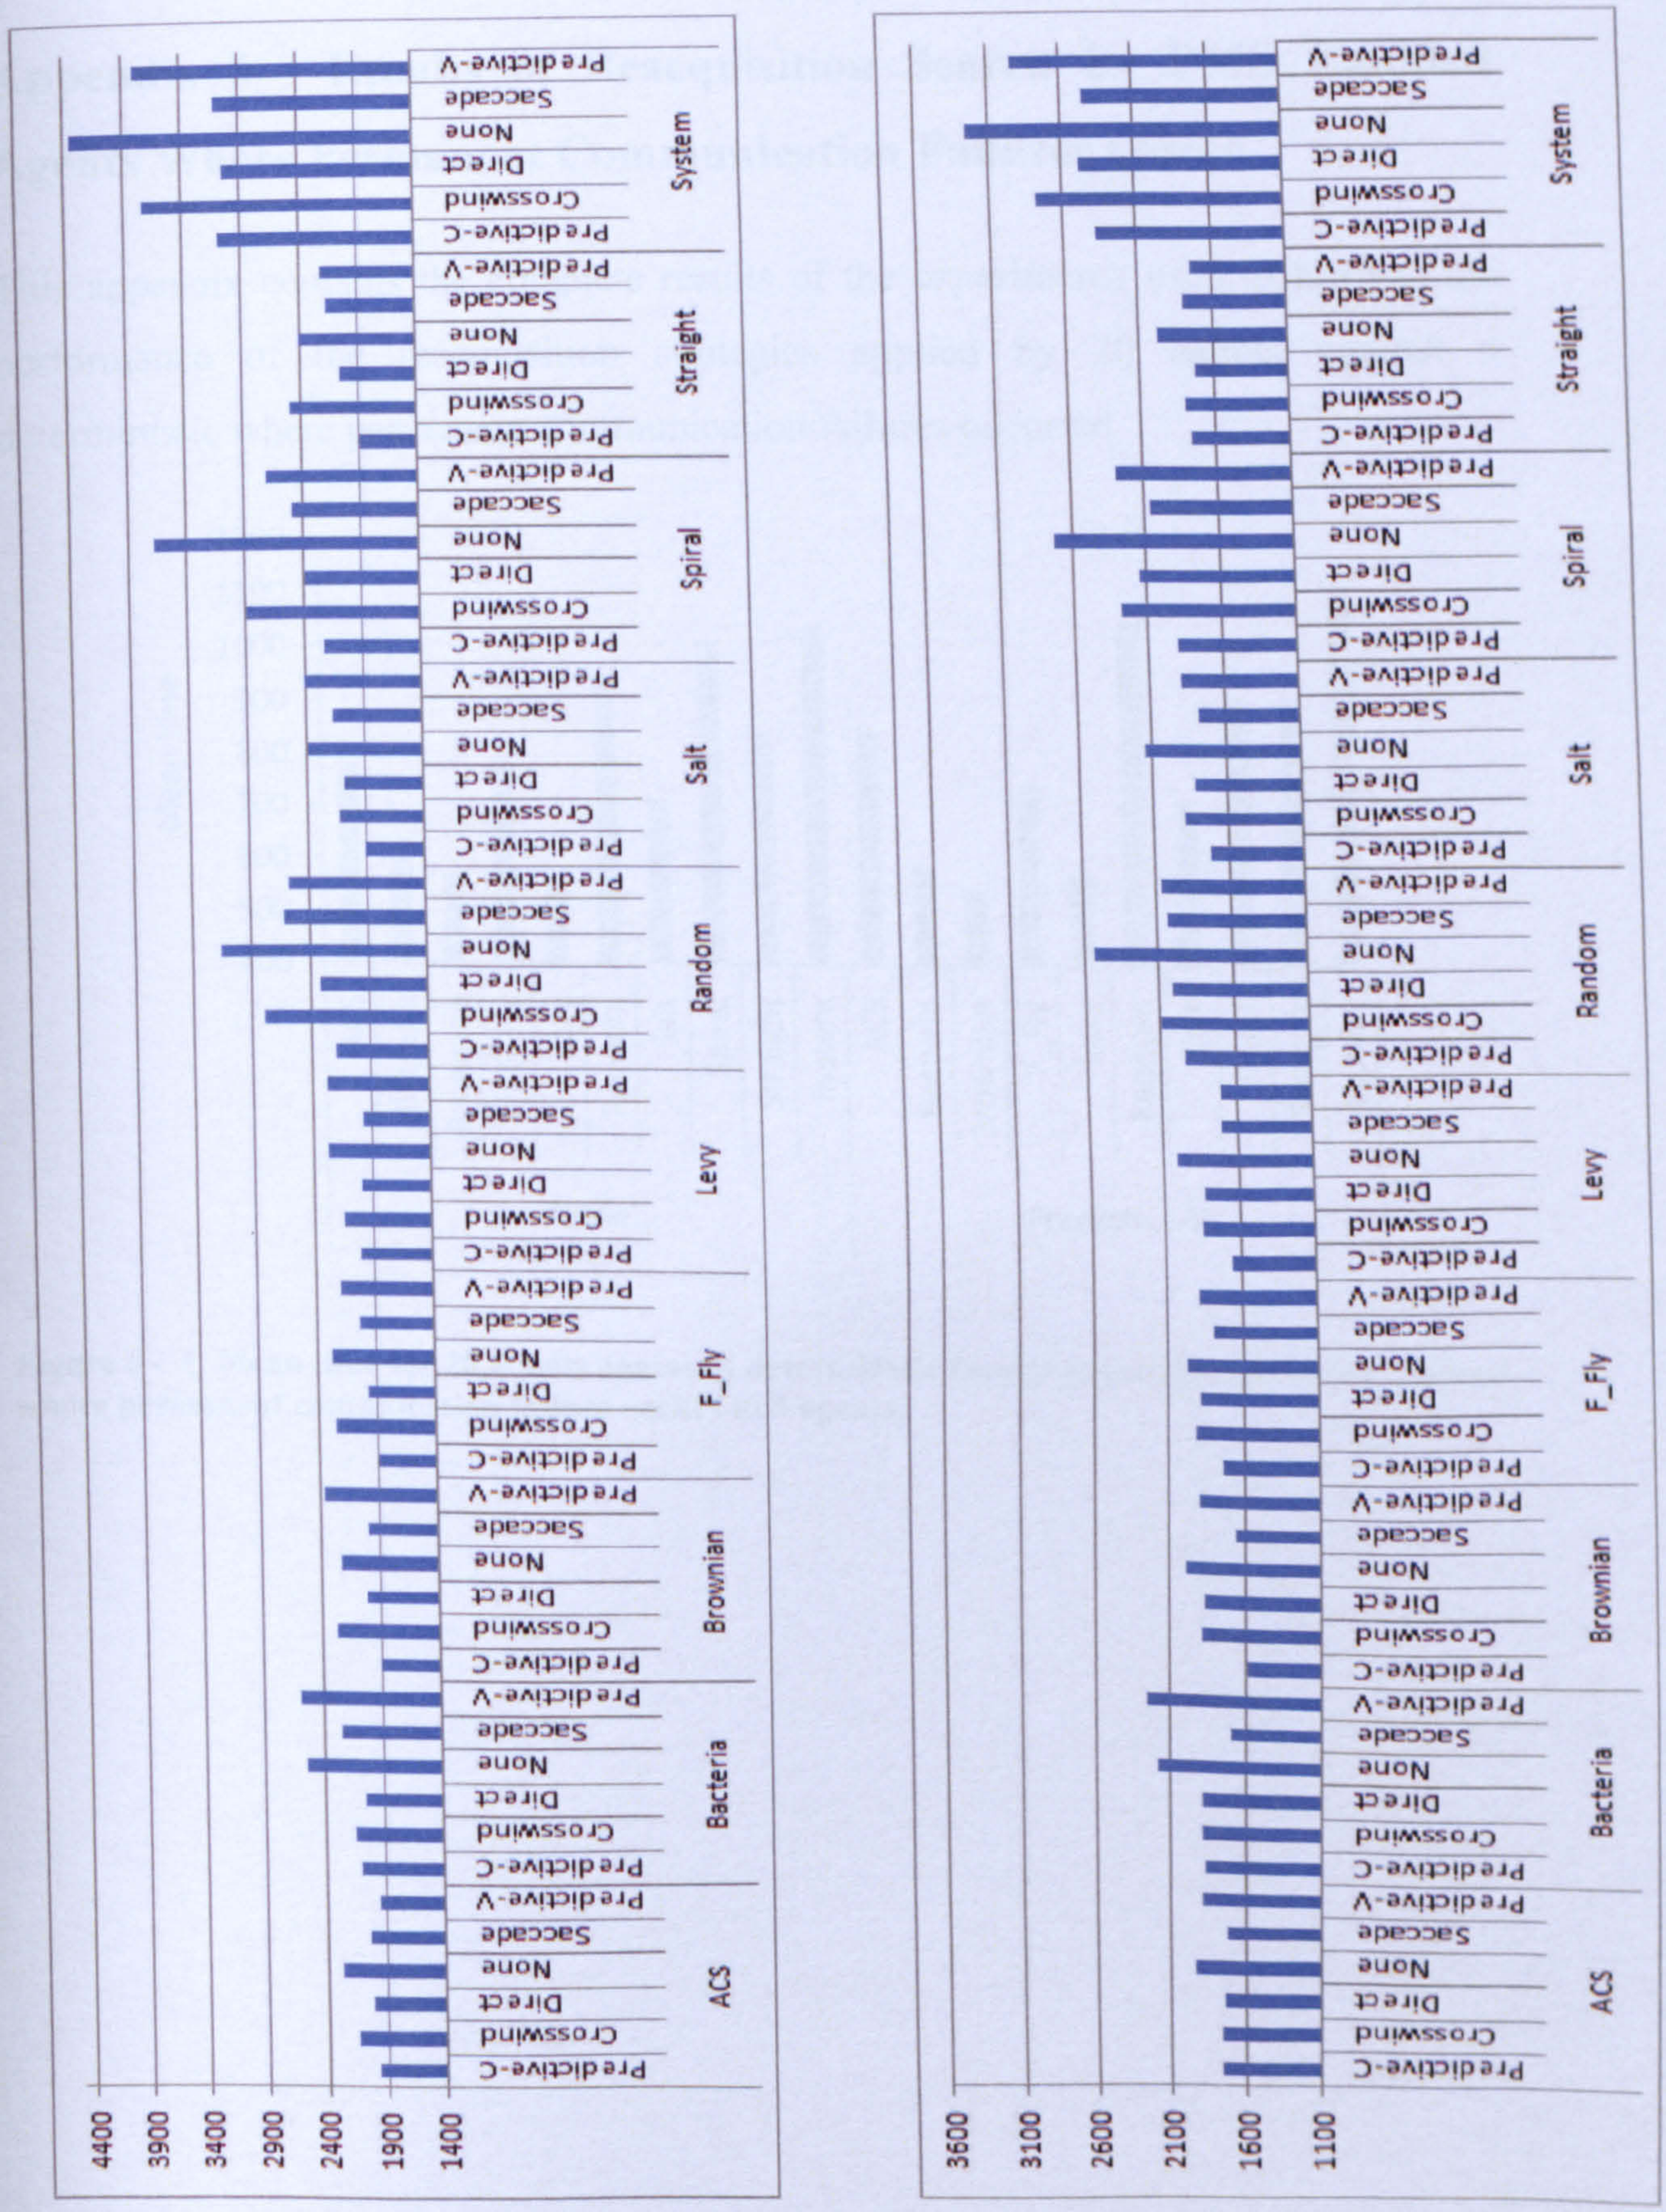

Figure 5 - 6. Mean time for 10 agents against 4 fruit fly targets travelling at 1.5 agent speed (left) and 2.0 agent speed (right).

Appendix 6 - Results of Reacquisition Search by PMS Enabled Agents Where Permanent Communication Failures Occur.

This appendix contains the complete results of the experiments used to analyse the performance of the reacquisition strategies applied by 20 agents against 4 deterministic where permanent communication failures occurred.

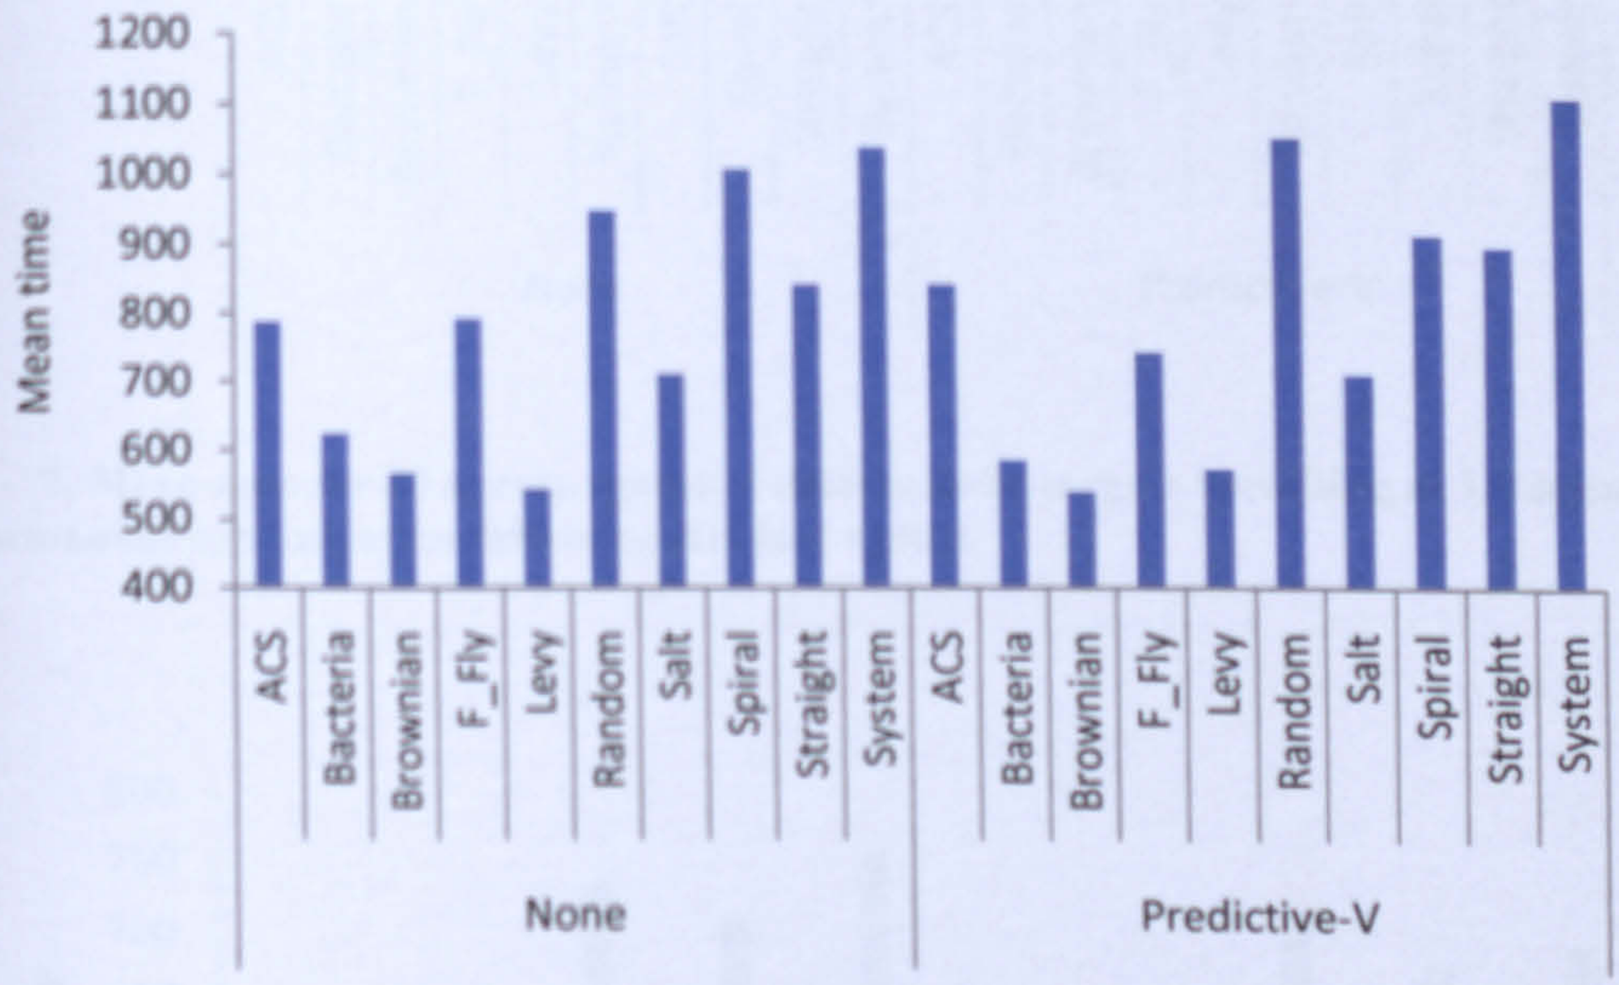

Figure 6 - 1. Mean time for 20 agents against 4 determinstic targets travelling at 0.5 agent speed where permanent comunication failure occurs in 4 agents.

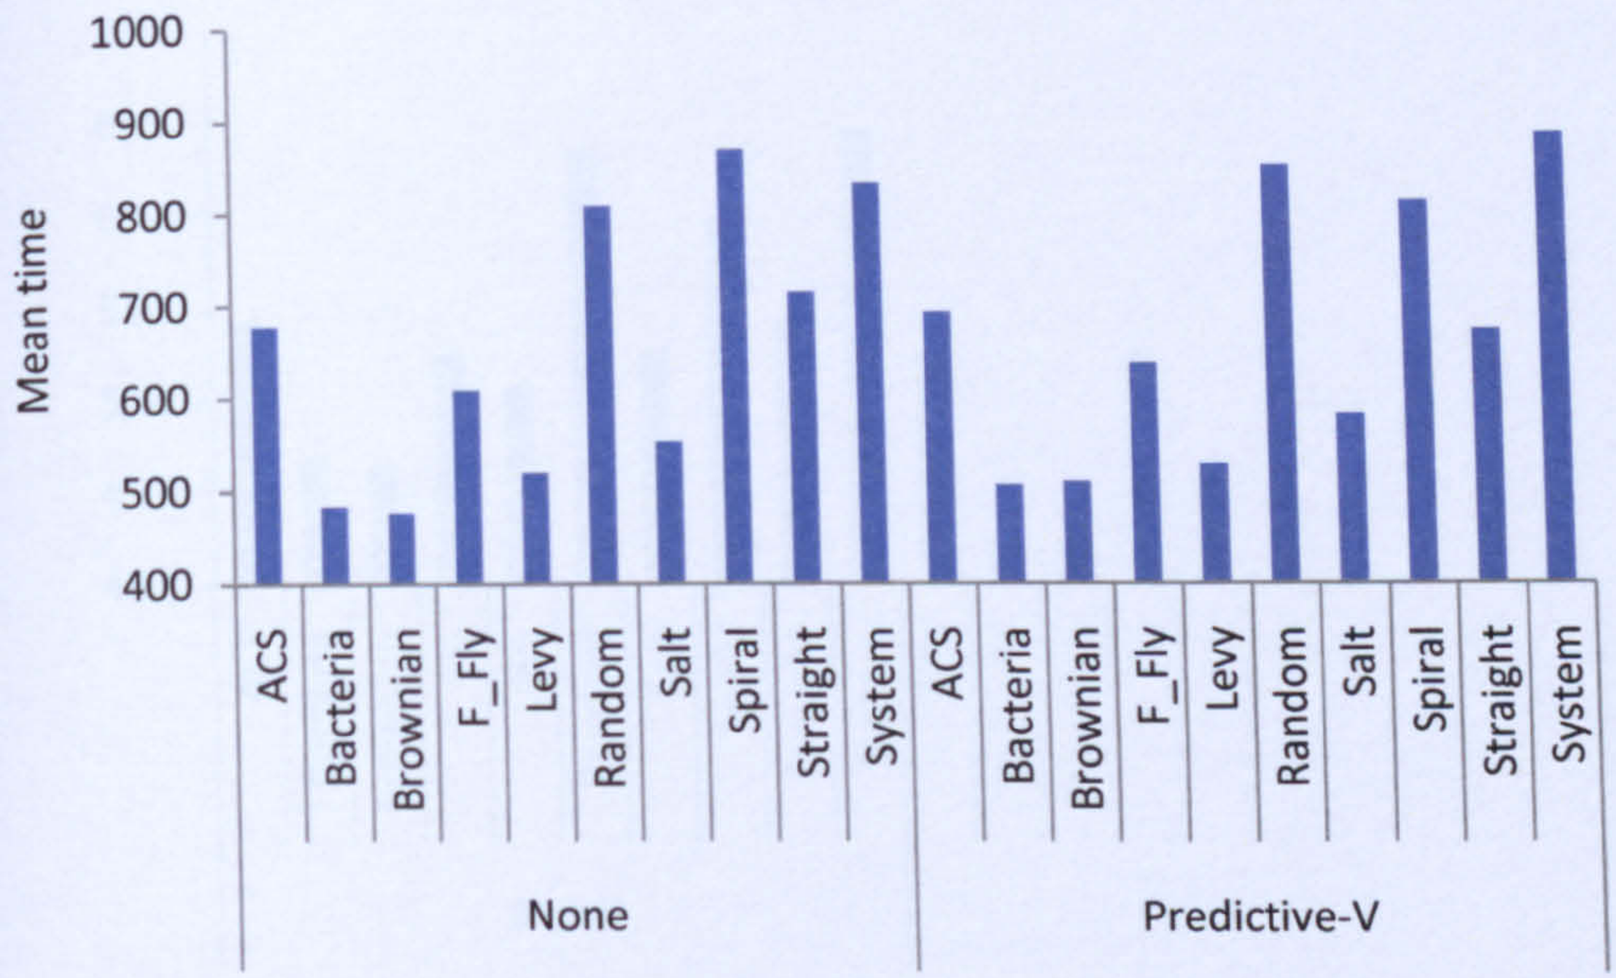

Figure 6 - 2. Mean time for 20 agents against 4 determinstic targets travelling at 1.0 agent speed where permanent comunication failure occurs in 4 agents.

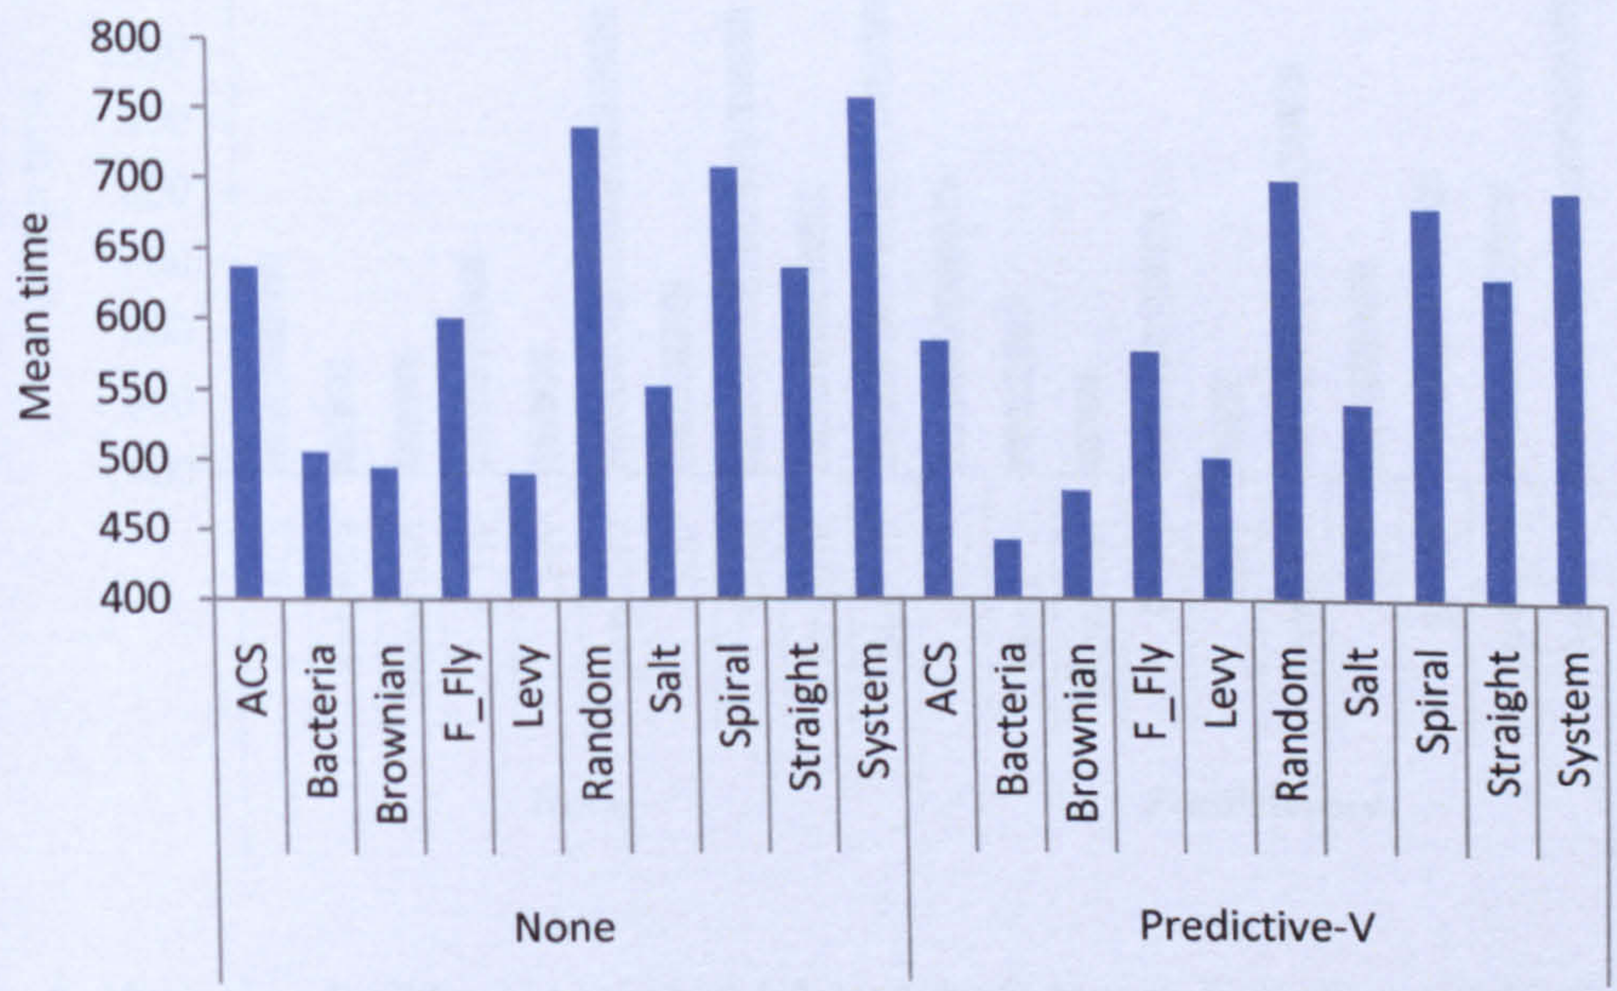

Figure 6 - 3. Mean time for 20 agents against 4 determinstic targets travelling at 1.5 agent speed where permanent comunication failure occurs in 4 agents.

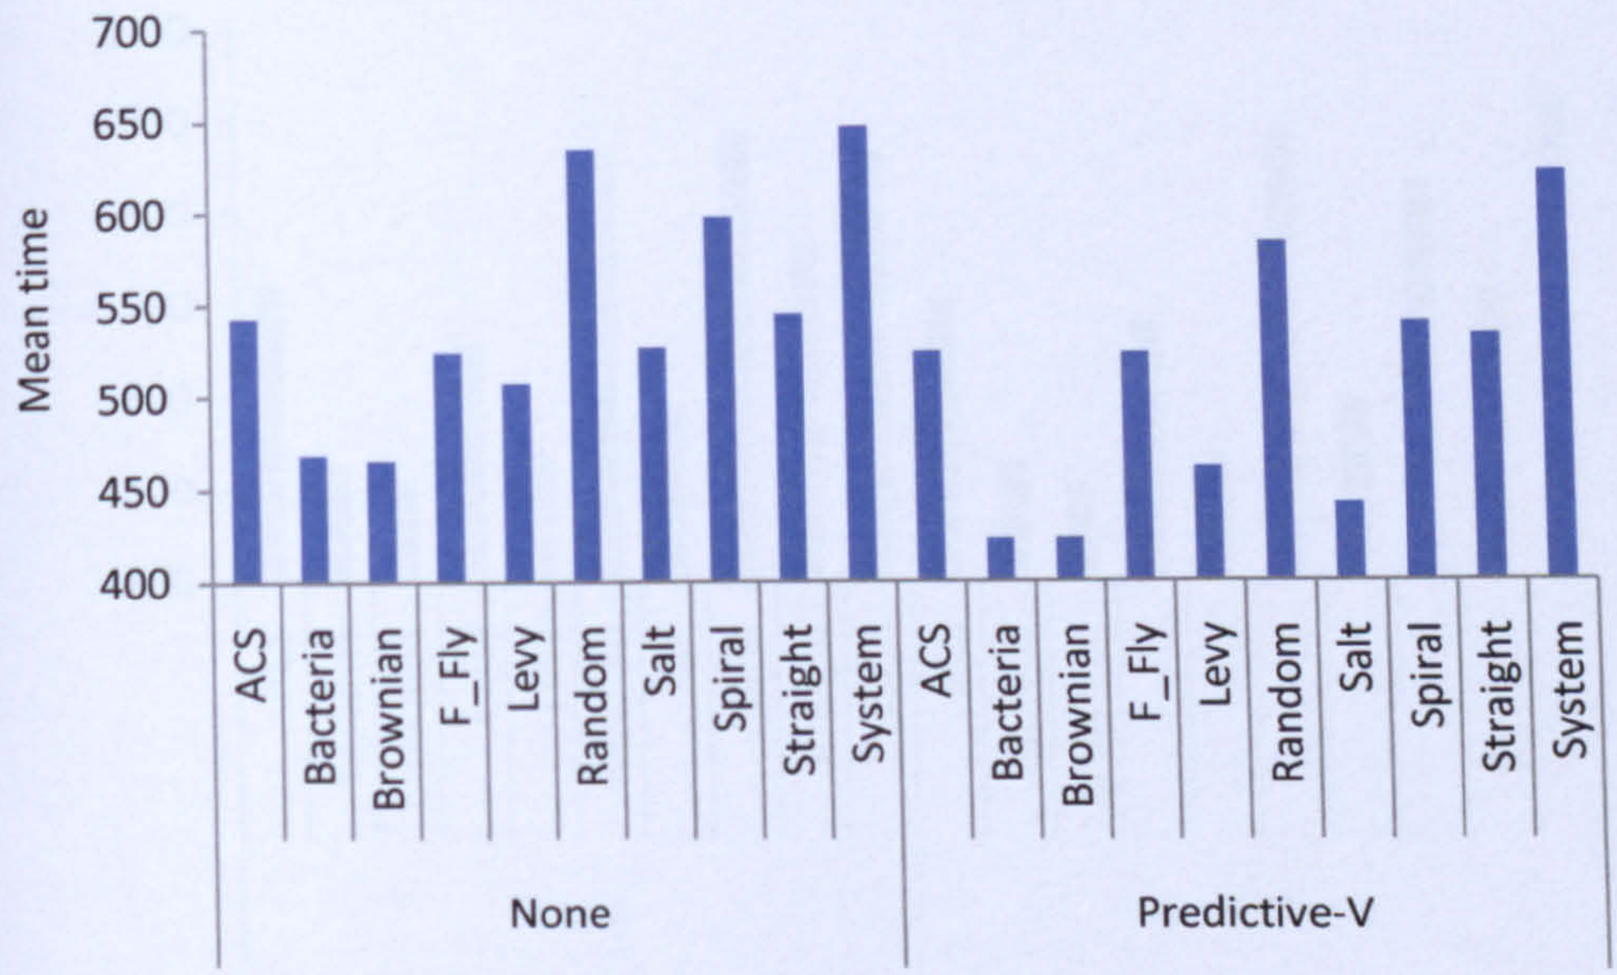

Figure 6 - 4. Mean time for 20 agents against 4 determinstic targets travelling at 2.0 agent speed where permanent comunication failure occurs in 4 agents.

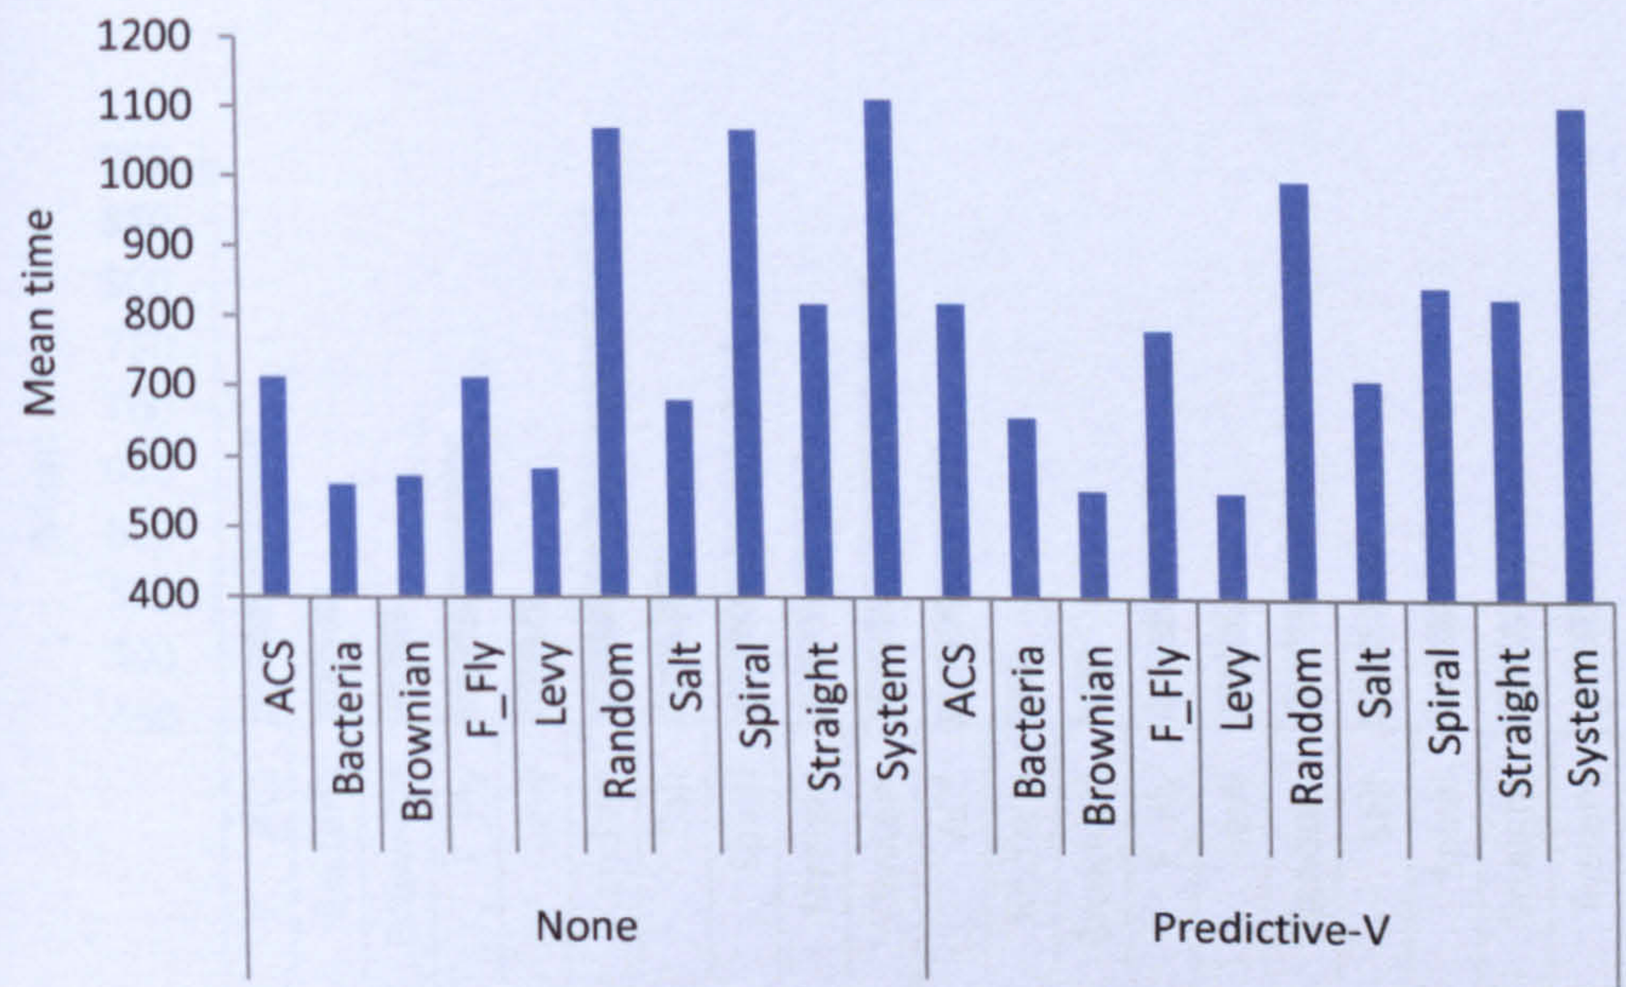

Figure 6 - 5. Mean time for 20 agents against 4 determinstic targets travelling at 0.5 agent speed where permanent comunication failure occurs in 8 agents.

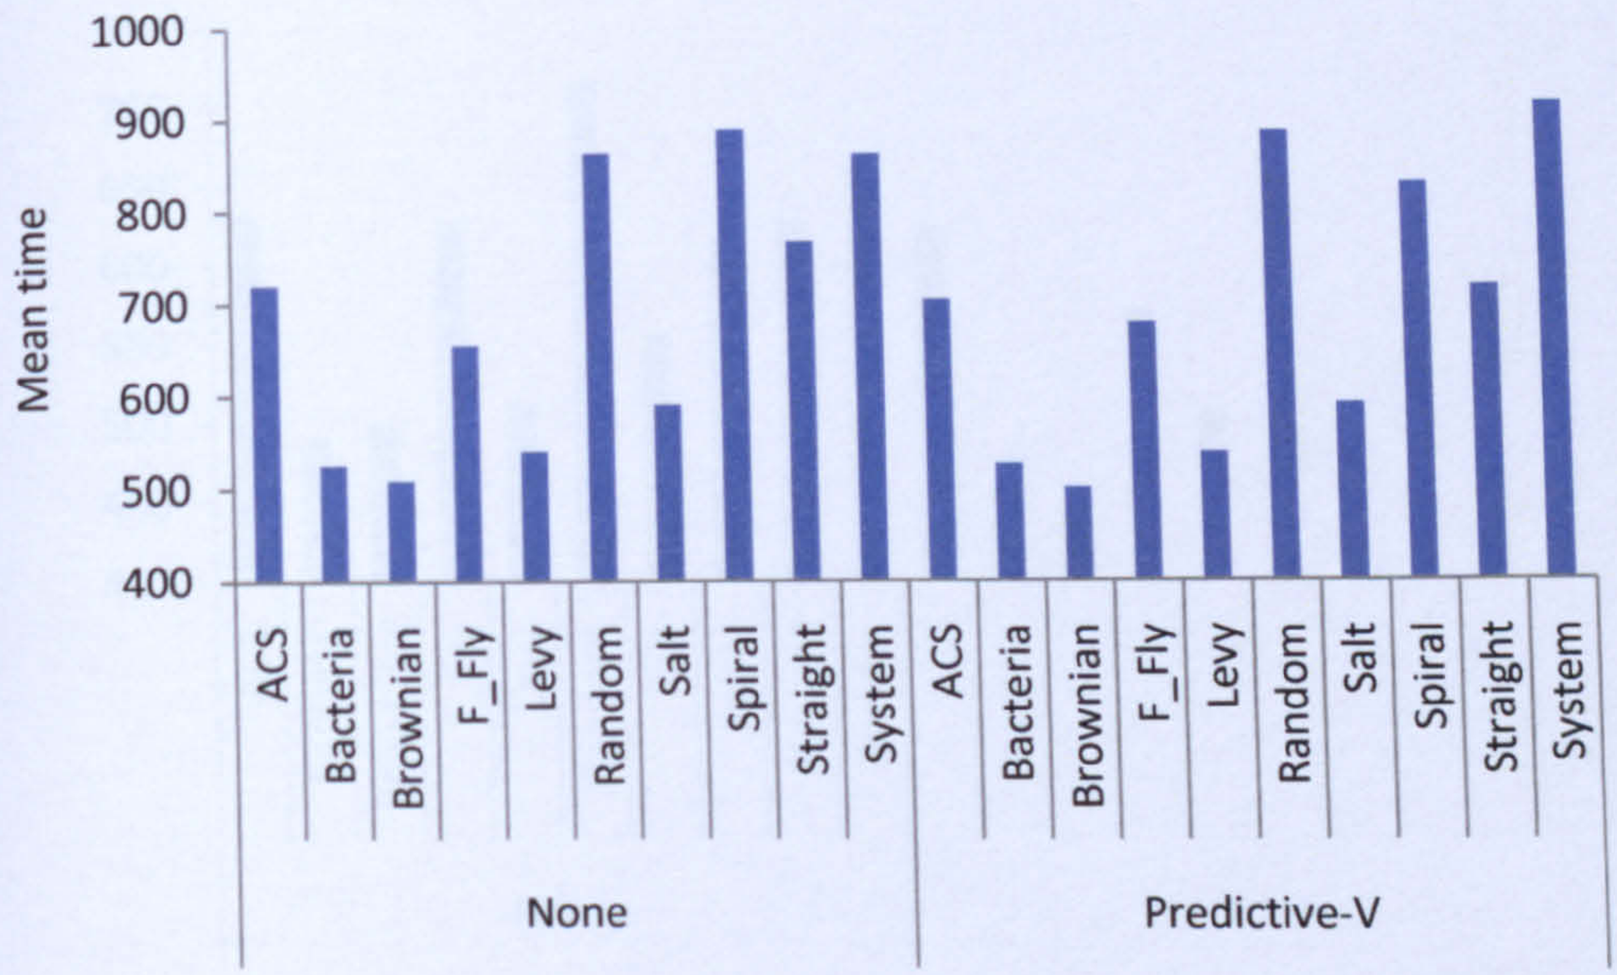

Figure 6 - 6. Mean time for 20 agents against 4 determinstic targets travelling at 1.0 agent speed where permanent comunication failure occurs in 8 agents.

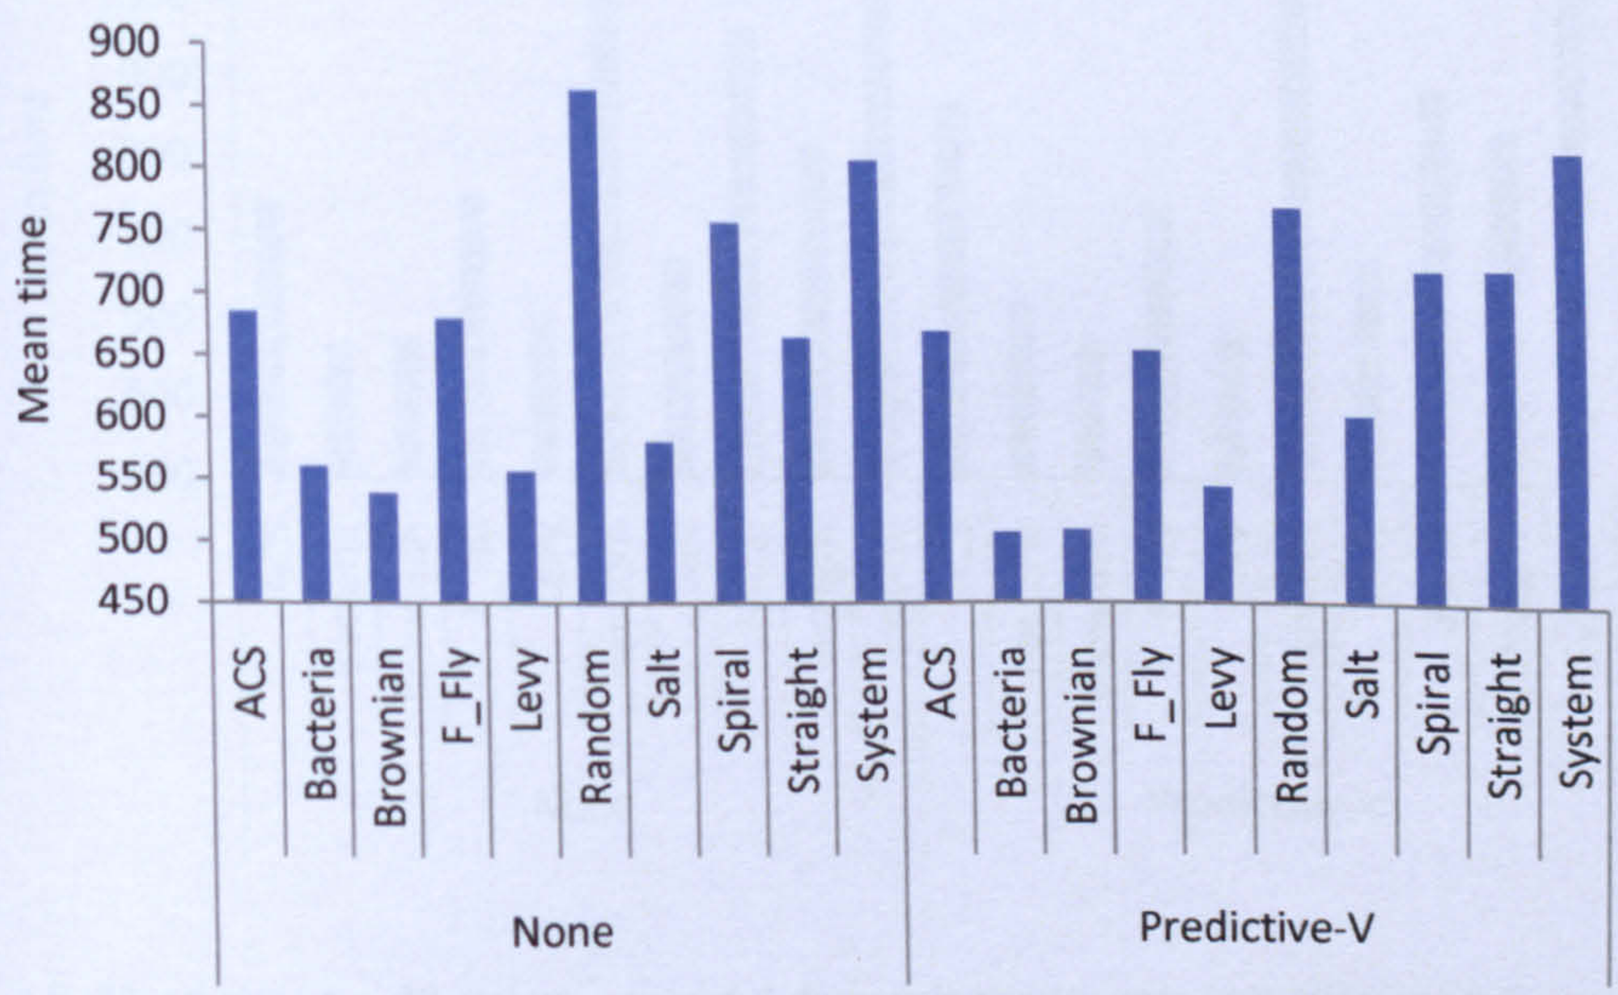

Figure 6 - 7. Mean time for 20 agents against 4 determinstic targets travelling at 1.5 agent speed where permanent comunication failure occurs in 8 agents.

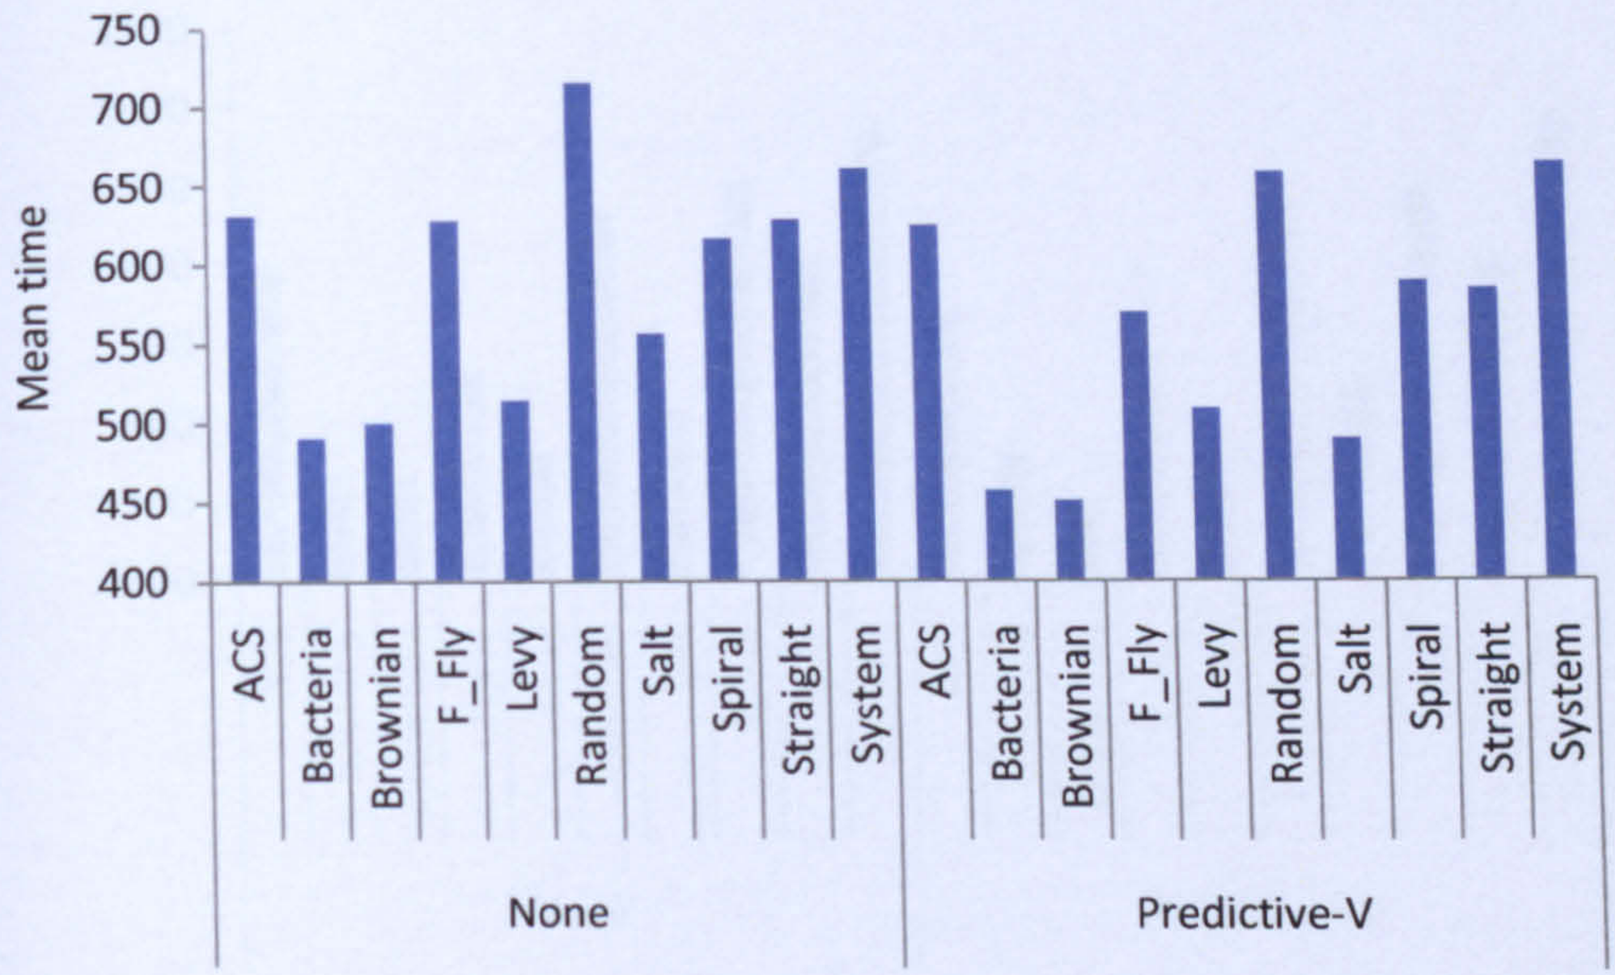

Figure 6 - 8. Mean time for 20 agents against 4 determinstic targets travelling at 2.0 agent speed where permanent comunication failure occurs in 8 agents.

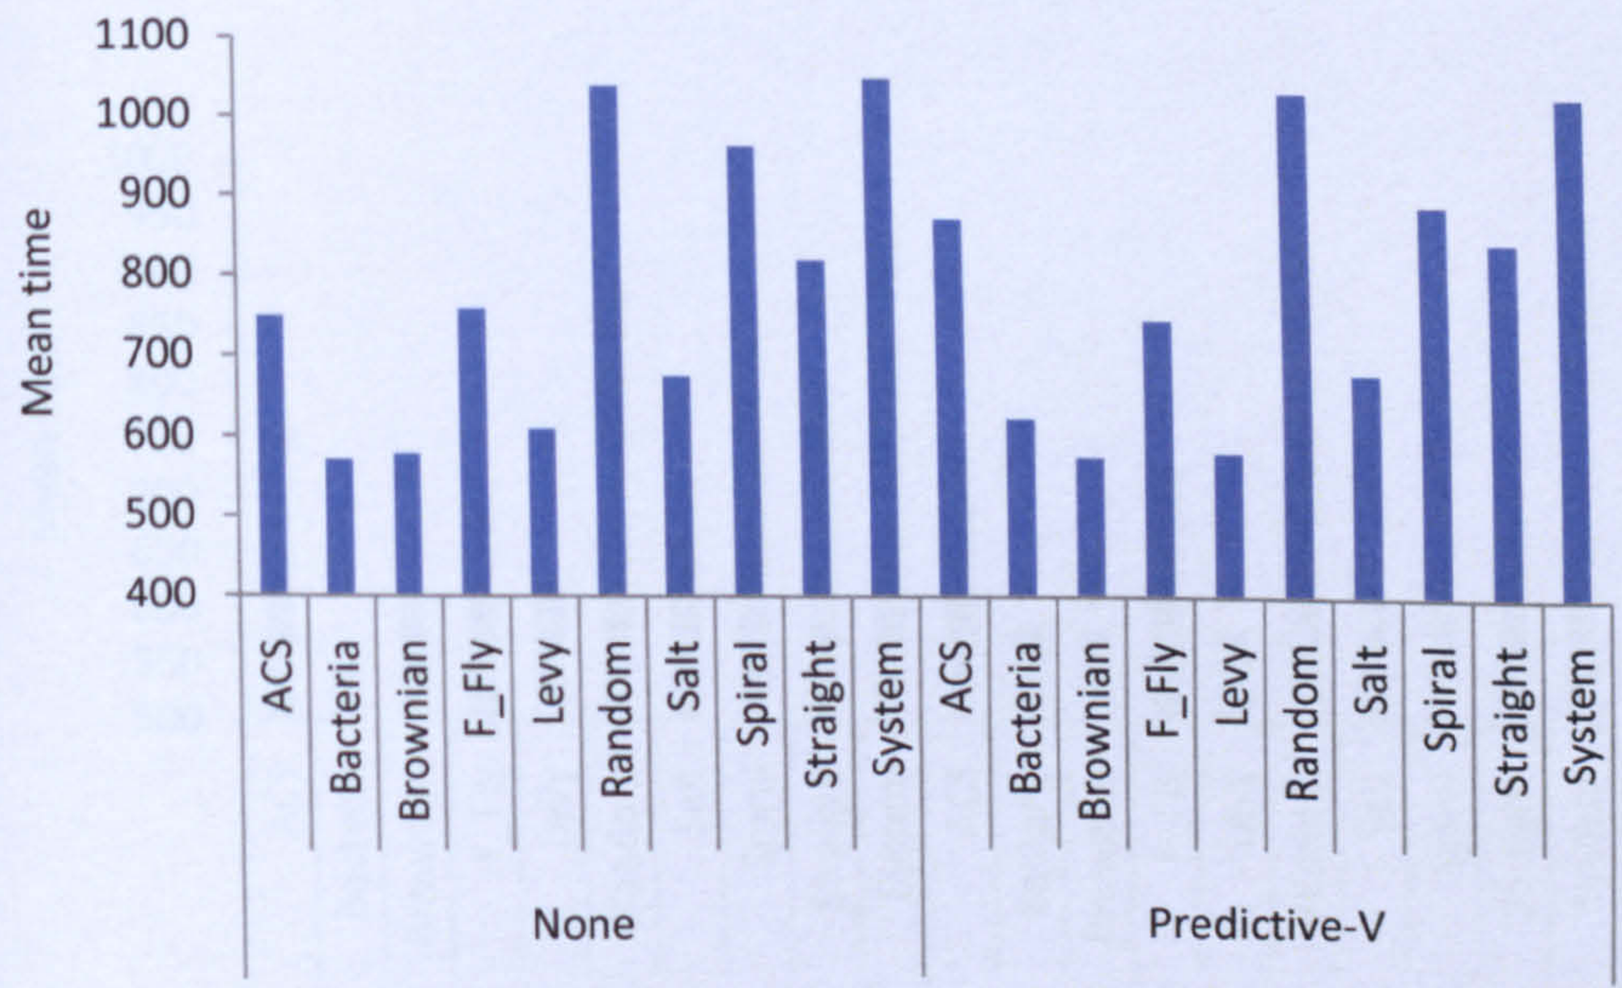

Figure 6 - 9. Mean time for 20 agents against 4 determinstic targets travelling at 0.5 agent speed where permanent comunication failure occurs in 12 agents.

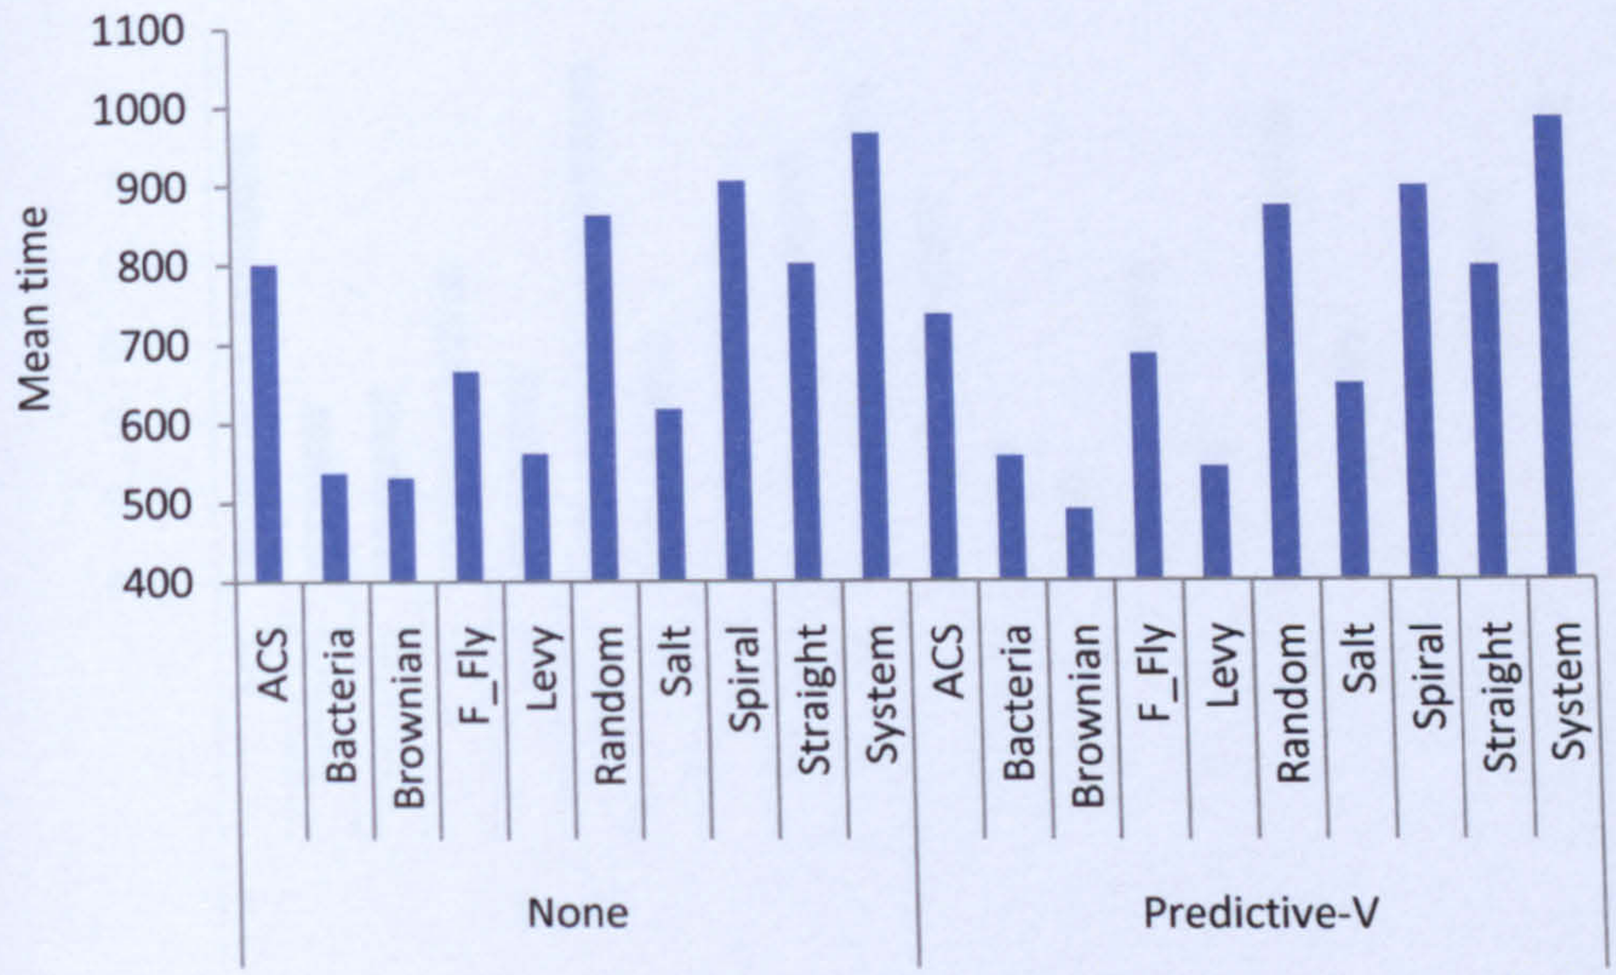

Figure 6 - 10. Mean time for 20 agents against 4 determinstic targets travelling at 1.0 agent speed where permanent comunication failure occurs in 12 agents.

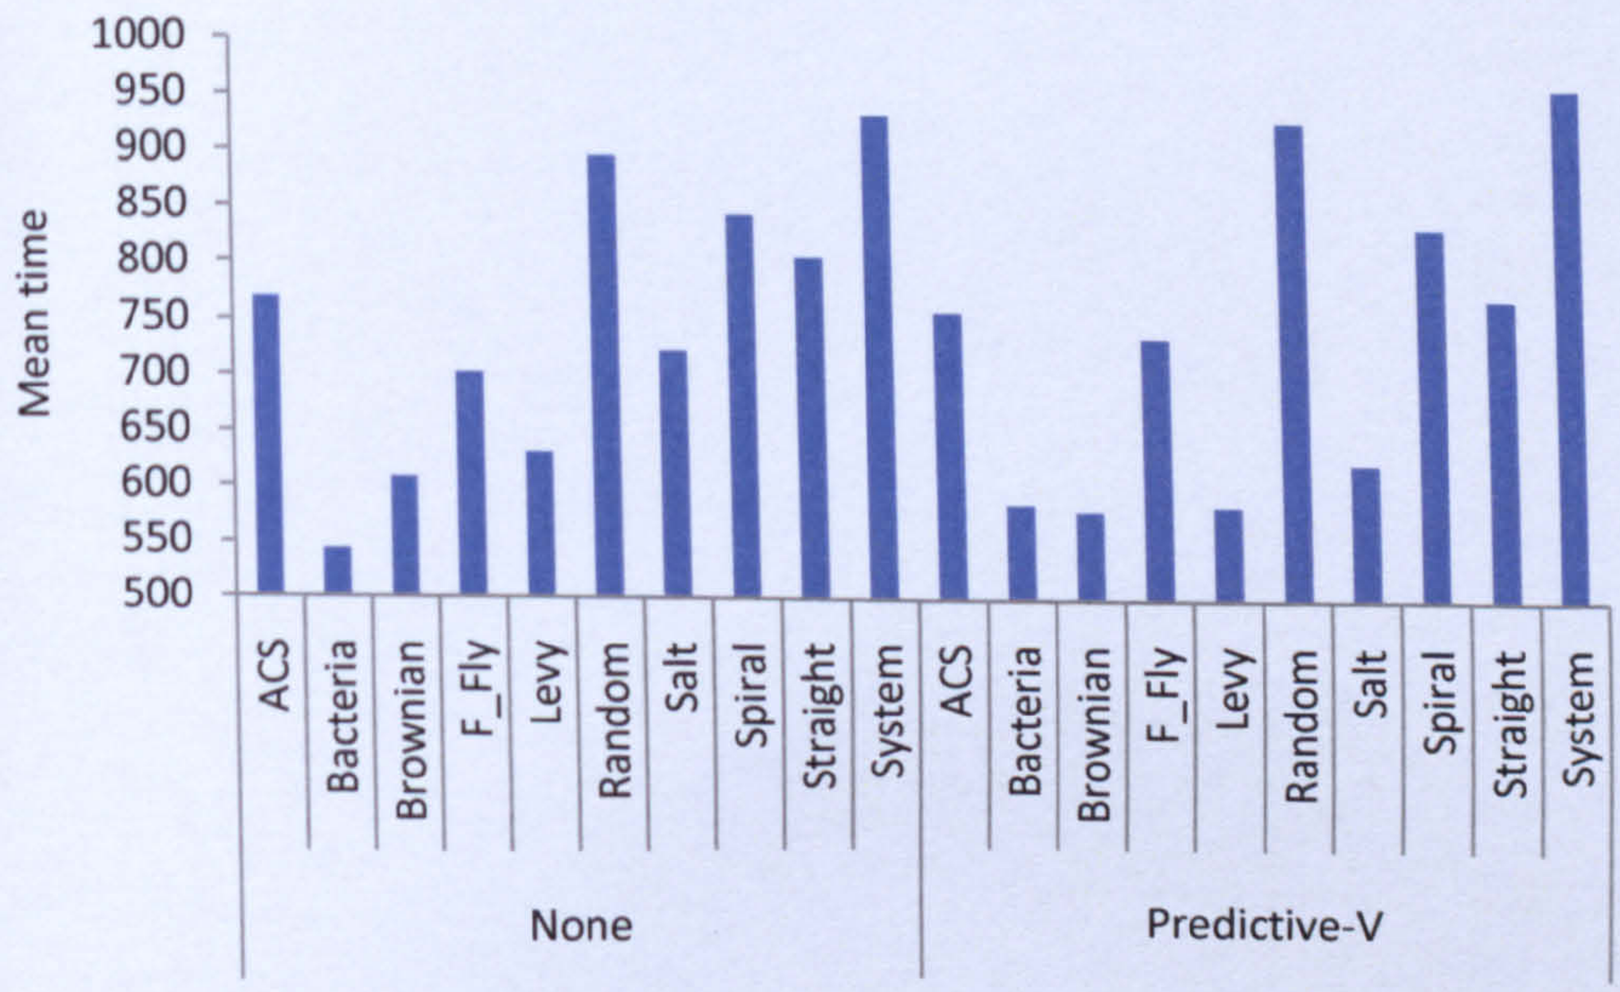

Figure 6 - 11. Mean time for 20 agents against 4 determinstic targets travelling at 1.5 agent speed where permanent comunication failure occurs in 12 agents.

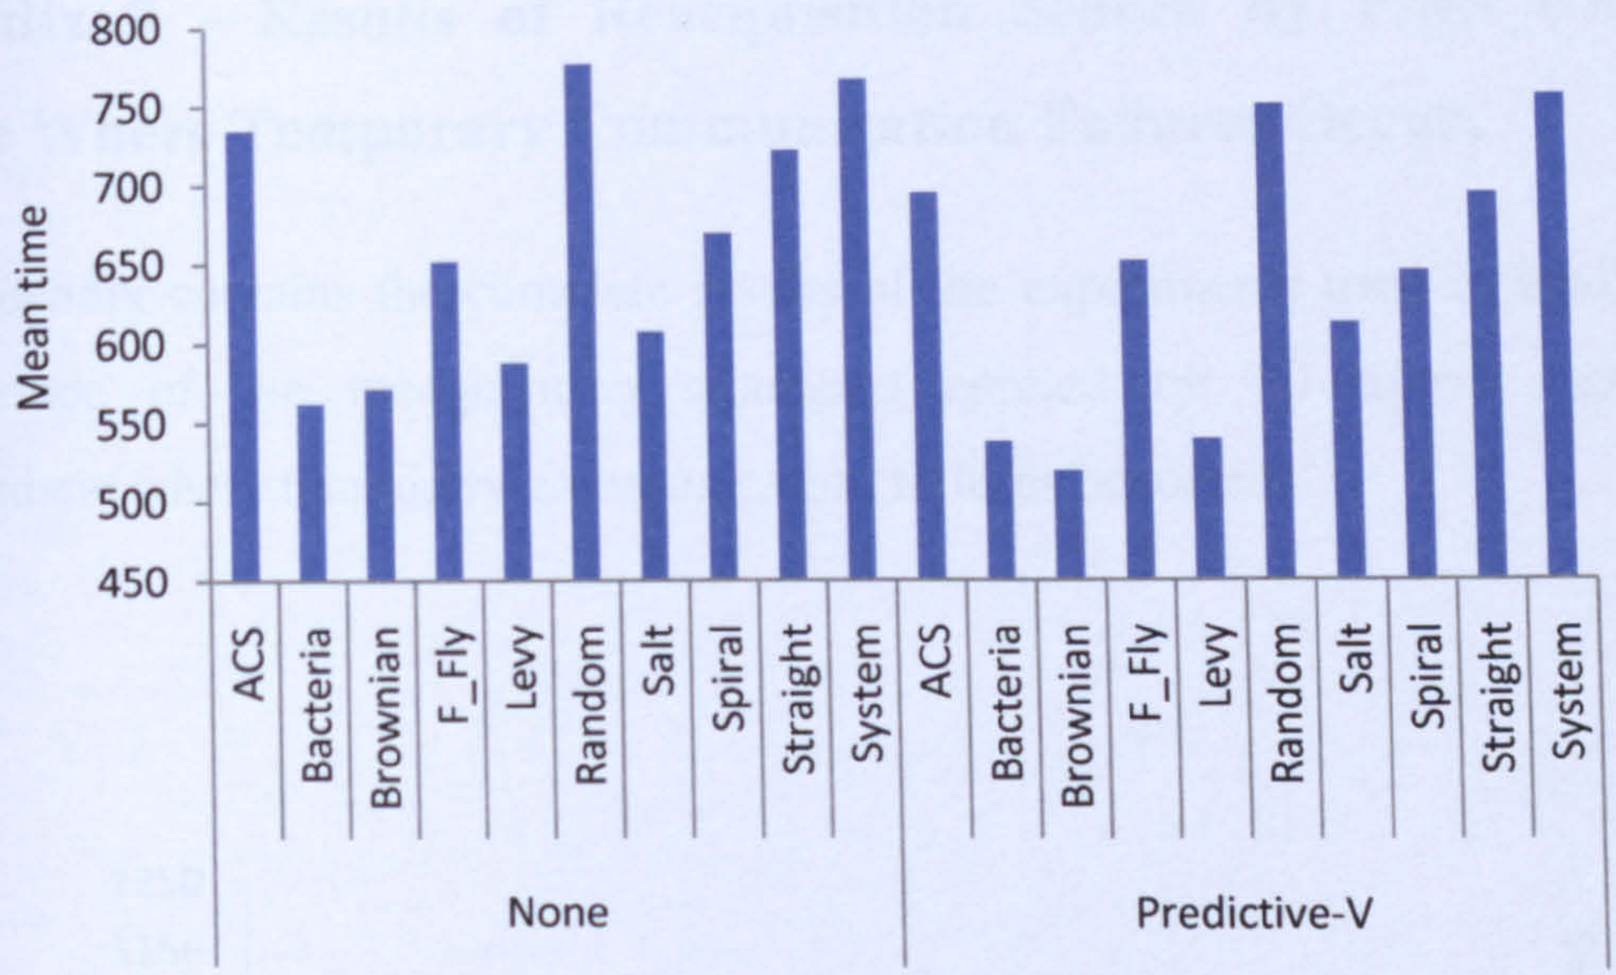

Figure 6 - 12. Mean time for 20 agents against 4 determinstic targets travelling at 2.0 agent speed where permanent comunication failure occurs in 12 agents.

Appendix 7 - Results of Reacquisition Search by PMS Enabled Agents Where Temporary Communication Failures Occur.

This appendix contains the complete results of the experiments used to analyse the performance of the reacquisition strategies applied by 20 agents against 4 deterministic where temporary communication failures occurred.

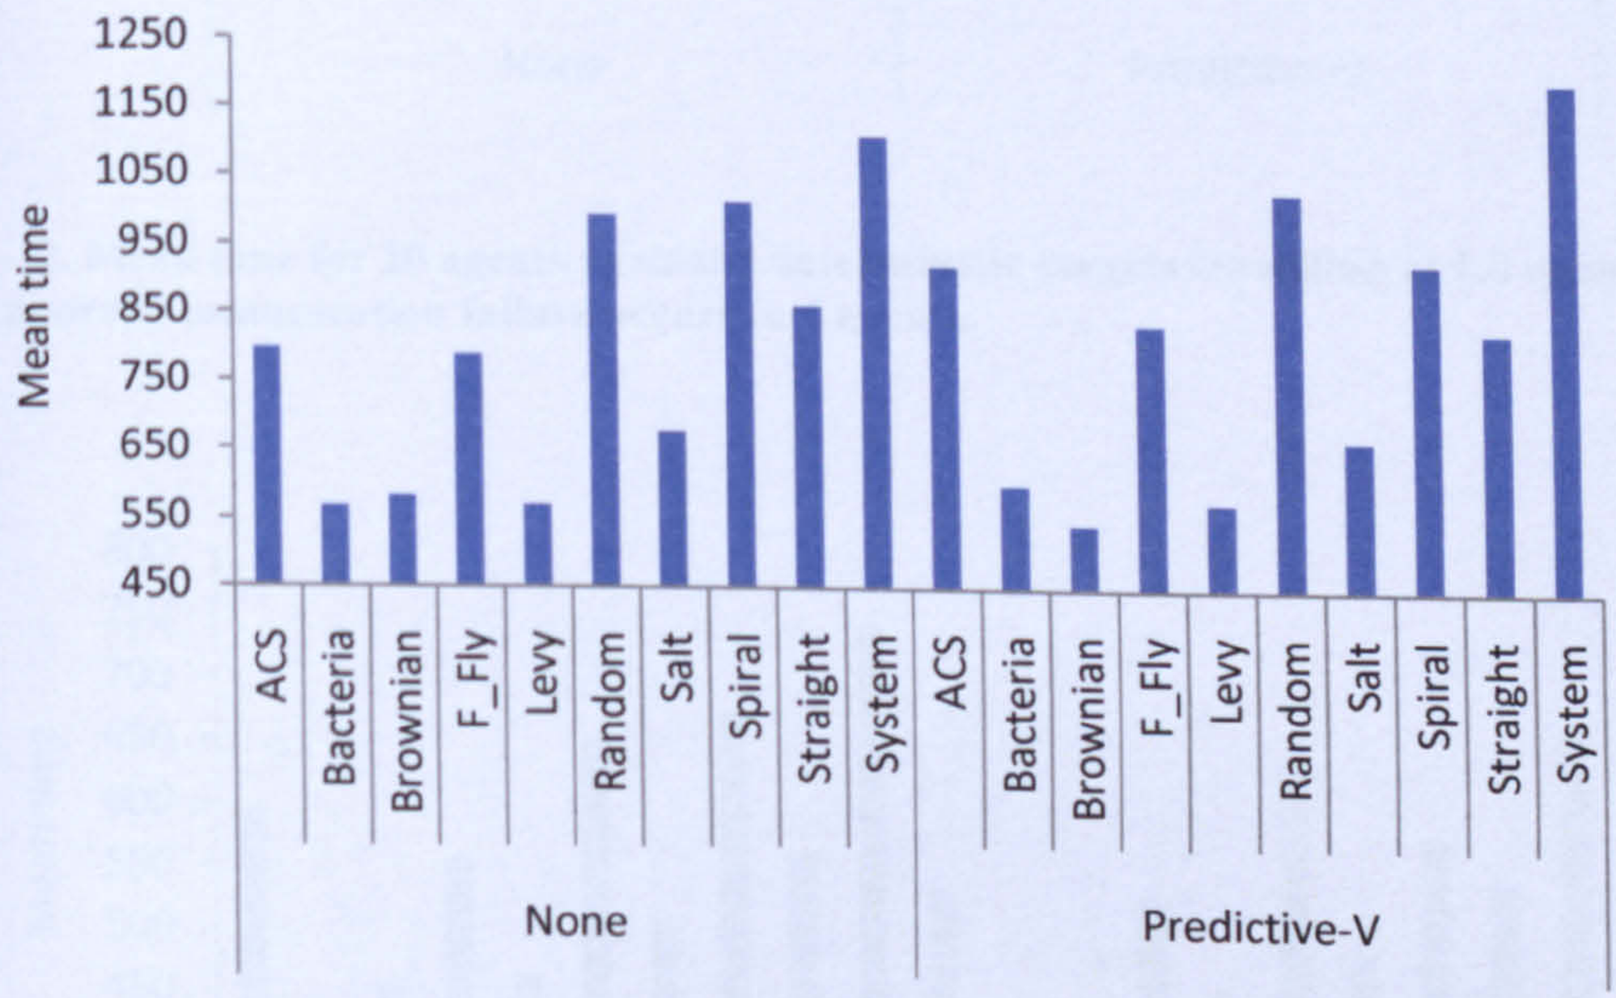

Figure 7 - 1. Mean time for 20 agents against 4 determinstic targets travelling at 0.5 agent speed where temporary comunication failure occurs in 4 agents.

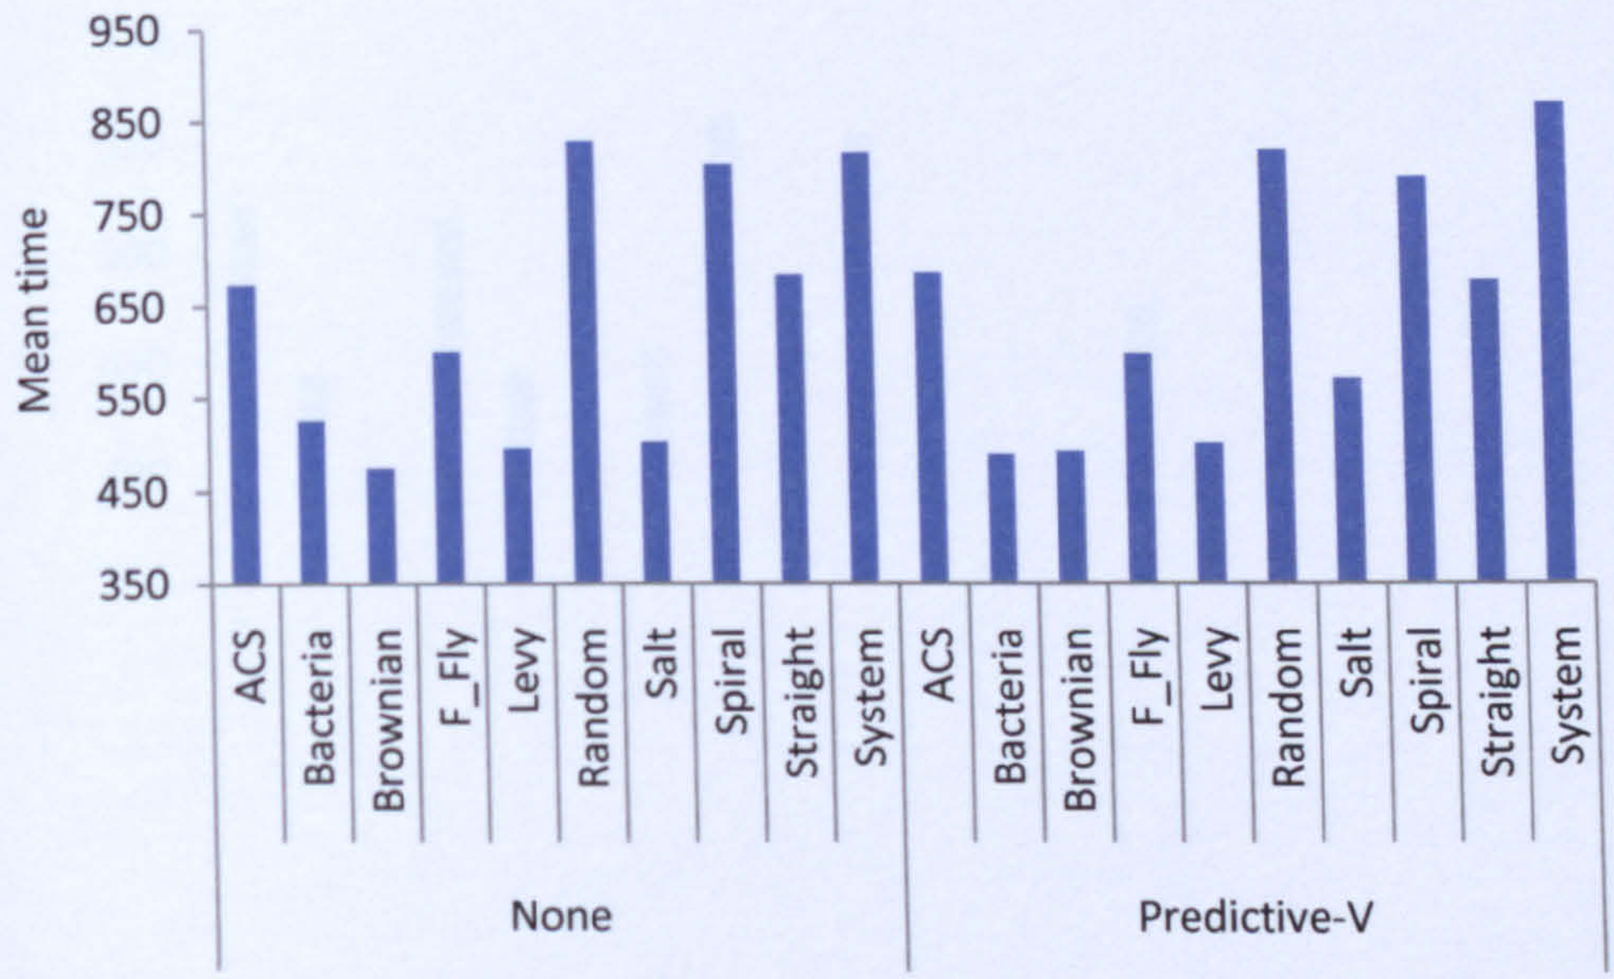

Figure 7 - 2. Mean time for 20 agents against 4 determinstic targets travelling at 1.0 agent speed where temporary comunication failure occurs in 4 agents.

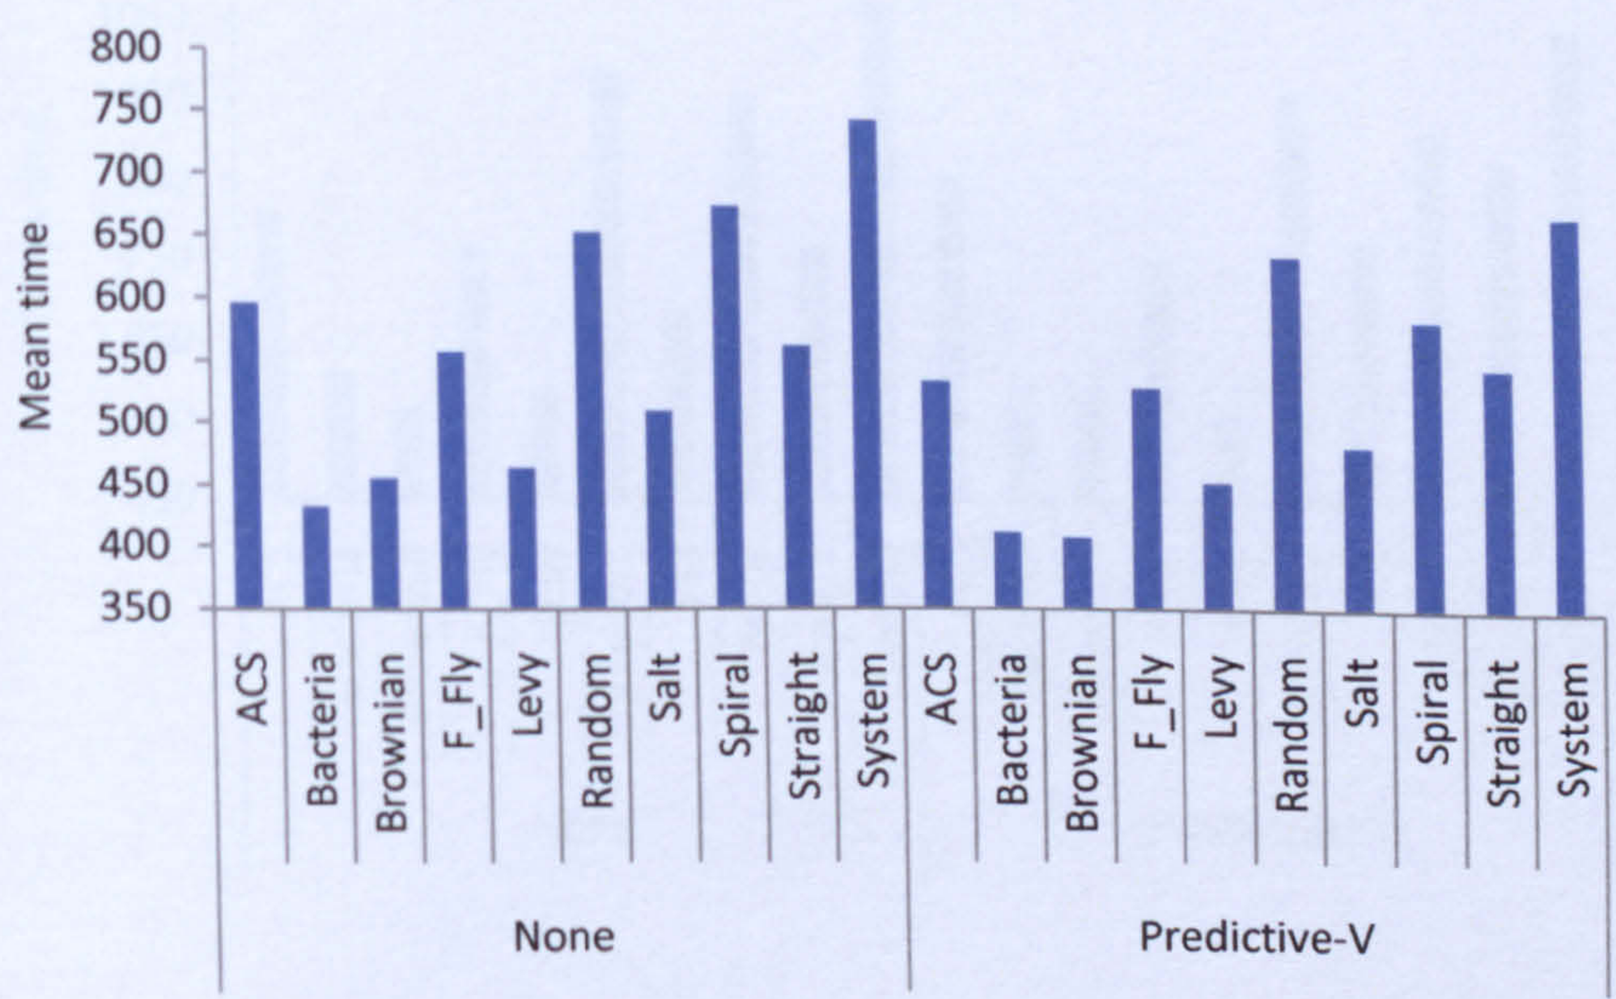

Figure 7 - 3. Mean time for 20 agents against 4 determinstic targets travelling at 1.5 agent speed where temporary comunication failure occurs in 4 agents.

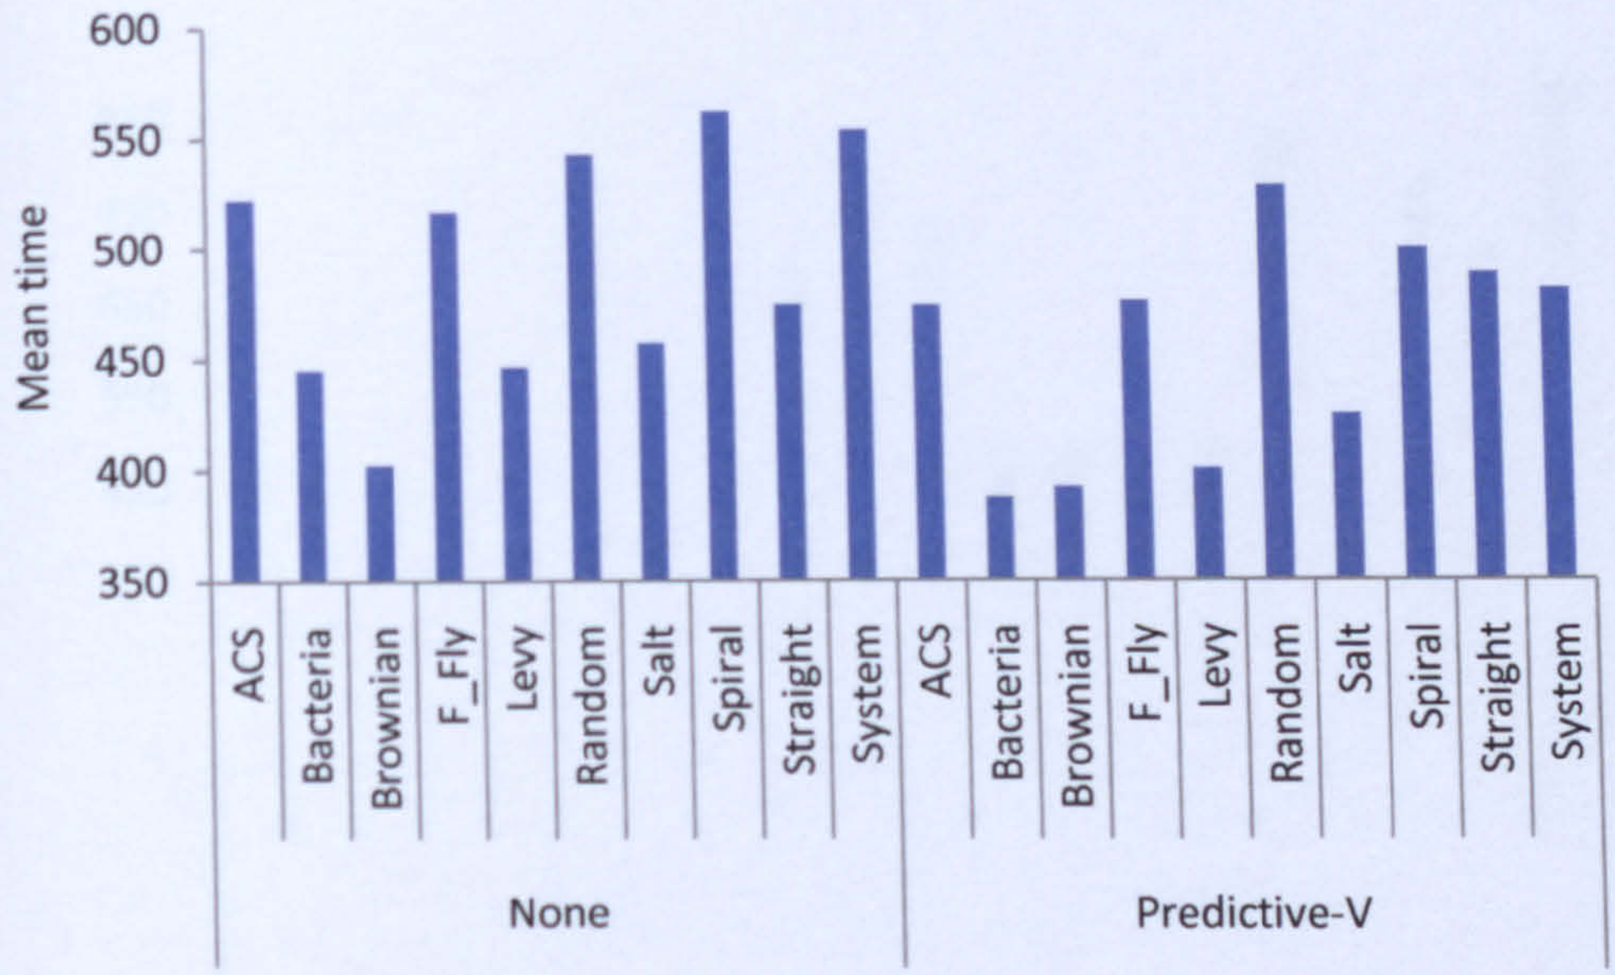

Figure 7 - 4. Mean time for 20 agents against 4 determinstic targets travelling at 2.0 agent speed where temporary comunication failure occurs in 4 agents.

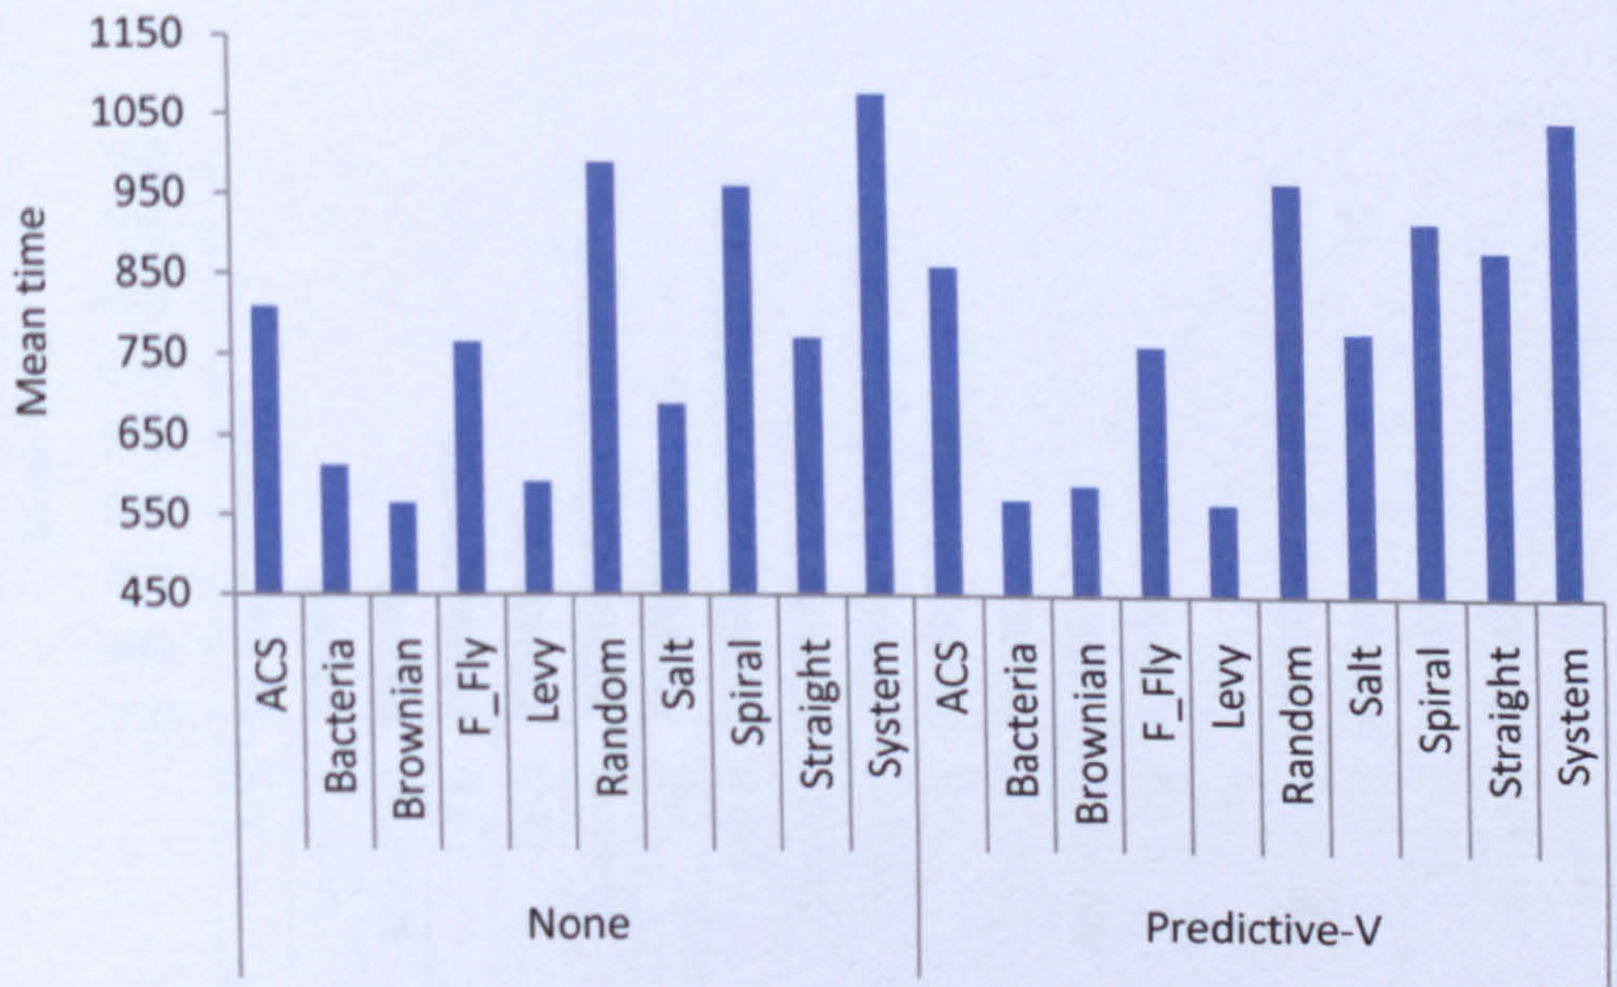

Figure 7 - 5. Mean time for 20 agents against 4 determinstic targets travelling at 0.5 agent speed where temporary comunication failure occurs in 8 agents.

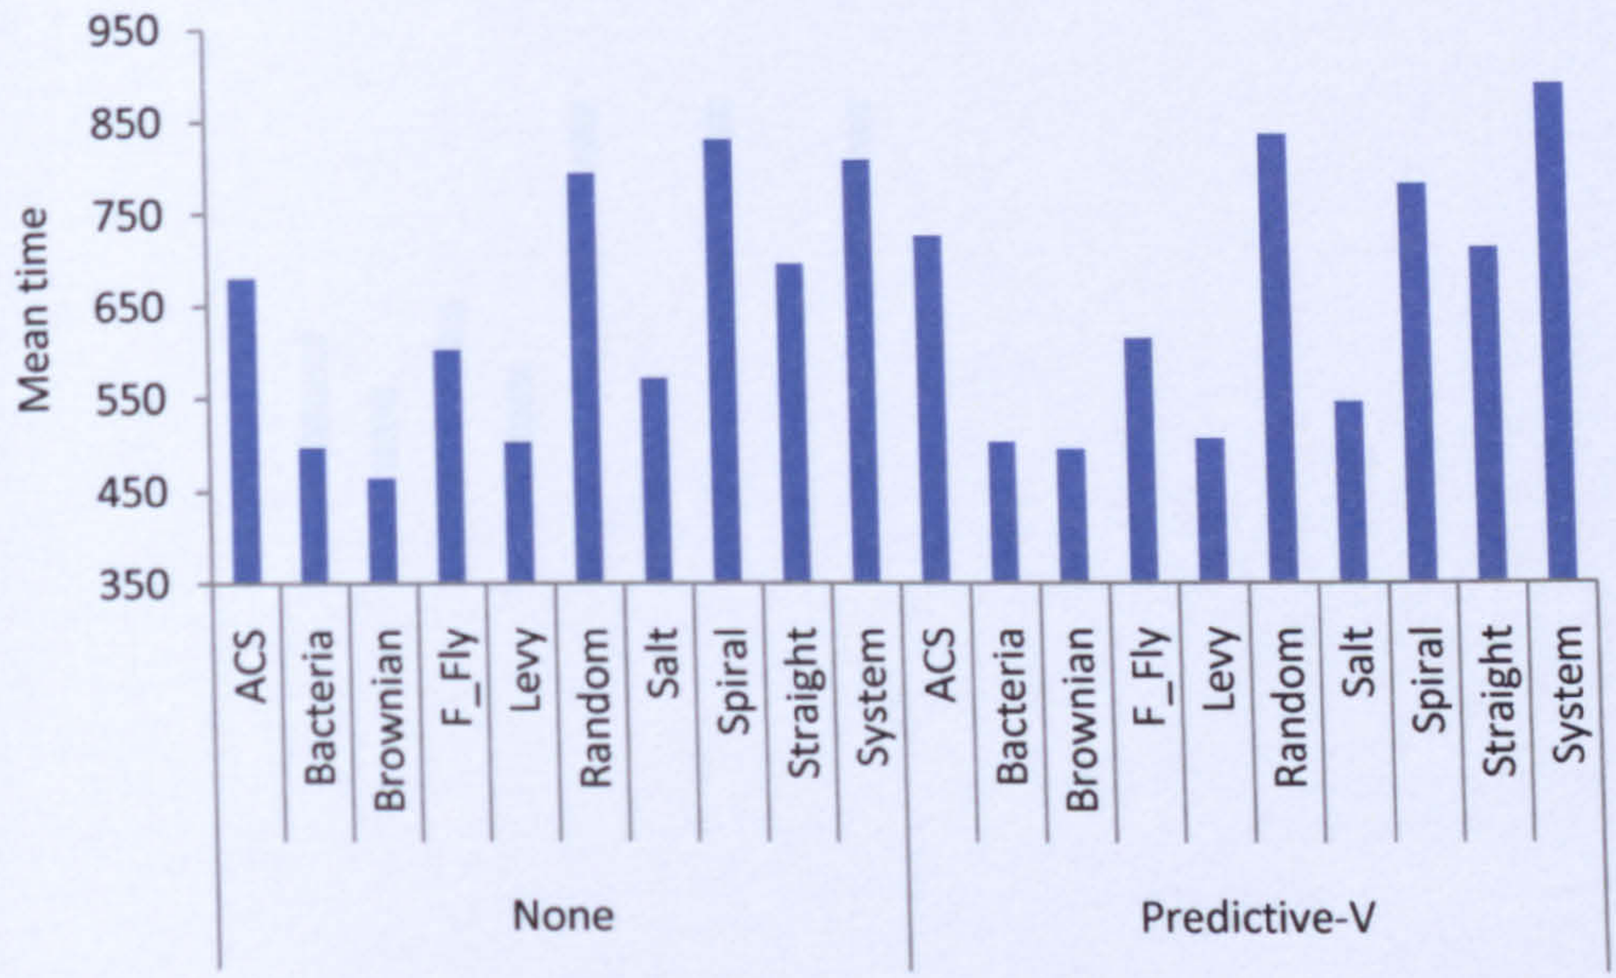

Figure 7 - 6. Mean time for 20 agents against 4 determinstic targets travelling at 1.0 agent speed where temporary comunication failure occurs in 8 agents.

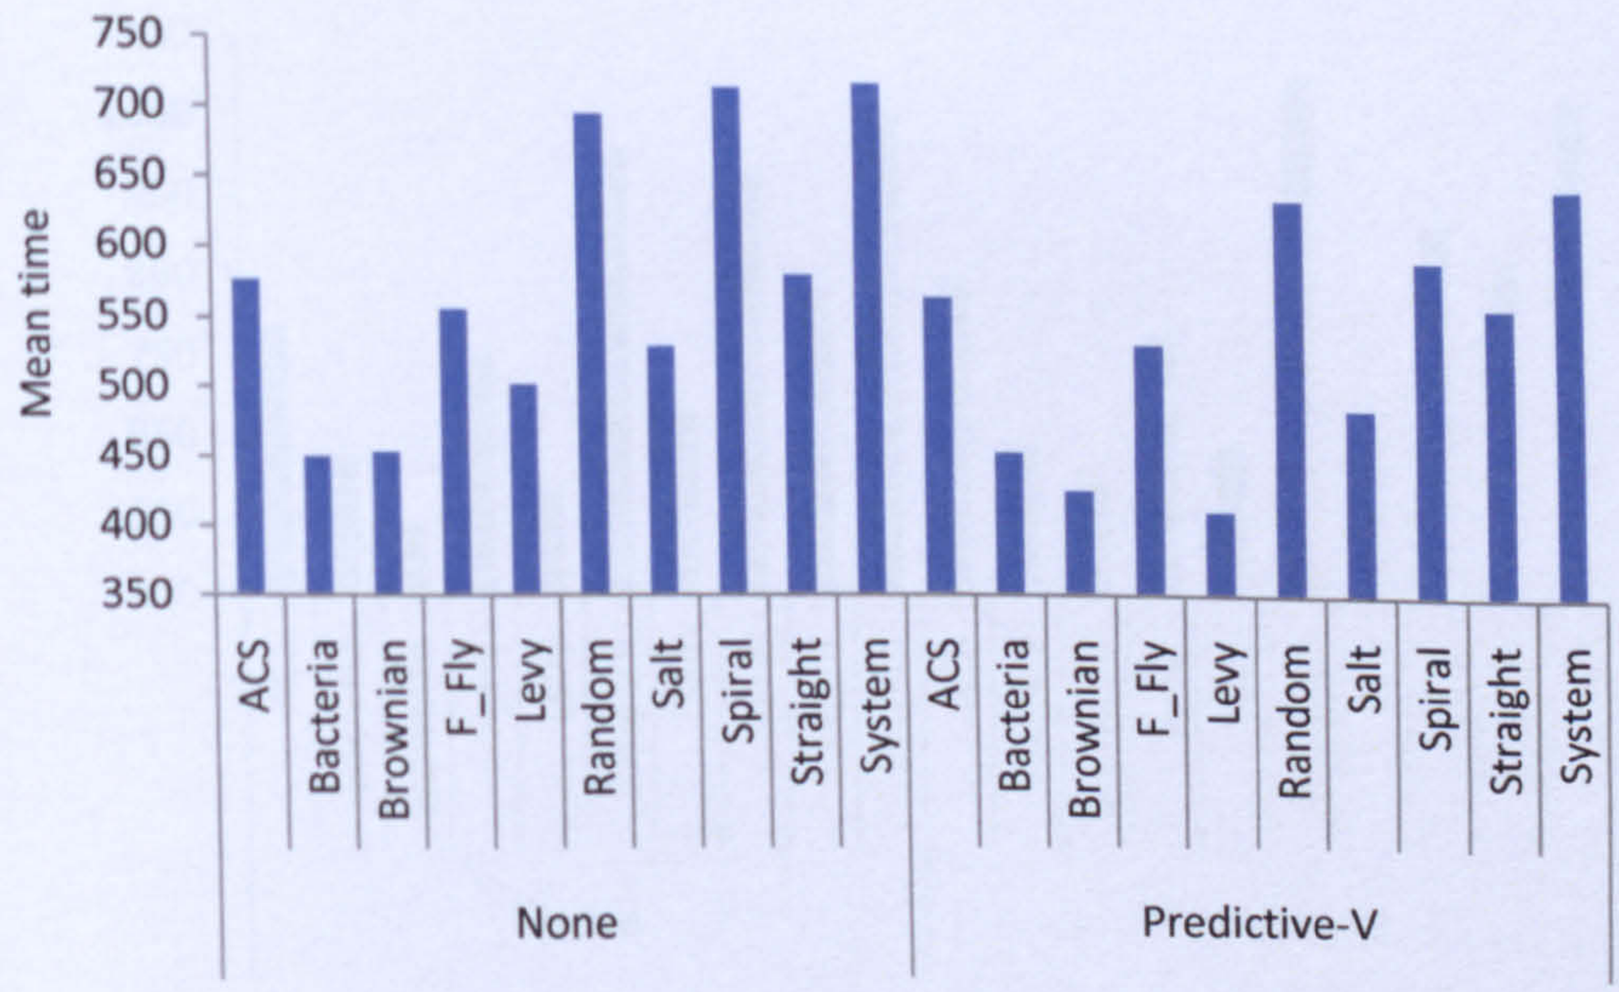

Figure 7 - 7. Mean time for 20 agents against 4 determinstic targets travelling at 1.5 agent speed where temporary comunication failure occurs in 8 agents.

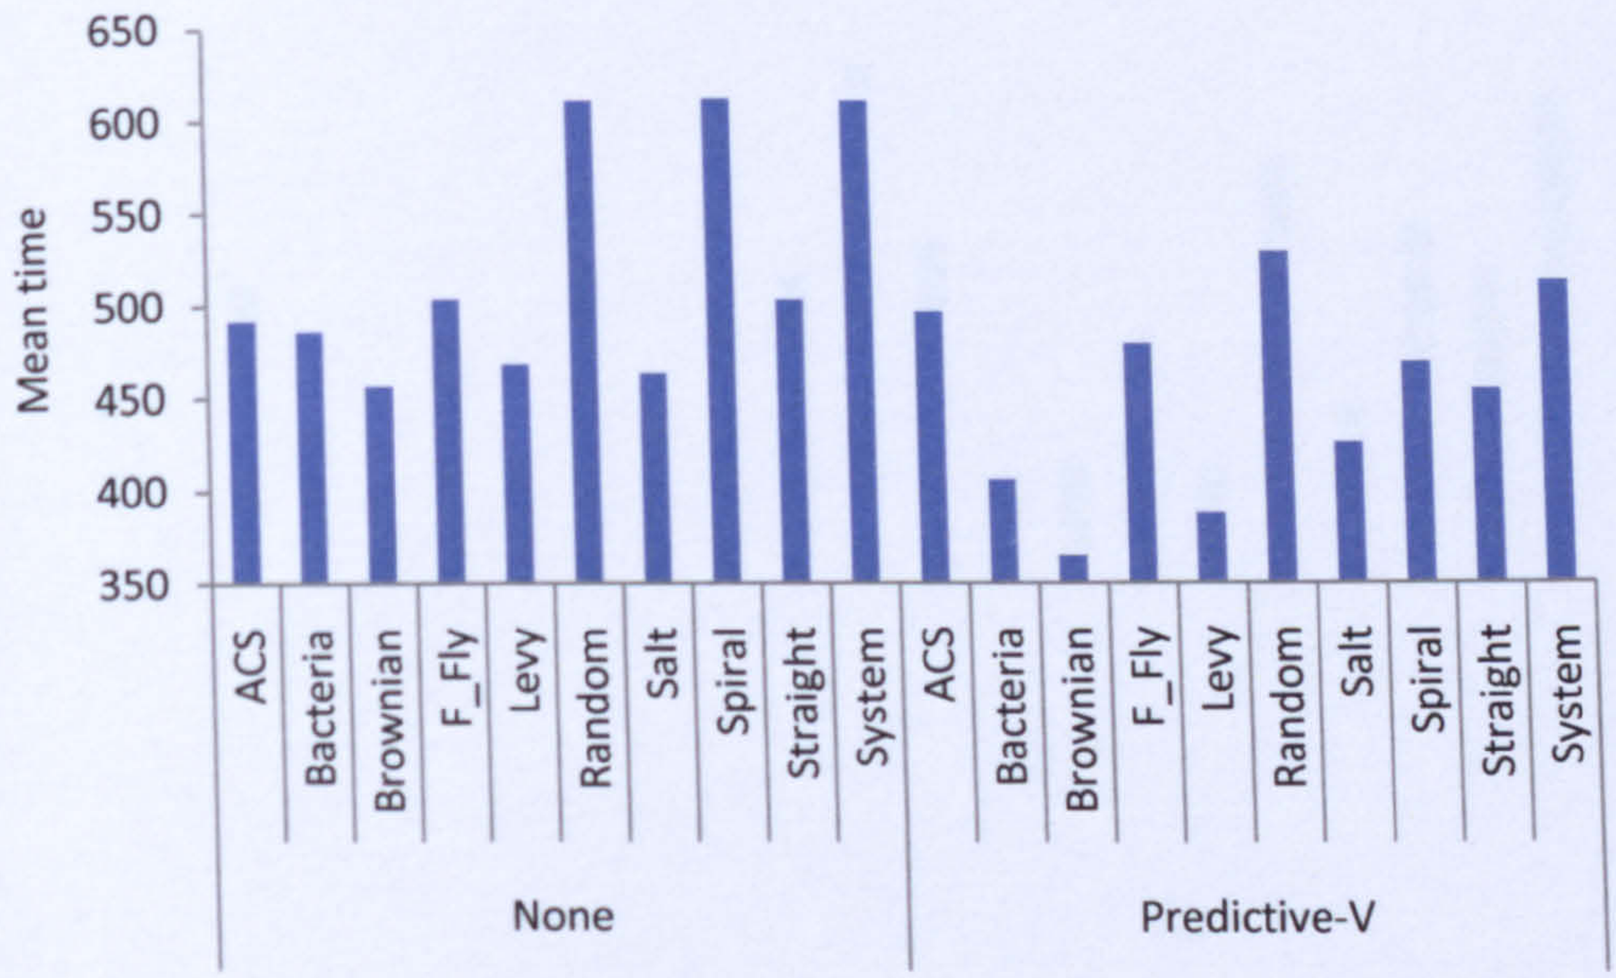

Figure 7 - 8. Mean time for 20 agents against 4 determinstic targets travelling at 2.0 agent speed where temporary comunication failure occurs in 8 agents.

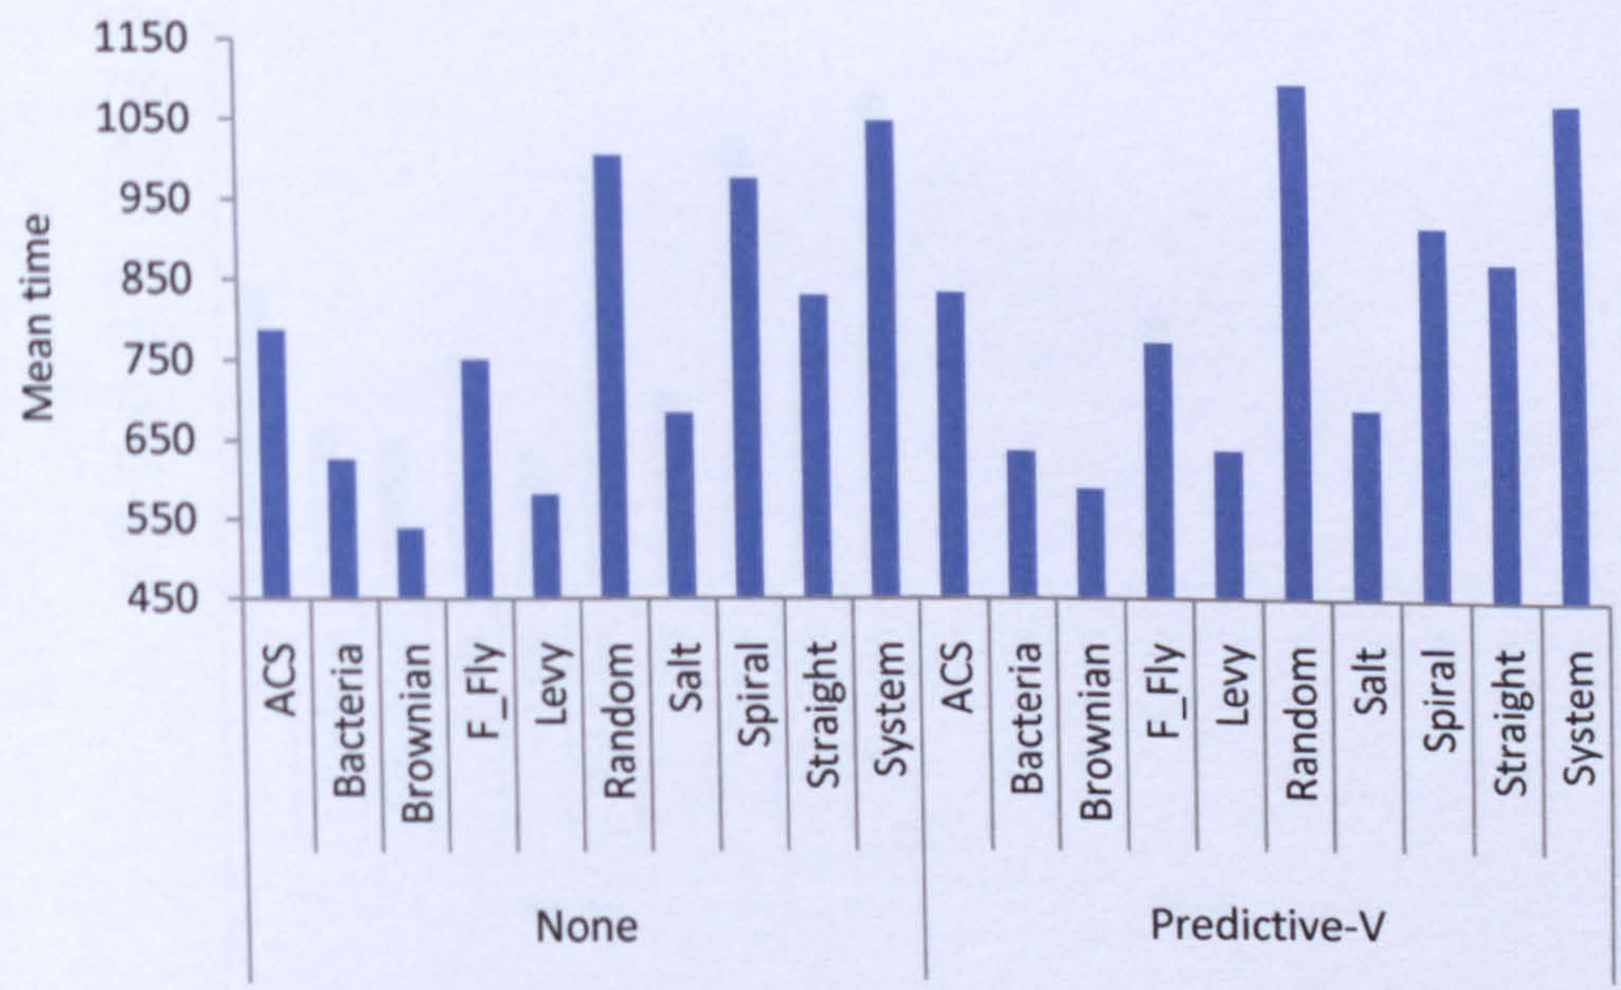

Figure 7 - 9. Mean time for 20 agents against 4 determinstic targets travelling at 0.5 agent speed where temporary comunication failure occurs in 12 agents.

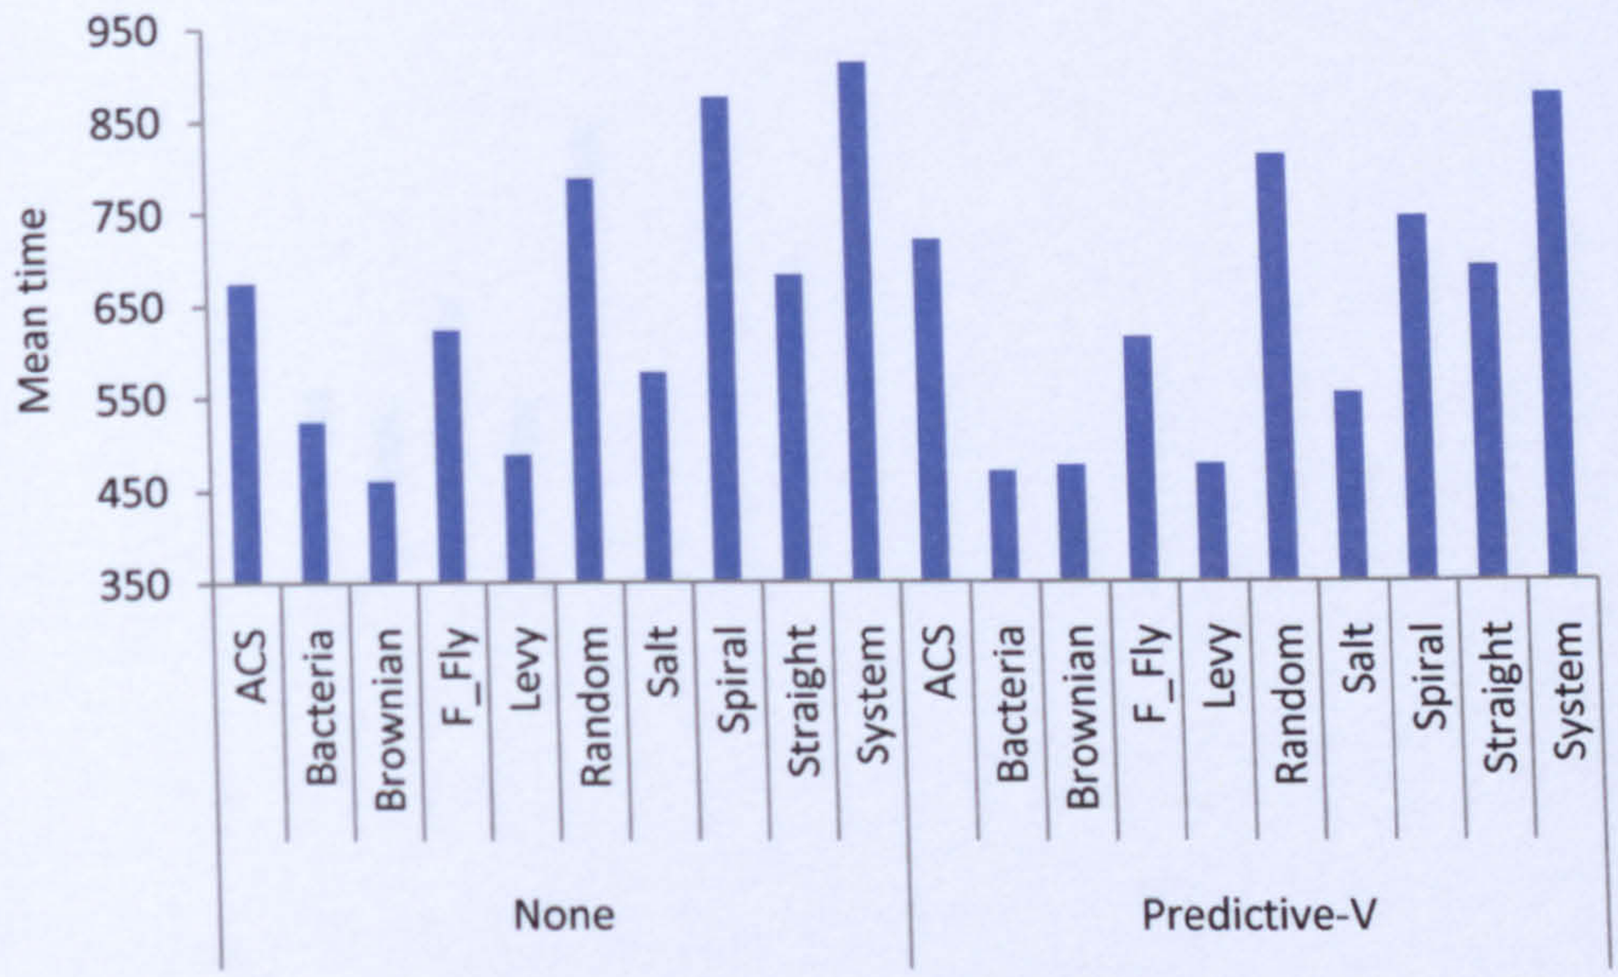

Figure 7 - 10. Mean time for 20 agents against 4 determinstic targets travelling at 1.0 agent speed where temporary comunication failure occurs in 12 agents.

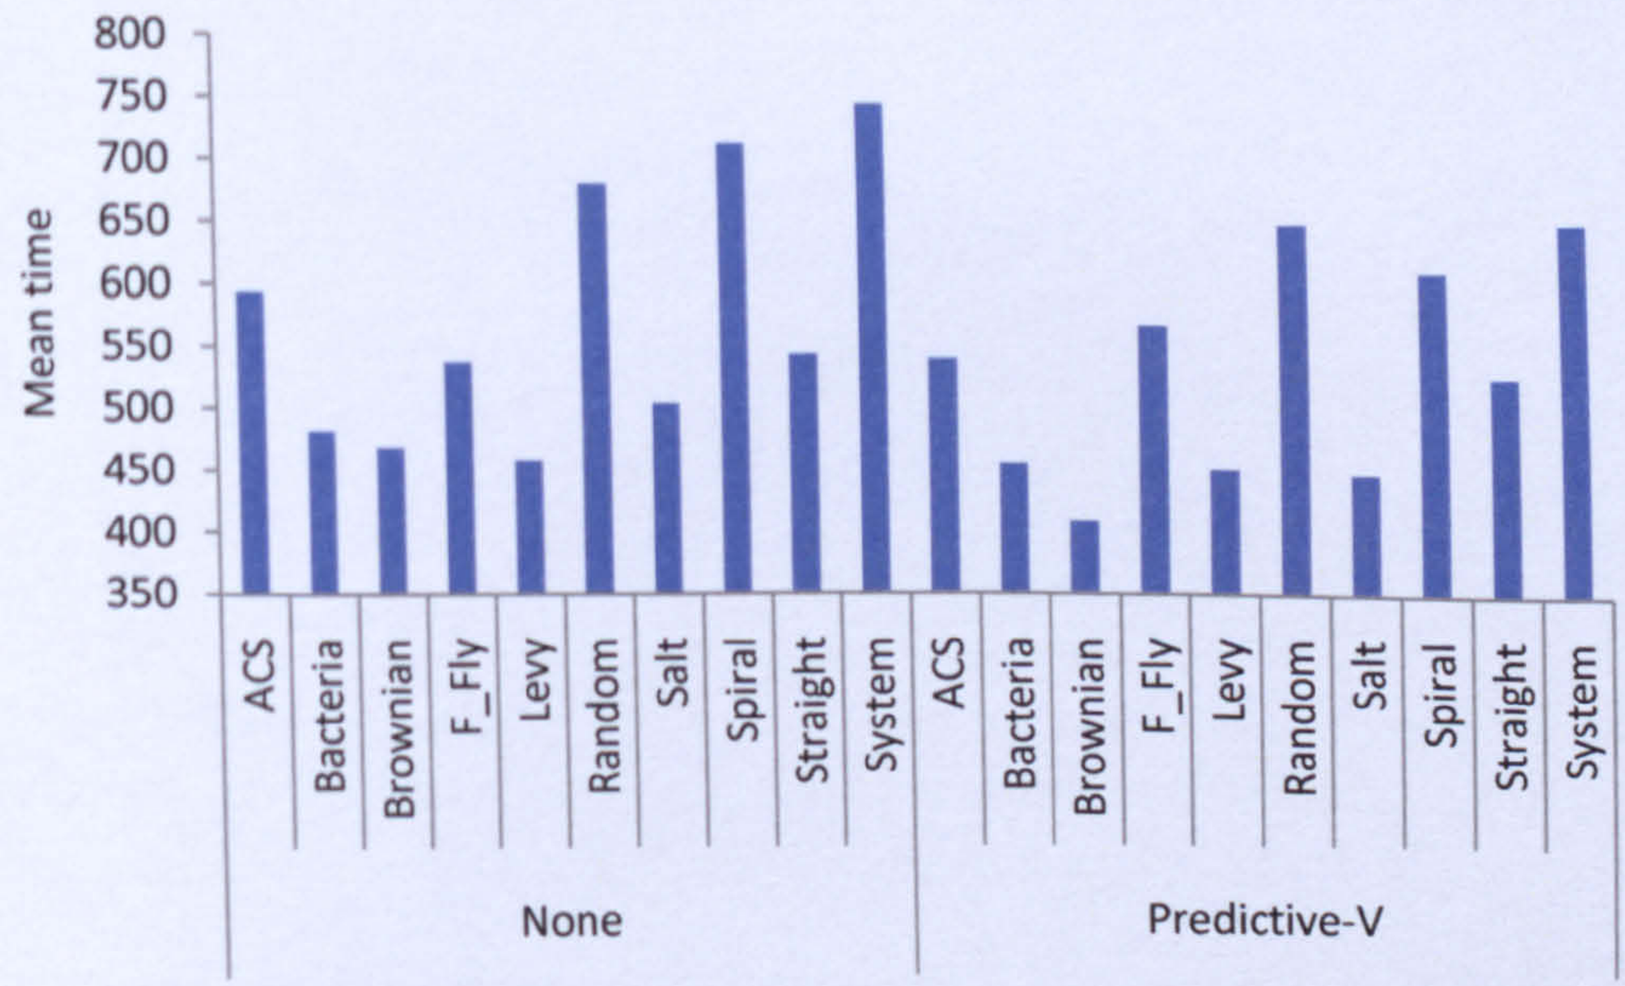

Figure 7 - 11. Mean time for 20 agents against 4 determinstic targets travelling at 1.5 agent speed where temporary comunication failure occurs in 12 agents.

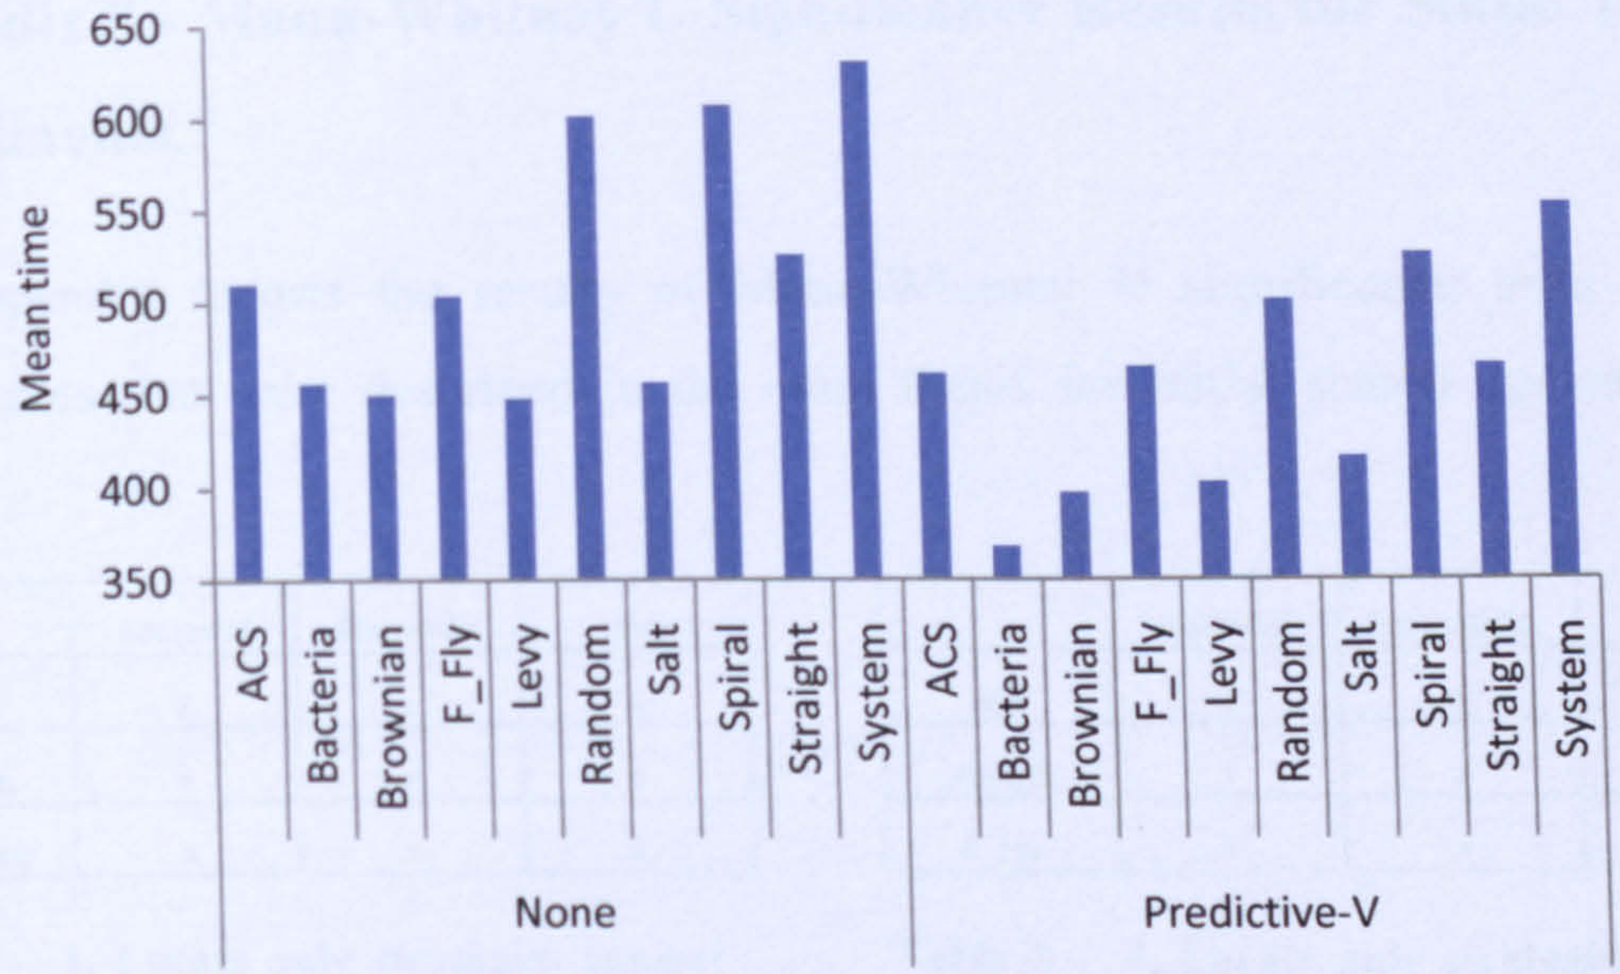

Figure 7 - 12. Mean time for 20 agents against 4 determinstic targets travelling at 2.0 agent speed where temporary comunication failure occurs in 12 agents.

Appendix 8 - Mann-Whitney U Significance Results for Static Target Experiments.

This appendix reports the results of Mann-Whitney U significance tests for all experiments that were described in the main thesis for initial search against static targets.

	Bacterial	Brownian	Levy
Spiral	x	x	x
Random	x	x	x
Systematic	x	x	x

Table 8 - 1. Locate only strategies against static uniformly distributed targets (5 agents against 4 targets).

	Bacterial	Brownian	Levy
Spiral	x	x	x
Random	x	x	x
Systematic	x	x	x

Table 8 - 2. Locate only strategies against static uniformly distributed targets (5 agents against 20 targets).

	Bacterial	Systematic	Spiral
Straight	x	x	x
F_Fly	x	x	x
ACS	x	x	x

Table 8 - 3. Locate only strategies against static uniformly distributed targets (20 agents against 4 targets).

	Bacterial	Brownian	Levy
Straight	x	x	x
ACS	x	x	x
Random	x	x	x

Table 8 - 4. Locate only strategies against static uniformly distributed targets (20 agents against 20 targets).

	Bacterial	Systematic	Brownian
ACS	x	x	x
Straight	x	x	x
F_Fly	x	x	x

Table 8 - 5. Locate only strategies against static clustered targets (5 agents against 4 targets).

	Levy	Brownian	Bacterial
Spiral	x	x	x
Random	x	x	x
F_Fly	x	x	x

Table 8 - 6. Locate only strategies against static clustered targets (5 agents against 20 targets).

	Spiral	Random	Systematic
Salt	x	x	x
Levy	x	x	x
Brownian	x	x	x

Table 8 - 7. Locate only strategies against static clustered targets (20 agents against 4 targets).

	Bacterial	Systematic	Brownian
ACS	x	x	x
Straight	x	x	x
F_Fly	x	x	x

Table 8 - 8. Locate only strategies against static clustered targets (20 agents against 20 targets).

	Bacterial	Brownian	Levy
Spiral	x	x	x
Random	x	x	x
Straight	x	x	x

Table 8 - 9. Locate-and-destroy strategies against static uniformly distributed targets (5 agents against 4 targets).

	Bacterial	Brownian	Levy
Spiral	x	x	x
Systematic	x	x	x
Random	x	x	x

Table 8 - 10. Locate-and-destroy strategies against static uniformly distributed targets (5 agents against 20 targets).

	Bacterial	Systematic	Random
ACS	x	x	x
Straight	x	x	x
Levy	x	x	x

Table 8 - 11. Locate-and-destroy strategies against static uniformly distributed targets (20 agents against 4 targets).

	Bacterial	Brownian	Systematic
ACS	x	x	x
Straight	x	x	x
Random	x	x	x

Table 8 - 12. Locate-and-destroy strategies against static uniformly distributed targets (20 agents against 20 targets).

	Random	Straight	Systematic
ACS	x	x	x
Saltatory	x	x	x
Levy	x	x	x

Table 8 - 13. Locate-and-destroy strategies against static clustered targets (5 agents against 4 targets).

	Bacterial	Brownian	Systematic
ACS	x	x	x
Spiral	x	x	x
F_Fly	x	x	x

Table 8 - 14. Locate-and-destroy strategies against static clustered targets (5 agents against 20 targets).

	Random	Systematic	Spiral
ACS	x	x	x
Brownian	x	x	x
Levy	x	x	x

Table 8 - 15. Locate-and-destroy strategies against static clustered targets (20 agents against 4 targets).

	Systematic	Random	Spiral
ACS	x	x	x
Levy	x	x	x
Brownian	x	x	x

Table 8 - 16. Locate-and-destroy strategies against static clustered targets (5 agents against 4 targets).

Appendix 9 - Mann-Whitney U Significance Results for Dynamic Target Experiments.

This appendix reports the results of Mann-Whitney U significance tests for all experiments that were described in the main thesis for initial search against dynamic targets.

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 1. Search strategies against deterministic targets travelling at 1.0 times maximum agent speed (10 agents against 1 target).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 2. Search strategies against deterministic targets travelling at 1.1 times maximum agent speed (20 agents against 1 target).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 3. Search strategies against deterministic targets travelling at 1.0 times maximum agent speed (10 agents against 4 targets).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 4. Search strategies against deterministic targets travelling at 0.9 times maximum agent speed (20 agents against 4 targets).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 5. Search strategies against random targets travelling at 1.3 times maximum agent speed (10 agents against 1 target).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 6. Search strategies against random targets travelling at 1.5 times maximum agent speed (20 agents against 1 target).

	Systematic	Random	Spiral
Brownian	x	x	-
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 7. Search strategies against random targets travelling at 1.3 times maximum agent speed (10 agents against 4 targets).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 8. Search strategies against random targets travelling at 1.3 times maximum agent speed (20 agents against 4 targets).

	Systematic	Random	Bacterial
Straight	x	x	-
Levy	x	x	-
Saltatory	x	x	-

Table 9 - 9. Search strategies against Brownian motion targets travelling at 1.0 times maximum agent speed (10 agents against 1 target).

	Bacterial	Systematic	Brownian
ACS	x	x	x
Straight	x	x	x
Spiral	x	x	x

Table 9 - 10. Search strategies against Brownian motion targets travelling at 1.1 times maximum agent speed (20 agents against 1 target).

	Random	Systematic	Spiral
Levy	x	x	x
Brownian	x	-	-
Saltatory	x	-	-

Table 9 - 11. Search strategies against Brownian motion targets travelling at 1.0 times maximum agent speed (10 agents against 4 targets).

	Systematic	Random	Bacterial
Levy	x	x	x
ACS	x	x	x
Straight	x	x	x

Table 9 - 12. Search strategies against Brownian motion targets travelling at 1.0 times maximum agent speed (20 agents against 4 targets).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 13. Search strategies against fruit fly targets travelling at 1.4 times maximum agent speed (10 agents against 1 target).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 14. Search strategies against fruit fly targets travelling at 1.4 times maximum agent speed (20 agents against 1 target).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	-

Table 9 - 15. Search strategies against fruit fly targets travelling at 1.3 times maximum agent speed (10 agents against 4 targets).

	Systematic	Random	Spiral
Brownian	x	x	x
Levy	x	x	x
Bacterial	x	x	x

Table 9 - 16. Search strategies against fruit fly targets travelling at 1.3 times maximum agent speed (20 agents against 4 targets).

Appendix 10 - Mann-Whitney U Significance Results for Efficiency Experiments.

This appendix reports the results of Mann-Whitney U significance tests for all experiments that were described in the main thesis for efficiency.

	F_Fly	Levy	Brownian
Bacterial	x	x	x
Saltatory	x	x	x
Systematic	x	x	x

Table 10 - 1. Efficiency tests.

Appendix 11 - Mann-Whitney U Significance Results for Predictive Memory System Experiments.

This appendix reports the results of Mann-Whitney U significance tests for all experiments that were described in the main thesis for PMS strategies during permanent occlusion, temporary occlusion and both permanent and temporary communications failures.

	Crosswind	None
Saccade	x	x
Predictive-V	x	

Table 11 - 1. Reacquisition search strategies against Brownian motion targets where no occlusion or communication failures occur.

	Predictive-V	Saccade
Direct	x	x
Predictive-C	x	x

Table 11 - 2. Reacquisition search strategies against deterministic targets where no occlusion or communication failures occur.

	Crosswind	Saccade
Predictive-V	x	x
None	x	x

Table 11 - 3. Reacquisition search strategies against fruit fly movement targets where no occlusion or communication failures occur.

	Crosswind	None
Direct	x	x
Predictive-V	x	x

Table 11 - 4. Reacquisition search strategies against Brownian motion targets where permanent occlusion occurs.

	Direct	None
Saccade	x	x
Predictive-V	x	x

Table 11 - 5. Reacquisition search strategies against deterministic targets where permanent occlusion occurs.

	None	Predictive-V
Saccade	x	x
Crosswind	x	x

Table 11 - 6. Reacquisition search strategies against fruit fly targets where permanent occlusion occurs.

	Crosswind	None
Direct	x	x
Predictive-V	x	x

Table 11 - 7. Reacquisition search strategies against Brownian motion targets where temporary occlusion occurs.

	None	Direct
Predictive-C	x	x
V_Shaped	x	x

Table 11 - 8. Reacquisition search strategies against determinsitic targets where temporary occlusion occurs.

	Crosswind	None
Predictive-C	x	x
Direct	x	x

Table 11 - 9. Reacquisition search strategies against fruit fly targets where temporary occlusion occurs.

	Predictive-V against None
0.5 Permanent	-
1.0 Permanent	-
1.5 Permanent	-
2.0 Permanent	x
0.5 Temporary	-
1.0 Temporary	-
1.5 Temporary	x
2.0 Temporary	x

Table 11 - 10. Reacquisition search strategies against deterministic targets where permanent and temporary communications failures occurs in 8 agents.