

A Hormone Inspired System for On-line Adaptation in Swarm Robotic Systems

James Stuart Wilson

DOCTOR OF PHILOSOPHY

UNIVERSITY OF YORK
ELECTRONIC ENGINEERING

October, 2019

Abstract

Individual robots, while providing the opportunity to develop a bespoke and specialised system, suffer in terms of performance when it comes to executing a large number of concurrent tasks. In some cases it is possible to drastically increase the speed of task execution by adding more agents to a system, however this comes at a cost. By mass producing relatively simple robots, costs can be kept low while still gaining the benefit of large scale multi-tasking. This approach sits at the core of swarm robotics.

Robot swarms excel in tasks that rely heavily on their ability to multi-task, rather than applications that require bespoke actuation. Swarm suited tasks include: exploration, transportation or operation in dangerous environments. Swarms are particularly suited to hazardous environments due to the inherent expendability that comes with having multiple, decentralised agents. However, due to the variance in the environments a swarm may explore and their need to remain decentralised, a level of adaptability is required of them that can't be provided before a task begins. Methods of novel hormone-inspired robotic control are proposed in this thesis, offering solutions to these problems. These hormone inspired systems, or virtual hormones, provide an on-line method for adaptation that operates while a task is executed. These virtual hormones respond to environmental interactions. Then, through a mixture of decay and stimulant, provide values that grant contextually relevant information to individual robots. These values can then be used in decision making regarding parameters and behavioural changes.

The hormone inspired systems presented in this thesis are found to be effective in mid-task adaptation, allowing robots to improve their effectiveness with minimal user interaction. It is also found that it is possible to deploy amalgamations of multiple hormone systems, controlling robots at multiple levels, enabling swarms to achieve strong, energy-efficient, performance.

Contents

List of Figures	19
List of Tables	23
Acknowledgements	25
1 Introduction	29
1.1 Hormone Inspired Systems For Robotic Control	30
1.2 Thesis Contributions	30
1.3 Thesis Structure	31
2 Background	35
2.1 Introduction	35
2.2 Features of Robot Swarms	35
2.2.1 Robot Swarms Vs Multi-Robot Systems	36
2.2.2 Applications For Robot Swarms	37
2.2.3 Task Allocation	42
2.2.4 Energy Consumption	45
2.3 Biomimicry In Robotics	52
2.3.1 Bio-Inspired Design	52
2.3.2 Genetic Algorithms	55
2.3.3 Hormone-Like Systems	58
2.4 Robotic Adaptation	71
2.4.1 Behavioural Adaptation	71
2.4.2 Morphological Adaptation	71
2.5 Summary	75

3	Virtual Hormones for Explicit Control	77
3.1	Introduction	77
3.2	Hormone Controller	78
3.2.1	Motor Control	78
3.2.2	Behaviour Switch	81
3.2.3	System 1 - Single Parameter Adaptation	83
3.2.4	System 2 - Multiple Parameter Adaptation	85
3.2.5	Base Line Test	85
3.3	Experiments	86
3.3.1	Experimental Set up	86
3.3.2	Parameter Values	89
3.3.3	Comparison	91
3.4	Analysis	96
3.4.1	Consistency	96
3.4.2	Sensitivity	99
3.4.3	Scalability	102
3.5	Chapter Summary	104
4	Virtual Hormones for Energy Efficient Task Allocation	107
4.1	Introduction	107
4.2	Hormone-Inspired Systems	108
4.2.1	Comparability Of Proposed System Vs Other Hormone Inspired Systems	109
4.3	Controllers	109
4.3.1	System 1: Hormone Arbitration	110
4.3.2	System 2: An Offline Optimised System	115
4.4	Experiments	115
4.4.1	Energy Efficiency	116
4.4.2	Simulation	116
4.4.3	Investigating The Effect Of Swarm Density	119
4.4.4	Investigating The Performance Of Large Swarms In A Highly Dynamic Environment	122
4.5	Chapter Summary	125

5	Virtual Hormones for Task Allocation by Self Identifying Traits	129
5.1	Introduction	129
5.2	Hormone-Inspired Behaviour Arbitration System (HIBAS)	131
5.3	Creating a Heterogeneous Swarm	136
5.3.1	Wheel Type: Wood Environment	136
5.3.2	Wheel Type: Grass Environment	137
5.3.3	Wheel Type: Conceptual	137
5.4	Experiments	138
5.4.1	Parameters	139
5.4.2	Simulation	139
5.4.3	Comparison System	140
5.4.4	Experiment 1 - Swarm Preference Between Two Environments with Static Features	140
5.4.5	Experiment 2 - Swarm Preference Between Two Environments with Mid Experiment Terrain Swap	140
5.4.6	Experiment 3 - Swarm Preference Between Three Environments With Static Features	141
5.5	Results	143
5.5.1	Experiment 1 - Swarm Preference Between Two Environments With Static Features	143
5.5.2	Experiment 2 - Swarm Preference Between Two Environments with Mid Experiment Terrain Swap	143
5.5.3	Experiment 3 - Swarm Preference Between Three Environments With Static Features	146
5.6	Analysis	146
5.6.1	Consistency	146
5.6.2	Scalability	149
5.7	Chapter Summary	152
6	A Multi-Hormone System for Arbitrating Traits and Behaviours in Dynamic Environments	153
6.1	Introduction	153
6.2	HIBAS Implementation for Control of a Foraging System with Deviating Motor Speeds	155

6.2.1	Energy Characteristics of Psi Swarm Robot Hardware	155
6.2.2	Hormone Interaction With Motor Speed	156
6.2.3	Comparison Systems	161
6.2.4	Analysis Of Systems Highlighting The Need For Adaptation	162
6.2.5	Results	164
6.2.6	Section Summary	167
6.3	Introduction Of The Sleep Hormone To A Foraging Swarm	169
6.3.1	Preliminary Tests For Sleep Hormone In A Demand Lead Foraging Task	170
6.3.2	Combining the Sleep Hormone With The Speed Deviating System . . .	173
6.4	Introduction Of Environment Selection Hormones With Sleep and Speed Regulating Hormone Systems	176
6.4.1	Environmental Setup	177
6.4.2	Effect Of Demand On Environment Selection When The Speed Hormone Is Combined With Environmental Preference Hormone	177
6.4.3	System combining Sleep, Speed and Preference Hormones.	178
6.4.4	Scalability of Final Amalgamated Hormone System	183
6.4.5	Reintroduction Of Obstacles For Amalgamated System	185
6.5	Robustness Analysis - Modelling Wear Over Time To Test System Robustness .	187
6.5.1	Wheel Wear Results	187
6.6	Chapter Summary	189
7	Conclusions and Future Work	191
7.1	Thesis Summary	191
7.1.1	Chapter 2 - Background	191
7.1.2	Chapter 3 - Virtual Hormones For Explicit Control	191
7.1.3	Chapter 4 - Virtual Hormones For Energy Efficient Task Allocation . . .	192
7.1.4	Chapter 5 - Virtual Hormones For Task Allocation By Self Identifying Traits	192
7.1.5	Chapter 6 - A Multi-Hormone System For Arbitrating Traits and Behaviours in Dynamic Environments	192
7.2	Concluding Remarks	193
7.3	Further Work	194

List of Figures

2.1	Cocoro swarm performance mid task. This group of robots exhibit true swarm-like behaviour as they conduct aquatic data collection. The members of the swarm are labelled into groups, executing different stages of the task: (a) robots exploring the seafloor, (b) robots creating a communication chain, (c) floating deployment station. Mintchev et al. (2014)	41
2.2	This photo shows a group of simple aquatic robots performing a homing task on a real aquatic surface. This provides an example of the effective implementation of a swarm-like system for a real world task. Duarte et al. (2016)	42
2.3	This figure illustrated how desired assignments transition from null or incorrect allocations. Desired task assignments are given in percentages of dark grey, green (medium grey) and light grey. The illustration here shows the goals of each algorithm; allocating unassigned or incorrectly assigned robots to the desired task (or colour) distribution. McLurkin & Yamins (2005)	43
2.4	This figure illustrates the convergence properties of each of the network algorithms for task assignment experiments conducted in McLurkin & Yamins (2005). In the test of each algorithm, robots started from a random initial assignment and the normalised distance to the goal allocation was measured as the robot groups performed allocation over time. This experiment was conducted with 8 trial repeats for each algorithm. McLurkin & Yamins (2005)	44
2.5	Simulation screen shot depicting an octagonal area containing a small swarm of robots foraging for food items(shown as the small red squares within the environment). In the most east corner of the environment is a light source, used by the robots to align themselves and return to the nest site upon finding food items.(Liu et al. (2007))	46
2.6	This figure shows three graphs illustrating how the Total Energy System Consumed (TESC) can reduce as the number of robots increases in multiple item collection scenarios. The environment dimensions for each of these graphs are as follows: (a) 40x40, (b) 50x50 and (c) 60x60. (Palmieri et al. (2017))	48

2.7	This figure shows how for almost all systems of various robot quantities, the Total Energy System Consumed (TESC) will increase as the number of target items increases. The graphs each represent experiments conducted in environments of the following sizes: (a) 40x40, (b) 50x50 and (c) 60x60. (Palmieri et al. (2017))	49
2.8	Snapshots of ATLAS, a bipedal humanoid robot, in simulation, running at 2m/s. Kuindersma et al. (2016)	53
2.9	Illustration of BigDog, a quadruped robot designed to travel in extreme environmental conditions. Raibert et al. (2008)	53
2.10	Wall Climbing robot hanging on a climbing track. This robot was capable of reproducing the gaits of cockroaches and gekos in order to examine vertical surface locomotion. Clark et al. (2007)	54
2.11	Soft actuators formed in a hand structure for prosthetic use, creating a tool that can be used by children while minimising the risk of injury to themselves or others that comes with typical hard prosthetic hands. Fras & Althoefer (2018)	55
2.12	Illustration of on-line adaptation and learning concept. The illustration features a birthing clinic, robots school, playground (in which developed robots will perform tasks) and recycling plant which would allow for a full robot life-cycle, ending in the reuse of old robot parts and rebirth into a new, slightly more developed, robot. Eiben et al. (2013)	57
2.13	Example of a GA evolved robot taking advantage of evolutionary criteria and using gravitational potential energy to produce speed by falling. Lehman et al. (2018)	58
2.14	This figure displays the operation of pheromone based guidance: (a) An ant follows a BHD path by chance, (b) Both paths are followed with the same probability and (c) A larger number of ants follow the shorter path. Ioannidis et al. (2011)	59
2.15	Image of Virtual pheromone trailing using foot-bots as foragers and eye-bots as virtual pheromone trackers. Eye-bots can be seen forming a path over the foot-bots, demonstrating their ability to create a guiding trail. (Ducatelle et al. (2011))	61
2.16	Obstacle avoidance trajectories from virtual hormone controlled robots with varied hormone parameters for hormone decay, velocity and sensor range (λ , v & r respectively) Stradner et al. (2009)	64

- 2.17 The figure shows a swarm of robots combining to create a new quadrupedal robot with a more specialised morphology. The combined robots are now capable of traversing more complex terrains than individual wheeled robots. Kernbach et al. (2008) 65
- 2.18 Group of robots (blue circles) using morphogen gradients to create a ring formation from an initial random deployment. The paths of these robots are highlighted in green as they move from their initial random positions to form a ring shape relative to one another. Jin et al. (2009) 66
- 2.19 Screen shot of an experiment in which a swarm of robots form a modular snakebot to navigate over small ledges to areas previously inaccessible to individual robots. Kuyucu et al. (2013) 68
- 2.20 Transformable wheel capable of swapping between rolling and leg-like movement based on the context of the terrain it is being used to cross. Kim et al. (2013). 73
- 2.21 Adaptive quad-rotor drone capable of expanding and shrinking in size when required . This allows it to take advantage of the additional stability associated with wider placement of rotors, while still having the capability to fit through tight spaces when required. Zhao et al. (2017) 73
- 2.22 Photographs showing frames of each step in the DALER's walking method as the robot re-purposes its wings to locomote on the ground. Daler et al. (2013). 74
- 3.1 Illustration of the signals received by a central robot and how these signals will effect the robot's hormone values. In the image the robot in the centre is receiving the total hormone values (the sum of the left and right hormones) of the robots within line of sight. It can be seen that the robots to the left of the central robot transmit values that will then only effect the hormone value associated with the left side of the central robot and the robots to the right will effect the hormone value associated with the right side in a similar manner. . 80
- 3.2 The Illustration shows separation hormone values effecting the movement of three robots. The close proximity of robots 1 and 2 has resulted in a larger hormone value in the hormone cells facing one another. This subsequently results in the robots turning away from one another. Robot 2 turns away from robot 1 regardless of the fact robot 3 is to its left as the greater distance between the robots means that the stimulant to the hormone cell facing robot 3 is not as great as that of the hormone cell facing robot 1. Robot 3 however, will still turn away from robot 2 as it has no other robot effecting its left hormone cell. The resultant trajectories are marked with dashed lines headed with arrows. . . . 82

- 3.3 State Diagram for context aware hormone system. Systems 1 and 2 differ only by the type of threshold required for transitioning between state S_0 and S_4 . As the baseline system contains no hormone mechanism it does not use this state diagram at all. Instead robots avoid one another and move randomly until a non-robot obstacle is detected, at which point the individual detecting an obstacle begins mapping. 83
- 3.4 Simple Obstacle Environment: A small environment, this is the most basic in terms of object complexity. This environment should be easy to navigate and map, putting minimal strain on the system. 87
- 3.5 Maze Environment: A large environment. This is the most complex environment to navigate and was designed as a compilation of all the environments, creating a problem that should be more challenging to optimise for. 87
- 3.6 Corridor Environment: The difficulties posted by this small environment are the narrow hallways the swarm will be required to navigate. This environment should show that smaller values of D are beneficial to wall following in enclosed spaces. 88
- 3.7 Office Environment: Another large environment, however, this one is made up of several compartmental rooms. This environment should make it difficult for robots in the same room to disperse and will test the swarms ability to operate with minimal communication. 88
- 3.8 A-Test results comparing the total area mapped at the final tick of simulation for each environment tested. For each environment the parameter for separation threshold was modified, ranging between 0 and 500 in steps of 10, to observe the systems sensitivity to said parameter and the potential ranges that should be explored when creating a dynamic separation threshold. The parameter sweep used to produce this graph was also used to identify optimal values of separation threshold for each environment. 90
- 3.9 Image capture of the Office environment for each system at 1000 ticks showing the area mapped within the environments at this time stamp (indicated by black dots seen on the floor of the environment). Image (d) displays the results of the office environment at 1000 ticks, showing that the single adaptation system completes no mapping at this stage of the experiment while the fully adaptive method and the random method are able to begin mapping early. . . 92

3.10 Image capture of the Office environment for each system at 16000 ticks showing the area mapped within the environments at this time stamp (indicated by black dots seen on the floor of the environment). Image (d) displays the results of the office environment at 16000 ticks, showing that the single adaptation system is able to recover from the poor mapping quantity early in the experiment to create a strong performance alongside the fully adaptive method against the random system.	93
3.11 Box plots of results comparing the mapping performance by the final time step of the experiment of the system 1 controller (Single Adapt), the random controller and the system 2 controller (Full Adapt).	94
3.12 A-test scores for data sets at 100 tick intervals for 5 experiment repetitions, showing all measured ticks reporting in large difference in results.	97
3.13 A-test scores for data sets at 100 tick intervals for 50 experiment repetitions, showing no large differences, but majority medium difference across all time steps.	97
3.14 A-test scores for data sets at 100 tick intervals for 75 experiment repetitions, showing some small difference and some time steps approaching a medium difference.	98
3.15 A-test scores for data sets at 100 tick intervals for 100 experiment repetitions, showing all ticks reporting only a small difference, confirming statistical consistency for this number of repetitions.	98
3.16 A-test results for parameter D in the office environment sweeping through a range of parameter values for D between 0 and 600, incrementing in steps of 10.100	
3.17 A-test results for parameter D in the obstacle environment sweeping through a range of parameter values for D between 0 and 600, incrementing in steps of 10.100	
3.18 A-test results for parameter D in the maze environment sweeping through a range of parameter values for D between 0 and 600, incrementing in steps of 10.101	
3.19 A-test results for parameter D in the corridor environment sweeping through a range of parameter values for D between 0 and 600, incrementing in steps of 10.101	
3.20 Scalability tests in the obstacle environment showing a linear increase to the number of points mapped as the number of robots increases from 5 to 55 in steps of 5 robots.	102
3.21 Scalability tests in the maze environment showing a gradual increase in number of points mapped as swarm size increases, performance can be seen starting to plateau after 20 robots.	103

3.22 Scalability tests in the office environment showing performance increasing but with diminishing gains as more robots are added, similarly to the maze environment.	103
3.23 Scalability tests in the corridor environment showing performance increase initially but remaining at a similar level between 20 and 55 robots.	104
4.1 State machine for foraging hormone system illustrating state names and the relative hormone ratios required for transition.	112
4.2 Graph showing the hormone values of a single robot across a selected time period in single experimental trial. The sleep hormone can be seen to take a much larger value than the other two hormones. Implemented to create longer periods of sleep when the swarm is substantially cluttered. Additional behaviours such as successful collection of items or obstacle avoidance can be seen in spikes in the respective hormone values.	113
4.3 Graph showing the cumulative mean (red line) and confidence intervals (dashed lines) as the number of replications increases for the 20 robot hormone system used in the first set of experiments. A vertical line at $n = 21$ marks where the deviation reaches an acceptable point (less than 0.05).	117
4.4 Graph showing the cumulative mean (red line) and confidence intervals (dashed lines) as the number of replications increases for the 20 robot hormone system used in the second set of experiments. A vertical line at $n = 183$ marks where the deviation reaches an acceptable point (less than 0.05).	118
4.5 Simulated ARGoS environment for experiment 1. The grey area represents the nest site, the black spots mark the location of food. 20 robots can be seen in the nest area awaiting the beginning of the experiment.	119
4.6 Box plots showing energy efficiency in the simple environment used for experiment 1 for both the Hormone inspired system and the timer-based system. The graph shows the results for swarm sizes of 5, 10, 15 and 20. The results show that the Hormone-inspired system is capable of obtaining competitive results versus the GA optimised system without the Hormone-inspired system having received bespoke optimisation itself.	121
4.7 Simulated ARGoS environment for experiment 2. This environment is larger at 8mx15m and features a wall, removed at the 500 second point in the experiment to future diversify the arena properties.	122

- 4.8 Box plot results showing the energy efficiency for swarm sizes of 5 to 20 in the larger, more complex, experimental environment. Even in this more complex environment, the Hormone-inspired system is capable of obtaining competitive results versus the GA optimised system without the Hormone-inspired system having received bespoke optimisation itself. 126
- 5.1 Two robots, A and B, operating in environment 1. Robot A has successfully discovered a food item and as a results its H_1 value is increasing. Robot B has entered the environment and is under performing relative to Robot A as indicated by its lower H_1 percentage. These robots transmit H_1 values and as Robot B receives an H_1 greater than its own H_2 is stimulated, as Robot A receives an H_1 value lower than its own no such stimulation occurs and H_2 is left to decay. The net result is an encouragement for Robot B to change preference to environment 2 and for Robot A to continue operating in the same environment. 133
- 5.2 A graph displaying the dynamics in the hormone values used in an experiment in which a swarm of three different robot types had an option of three varied environments to forage in. Curved increases in hormone value seen in this graph typically represent the collection of an item, while more abrupt increases will be caused by other robots transmitting hormone values. 135
- 5.3 Wheel designed to specialise in the environment with a wooden floor. The wheel has a large diameter, creating a higher top speed as well as a lightweight, thin and hollow design to allow for quick acceleration. The wheel affixes to the robot via bolts attached through the three holes in the middle of the design. The wheel was designed in Autodesk Inventor Professional 2018. 137
- 5.4 Wheel designed to specialise in the environment with a grass floor. The spoke-like outer rim of the wheel provided traction for the wheel, preventing slippage when travelling over uneven surfaces. The substantially wider wheel also provided stability, preventing the robots from moving off course when encountering small bumps. The wheel affixes to the robot via bolts attached through the three holes in the middle of the design. The wheel was designed in Autodesk Inventor Professional 2018. 138
- 5.5 Screenshot of simulation environment used in experiment 1. Black dots represent food items, white ground represents wooden flooring, green represents grass and the grey area in the centre represents the nest site. Robots of different types (represented by red and black wheels either side of the robots) can be seen moving to and from respective environments, some carrying food items (represented by black rings above them). 141

- 5.6 Screenshot of simulation environment used in experiment 3. Black dots represent food items, white ground represents wooden flooring, green represents grass, red represents very rough terrain and the grey area represents the nest site. Robots of varying type (indicated by the presence of red, green and black wheels) can be seen moving to and from respective environments, some carrying food items (represented by black rings above them). 142
- 5.7 Graph showing the mean results across 150 trials in experiment 1. The performance of both the random and hormone-inspired system is shown in terms of correct categorisation and food items foraged. The small difference between food collection rate verses the large disparity in allocation type shows that, even given a small difference in collection capability, categorisation is still possible using the hormone system in this environment. 144
- 5.8 Graph showing the mean results across 150 trials in experiment 2. The performance of both the random and hormone-inspired system is shown in terms of correct categorisation and food items foraged. The large flip in correct categorisation percentage seen in the hormone system illustrates the point at which the environment types swapped. The hormone system is subsequently able to recover from this, eventually outperforming the random system once more. 145
- 5.9 Graph showing the mean results across 180 trials in experiment 3. The performance of both the random and hormone-inspired system is shown in terms of correct categorisation and food items foraged. With the introduction of the red environment there is a greater penalty to assigning robots to search an arena area they are not suited to. As a result, the effect of correct categorisation can be seen more clearly in the rate at which food is collected by the hormone system versus that of the random. 147
- 5.10 2 robot types - Cumulative mean and confidence intervals for experiment with two terrain types and time step requiring the largest number of replicates. The vertical line at $n = 150$ marks where the deviation reaches an acceptable value (less than 0.05). 148
- 5.11 3 robot types - Cumulative mean and confidence intervals for experiment with three terrain types and time step requiring the largest number of replicates. The vertical line at $n = 180$ marks where the deviation reaches an acceptable value (less than 0.05). 149

- 5.12 Correct categorisation percentages recorded 500 seconds into the experiment, at a variety of swarm sizes in an experiment containing two robot types and two terrain types to chose from. The horizontal line through the graph displays the expected median percentage performance for a system categorising at random between two options. 150
- 5.13 Correct categorisation percentages recorded 500 seconds into the experiment, at a variety of swarm sizes in an experiment containing three robot types and three terrain types to chose from. Statistical expected median percentage is marked with a dark solid line, the expected median percentage based off of previous results is marked with a dark dashed line. 151
- 6.1 Behavioural hierarchy for the behaviours investigated within this Thesis. . . . 154
- 6.2 Graph displaying the results of the power consumption of a Psi Swarm robot increasing motor speed gradually. 10 repetitions were taken for these results and a trend line has been fit to the mean of these results and is shown in red. The equation forming the trend line is shown in Equation 6.1. 156
- 6.3 Graph showing average food gathered per energy unit consumed in a swarm of 20 foraging robots across 50 trials. The speed showing the greatest food collection per energy consumed is marked at 35cm/s. 157
- 6.4 A graph displaying the dynamics in the hormone values used in an experiment in which a swarm with hormone regulated speeds conducted a foraging behaviour. The graph displays the hormone values of an individual robot within this swarm across 500 experimental ticks. 160
- 6.5 Screenshot of first simulated environment used in initial experimentation with the new speed regulating systems. Food items are shown as black circles in the white environment, puck robots can be seen waiting in the nest area (light grey).162
- 6.6 Screenshot of second simulated environment used in initial experimentation with the new speed regulating systems. Food items are shown as black circles in the white environment, puck robots can be seen waiting in the nest area (light grey). Obstacles creating corridors are illustrated in dark grey. 163
- 6.7 Results for the three systems tested in in environment 1. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 165

- 6.8 Box plot results for the three systems tested in in environment 2. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 168
- 6.9 State machine for foraging hormone system illustrating state names and the relative hormone ratios required for transition as first displayed in Chapter 4. H_S is now replaced by H_σ 171
- 6.10 Box plot results for the hormone inspired sleep system tested in environment 1. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 172
- 6.11 Box plot results for the hormone inspired sleep system tested ed in in environment 2. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 173
- 6.12 Results for the combined hormone sleep and speed system tested in both environments. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 174
- 6.13 Effect of item target values driving different demands in the speed hormone on the percentage of robots taking preference to their optimal environment. The categorisation system running with no speed hormone present is marked as 'no adapt'. Results show a fairly consistent percentage of categorisation relative to item target, indicating that the speed hormone has minimal effect on the categorisation abilities of the combined hormone system. 178
- 6.14 Box plot results for the hormone preference system tested in the environment containing two difference terrain types. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 179
- 6.15 Box plot results for the hormone preference system, combined with the sleep hormone system tested in the environment containing two difference terrain types. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 180

- 6.16 Box plot results for the hormone preference system, combined with the speed hormone system tested in the environment containing two difference terrain types. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 181
- 6.17 Box plot results for the hormone preference system combined with both the sleep and speed hormone system, tested in the environment containing two difference terrain types. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 182
- 6.18 Results for energy efficiency as the number of robots in the swarm increases from 12 to 60. The introduction of additional robots can be seen to decrease energy efficiency in an almost linear fashion. This is to be expected when increasing swarm size but confining operation to a confined environment. With more robots present traffic when collecting items will slowly build, making it more difficult for the swarm to collect items efficiently. 184
- 6.19 Test environment for hormone amalgamation containing two areas of different terrain, highlighted in white and green as well as funnel like obstacles designed to increase swarm density for a section of both terrains. 185
- 6.20 Results from the hormone amalgamated system in an environment with two directions to forage in, each with different terrain types and both containing obstacles. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%). 186
- 6.21 Graph showing the A-Test scores of the combined hormone system from trial to trial as Zeta increases in effect, degrading the motor system more quickly over time. The Zeta value can be seen starting to have a greater effect on the system as it approaches 100%. 188

List of Tables

2.1	Descriptions of various swarm robot behaviours commonly used in testing. . .	40
2.2	Summary Table for Section 2.2 Features of Robot Swarms	51
2.3	Summary Table for Section 2.3 Biomimicry in Robotics.	70
2.4	Summary Table for Section 2.4	75
3.1	Parameter values selected for hormone mapping system.	89
3.2	P values for comparisons between three controllers for each environment. Results that reject H_0 or H_1 are highlighted in bold.	95
4.1	Parameter values selected for hormone arbitration through the mathematical reasoning explained at the end of Section 4.3.1	114
4.2	System 2 parameters optimised via genetic algorithm across 30 generations for experiment 1. These results show that, as would be expected, the smallest group of robots require low levels of sleep with a large resistance to avoidance, allowing for continual operation. Groups larger than this required relatively low avoidance counters, allowing robots to return to the nest after brief collisions. In terms of sleep counter, times gradually increased with swarm size, reducing traffic outside of the nest to account for the increased population.	120
4.3	Wilcoxon rank sum tests comparing the two systems for each swarm size with results for energy efficiency taken from the first, simpler, experiment environment. Significant difference can be seen at all swarm sizes with the exception of 15.	121

4.4	System 2 parameters optimised via genetic algorithm across 30 generations for experiment 2. The second experiment shows a clearer gradient of values for avoidance counter. It is clear that robots get less resistance to collision as the swarm size increases, promoting the more frequent return to the nest of robots in more dense populations. Conversely, the sleep counter shows little difference within larger groups of robots, taking a value of 0 for the smallest. This implies that robots in the smallest group should never sleep to get the best performance. However, in groups of more than 5 robots they will sleep for approximately a fifth of the experiment time if they enter the avoidance state for long enough to return to the nest.	124
4.5	Wilcoxon rank sum tests comparing the GA optimised and Hormone based systems for each swarm size with results for energy efficiency taken from the second, more dynamic, experiment environment. A significant difference can be seen in all tested swarm sizes except 20.	124
5.1	Key for the symbols used in the hormone equations.	132
5.2	This Table lists the speeds recorded when each of the wheels were tested in hardware examples. The figures listed represent the average speed recorded across 10 repetitions.	136
5.3	Table listing the parameters used in all software experiments. Parameter symbol meanings are described in Table 5.1	138
6.1	Parameter values for the Return and Speed Hormones. Where $H_s(t)$ is the Speed Hormone, λ_s is the decay rate for the system, γ_{s1} is the weighting for the stimuli and γ_{s2} is the weighting for the inhibitor.	161
6.2	Environment 1: Wilcoxon rank sum tests comparing the three systems for the tested item collection targets between 10-150 in terms of energy efficiency. Significant differences (indicated by a p value of < 0.05) are highlighted in bold .	167
6.3	Environment 2: Wilcoxon rank sum tests comparing the three systems for the tested item collection targets between 10-150 in terms of energy efficiency. Significant differences (indicated by a p value of < 0.05) are highlighted in bold .	169
6.4	Parameter values for the hunger hormone and new sleep Hormone. Parameter descriptions can be seen alongside Equations 6.7, 6.8 & 6.9.	170
6.5	Wilcoxon rank sum tests comparing the combined hormone system with both the speed hormone system and the sleep hormone system, in both of the previously established environments. Tests were conducted for item collection targets between 10-150 in terms of energy efficiency. Significant differences (indicated by a p value of < 0.05) are highlighted in bold	176

-
- 6.6 Speed Coefficients for Heterogeneous Robot Wheels on Different Terrains. These values were used to simulate different capabilities amongst a heterogeneous swarm of robots. This created a requirement to assign robots to appropriate environments in order to obtain efficient performance from the swarm. . . . 177

Acknowledgements

I would like to start by thanking my parents who have not only provided the care and support that has given me the opportunity to pursue academia, but have also kept me highly motivated through their pride in my accomplishments and constant encouragement. None of this work would be possible without them.

Next, I would like to thank both of my PhD supervisors Jon Timmis and Andy Tyrrell for their support and advice throughout the past three years. They have helped me focus and develop my research and provided me with the chance to share my work both nationally and internationally on multiple occasions.

I would also like to thank my examiners Martin Trefzer and Sabine Hauert for their time and guidance in the last stages of this work.

Finally, I would like to thank my friends and colleagues in the department, Naomi, Richard, Kieran and James who have consistently given me good advice, assistance and criticism when required. Across the past few years you have all been there for me, not just at work but also in our downtime.

You have all made this experience the best it could be and I hope that my appreciation for your time and company is reflected in the dedication and hard work put into this thesis.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References.

Aspects of the work presented in this thesis have been published and are as follows:

- Wilson, J., Timmis, J. & Tyrrell, A. (2018), 'A hormone arbitration system for energy efficient foraging in robot swarms', in 'Annual Conference Towards Autonomous Robotic Systems ', pp. 305–316, Springer.
- Wilson, J., Timmis, J. & Tyrrell, A. (2018), 'A Hormone-Inspired Arbitration System For Self Identifying Abilities Amongst A Heterogeneous Robot Swarm', IEEE Symposium Series on Computational Intelligence (SSCI). pp. 843–850, IEEE.
- Wilson, J., Timmis, J. & Tyrrell, A. (2019), 'An Amalgamation of Hormone Inspired Arbitration Systems for Application in Robot Swarms', Applied Sciences, 9(17), 3524, MDPI.

Chapter 1

Introduction

Individual robots, while providing the opportunity to develop a bespoke and highly specialised system, suffer in terms of performance when it comes to executing a large number of tasks. By introducing additional agents to a system in which multiple tasks must be completed, it is reasonable to assume that the tasks as a whole would be completed faster. With the groups of robots operating in parallel with one another to increase the productivity of the system. However, the introduction of additional robots comes at a cost. More robots must be produced and, to abide by the same budget as the initial robot, the subsequent robots must be made more simple. Swarm robotics is an example of this process taken to its furthest state. A large quantity of robots with simple functionality are produced but, given the larger number, they are able to provide a powerful method for parallel task execution (Beni (2004)).

Producing a large quantity of robots that will be able to operate effectively in the long term comes with its own set of problems, namely the coordination of such a large group of agents. A simple solution to this would be to introduce a central coordinator which could then delegate tasks amongst the entire group of robots. However, this would contradict one of the key features of a robot swarm; a swarm should be able to perform successfully even if any given robot were to fail at any point during a task. This feature allows for swarms to be inherently robust and suitable for volatile tasks in constantly changing environments (Şahin (2005)).

Conversely, with a central computer or leader providing instruction, large groups of robots can gain the benefit of global knowledge, as each member of the group can contribute information to a centralised agent. In turn, this information can be used to create an efficient strategy for effective utilisation of individual robot capabilities. If executed properly this can be a very powerful methodology for multi-robot control. However, for this system to work, individual agents must remain connected to the central agent. Without the ability to react or coordinate themselves without this centralised agent, a lack of instructions from the leader may result in disconnected members becoming lost to the group. Lost swarm members may get locked into a repetitive task that is unhelpful to the goals of the multi-robot system or may simply be

unable to navigate back to a centre of operations, resulting in the permanent loss of the agent. In an attempt to avoid situations such as this, most work involving swarm robotics aims to allow individuals within a swarm to self regulate, creating an emergent coordination which requires no designated leader for the swarm to function effectively. The difficulty with this method comes when a robot must have some understanding of the context in which they are operating in. In a uniform environment with a well understood task, robots operation can be programmed discreetly with a user planning and optimising each step in their behaviour. However, with multiple agents operating in real environments, agents must have some of their own understanding to react to unforeseen events. This understanding can then be actualised in the alteration of behaviour or parameters.

Hormone inspired systems look to amend this problem by providing a computationally inexpensive manner for evaluating an environment over time, using this information to alter the temperament of individual robots amongst a swarm.

1.1 Hormone Inspired Systems For Robotic Control

Examples of hormone regulation can be seen throughout nature (HARANO et al. (2008); Watanabe & Kuczaj (2012)). In these examples they can be seen working as an alternative to neural signalling, which typically acts as a direct and deliberate control signal to an organisms limbs, instead operating over greater periods of time in response to stimuli. The changes produced from hormone signals can vary between substantial morphological changes (as seen in tadpoles Riddiford & Truman (1993)) to changes in how the organisms might interact at a behavioural level (such as levels of aggression Watanabe & Kuczaj (2012)).

Drawing from these natural examples of hormone function, virtual hormone systems take stimuli in the form of a robots interaction with the environment, measured via simple sensors, creating hormone values which rise as interaction increases and fall at a rate imposed by a decay value attached to the virtual hormone equation. By monitoring these values over time and comparing similar hormone equations to one another features such as success rate, swarm density and locomotive efficiency can be extrapolated. These features are discussed in much greater detail throughout Chapters 3-6.

1.2 Thesis Contributions

Chapter 3 By re-purposing a technique previously designed for obstacle avoidance a novel bio-inspired method for robot deployment and dispersion is proposed, creating an effective method for mapping complex environments with no prior knowledge of environmental features.

Chapter 4 presents a method of hormone inspired behaviour state regulation not previously explored. While other studies have explored the use of hormones to define the boundary between two behaviour states, these have had to include virtual pheromones and pre-defined thresholds. While this system shows an original use in which the relative balance between multiple hormones generate adaptive thresholds capable of successfully reacting to a greater range of environmental dynamics.

Chapter 5 introduces a system of robotic self evaluation not previously studied. Using the transmission of hormone values from robot to robot within a swarm along with various environmental stimuli robots are able to create preference for a particular environment, using a computationally inexpensive method of ranking their performance against other robots encountered during a task.

Chapter 6 first presents a novel system of speed regulation, using a hormone system to directly control motor speeds during demand lead following. Later in the chapter, a first case combination of multiple hormone types is introduced. The combination includes hormones for speed regulation, behaviour state control and environmental preference. The simplicity in combining these systems is shown and an argument is made for the development of multi-hormone systems for real world swarm robot applications.

1.3 Thesis Structure

This thesis concentrates on the design, innovation and implementation of multiple virtual hormone types, looking at different levels of behavioural control and the benefits each level provides to a swarm robot system. Through the examination and experimentation with multiple virtual hormone systems, this thesis is guided by the following general hypothesis.

Hypothesis: *A swarm robotic system can obtain a greater efficiency or effectiveness against a comparison technique through the implementation of a hormone inspired system. Hormone inspired systems will help agents within the swarm adapt over time, without prior knowledge of the environment properties. Adaptation provided by the hormone systems will regulate either robot features or behaviour states.*

The content in the chapters of this thesis are summarised as follows.

Chapter 2 This chapter explores the background work to the thesis. Examining past literature to provide context to the experimental work it precedes. This literature includes sections which discuss the fundamentals of swarm robotics, establishing what makes a group of robots truly swarms like, with examples of previous implementations and potential applications, a background to other bio-inspired systems, placing the innovations

in swarm and hormone based system within the context of a larger field and the various methods of adaptation available to robotic systems along with their potential shortcomings.

The content within each of these sections is then evaluated and a summary provides instruction on which elements may be applicable to the content of each experimental chapter.

Chapter 3 The first experimental chapter, this examines the potential of a virtual hormone system capable of directly controlling the motors of a robot. The virtual hormone in this example is stimulated by a combination of transmitted hormone values from other swarm members and the distance at which they are transmitted. The goal of the system was to increase the mapping capabilities of the swarm by finding appropriate dispersion's after deployment and maintaining effective distance from other swarm members for strong mapping capabilities in a variety of complex environments. The chapter concludes that virtual hormone systems can be effectively implemented for this purpose. However, it is noted that hormone systems are not ideal for actuator control as, implementing a hormone system for the direct control of anything more complex than a wheel based system would be heavily time consuming and would be achieved more effectively with traditional control methods.

Chapter 4 This chapter develops a new hormone inspired behavioural arbitration system to regulate the sleep states of a swarm of foraging robots. Sleep states have been shown to increase the energy efficiency of item collection while foraging as they allow for de-cluttering in environments by momentarily preventing less successful robots from participating in the foraging task. The experiments in this chapter compare the virtual hormone regulation with a genetic algorithm optimised sleep system, which selects a single sleep time for each robot as they return to the nest. The chapter concludes that the variability in sleep time afforded by the hormone system makes it at least as effective a method as the optimised case in the simple environment while also not requiring the same level of information as the system it is compared against. In the more dynamic environment, the hormone system consistently outperforms the optimised system which has to select a compromised value for before and after the change in the environment, highlighting one of the weaknesses of a system tuned by genetic algorithm.

Chapter 5 This chapter continues to investigate hormone inspired behavioural arbitration. However, rather than coordinating states, the hormone system presented in this chapter looks to create preference amongst a heterogeneous swarm. Allowing robots to select an environment they are best suited for working in without having exact knowledge of the environments they are presented with or their own capabilities. The chapter concludes that substantial improvements can be made compared to a random system when it

comes to correct allocation percentage when given the option of two environments. When given the option of three, increases can still be made, though not as large due to the difficulty of the problem.

Chapter 6 This chapter combines all of the levels of hormone control presented in Chapters 3-5 to identify if the positives each type of system provide can be compiled into one, powerful, virtual hormone system without significant drawback. To begin this chapter a new wheel-motor controlling hormone system suitable for a foraging task is developed to replace the system worked on in Chapter 3. This new system improved the energy efficiency of a swarm by fluctuating the sleep based on a specified demand, keeping energy consumption low but also attempting to retain a good collection rate verses the stated demand. Once this is tested, the systems are combined, testing and identifying the potential negatives as each new system is added. The chapter concludes that the full combination of the hormone system experiences minimal negative interference between hormone types, producing the best combination of results of the systems tested.

Chapter 7 This chapter concludes the Thesis, providing a summary of the findings of each chapter and returns to the general hypothesis, examining it's validity and providing examples of future work.

Chapter 2

Background

2.1 Introduction

The information provided within this chapter introduces the fields of swarm robotics, robotic adaptation and biomimicry. This provides a basis to the understanding of each field's fundamental principles. Additionally, the chapter gives an insight to the shortcomings of each subject and recent advancements that seek to reduce issues that have been identified within these fields. By analysing the previous advancements of each area, the chapter will highlight features that should be considered when designing swarm robotic systems with elements of adaptation and biomimicry.

Through the review of these fields, this chapter will supply the reader with the fundamental knowledge required to adequately interpret the systems developed as a part of this thesis.

2.2 Features of Robot Swarms

Robot swarms attempt to provide a solution to complex robotic tasks, not through intricate and complex designs of individual robots, but through the use of numerous robots, each with basic features. Swarms of robots are capable of creating robust systems, in part due to their inherent redundancy. In an ideal robot swarm, should any one robot fail, the effect to the entire system's performance should be negligible. The ideal robot swarm will have no system wide vulnerabilities or dependency on an individual (Winfield & Nembrini (2006)), allowing the rest of the swarm to continue as normal in the event of a fault, performing the tasks required of the collective. Furthermore, as a collective, the swarm can perform multiple tasks and functions in parallel, providing a large benefit to the efficiency of exploration and logistical tasks over that of a single, complex robot. However, a swarm of robots is not purely defined by a large group of robots operating at the same time. There are distinctions that need to be made in terms of decentralisation, individual capability and simplicity in design.

The next section discusses these stipulations and elaborates on the difference between a true swarm and a multi-robot system.

2.2.1 Robot Swarms Vs Multi-Robot Systems

Şahin (2005) provides a set of criteria for a system to be defined as a swarm of robots. These criteria are as follows:

1. The ability to autonomously interact with the environment as individual members.
2. Groups of robots should be at a minimum of 10-20 robots in size.
3. Robots within groups should be relatively homogeneous with groups of robots that contain more heterogeneity being less swarm-like.
4. Robots making up a swarm should be incapable or inefficient as individuals, only capable of performing complex actions as a group.
5. Coordination throughout the swarm should be distributed, with members having only local communications and sensing available to them.

To elaborate on these points, the homogeneity referenced in point 3 can be interpreted as the requirement for each robot within a swarm to have a similar ability to execute any simple task presented to the swarm. It is possible for heterogeneity amongst a group of robots to avoid reducing its swarm-like nature, if individual specialities only improve an individuals ability to complete a task, but do not give the agent the capability to perform a task which differs from those available to other swarm members. If small individual differences did in fact reduce the swarm-like qualities of a system, a realistic swarm of robots would be at risk of becoming less swarm-like over time, as actuator degradation or other damage experienced over time create a difference in capability amongst a swarm. As such, the homogeneity in point 3 must refer to homogeneity as a lack of bespoke ability. For example, disallowing a sub-group of robots in a swarm the ability to lift an item or having a select few swarm agents with the ability to fly.

This idea of speciality extends to point 4. Robots should only be incapable or inefficient relative to an individual robot with a complex bespoke designed for a task. Swarm robots should be able to perform simple tasks individually, without having such a complex design as to be indispensable. When it comes to more complex tasks, it is acceptable for individual robots to be incapable of execution. However, with the help of additional members said task should be performable. For example, an individual robot may not be able to move a large, heavy object. But, with the assistance of a few additional robots the swarm should be able to complete the task, rather than the swarm being incapable of performing the task at all.

To address point 5, the distribution of a swarm prevents the reliance on any individual swarm agent. No robot within the swarm should be fully dependant on another to function. This

means that the swarm is able to continue operation should any member of the swarm be damaged or otherwise compromised. With this in mind, general communication between robots, or the sharing of knowledge through a network would still be considered swarm-like. With the points 3,4 and 5 interpreted as above, the criteria defined here align with the hormone inspired system implementations presented in later experimental sections. Thus, the systems presented in this thesis can be considered to be swarm-like.

These give a set of features that can be used to separate swarms of robots from multi-robot systems. While these constraints mean that individuals amongst a swarm are very simple in design, the uniformity of each robot means that typically, any member of the swarm is capable of swapping function with any other member, allowing for greater flexibility in tasks. A multi-robot system may be well suited to forming a production line from a group of specialist robots, each capable of completing a single complex task. But such a system would struggle should something unexpected happen during the construction process. In such an unexpected instance, task allocation or roll-swapping may be difficult if not impossible due to the level of specialisation within the multi-robot system. Therefore, robot swarms in their relative simplicity provide a good platform for adaptation, reallocation and reconfiguration. Though the tasks they are instructed to enact should be more general and require minimal specialisation.

The versatility of a swarm of robots comes in part from their potential fault tolerance, with redundancy amongst the large group of robots allowing the possibility for multiple failures during the performance of a task. The concept of this swarm fault tolerance has been refuted by Bjercknes & Winfield (2013) pointing out exceptions to this tolerance. They highlight a susceptibility to failure as swarm size increases, in cases where swarms are not aware of faults within their own members. However, there have been several proposed methods for fault detection within robot swarms including the use of a behavioural feature vector system (O’Keeffe et al. (2017)), artificial immune systems (Ismail et al. (2015)) or visual indication (Christensen et al. (2009)). These fault detection systems, when implemented, address the issue of robot fault ignorance and allow additional upscaling of the swarm size.

2.2.2 Applications For Robot Swarms

Swarm robotics, though studied by many (with experiments including: McLurkin & Yamins (2005); Nam & Shell (2018); Jones et al. (2006); Liu et al. (2007)), is still in a developmental phase and most swarm robotic systems are used only in lab environments. It is very rare for a true swarm robot system to form a solution to a real world problem. At present, applications for swarms of robots are typically considered theoretically. Şahin (2005) suggests that robot swarms would be well suited to a multitude of tasks, including those involving: the monitoring of a large space, dangerous tasks where swarm members can afford to be expendable, tasks

that require scalability, calling on more or fewer robots to solve a problem, or tasks that require redundancy i.e. a robust communication network. A review conducted in Brambilla et al. (2013) furthers these points, suggesting four key categories for swarm behaviour currently studied in literature:

1. Spatially organising behaviours
2. Navigation.
3. Collective decision making
4. Other collective behaviours that do not fit into any one of the categories mentioned prior.

These collective behaviour categories, and the behaviours that comprise them, are captured in the section on testing applications that follows.

Testing Applications

When studying swarm robot systems there are several basic behaviours that can be used to test the proficiency of a robot swarm and are frequently used as an experimental foundation of a newly developed system. These behaviours, as identified by Navarro Oiza & Matía Espada (2013) and Brambilla et al. (2013) are shown in Table 2.1.

Behaviour	Description
Aggregation	The gathering of a robot swarm, typically a basic requirement for the implementation of a more complex task.
Dispersion	Effective distribution amongst a space. Used to create networks, explore environments, form a distributed sensor etc...
Pattern Formation	Producing a shape or morphology as a swarm. This can be used for task allocation or self repair.
Chain Formation	Robots position themselves to connect two points, these chains of robots can then be used as navigational guides. This process is not dissimilar to robot pheromone use, discussed further in Section 2.3.3.

Self Assembly	A method in which individual robots are able to connect to one another, forming a new morphology. The new configuration of robots should then allow the robot group to perform a task more effectively or perhaps perform something it was previously unable to do. Examples of this can be seen in Christensen et al. (2008) where a group of simulated robots are able to connect to one another, enabling them to traverse large gaps or steep inclines which were previously impassable as individuals.
Clustering and Assembly	Robots move objects scattered around an environment, for either the purpose of assembling a structure or simply to gather items to a predefined location. Item collection such as this has been used in the testing of multiple swarm behaviour implementations (Liu et al. (2006); Campo & Dorigo (2007); Krieger et al. (2000)) as the task involves several sub behaviours (searching, collecting, retrieving, etc.) and collected items can be easily abstracted to objects of value, energy sources or even people requiring rescue in disaster situations, allowing the task to be likened to real world examples.
Collective Moment	Within swarm robotics group coordination is important to achieve a task that requires large amounts of co-operation members of a swarm must be able to cohesively move without collision or significant interruption to direction.
Source Search/Collective Exploration	Swarms lend themselves well to search tasks due to their capability to explore in parallel. This behaviour acts well as a test of multiple levels of behaviour, typically requiring aggregation, dispersion, collective movement and potentially collective mapping.
Collective Transport of Objects	Another more complex behaviour, this task typically performed from a combination of all other previously mention behaviour types. The swarm may have to deliver objects that no single member could carry, or the large swarm size could be used to increase the rate at which objects are discovered and transported.
Task Allocation	With multiple robots and multiple tasks labour must be divided efficiently. In experimentation, the specifics of the tasks can be abstracted while testing and focusing on the methods of allocation instead. This is discussed in greater detail in Section 2.2.3.

Consensus Achievement	With no centralised unit making decisions for a swarm, if a swarm-wide decision needs to be made the individual agents must come to a consensus with one another. This is usually difficult due to the dynamics that can be expected while a swarm is operating; the best choice may be unclear or may change over time. As such there have been several explorations into networked decision making or decision making between multiple agents.
Fault Detection	Swarms executing a task for significant periods of time will experience wear and eventually faults. In order to reduce the effect of swarm agent degradation it can be beneficial for swarms to identify faulty members. This can be achieved by comparing the performance of swarm members or by looking for key features that might indicate faults. Examples of such fault detection can be seen in O’Keeffe et al. (2017) and Tarapore et al. (2019)
Group Size Regulation	This is the capability of regulating a group of robots to be a particular size. Too many agents in one location can reduce swarm performance due to heavy traffic (Lerman & Galstyan (2002)) and too few can lead to a task having insufficient resources to complete when required. Thus, finding the balance between these factors is important for the effective implementation of swarm systems.
Human Swarm Interaction	Due to the autonomy that is associated with a swarm of robots, they can be difficult to influence by an external agent i.e. a Human. Human-swarm interaction looks to understand how a human operator might be able to control a swarm or receive live performance information as the swarm executes its task.

Table 2.1: Descriptions of various swarm robot behaviours commonly used in testing.

Of the behaviours listed in Table 2.1, object clustering will be the most relevant to the contributions presented in later chapters. This is due to the fact that almost all of the experiments conducted with the systems proposed in this thesis involve foraging. Group size regulation also has some relevance to Chapters 4 and 6 which involve the implementation of a hormone sleep system, removing individual agents from a task for a period of time.

Amongst others, work involving foraging behaviours has included: providing a test bed for examining errors as robots cover more distance Buchanan et al. (2016), improving the efficiency of search behaviour Schroeder et al. (2017), exploring methods of task allocation, creating periods of inactivity when workload is low Charbonneau & Dornhaus (2015) and increasing the efficiency of motion, reducing congestion by letting swarm members rest Liu




Figure 2.1: Cocoro swarm performance mid task. This group of robots exhibit true swarm-like behaviour as they conduct aquatic data collection. The members of the swarm are labelled into groups, executing different stages of the task: (a) robots exploring the seafloor, (b) robots creating a communication chain, (c) floating deployment station. Mintchev et al. (2014)

et al. (2007). These have provided an insight into what makes a foraging swarm and will be used in later sections to aid in the construction of experiments.

Practical Applications

Though presently rare for swarm systems to be implemented to solve real world issues, there have been a small number of studies that attempt to design and implement a swarm of robots for practical use. The study in Schmickl et al. (2011) is an example of one such system that can be used to exhibit true swarm-like behaviour, the experiments conducted in this study look at aquatic data collection in a swarm of robotics with self-awareness (Shown mid operation in Figure 2.1). The proposed system uses bioinspired taxis behaviour to dictate the movement of the swarm and decentralised wireless communication to find quorum when deciding communally within the swarm if a task was completed. The project later went on to produce a swarm of hardware robots capable of performing such behaviour (Mintchev et al. (2014)).

Similarly Duarte et al. (2016) successfully performed a swarm-like environmental monitoring task using a set of simple robots. By using evolved controllers to perform a variety of tasks requiring self-organisation (homing, clustering, dispersion and area monitoring), the swarm was capable of performing tasks effectively, despite the unpredictable conditions associated with real aquatic surfaces (shown in Figure 2.2). While the number of robots used in the hardware experiment was small for a swarm (a maximum of ten robots were used)



Figure 2.2: This photo shows a group of simple aquatic robots performing a homing task on a real aquatic surface. This provides an example of the effective implementation of a swarm-like system for a real world task. Duarte et al. (2016)

the autonomous interaction, homogeneity, simplicity of individual design and the lack of a centralised point of instruction qualified the system as swarm-like by the definitions in Şahin (2005).

Outside of these examples most swarm systems are usually either simulated behaviours where experiments take place virtually or hardware experiments that only deal with an abstracted problem in a heavily controlled environment, testing the fundamentals of robot to robot interaction. At present, due to a combination of hardware availability and the algorithmic capabilities available for regulating swarms system behaviour, swarm robotic systems are almost at a stage ready for effective real world deployment. To make the leap into solving real world problems well-defined formal notation and verification techniques will need to be put in place, such as those proposed by Ribeiro et al. (2017). Once implemented these techniques will provide guarantees on the behaviour of swarms in complex environments, with controller code being automatically generated from mathematically verified definitions. At such a point, swarm robot systems will become a more attractive prospect for wide spread use. As such, the experiments shown in later sections have been developed to empirically prove that systems presented can provide utility to the field of swarm robotics, with testing taking place primarily in simulated environments.

2.2.3 Task Allocation

The size and flexibility of robot swarms comes with the issue of having to organise the members efficiently. The following section explores methods that offer solutions to task allocation. These methods will look at both task allocation performed co-operatively by an entire swarm, forming networks that can then assign and trade tasks explicitly, and task allocation performed without



Figure 2.3: This figure illustrated how desired assignments transition from null or incorrect allocations. Desired task assignments are given in percentages of dark grey, green (medium grey) and light grey. The illustration here shows the goals of each algorithm; allocating unassigned or incorrectly assigned robots to the desired task (or colour) distribution. McLurkin & Yamins (2005)

robot to robot communication, with the individuals that comprise a swarm making personal decisions on task acceptance. In this latter case, organisation becomes emergent from these individual decisions.

Task Assignment Through Swarm Networks

McLurkin & Yamins (2005) compares the performance of different network algorithms capable of broadcasting tasks to an entire swarm of robots. In the analysis presented the robots were capable of communicating with their immediate neighbours and the network algorithms under investigation were used to create desired distributions of robot modes (three states represented by the different colours shown in Figure 2.3).

The three network algorithms presented are as follows:

1. Extreme-Comm: An algorithm requiring a large amount of robot to robot communication, as messages are propagated through the swarm neighbour to neighbour seeking to deterministically converge to the correct task distribution.
2. Card-Dealer: This algorithm is much slower than the Extreme-Comm system but also requires very little memory and minimises inter-robot communication. The robots choice in task is educated by a system that selects the next task type to most closely approximate a target distribution.
3. Tree-Recolour: This algorithm creates a spanning tree of a network. Instructions are passed down the tree layer by layer as a gradient message, containing the information required to swap robots at each level from task to task.

This study showed that the Extreme-Comms algorithm was by far the fastest solution to the task allocation problem, though requiring a relatively large amount of memory to perform. The most consistent and reliable solution however, was the card-dealer's algorithm with the




Figure 2.4: This figure illustrates the convergence properties of each of the network algorithms for task assignment experiments conducted in McLurkin & Yamins (2005). In the test of each algorithm, robots started from a random initial assignment and the normalised distance to the goal allocation was measured as the robot groups performed allocation over time. This experiment was conducted with 8 trial repeats for each algorithm. McLurkin & Yamins (2005)

added benefit of requiring minimal memory to perform. Both algorithms are examples of heavily engineered systems, providing exhaustive solutions to task allocation through group made decisions in a decentralised manner.

This study shows that in order to exhaustively task allocate a swarm quickly, a large amount of memory is required. In the contributions shown in later experimental chapters, task allocation is attempted without a large amount of memory, instead the swarm will regulate tasks through behavioural control; finding the appropriate tasks through emergent means. In addition to this, in order to keep robot composition simple, many robot swarms have limited communicative abilities. It is therefore important to explore methods of task allocation that allow robots to decide on their own task with no knowledge other than the stigmergic information taken from interacting with their local environment.

Natural Task allocation

Task allocation can also be found in natural examples, it has been shown that some species of ants are able to adjust the number of individuals working on a task based on encounters, switching to appropriate tasks without knowledge of swarm size, number of ants involved in a task or making an informed assessment of which task should be performed Gordon & Mehdiabadi (1999). The study observing this behavioural allocation inspected a swarm of harvester ants, paying attention to the interaction between dedicated midden ants (ants sorting and carrying the colonies refuse) and ants engaged in other tasks. The study found that there was a positive correlation between the number of midden ants encountered and the likelihood of an ant swapping to perform midden work. This work has since been used to educate the construction of robotic systems, making effective task allocation methods for groups of robots by mimicking harvester ant behaviour (Zhao (2013)). The contributions featured in the later experimental chapters of this thesis have kept the idea of these natural examples in mind,

using a naturally inspired system as an approach for robot coordination. The work contributed in this thesis attempts to step away from explicit coordination and organisation and instead attempts to allow robots to settle into tasks in an emergent, reactive manner. This copying of natural systems to produce effective solutions is referred to as 'biomimicry', discussed in further detail in the next section.

2.2.4 Energy Consumption

Energy consumption is an interesting aspect of swarm robotics as, while multiple robots may be able to complete a task much more quickly than a single robot, they may not complete that task in an energy efficient manner. Aspects such as collisions and redundancy can mean that as swarm density increases, robots amongst the swarm use energy more wastefully.

Studies have been conducted to look at mechanisms in which the energy efficiency of robot redundancy can be reduced. Liu et al. (2007) investigates a mechanism in which labour can be divided amongst the swarm by invoking a low energy sleep state in robots after they have completed their task. These periods of sleep post task reduced the number of robots actively performing a task (in this case foraging for food items) to reduce environmental clutter. The length of the sleep period was adapted in each robot across the length of the experiments, with the successful retrieval of food reducing sleep time in successful robots and their peers, increasing the number of robots in the environment, and failure to retrieve food increasing the sleep period, reducing the number of robots in the environment (simulation screen shot shown in Figure 2.5). By adapting sleep times through individual success and social cues the system was found to be capable of effectively guiding the swarm to energy optimisation. The results from these experiments were promising, though the swarm sizes used did not exceed 10 meaning that the behaviour tested was only borderline swarm-like. These experiments formed a foundation for the system proposed in Wilson et al. (2018) in which sleep and foraging states were arbitrated by virtual hormones to attempt to further optimise energy efficiency. The hormone system and the benefits of behavioural adaptation it provided are discussed in detail in Chapters 4 and 5.

Other works such as Palmieri et al. (2017) have investigated the energy consumption of bio inspired robotic coordination procedures. The paper explored strategies of recruitment including a Firefly-based team strategy, a particle swarm optimisation technique and an artificial bee colony algorithm. The comparison performed experiments with each type of energy saving procedure in which a swarm would have to explore an environment, identify a number of targets and then perform a task requiring a specific number of robots. The fundamental features of each of these procedures are as follows:

1. Firefly-based team strategy: This strategy, initially proposed by Yang (2009), mimics the flashes exhibited naturally by fireflies to attract and coordinate with one another,




Figure 2.5: Simulation screen shot depicting an octagonal area containing a small swarm of robots foraging for food items (shown as the small red squares within the environment). In the most east corner of the environment is a light source, used by the robots to align themselves and return to the nest site upon finding food items. (Liu et al. (2007))

the flashes in this case being replaced with the broadcast of a short range signal from each robot. This signal was used to draw robots together. By combining this attraction with a dynamic leadership system (i.e. robots that discovered target locations were able to act as a primary coordinator and attract other swarm members) a method of formal coordination was produced that allowed the swarm to reach target locations as a group.

2. Particle swarm optimisation: This optimisation technique, proposed initially by Eberhart & Kennedy (1995), used the social observations of neighbouring robots velocity and position, along with a set of best previous positions and used this information to decide on the direction individual robots within the swarm would next move. This method of coordination is only called upon when a robot discovered a target location and requested recruits, while not recruited all of the swarm members explored the environment.
3. Artificial bee colony algorithm: As an algorithm proposed by Karaboga & Akay (2009) that took inspiration from honey bees, the behavioural elements of this algorithm are split into bee-like roles: scout, onlooker and employed bees. Scouts identify new food positions (or in the case of Palmieri et al. (2017)'s comparison, target locations) and update the stored list of available food coordinates. Employed bees travel to food locations and provide information about the target, while onlooker bees use this information to decide which position to travel to.

The experiments performed with these three strategies highlighted that when increasing the swarm size and arena size simultaneously while keeping the number of targets within the area the same, the energy consumption for every strategy decreased for almost every increment of scale between 10 robots and 60 (incrementing in steps of 5 robots, graph shown in Figure 2.6). Conversely, when increasing swarm size, arena size and the number of tasks, energy consumption typically increased as the parameters did (shown in Figure 2.7). This emphasises the need to understand the task and environment a swarm will be performing in and the importance of identifying the correct swarm size if a swarm is to operate in an efficient manner. It also strengthens the need for self-allocating systems such as Liu et al. (2007); Wilson et al. (2018) that can modulate the number of active robots performing a task. These self-regulating systems reduce the need for a centralised decision on swarm size and instead means that swarms perform multiple different tasks or in a series of environments, without needing to return and redeploy.

Even if the energy efficiency of a task is not of concern, for a swarm to execute tasks productively and over long periods of time, the recharging protocol of the swarm must be designed so as to not fall victim to the cluttering associated with high swarm density. Kernbach & Kernbach (2011) showed a method of coordination that attempted to reduce clutter within the swarm and prevent bottlenecks at charging stations to allow the system to move towards energy homeostasis. By broadcasting numerical information from robot to







Figure 2.6: This figure shows three graphs illustrating how the Total Energy System Consumed (TESC) can reduce as the number of robots increases in multiple item collection scenarios. The environment dimensions for each of these graphs are as follows: (a) 40x40, (b) 50x50 and (c) 60x60. (Palmieri et al. (2017))



(a)



(b)



(c)

Figure 2.7: This figure shows how for almost all systems of various robot quantities, the Total Energy System Consumed (TESC) will increase as the number of target items increases. The graphs each represent experiments conducted in environments of the following sizes: (a) 40x40, (b) 50x50 and (c) 60x60. (Palmieri et al. (2017))

robot, a network was created within the swarm that allowed for a collective calculation of energy consumption. The value for swarm energy consumption then provided individual high and low energy robots with the information required for behaviour switching. As a result, robots amongst the swarm were able to switch between the swarm's collective task and recharging when appropriate. The primary drawback to the proposed system was the configuration of the charging stations. If a robot were to run out of power just before reaching a charging station, said station would be blocked by a robot unable to move. If implemented in a system with a slightly more intelligent group of robots, this would most likely not be an issue. If the system was able to identify a blocked docking station, fully charged robots could be used to clear the path to the docking stations before beginning their primary tasks.

Subsection Title	Key Points
Robot Swarms Vs Multi-Robot Systems	<p>Five criteria points for robotic swarm definition provided by Şahin (2005):</p> <ol style="list-style-type: none"> 1. The ability to autonomously interact with the environment as individual members. 2. Groups of robots should be at a minimum of 10-20 robots in size. 3. Robots within groups should be relatively homogeneous with groups of robots that contain more heterogeneity being less swarm-like. 4. Robots making up a swarm should be incapable of inefficient as individuals, only capable of performing complex actions as a group 5. Coordination throughout the swarm should be distributed, with members having only local communications and sensing available to them. <p>The section also emphasises the versatility achieved by swarms through members performing tasks in parallel and redundancy amongst each task as the primary benefits of swarm-like implementations over a multi-robot systems.</p>
Applications For Robot Swarms	<p>Testing Applications: While there are many behaviours that can be used as a test bed for swarm robotic systems, the focus of the experiments presented in this thesis will be on foraging. Foraging behaviours provide a complex, yet easily explainable proving ground to test novel swarm systems.</p> <p>Practical Applications: Though there are not many real world implementations of true swarm robotic systems, experiments have shown the benefits large groups of decentralised robots offer, examples including: environmental exploration, data collection and logistical tasks.</p>
Task Allocation	<p>Task allocation plays an important role amongst swarms, whether due to social assignment or individual decisions. By observing the methods investigated within the section, it is clear that allocation can be achieved through both: heavily engineered methods and high level, naturally inspired behaviour changes. The variety of coordination methods explored within the section showed that the split of tasks amongst a swarm was typically the root of its effectiveness, no matter the method.</p>
Energy Consumption	<p>This subsection reviews several papers investigating the rate at which swarms use power to maximise the efficiency of energy depending on the type of task presented. Identifying a group of the correct size for a task was shown as a common theme amongst work looking for energy optimisation. Systems capable of automating this process mid task, were highlighted as a valuable asset to robot swarm performance.</p>

Table 2.2: Summary Table for Section 2.2 Features of Robot Swarms

2.3 Biomimicry In Robotics

While biomimicry has been used for centuries, the term itself was made popular in Benyus (1997) in which biomimicry is defined as 'Innovation inspired by nature'. The book discusses various projects in which natural discoveries have led to a design or made a design possible. Examples the book provides include the manufacture of crystal wind shields using SAM (self assembled monolayers), taking inspiration from organic templating seen in sea shells, the discovery of life saving drugs using a method referred to as 'biorational drug prospecting' which pertains to following various primates through jungles and observing their dietary habits to identify plants with potential medical value. More relevantly to robotics, the book discusses naturally inspired computation and the various issues at present in bridging the gap between how human brains function and how this can be achieved by a silicone based computer. The literature in this section attempts to demonstrate the designs and behaviours that have been produced through biomimicry and how beneficial they can be to the world of robots. These examples have been showcased in an attempt to demonstrate the reasoning behind pursuing ideas inspired by natural systems, as has been done to produce the systems contributed within this thesis.

2.3.1 Bio-Inspired Design

The developments discussed in Benyus (1997) are not the limit of Bio-inspired design for robotics. However, it is very common for nature to influence robotic design. This can clearly be seen in the development of humanoid-like robots, namely the ATLAS robot created by Boston Dynamics (illustrated running in Figure 2.8). This full scale humanoid robot was designed to operate dependably in varying terrains and it is hoped that it will be able to perform complex human-like tasks such as tool manipulation Kuindersma et al. (2016). Boston Dynamics have also created an autonomous quadruped robot named BigDog, illustrated in Figure 2.9, designed to travel in extreme environmental conditions, carrying loads of up to 154Kg Raibert et al. (2008). BigDog achieves animal-like mobility and is not just inspired by pack animals but manages to replace them as a more reliable, controllable and arguably ethical system.

Bio-Inspired methods of movement are not restricted to horizontal surfaces, Clark et al. (2007) analysed the leg coordination in cockroaches and Geckos and used the knowledge gained from these creatures in the testing of a two legged wall-scaling robot (Figure 2.10). The wall-scaling robot produced was capable of reproducing the examined gaits, creating a fast moving climber and also creating a testbed for analysing animal movements.

In addition to locomotion, bio-inspired actuators have been produced that allow for soft interaction. These manipulators have taken inspiration from mammalian tongues, elephant trunks and octopus arms to create tools capable of performing delicate tasks in congested environments (Trivedi et al. (2008)). Soft actuators also have strong applications within the




Figure 2.8: Snapshots of ATLAS, a bipedal humanoid robot, in simulation, running at 2m/s. Kuindersma et al. (2016)




Figure 2.9: Illustration of BigDog, a quadruped robot designed to travel in extreme environmental conditions. Raibert et al. (2008)




Figure 2.10: Wall Climbing robot hanging on a climbing track. This robot was capable of reproducing the gaits of cockroaches and gekos in order to examine vertical surface locomotion. Clark et al. (2007)




Figure 2.11: Soft actuators formed in a hand structure for prosthetic use, creating a tool that can be used by children while minimising the risk of injury to themselves or others that comes with typical hard prosthetic hands. Fras & Althoefer (2018)

development of prosthetic hands (example shown in Figure 2.11), providing a solution to children who are likely to injure themselves on traditional hard prosthetic limbs, while also remaining replaceable and easily scalable as the child grows (Fras & Althoefer (2018)).

2.3.2 Genetic Algorithms

Currently, one of the most common, naturally inspired, methods for computational learning is the genetic algorithm (GA). A process which copies natural selection by breeding features of a system based on a decided fitness. By allowing only the fittest to survive generation to generation, systems result in an increased performance in what the designer has selected as the 'fitness'. Davis (1991) provides an extensive review of various GA's discussing methodology, their accuracy and use to solve real world issues.

Off-line

The most common use of GA optimisation is off-line. The optimisation takes place in some simulated environment pre-task and parameters discovered from this optimisation can then be used when the task is enacted. Multitudes of this type of GA are presented in Davis (1991) which are still used presently to effectively solve problems. However, within the field of robotics, systems operating in dynamic environments will require a level of adaptability that an off-line GA cannot offer. For repetitive tasks this is not an issue but as the tasks robots are expected to perform become more complex, an off-line GA may only be capable of producing compromise parameters for the system. Additionally, simulation of the real world task may not be accurate enough, a problem identified as the 'reality gap' (Jakobi et al. (1995)). This is an issue that can be addressed by adding the appropriate amount of noise to a simulation. The type of noise required and to which parameter the noise should apply may not always be clear when modelling an environment. Thus, the parameters generated by an offline simulation when applied to a real system may not be truly optimal. This can be seen in Hecker (2015)

where parameters evolved with simulated error and evolved with no error are ported into a real system producing significantly disparate results.

On-line

Bredecche et al. (2009) addresses the lack of mid-task adaptation typically present in GA's by combining a genetic algorithm system with a swarm performing a live task. The paper studied the implementation of a virtual genome consisting of neural network weightings that explicitly controlled the robot. These genomes were then evolved using a (1+1)-Online evolutionary algorithm in which a singleton population was able to dynamically trail new genomes to attempt to defeat a current champion neural network weighting. The generational trials were made dynamic by changing the size of the parameters sets based on performance change. If a small change to parameters yielded improvement, the next generation would continue to make small changes. However, if the small change to parameter value produced minimal change, the next change would be much larger so as to avoid settling at a local optimum.

Similar work was conducted in Bredecche et al. (2012) though in this study a swarm of robots were introduced, capable of sharing genome information between one another through the use of the MEDEA algorithm. With this system implemented it was found that large populations were capable of emerging to a consensus, displaying unique behavioural strategies formed in a computationally lightweight manner.

In addition to on-line behavioural adaptation there have been proposals for the development of on-line morphological adaptation. Projects such as those detailed in the works of Eiben et al. (2013) have suggested providing swarms of robots with a hub or 'birthing clinic' in which they are capable of reproduction by either recycling the parts of parent robots or constructing new robots from scratch. These new robots would take on a combination of morphological traits from parent robots that had already been evaluated to be successful in their task performance and then be set free to perform the same task. Over time this process should evolve a progressively more effective system. However, the proposing paper highlights that for the evolution to be brought on-line, firstly the generated structures will need to be inspected for viability to succeed prior to their birth and a robot nursery will be required for 'infant' robots to learn the functions of their newly located sensors and actuators. This process is illustrated in Figure 2.12.

Issues presented by GA adaptation and optimisation

While GA's can be a powerful method for improving a system, they do suffer from some drawbacks. If the correct rules are not put in place, the wrong measure of fitness is requested or errors exist within a simulation, genetic algorithms are likely to exploit this system. Lehman et al. (2018) lists several examples of genetic algorithms behaving unexpectedly. These




Figure 2.12: Illustration of on-line adaptation and learning concept. The illustration features a birthing clinic, robots school, playground (in which developed robots will perform tasks) and recycling plant which would allow for a full robot life-cycle, ending in the reuse of old robot parts and rebirth into a new, slightly more developed, robot. Eiben et al. (2013)




Figure 2.13: Example of a GA evolved robot taking advantage of evolutionary criteria and using gravitational potential energy to produce speed by falling. Lehman et al. (2018)

examples include the attempted evolution of interesting locomotion methods. However, the fitness of the evolving creatures was defined as the ‘average ground velocity during its lifetime of ten simulated seconds.’ This caused the simulated creatures to prioritize tall structures that would simply fall over to reach high velocities, taking advantage of potential energy rather than creating legs or other propulsion methods, shown in Figure 2.13 (Sims (1994)).

The ability for GAs to exploit the underlying fitness metric can sometimes make it difficult to understand how or why the generated solution performs the way it does. This can be problematic when it comes to extending the system, requiring long periods of time relative to that of an engineered system when it comes to studying and understanding the evolution and the system's viability. In addition to this, behaviours evolved in simulation can be untrustworthy. If the simulation design suffers from the flaws previously mentioned, the evolved behaviour could be unexpected or dangerous when reproduced outside of the simulation and expecting reality to mirror the flawed virtual environment.

2.3.3 Hormone-Like Systems

In nature, hormones exist as a live adaptation technique in the form of chemical signalling providing behavioural changes. As stimuli reach cells or organs, hormone chemicals are produced and diffused throughout the anatomy of an individual. The build up and gradual decay of these hormones as they are metabolised gives an individual information on how frequently various stimuli are received. The balance and concentration of various hormones can then affect the behaviour of the individual to react beneficially to the stimuli they are receiving.

The human ‘fight or flight’ response described initially by Cannon (1929) provides an example of a hormone response. In the presence of perceived danger a neural signal is sent to the

adrenal glands which in turn emit the hormone epinephrine, as the epinephrine diffuses throughout the body organs react to it in different manners, the heart pumps harder, the liver produces more glucose and blood vessels dilate to allow more blood to flow to muscles. This prepares the body to run or fight by providing ample energy to the required muscles and thus effectively reacting to the presented stimuli.

Pheromones

Pheromones act very much like hormones though they are secreted and left to exist outside of an individual's body, they are dispersed according to stimuli and decay over time just as hormones would be dispersed internally to a system. Pheromones are then sensed by other individuals, conveying messages and information. By laying down pheromone chemical signals as they move, several species of insects are able to communicate through stigmergy to identify the best way to navigate Garnier et al. (2007). The identification of the shortest route can be implicitly discovered through a gradual build up of deposited pheromones. This is due to the fact that it takes less time to travel shorter routes to and from a food source and as a direct result of this, insects travel along the shorter route more frequently. With greater frequency, more pheromone is deposited and thus the optimal path is identified through the pheromone concentration, higher pheromone concentration indicating a shorter route. The deposited pheromones also decay over time meaning that the less travelled paths eventually have no trace of pheromone at all. This process is shown in Figure 2.14.




Figure 2.14: This figure displays the operation of pheromone based guidance: (a) An ant follows a BHD path by chance, (b) Both paths are followed with the same probability and (c) A larger number of ants follow the shorter path. Ioannidis et al. (2011)

The largest issue with the currently presented pheromone techniques is viability for hardware experimentation. Simulated systems can create pheromone path finding very easily as the pheromone can be abstracted virtually. Unfortunately, creating a synthetic pheromone for a swarm to track in a real world situation can be problematic. In most tasks placing physical

medium in an environment is deemed unacceptable, studies involving the guidance through the use of chemical sensors Devezza et al. (1994); Fujisawa et al. (2008) highlight the importance of an appropriate medium to be sensed and the difficulties of laying chemical trails in real environments with chemicals that could potentially interfere with robot sensors, leading a robot in the wrong direction.

An interesting attempt at trail following with a physical medium is shown in Schmickl & Crailsheim (2006) which appropriates a foraging example to complete a load moving task. In this study the object source is a mound of dirt that must be transported to a dump site (to extend the parallels to foraging this can be thought of as the nest). As the robots travelled and passed the dirt between one another small amounts of dirt were deposited on the ground. This dirt contrasted with the clean environment and was identifiable by the robots in the swarm. This allowed the dirt to act as a stigmergic marker, simultaneously acting as a pheromone abstraction and the object to be foraged. This method allowed robots to successfully identify short routes, sometimes multiple if two equidistant paths were presented, and provided an acceptable guiding trail given the nature of the task.

This dirt moving task is an exception to most modern studies, with the majority of pheromone based research moving towards potential virtual pheromones that remove the need for physical medium. One such example is presented in Ducatelle et al. (2011) on cooperative robot swarms. This method utilised a homogeneous swarm built from aerial eye-bots and ground based foot-bots. The pheromones in this example were emulated by the eye-bots as they formed an array above the environment in which the foot-bots would traverse. These eye-bots monitored the movement of the foot-bots, repositioning themselves to monitor the maximum number of foot-bots. Repositioning of the eye-bots was controlled through a simple algorithm; If there were more foot-bots observed in one direction than in others the eye-bot shifted towards that direction.

The instructions given to the foot-bots indicated which direction they should move next. The instructions were initially given at random, though as the eye-bots observed foot-bots avoiding obstacles, the system learned about the positioning of objects. As a result, a larger weighting was given to the probability of instructions being chosen that would lead foot-bots in directions away from known obstacles. In addition to these rules, foot-bots would ignore any instruction given to them that would cause them to return to a direction they had just been instructed to move away from in order to prioritise dispersion.

The combination of these behaviours created a pheromone-like response in which the foot-bots, without directly communicating with one another, found the shortest route to an objective and created a trail (in this case a trail of eye-bots) along the optimal path to their goal to inform other members of the swarm which direction would be the best to take.

Another example of robot pheromone implementation looks at a flying swarm capable of virtually depositing pheromone signals via short range communication (Hauert et al. (2008)).




Figure 2.15: Image of Virtual pheromone trailing using foot-bots as foragers and eye-bots as virtual pheromone trackers. Eye-bots can be seen forming a path over the foot-bots, demonstrating their ability to create a guiding trail. (Ducatelle et al. (2011))

The proposed work took inspiration from army ant foraging patterns. Exploring aerial vehicles began their task by using ant-like behaviour to navigate their environment. To do this aerial vehicles already deployed acted as nodes, creating communication links and transmitting a pheromone value. These pheromone values were used by ant-state robots to select their next target nodes to efficiently navigate the network.

Should an exploring robot break the communication link with its reference node, it was assumed that an area unoccupied by any robots within the swarm had been discovered. The exploring robot at this point entered the node-state for other robots to communicate with. Once an aerial vehicle swapped from an ant-state to a node-state it received an initial pheromone value. Interactions with adjacent node robots and passing robots exhibiting the ant behaviour increased the stored pheromone value. This resulted in well travelled routes experiencing large pheromone values, indicating an effective path, and infrequently travelled routes experiencing a deficiency in pheromone values. In addition to this, pheromone values would decrease over time, emulating evaporation seen in natural pheromone examples. Once a pheromone value had decreased to 0, the node robot storing the value would swap back to the ant behaviour, taking advantage of the fact that its current location was not in use by the rest of the swarm.

Through this searching method the system eventually identifies a target location, the target being a user in the case of the experiments featured in the work, and collapses the remaining redundant branches to the established network.

This system was found to efficiently produce communication networks in a robust and scalable manner. What's more, with no requirement for physical medium, as typically used for pheromone mimicking behaviours, the produced system was very viable for real world implementation.

The lack of physical medium in Hauert et al. (2008) made the system very similar to that of a virtual hormone system. Systems that likewise take advantage of chemical signalling seen in nature but do not place these markers stigmergically. The features of virtual hormones will be discussed in greater detail in the following section.

Virtual Hormones

Fundamentally virtual hormones are constructed from a decay and a stimulus. The decay reduces the level of the hormone value over time and the stimulus exists as a condition, which when met, increases the level of the hormone value. Stimuli can take the form of an interaction with the environment or another robot. Examples of these interactions might include discovering a point of interest, colliding with another robot or the presence of another robot's hormone value. Some systems might also use inhibitors, triggered by interactions in the same way as stimuli, but instead decreasing the hormone level. Hormone values constructed

in this manner are in accordance with the properties highlighted as intrinsic to hormone messages in Shen et al. (2000) i.e. hormone messages must:

1. Float in a distributed system with no particular destination.
2. Have a life time i.e. have a decaying element.
3. Be capable of triggering different actions at different receiving sites.

Early work related to virtual hormones for robotic control was conducted by Neal & Timmis (2003) with the development of an artificial endocrine system. The presented work established a combination of a neural network and artificial endocrine system for the control of a single robot in a office environment. The neural network was designed to control robot motors and avoid obstacles within the chosen environment, however the selected weightings for the neural network were found to be inadequate for effective obstacle avoidance. Specifically, as the environment varied, robots began to react poorly to larger amounts of clutter, unable to navigate the small spaces successfully with the standard avoidance distance established prior to the experiment. The endocrine system was introduced to regulate the robots avoidance distances, implemented as a 'gland' which would influence synapses within the neural network. The hormone levels produced by the virtual gland was designed to create a 'more expeditious retreat' as obstacles were encountered closer to the robot. This mechanism was found to produce behaviour in the robot which may be beneficial to exploration in dynamic environments. This work was furthered in Vargas et al. (2005) making the system more biologically plausible by introducing a hormone level repository and hormone production controller. These were respectively responsible for monitoring and secreting hormones when appropriate to achieve homeostasis (in this context regarding a stable equilibrium within the robot controller allowing for orderly control of a robot). This work demonstrated the process by which external environmental factors can be used to regulate a robotic system, demonstrating the foundation upon which most of the virtual hormone systems in this thesis will be built upon.

Explicit Control Virtual hormones and hormone-inspired systems have also previously been used more directly, controlling the motor functions of a single robot. In Stradner et al. (2009) the authors presented a method that modelled a robot as two cells controlling the left and right motor of a puck robot, each motor was driven by their own hormones H_r and H_l with wheel speed changing proportionately with the magnitude of hormone value. The hormones for each cell were stimulated by a proximity sensor and were capable of diffusing between cells, acting as an inhibitor to the apposing hormone when present in the neighbouring cell. With the hormone values corresponding to the wheel speeds on the respective sides of the robot, this produced an effective hormone controlled method for obstacle avoidance, displayed




Figure 2.16: Obstacle avoidance trajectories from virtual hormone controlled robots with varied hormone parameters for hormone decay, velocity and sensor range (λ , v & r respectively) Stradner et al. (2009)

in Figure 2.16. The study found that this system could be successfully implemented in hardware and could be well studied with an exhaustive parameter sweep at ‘reasonable computational cost’. Through the implementation of this first case example, evidence was given for the viability of a virtual hormone controlled robot system.

Similarly, Kernbach et al. (2008) produced a system which allowed hormones to regulate the movement of individual robots in a similar manner to Stradner et al. (2009). This work added additional function to the virtual hormone, using the same hormone to regulate an additional behaviour state. In this new behaviour state the robots conjoined to produce a larger, specialised morphology (depicted in Figure 2.17). The hormone in this state was re-purposed to create a hormone gradient, regulating the size of the newly formed conjoined organism. The proposed hormone system in this case, was to have its parameters stored within the virtual genome of the robot system and allowed to evolve as a control method for the system.




Figure 2.17: The figure shows a swarm of robots combining to create a new quadrupedal robot with a more specialised morphology. The combined robots are now capable of traversing more complex terrains than individual wheeled robots. Kernbach et al. (2008)




Figure 2.18: Group of robots (blue circles) using morphogen gradients to create a ring formation from an initial random deployment. The paths of these robots are highlighted in green as they move from their initial random positions to form a ring shape relative to one another. Jin et al. (2009)

Behavioural Control While explicit control over a robot is attainable with a virtual hormone system, it has been shown that virtual hormones are very effective at arbitrating behaviour states. Hormone-inspired controllers have been successfully implemented to adapt swarm morphology, giving context to environments via stimuli and then constructing appropriate formations Jin et al. (2009); Kuyucu et al. (2013). These studies show that hormone-inspired systems can be engineered to provide an effective, computationally inexpensive method for robot control.

Jin et al. (2009) shows a hormone inspired method for shaping formations of robot groups using a hormone diffusion model to create a morphogen gradient. The presented system treats each robot as a single cell, each only storing local information and interacting through virtual proteins and the morphogen gradient. The virtual proteins emitted from each robot act as a way to avoid colliding with one another, with robots moving away from proteins that are detected to be coming from other robots. The morphogen gradient is embedded with a predefined shape to regulate the dynamics of the system and allow them to form the desired structure. The resultant system was capable of consistently self organising into predefined shapes and was robust to system and environmental changes. Figure 2.18 shows a group of robots forming a ring (as predefined) using hormone gradients to guide them.

Kuyucu et al. (2013) presents a hormone method regulating a swarm of robots between two behaviour states. In the first state a swarm of homogeneous robots explore an environment

individually, quickly surveying the area they are restricted to and placing a virtual pheromone as they explore. Once they have laid the pheromone marks areas that have already been explored. In the second state, the robots form a 'Modular Snakebot' a robot formed from a chain of the individual swarm members, allowing for better stability, better traction and ability to operate in difficult terrain. The hormone in this method represents a value of "impatience" which is stimulated when the robot detects a high concentration of pheromone. Once an individual's hormone level reaches a predefined hormone threshold the robot seeks to dock with other robots that have already docked or are searching to dock. Once docked and formed, the 'Modular Snakebot' allows the swarm to escape the confinement of their current environment, passing over walls previously impassable by robot individuals, shown in Figure 2.19. In the new unexplored environment the snakebot detects the low value of pheromone and then disassembles, returning to the swarm searching state. This method provided a good solution to exploring an unknown environment, though it did heavily depend on a placed pheromone to convey information about the environment. As discussed earlier in this chapter, a fully applicable solution to placed pheromones for robotics has not yet been proposed outside of a heavily controlled environment. Which means, while this hormone regulated behaviour switching performed well in simulation, implementing the system to explore genuine unknown environments is not yet possible.

Other hormone-like behavioural control systems include the methods proposed by Neal & Timmis (2005); Vargas et al. (2005), developing an artificial endocrine system to regulate an individual robot. The studies present a method for modelling hormone secretion glands as they exist within a mammalian body. This method, used a combination of stimuli to induce the increase of hormone values, which in turn decayed over time. An implementation that aligns with the previously mentioned defining traits of hormone messages. Vargas et al. (2005) introduced the aforementioned artificial endocrine system to a simple robot, carefully selecting the stimuli to the hormone values to manage the internal states of the robot. In the case, the endocrine system managed the desire to recharge the batteries of the navigating robot (both in simulation and in hardware experimentation). The study identified an adaptive method for autonomous navigation in which a robot had a basic understanding of its energy limitations. Subsequently the robot was able to decide when it was appropriate to explore an environment and when it was time to return to a charging station.

This work has since been built upon and used to regulate the behaviour of autonomous sailing robots capable of recharging via solar panel (Sauzé et al. (2010)). Rather than conducting experiments in lab-like conditions with predefined charging point such as the experiments in Vargas et al. (2005), this endocrine system regulated the time of day that sailing robots would consume energy, using battery level and available sunlight as stimulating factors. The work found that it was possible to use an artificial endocrine system to improve the power usage from solar sources based on the availability of daylight. With the hormone system




Figure 2.19: Screen shot of an experiment in which a swarm of robots form a modular snakebot to navigate over small ledges to areas previously inaccessible to individual robots. Kuyucu et al. (2013)

implemented the sailing robots performing the majority of energy-expensive tasks during an abundance of light and reduced the likelihood of ever reaching a battery level of 0.

There are also examples of hormone-like behavioural control systems forming hybrid systems with existing methodologies. Timmis et al. (2010) demonstrates a behavioural controller formed from a combination of a neural network and an endocrine system used to manage the behavioural processes of a foraging swarm. The neural-hormone system was specifically used to modify the movement-based actions of individual robots. Examples of movements examined in the study include: obstacle avoidance, separation (the preference for robots to move away from other robots), cohesion (the preference for robots to move towards other robots), seeking items for foraging, seeking charging stations and seeking the bin (or nest site) to deposit found items. The movement of the robots was decided by summing the output values of the neural endocrine system. Subsequently, the behaviours with the largest current value according to the neural endocrine system would have the greatest effect on the desired movement. The results of the study found that simple neural-endocrine systems could be easily used for the development of foraging swarm systems. Though, as frequently found in the development of swarm systems, the performance within the proposed system began to suffer as more robots were introduced to the environment.

Subsection Title	Key Points
Bio-Inspired Design	<p>The design of robotics systems can be heavily guided by natural examples. This is clear in the common production of human-like, dog-like or other mammalian based foundations for robotic anatomy.</p> <p>Bio-Inspired design extends past pure structure and can also influence methods of locomotion, as seen in geko or insect like robots capable of travel upon vertical surfaces.</p> <p>Invertebrates have also provided guided the design of new forms of actuators. Leading to the development of soft robotic grippers with applications in surgery or as components in prosthetic limbs.</p>
Genetic algorithms	<p>Genetic algorithms can be used as a method for innovating or optimising a system. This is typically achieved off-line in simulation prior to a task, though it is also possible to create a system capable of synthetically 'breeding' and exchanging genomes based on fitness.</p> <p>While GA's can provide a powerful option for improving a system, the unpredictability and heavy handed changes associated with evolved designs and behaviours could prove to be a cause of concern, leading to poor performance in tasks or the presentation of abnormal behaviours, especially if a system is allowed to adapt during a task with no human intervention.</p>
Hormone-Like Systems	<p>The defining traits of hormone messages were established as:</p> <ol style="list-style-type: none"> <li data-bbox="531 1182 1324 1249">1. Hormone messages must float in a distributed system with no particular destination. <li data-bbox="531 1279 1174 1308">2. Hormone messages must have a finite life time. <li data-bbox="531 1337 1324 1404">3. Hormone messages must be capable of triggering different actions at different receiving sites. <p>Differences between general virtual hormones and pheromones were distinguished. The most noteworthy of the differences being that pheromones exist outside of an individuals body, secreted or left in a location, to be sensed by another individual. Moreover, works in which pheromones had been successfully used where highlighted, though it was pointed out that most pheromone systems are fundamentally flawed by the fact that placing physical medium amongst a working environment is generally required.</p> <p>As well as pheromone systems, examples of virtual hormone systems for the direct control of motors and for behavioural coordination where given. These emphasised the advantage virtual hormones provide in the adaptability they give a system.</p>

Table 2.3: Summary Table for Section 2.3 Biomimicry in Robotics.

2.4 Robotic Adaptation

2.4.1 Behavioural Adaptation

Robotic behavioural adaptation can be constituted as a robotic system that modifies its behaviour based on environmental interaction to enhance its performance in a task. This description gives a broad scope for what might be classified as behavioural adaptation, most of the systems described in Section 2.2.3 on task allocation and Section 2.2.4 on energy consumption can be fitted within this category.

Due to the breadth of examples that exist within the field of behavioural adaptation, the existing literature reviewed within this chapter will be limited to Sections 2.2.3 and 2.2.4 or highlighted as they are discussed in future sections.

The experimental work presented in later chapters will focus on adaptation within task allocation and energy consumption as they are representative of the broader scope of behavioural adaptation. Adapting energy consumption (which can be achieved through task allocation) creates a complex challenge which is suitable for optimisation and has a grounding in real world applications. Observing energy consumption also give the opportunity to investigate system trade offs, some times the most energy efficient solution is not always the most effective for a task, especially if it is critical for tasks to be completed urgently.

Task allocation is probably the most important aspect of swarm robotic behaviour, the benefits of good allocation being shown in Section 2.2.3 to have considerable affect on task performance. Having robots performing the correct actions at the correct time defines the fundamental principle of an effective swarm, whether the allocation was achieved explicitly or through emergence. Due to this and the large variety of methods and behavioural strategies that task allocation encompasses, investigating task allocation as a primary behavioural adaptation will provide an expansive and complex testbed for proposed contributions.

2.4.2 Morphological Adaptation

In this section, morphological adaptation refers to changes to shape or design a system is able to make to itself without an evolutionary or generational element assisting the process whether online or offline (evolutionary and genetic algorithms are discussed in more detail in Section 2.3.2).

Morphological adaptation exists as a much smaller field than behavioural arbitration due to the hardware constraints in having robots change their shape. Some swarm robotic systems attempt to create adapting morphologies through the change in behaviour of individual robots. By creating a desire to dock with other robots in the swarm, multiple robots are able to construct larger shapes capable of different types of locomotion i.e. crawling, walking or

rolling. Some examples of these types of morphological adaptation are discussed in some detail within Section 2.3.3, the works of Kuyucu et al. (2013) providing a good simulated example of the morphology switching. These multi-robot structures have also been produced in physical experiments. In the REPLICATOR and SYMBRION projects (Kernbach et al. (2008)) prototype modular robots with linking capabilities were produced which were able to form a four legged multi-robot organism. This organism was designed as a potential solution to passing a barrier separating the swarm in the test environment.

There are also examples of robots that are capable of adapting their morphology without relying upon other robots to form structures. Kim et al. (2013) developed a transforming wheel capable of traversing obstacles of heights greater than the wheel radius (Figure 2.20). The wheels in question were able to unfold to create three legs, these legs allowed the robot to climb up to 2.6 times the radius of the wheels. These legs were passively deployed whenever a vertical surface was encountered, adapting to the encountered problem. This adaptation allowed the robots to access difficult terrains that would typically only be accessible by robots with actuated legs, while maintaining the high performance and mobility while operating on a flat surface, gained by having wheels. However, this additional option of locomotion may act as a trade off between versatility and robustness, with the additional features of the wheel creating new points of potential failure.

Morphological adaptations have also been used to increase environmental versatility in aerial robots. Zhao et al. (2017) designed an adaptive quad-rotor drone capable of expanding and contracting airframe size to trade between flight stability and agility, capable of navigating spacial challenges created by complex environments (drone shown in Figure 2.21). The scissor-like folding structure designed to allow for this adaptive airframe was found to provide ‘excellent obstacle surmounting performance, minimal aerodynamic influences, and great flight adaptability.’

In some cases robots may be designed to dramatically switch between locomotion types. For example Daler et al. (2013) produced a prototype ‘Deployable Air Land Exploration Robot (DALER)’. The robot was capable of fixed wing flight but was also capable of using its wings as ‘whegs’ to move over rough terrains (illustrated walking in Figure 2.22). Repurposing the wings for ground based transport reduced the weight of the platform and enabled effective movement on carpet, snow, grass, road and parquet while maintaining the ability to fly at roughly 14m/s. While this platform was capable of successful movement on land and through the air, the use of wings for legs does produce the concern that delicate control surfaces or the aerodynamic characteristics of the wing may be damaged over time.

From the examples given in this section it is clear that the added versatility platforms gained from morphological adaptation can be a large positive. However, due to physical constraints, it is clear that effective adaption of morphology requires the design of bespoke platforms. The design of these platforms must take into account the multiple features that may be



Figure 2.20: Transformable wheel capable of swapping between rolling and leg-like movement based on the context of the terrain it is being used to cross. Kim et al. (2013).




Figure 2.21: Adaptive quad-rotor drone capable of expanding and shrinking in size when required. This allows it to take advantage of the additional stability associated with wider placement of rotors, while still having the capability to fit through tight spaces when required. Zhao et al. (2017)




Figure 2.22: Photographs showing frames of each step in the DALER's walking method as the robot re-purposes its wings to locomote on the ground. Daler et al. (2013).

required for adaptation and as a result, the platforms are typically quite complex or form a compromise system due to their non-specialist nature. Additionally, without the ability to perform genuine self-reconstruction, the platform will never be able to truly and universally adapt its morphology, though this is where genome-like adaptation such as that proposed in Eiben et al. (2013) will come into play (discussed previously in Section 2.3.2). For these reasons when versatility is required it may be worth considering the use of simple robots, incapable of morphological adaptation, but each with different capabilities and as a group forming a diverse multi-robot system.

Subsection Title	Key Points
Behavioural Adaptation	This type of adaptation covers most complex robotic systems. If robots in a system are able to react to environmental features and change their behaviour as a result of this, it could be argued that said robot is capable of behavioural adaptation. However, this section highlights that the research contributions within this thesis will be focused upon behavioural adaptation for task allocation and to affect energy consumption.
Morphological Adaptation	Morphological adaptation is seen less within robotics and is mainly constrained to modular systems, capable of reconstructing their morphology to react to a problem, or robots designed with hybrid locomotive capabilities, with examples of wheel-leg and wheel-wing combinations given within the section.

Table 2.4: Summary Table for Section 2.4

2.5 Summary

The content of this literature review will guide the experimental work undertaken in later chapters. Each of the following chapters will involve the content discussed in Subsection 2.3.3 on Virtual Hormones and it will also be important to take into account the properties and abilities intrinsic to robot swarms as reviewed in Section 2.2. While the sections cover the fundamentals, individual chapters will require additional aspects of the reviewed literature, these cases will be highlighted as follows:

Chapter 3 - Virtual Hormones for Explicit Control These experiments are based on the use of virtual hormones to directly control motor functions. The design of the virtual hormone controller attempted to create a fast dispersion system for mapping an environment, adapting based on swarm density. This work represents the first step towards adapting biological hormone examples to engineer a solution and as a result, only requires a fundamental understanding of swarm and virtual hormone systems.

Chapter 4 - Virtual Hormones for Energy Efficient Task Allocation The second series of experiments presented in this thesis, explore behavioural adaptation in greater depth, along with the viability of Hormone-like signals as state arbitrators. For this section it is important to have a strong understanding of the problems involved in creating an energy efficient swarm, principally work on the energy consumption of scaling swarms and works that provide dynamic methods for changing swarm sizes. Referenced in Subsection 2.2.4.

Chapter 5 - Virtual Hormones for Task Allocation by Self Identifying Traits The work in this chapter extends the behavioural control of virtual hormone system, though this

time looking at the regulation of a heterogeneous swarm. The different features and traits of the robots in this swarm are limited, but Subsection 2.2.3 does provide a good insight to systems containing members of different capabilities and previous methods of arbitrating allocation.

Chapter 6 - Virtual Hormones for Creating Dynamic Traits The discussion in Section 2.2.1 mentions versatility as one of the greatest strengths of a swarm robot system. Capitalising on this, Chapter 6 looks to explore the adaptability granted by hormone systems, in complex and realistic situations, combining work from previous chapters to identify what is required to create an adaptable, but sustainable Virtual Hormone controlled robotic swarm. Attributing to this, Section 2.4 contains important background to this system, providing context to the adaptability of the system.

Through each of these chapters, virtual hormone systems will be tested and analysed with the goal of presenting a validated method for engineering hormone inspired systems. These chapters will also identify the appropriate applications for such systems, showing the benefit and disadvantages encountered through testing.

Chapter 3

Virtual Hormones for Explicit Control

3.1 Introduction

As discussed in Section 2.3 behaviours based on biological systems can provide inspiration in solving complex engineering problems. In recent years work has been undertaken to mimic the behaviour of cellular systems. More specifically, work that uses the natural interactions seen in proteins, peptides, steroids or other hormones to inspire robot control (Neal & Timmis (2003); Levi & Kernbach (2010)).

As previously described, hormone inspired controllers have been used predominantly for the control of swarm structure. In these cases it is not uncommon for hormones to indicate when it is appropriate to change morphology. In these structure and morphology changing systems, a hormone value is typically used to select behaviour states rather than to directly control robot actuators. Hormones within these systems are used primarily for coordination. As hormone values build they provide context to the system based on their stimuli and thus allow systems to coordinate the transformation into pre-planned structures at the appropriate time. Along side this, hormone inspired systems have also been used for the direct control of robot movement, previous works have represented wheel motors of a robot as a cell to be hormone driven (Stradner et al. (2009)).

The work presented in this chapter looks to verify the viability of virtual hormone based control for robot swarms by investigating a simple arena mapping example. The information gained from the experimentation in this section is used to educate investigation in ensuing chapters, identifying areas that should be avoided, areas worth exploring and how a hormone inspired system should be designed in regard to a swarm of robots.

This chapter introduces a hormone inspired system which combines multi-cell movement controllers and behaviour swapping systems to produce an adaptive dispersion controller. The controller directs a swarm capable of inducing either dispersion or attraction behaviour when beneficial to the mapping of an environment. Effective use of these behaviour states allows

for concentrated deployments of robots to disperse efficiently, improving the consistency and speed that the swarm is able to map various environments versus a system exploring at random. The performance of the controllers proposed in this chapter will be based upon the number of unique coordinate points the swarm is able to store, these points are only recorded when a robot identifies an obstacle so by collecting the points stored by every member of the swarm a map of walls and objects within the environment will be produced. For the purpose of measuring performance, duplicate coordinate points within 15 cm of one another will be ignored when counting the number of points recorded. The individual swarm members are not aware of coordinate points already mapped by other robots due to their distributed nature, therefore the dispersion and coordination provided by the proposed systems will have a large bearing on the measured performance.

Two sets of hormone equations comprise the proposed controller. One controller adapts a single parameter based on environmental context and optimises a second parameter to find the best value for each environment. The second controller self modifies all parameters subject to environmental factors in an attempt to eliminate the need for optimisation prior to encountering new scenarios.

3.2 Hormone Controller

The proposed hormone controllers operate at two levels. The first directly controlling the wheel motors of the robots within the swarm and the second controlling their behaviour state. In this section the specifics of the controller operation are detailed for systems 1, 2 and the base line test system.

3.2.1 Motor Control

Similar to previous work (Stradner et al. (2009); Kernbach et al. (2008)) the system described in this chapter divides a robot into two cells, each capable of storing a hormone value. These cells each activate the motor on their corresponding side.

Rather than using virtual hormones to proportionality modify wheel speed, as would be similar to systems discussed in earlier works, the systems presented in this chapter have the motor cells on-board each swarm member compete directly with one another. The cell with the highest hormone value activating the respective cell's wheel while the other wheel remains stationary. This means that the speed of the robot is not proportionate to the hormone value stored in each cell, allowing the robots to always disperse or map at full speed. It is also important to note that, in this chapter, the hormone values effecting each of the motor cells associated with a robot do not effect one another. As a result, the only factor reducing these hormone values will be the decay present in the equations. Thus the controller will rely purely

on decay to reach equilibrium. Equilibrium in this case referring a state in which hormone values on board a robot show no substantial difference between one another, allowing the robots to travel without hormone values having an effect on behaviour.

The primary hormone stimulant in this new controller is the presence of a foreign robot. The presence is communicated as robots transmit and receive up to date records of their current hormone values at each time step. This approximates natural cell to cell hormone communication without the requirement of a physical medium in the operating environment as is required in most pheromone communication tasks (the negatives of which have been discussed previously in Section 2.3.3).

The transmitted hormone values were detected via line of site range and bearing sensors. The sensors were capable of both detecting transmitted hormone values and reporting the direction of a transmitting robot relative to the detecting robot. The direction of the received signal was used to decide whether the hormone value of the left or right motor cell was affected by the detected hormone. With signals coming from the right of the robot making changes to the right cell and signals coming from the left making changes to the left cell (illustrated in Figure 3.1).

The hormones in the proposed controller are stimulated or inhibited by the detected presence of other robots. The magnitude of the hormone value in the detected robot relative to the hormone value of the detecting robot determines whether an individual is inhibited or stimulated by the presence of another robot. Receiving higher hormone values stimulates the hormones in the detecting robot and receiving lower hormone values inhibit the hormones in the detecting robot (this has been illustrated in Equation 3.1).

$$\begin{aligned}
 &\text{while: } 1 < H_L(t) \& H_R(t) < 250 \\
 H_L(t+1) &= \alpha + \lambda H_L(t) + \sum_{i=0}^{n_g} \frac{D}{d_i} - \sum_{i=0}^{n_l} \frac{d_i}{D} \\
 H_R(t+1) &= \alpha + \lambda H_R(t) + \sum_{i=0}^{n_g} \frac{D}{d_i} - \sum_{i=0}^{n_l} \frac{d_i}{D}
 \end{aligned} \tag{3.1}$$

$H_L(t)$: Left hormone value at previous time step (This must be positive and saturates at 250).

$H_R(t)$: Right hormone value at previous time step (This must be positive and saturates at 250).

α : Regular increase in hormone value.

λ : Hormone decay rate.

n_g : Number of robots connected with a greater Hormone Value than current robot.

n_l : Number of robots connected with a lesser Hormone Value than current robot.




Figure 3.1: Illustration of the signals received by a central robot and how these signals will effect the robot's hormone values. In the image the robot in the centre is receiving the total hormone values (the sum of the left and right hormones) of the robots within line of sight. It can be seen that the robots to the left of the central robot transmit values that will then only effect the hormone value associated with the left side of the central robot and the robots to the right will effect the hormone value associated with the right side in a similar manner.

d_i : Distance between current robot and connected robot 'i'.

D : Scaling value defining the relative hormone change based on robot concentration.

$H_R(t)$ and $H_L(t)$ were set to saturate at 250 and were not able to take a value lower than 0. This limits the total hormone value in each robot ($H_{Total} = H_L(t) + H_R(t)$) between 0 and 500. The importance of this is discussed further in Section 3.2.3.

It is worth noting that while distance is not considered for the comparison of hormone values, distance does change the effect a foreign robot's hormone value has in stimulating or inhibiting a motor cell, though not by the magnitude of the foreign hormone. In the stimulating case, the rate of increase in hormone value is reduced as distance between the detecting and transmitting robot increases. In the inhibiting case, the effect a foreign hormone has on a motor cell increases proportionately with distance. The result of this distance based change means that robots with a large distance between them have a minimal influence on dispersion, but a large influence on attraction.

This effect was engineered so that robots with high hormone values, which imply the near presence of other robots, would move away from clusters and towards robots with low hormone values. These low hormone values in turn imply a location in the environment with few close neighbours. This produces an effective push and drag dispersion effect with the hormone values dictating whether a robot should travel towards to away from a robot based on the current environmental context.

With this arrangement, each member of the swarm builds up hormone values in the cell facing the highest concentration of robots. On deployment this caused the robots to turn away from the cluster and disperse (this is exhibited in robot 1 and 3 in figure 3.2) reducing chance of robots remaining in an already mapped area. This system also ensured that later in the simulation, should a cluster form within an enclosed section of the map, swarm members in a cluster were able to connect to low H_{Total} robots outside of the enclosed space, attract towards them and navigate openings without explicit knowledge of the environment (attraction behaviour is shown in robot 2 in figure 3.2). This attraction could only take place to guide robots away from obstructions due to the nature of the line of sight sensors; robots were not attracted to robots with obstructions between them and as a result, were only attracted through clear paths.

3.2.2 Behaviour Switch

In addition to the direct control of the motors, the artificial hormone also provided important information about the swarms' environment. As stated H_{Total} will be high when robots are clustered together and low when the robots are dispersed. Due to this trait, each member of the swarm could use the hormone value to decide which behaviour they should exhibit.




Figure 3.2: The Illustration shows separation hormone values effecting the movement of three robots. The close proximity of robots 1 and 2 has resulted in a larger hormone value in the hormone cells facing one another. This subsequently results in the robots turning away from one another. Robot 2 turns away from robot 1 regardless of the fact robot 3 is to its left as the greater distance between the robots means that the stimulant to the hormone cell facing robot 3 is not as great as that of the hormone cell facing robot 1. Robot 3 however, will still turn away from robot 2 as it has no other robot effecting its left hormone cell. The resultant trajectories are marked with dashed lines headed with arrows.




Figure 3.3: State Diagram for context aware hormone system. Systems 1 and 2 differ only by the type of threshold required for transitioning between state S_0 and S_4 . As the baseline system contains no hormone mechanism it does not use this state diagram at all. Instead robots avoid one another and move randomly until a non-robot obstacle is detected, at which point the individual detecting an obstacle begins mapping.

Along side the hormone motor control state detailed in Section 3.2.1 the robots comprising the swarm are also able to swap into a mapping behaviour state. In the mapping state, the motor cells still store hormone values and interact with one another as described previously. However, the hormone values no longer have control over the motors. Instead the robots use their short range proximity sensors to identify obstacles and attempt to maintain a constant distance from said object as they move forward, tracing the object they have found. In this mapping state the range and bearing sensors are used to identify if two robots are within close proximity of one another, preventing mapping if two robots are within a short range (1 robot length diameter) of one another. This ensures that the only environmental features are mapped. The state transitions for this system are detailed in Figure 3.3.

While tracing around obstacles, robots record and store coordinate points to form a map of environmental discoveries. If at any point during mapping any individual robot's H_{Total} value exceeds the behaviour switch threshold (Th_b) that individual swaps back into the diffusion state, with its motors once again controlled in the manner described in Section 3.2.1.

3.2.3 System 1 - Single Parameter Adaptation

The parameter D , acts as the scaling factor for both stimuli and inhibitors and has the largest effect on performance (discovered through preliminary analysis of parameters, a deeper review

of hormone parameters is given in Section 3.4). After observing the behaviour of the swarm with a static D value repelling other mapping robots in cramped areas, often to the detriment of the systems mapping performance. It was theorised that it may be beneficial for this value to adapt based on the context of the robots' environment. The adaptive replacement of D further regulated hormone fluctuation to be appropriate for both cramped and spacious environments, encouraging a switch in behaviour relative to the current density of the swarm. This approach was intended to outperform a static value in complex environments and required minimal optimisation prior to deployment.

$$\begin{aligned}
 & \text{while: } 0 < H_R(t) \& H_L(t) < 250 \\
 H_L(t+1) &= \alpha + \lambda H_L(t) + \sum_{i=0}^{n_g} \frac{(H_{sat} + 1) - H_L(t) - H_R(t)}{d_i} - \\
 & \sum_{i=0}^{n_l} \frac{d_i}{(H_{sat} + 1) - H_L(t) - H_R(t)} \\
 H_R(t+1) &= \alpha + \lambda H_R(t) + \sum_{i=0}^{n_g} \frac{(H_{sat} + 1) - H_L(t) - H_R(t)}{d_i} - \\
 & \sum_{i=0}^{n_l} \frac{d_i}{(H_{sat} + 1) - H_L(t) - H_R(t)}
 \end{aligned} \tag{3.2}$$

In order for each robot to adapt effectively, a measure of the current swarm density within communication range was required. Based on this measure, the robot was able to increase or decrease the rate at which hormone values were stimulated or inhibited, modifying the effect that each connected robot had on the hormone value.

To implement an adaptive element to the D value, the system was designed so that greater H_{Total} values of the detecting robot produced smaller D values. The D values within the equations were subsequently replaced with $(H_{sat} + 1) - H_L(t) - H_R(t)$ (where H_{sat} is the maximum value H_{total} can take), designed to range between 1 and H_{sat} . In the experiments in this chapter H_{sat} was set as 500 for ease of calculation, due to the fact 500 was the largest value d_i could take. The values available for d_i were restricted by the maximum communication range robots could achieve with the equipped range and bearing sensor (RAB). Across the available communications range a value of 500 would be returned at the maximum distance and 1 at the minimum.

The new D value's lower limit was restricted to 1 to avoid an undefined result in the case of a division by 0. By keeping the denominator and numerator limited to the same maximum and minimum values, the stimulant per connected robot was limited to a value between 0 and 1. Preliminary tests found that these equations produced acceptable changes to the hormone values, rarely saturating or remaining at 0. The behaviour emerging from this system prevented concentrations of robots forming in the environment. It also prevented robots from

interrupting the mapping of one another at inappropriate times, should the environment force the swarm to stay close together.

With an adaptive D value in place, the first contextually aware hormone controller is presented in Section 3.3.

3.2.4 System 2 - Multiple Parameter Adaptation

In order for the system to be viable for hardware application, optimisation of the swarm on an environment to environment basis would be laborious and time consuming. Therefore, to simplify use while still maintaining a high performance, all parameters that have a significant effect on the system can be either: optimised for a general case (thus not perform optimally in all environments) or adapt to the requirements imposed by the environment. In System 2, the latter will be achieved by adding an adaptive behaviour switch threshold (Th_b).

The Th_b value was configured to change with an inversely proportionate ratio to the density of the swarm, tuned through parameter investigation. This was achieved by utilising the A test as described in Alden et al. (2014) (an effect-magnitude test which quantifies the difference between data groups, discussed in greater detail within section 3.3.2) which produced results for the single parameter adaptation that clearly showed the behaviour switch threshold should range between 100 and 300 to provide the maximum benefit across each environment tested while not branching into values of limited effect on performance.

Taking into account the required proportionality and value limits, the following adaptive parameter was created:

$$Th_b = U_{lim} - \frac{(U_{lim} - L_{lim}) * n}{N} \quad (3.3)$$

Where N is the total number of robots in the swarm, n is the number of robots connected to the controlled robot, U_{lim} is the upper limit for Th_b (300) and L_{lim} is the lower limit for Th_b (100). With this additional adaptive parameter the aim was that this system should be able to achieve a similar or greater performance than a static optimised threshold.

3.2.5 Base Line Test

In order to obtain a set of results to act as a baseline in the mapping comparison, a simple controller was developed. In this system the robots, once deployed, would move randomly until they discovered a wall or obstacle. While exploring, robots were still able to avoid one another using the same short range distance sensors available to the robots utilising systems 1 and 2 as control methods. Non-robot features in the environment were also detected with these sensors (at a maximum range of 50 cm). In order to ensure that detected items were not robots, simple robot-to-robot messaging was deployed. Throughout the experiment robots

communicated through a separate sensor capable of sending and receiving small packets of data from other robots, as well as detecting the distance between them. If this sensor identified another robot within 50cm of the robot detecting an item, it was assumed that the discovered item was a robot and not an obstacle worth mapping. On discovery of a non-robot feature, the robots would begin to trace around the object, using the short range sensors to maintain a constant distance from the detected obstacle. As previously mentioned, mapping was prevented when a robot was detected within 50cm of another robot. This actually created an additional benefit to robots already mapping as, if a wall following robot collided with another robot, it most likely meant that the section of wall they were about to travel along had already been mapped. As a result of identifying the near-by robot, a mapping robot would stop the wall following behaviour, and begin to move away from both the wall and the presented robot. This would allow the robot to explore a new area that had potentially received no mapping as of yet, rather than circling an already mapped obstacle indefinitely.

Additionally, other than identifying robots within close proximity, the base line random movement system would have no restrictions on when robots were able to begin mapping. Instead of waiting to reach a threshold point as in systems 1 and 2, as soon as a robot using the random controller identified an obstacle it was able to record a coordinate (assuming it was previously undiscovered). As a result, the swarm using this system was able to begin mapping in the early stages of the experiments, gaining an advantage over the robots using system 1 and 2. This was implemented by design to highlight how the other systems would have to provide an adaptive method capable of making up for lost mapping time in the beginning of the experiment.

3.3 Experiments

3.3.1 Experimental Set up

The simulations performed in these experiments were produced using ARGoS (Pinciroli et al. (2012a)), a system capable of simulating arenas with multiple robots. For all experiments the controllers detailed previously were implemented on a swarm of 10 foot-bots (designed for the Swarmonoid project Ducatelle et al. (2014)).

To test the performance of the proposed controllers the total number of unique coordinates plotted by the swarm was summed. This number was recorded every 100 time steps up to 8000 for small environments and 16000 for larger environments. This provided a good measure for mapping ability throughout the experiments and made behaviours easy to monitor.

To obtain conclusive results a variety of environments were used in the simulations. Each of these environments were produced to test different capabilities of the swarm controller. The environments are shown in Figures 3.4, 3.5, 3.6 and 3.7. Of these environments the Obstacle




Figure 3.4: Simple Obstacle Environment: A small environment, this is the most basic in terms of object complexity. This environment should be easy to navigate and map, putting minimal strain on the system.




Figure 3.5: Maze Environment: A large environment. This is the most complex environment to navigate and was designed as a compilation of all the environments, creating a problem that should be more challenging to optimise for.




Figure 3.6: Corridor Environment: The difficulties posted by this small environment are the narrow hallways the swarm will be required to navigate. This environment should show that smaller values of D are beneficial to wall following in enclosed spaces.




Figure 3.7: Office Environment: Another large environment, however, this one is made up of several compartmental rooms. This environment should make it difficult for robots in the same room to disperse and will test the swarms ability to operate with minimal communication.

λ	α	H_{sat}	U_{lim}	L_{lim}
0.94	2	500	300	100

Table 3.1: Parameter values selected for hormone mapping system.

environment offers a simple, open and easy environment to map, the Maze environment creates a variety of problems, narrow openings, large spaces and complex obstacle shapes, creating a challenging environment to map that will be difficult to optimise a static value for, the corridor environment has been designed to observe the effectiveness of the swarm systems in long, narrow spaces and the office environment creates several cubicles with small openings that should make diffusing for effective mapping difficult.

The investigated data for each trial is the result of 100 repeated experiments as indicated by the consistency analysis conducted. This is discussed in greater detail in Section 3.4.1.

3.3.2 Parameter Values

The parameters selected for the experiments are shown in Table 3.1. The hormone specific variables λ (decay rate) and α (base stimulant increment) were chosen to reflect an appropriate relevance period (the time during which a hormone should maintain value before decaying below a value deemed relevant) and settling point for the hormone values. To do this, λ was the first parameter considered, chosen to decay the hormone from saturation (H_{sat} , in this case 500) to the lowest value deemed relevant (H_{fin} , in this case 1) within 100 time steps when no stimulant was present. This number of time steps (n) was then converted to a decay rate using Equation 3.4. This was approximately the amount of time taken to perform a long manoeuvre when avoiding a wall or robot in close proximity.

$$\lambda = \sqrt[n]{\frac{H_{fin}}{H_{sat}}} \quad (3.4)$$

The value for decay could then be used to find α using the equation: $H_{set} = \frac{2\alpha}{(1-\lambda)}$ where H_{set} is the decided settling point.

Due to the nature of the small environments the robots were exploring, in combination with the relatively short range of their communications devices (5 meters), long range transmission of hormone values would be fairly uncommon. It was therefore assumed that robots capable of communicating hormone values (i.e. robots with a clear line of sight of one another) were more likely to be in close proximity than far apart from one another. As a result this would induce hormone values of greater magnitude.

To cater for this, the setting point for the hormone value was designed to give a greater range of values above than below, allowing for greater resolution when comparing hormone values and reducing the likelihood of hormone saturation. It was therefore decided that the settling




Figure 3.8: A-Test results comparing the total area mapped at the final tick of simulation for each environment tested. For each environment the parameter for separation threshold was modified, ranging between 0 and 500 in steps of 10, to observe the systems sensitivity to said parameter and the potential ranges that should be explored when creating a dynamic separation threshold. The parameter sweep used to produce this graph was also used to identify optimal values of separation threshold for each environment.

point should be establish at approximately one eighth of the saturation point.

After initial parameter adjustments, the Spartan package A-Test was applied to analyse the effect of the behaviour switch threshold on the system. This was performed across a range of values between 0 and 500 in steps of 10 for each of the four environments tested. The trials for each parameter value of the set were formed from 100 experimental runs. Each data group was compared to the $Th_b = 150$ results giving an A-test score relative to the performance achieved with this value. The results for each environment are shown in figure 3.8.

Before obtaining the results from this parameter robustness test the following predictions were made:

1. Low values of Th_b (under 100) would produce a large change in results compared to the 150 set as the robots would never be able to switch into their mapping state.
2. Very high values would produce largely different results due to minimal dispersion.

3. The best results would be achieved between 100 and 300 for each environment.

The first and third predictions were confirmed by the results showing large changes in performance with values lower than 100 in most cases and optimal performance identified at 270, 130, 250 and 90 for the Obstacles, Maze, Corridor and Office environments respectively. The office result is somewhat anomalous to the expected parameter value, although this small value is most likely due to the large separation required by such a complex environment.

Contradicting the second prediction, the results showed that only the obstacle environment had largely different results for high values of Th_b although there were no optimal performance values found above 300, suggesting that the minimal dispersion of robots does decrease performance.

The highest performing parameters were taken from this parameter analysis and used in the comparison as an optimal case for system 1.

3.3.3 Comparison



For the proposed controllers to be proven successful they were tested against the following null hypothesis:

H_0 : System 1 will have a mapping performance no different to a random wall follow controller.
(rejected with 95% tolerance)

H_1 : System 2 will have a mapping performance no different to system 1. (rejected with 95% tolerance)

The results of the controller comparison can be seen in figures 3.9 and 3.10 showing the performance of the office environment near the start and at the end of the experiments. Additionally figure 3.11 shows box plots of the mapping performance for each controller at the final tick of each simulated environment.

Observing the performance of system 1 versus the random controller it can be seen that in all cases except the obstacle arena there is an initial delay in mapping from system 1. During this time it is outperformed by the random controller. This effect is most obvious in the office environment where the hormone controller does not start mapping until roughly 1000 ticks into the simulation (figure 3.9). The random system has this early advantage in mapping due to the fact that the random controller starts plotting points as soon as an object is identified. Meanwhile, system 1 waits until the hormone reaches a level that exceeds the static threshold value before plotting points. The initial wait before mapping creates some inefficiency in System 1, this can be explained by the high Th_b value identified as being optimal for maximum area mapped by the end of the experiment. While this value gave the best mapping by the end of the simulated time, it was clearly required to maintain a good



(d) Box plot of results for System 2 (Full Adapt), Random System and System 1 (Single Adapt)

Figure 3.9: Image capture of the Office environment for each system at 1000 ticks showing the area mapped within the environments at this time stamp (indicated by black dots seen on the floor of the environment). Image (d) displays the results of the office environment at 1000 ticks, showing that the single adaptation system completes no mapping at this stage of the experiment while the fully adaptive method and the random method are able to begin mapping early.







Figure 3.10: Image capture of the Office environment for each system at 16000 ticks showing the area mapped within the environments at this time stamp (indicated by black dots seen on the floor of the environment). Image (d) displays the results of the office environment at 16000 ticks, showing that the single adaptation system is able to recover from the poor mapping quantity early in the experiment to create a strong performance alongside the fully adaptive method against the random system.



(a) Maze environment



(b) Corridor environment



(c) Obstacle environment

Figure 3.11: Box plots of results comparing the mapping performance by the final time step of the experiment of the system 1 controller (Single Adapt), the random controller and the system 2 controller (Full Adapt).

Controller		Obstacles	Maze	Corridor	Office
Random System 1	Vs	$P < 0.001$	$P < 0.001$	0.523	$P < 0.001$
Random System 2	Vs	$P < 0.001$	$P < 0.001$	$P < 0.001$	$P < 0.001$
System 1 System 2	Vs	0.013	0.494	$P < 0.0039$	$P < 0.551$

Table 3.2: P values for comparisons between three controllers for each environment. Results that reject H_0 or H_1 are highlighted in bold.

level of separation later in the experiment, offsetting the initial time taken to begin mapping considerably. The selected optimal Th_b value is most likely larger than that required for the initial dispersion, which highlights the need for an adaptive threshold value which system 2 attempts to amend.

Even though the random controller outperforms system 1 early in the simulations, as time and mapping progresses, system 1 overtakes the random controller in terms of performance. By the end of the simulations, the mean area mapped across all trials was higher than the random controller in every environment tested.

Using a Mann Whitney test, the difference between these data sets was confirmed to be significant in three of the four environments (results table 3.2). Each of the results comparing the Random System and System 1 reject H_1 with the exception of the Corridor environment. The similarity of results in the corridor environment could be due to the fact that almost all of the environment would have been mapped by both controllers by the final time-step.

Through the observation of System 2's performance it can be seen that the issues of the static threshold values are resolved to some extent. In the office environment System 2 takes advantage of the initial high obstacle density, starting mapping early with a similar mapping approach to the random dispersion. Post-dispersion System 2 behaves more similarly to System 1, keeping distant spacing between robots. Comparing the results for System 1 and 2 in each of the box plots displaying final area mapped (Figure 3.11), it can be seen that system 2 provides a greater average mapping in every environment. However, the Mann-Whitney tests only confirmed the increase as significant in Corridor environment. Even without a consistent significant increase, the results of System 2 are still encouraging as no parameter optimisation has taken place prior to the swarms deployment and System 2 is still capable of achieving a statistically indistinguishable performance from System 1 by the end of the simulation.

3.4 Analysis

This section details supplementary analysis of the proposed systems providing: statistical evidence supporting the number of replicate trials required for consistency amongst results, insight to the decisions made when deciding which parameters to adapt and how to adapt them, as well as the results from tests on scalability, undergone to ensure the systems function in a truly swarm-like manner.

3.4.1 Consistency

In order to determine the number of trials needed to ensure statistical consistency an A-test based consistency analysis was performed on the mapping example data sets as directed in Alden et al. (2013). The number of repeats required was identified by comparing distributions of simulated data at each recorded time point (100 ticks of simulation, resulting in 10 seconds of simulated time). As data sets get larger, the probability of creating distinguishable data sets reduces, with the difference in distribution typically getting smaller as data size increases. In this case the size of the data set was dependant on the number of repeated experiments. Once the difference between distributions became negligibly small it was implied that the stochastic nature of the simulated results has been mitigated by the number of trials. Figures 3.12, 3.13, 3.14 and 3.15 show the need for 100 repeated trials as there are still time points showing a large or medium difference in A-Test score in repeat sizes of 5, 50 and 75, indicating that a greater number of repetitions are required.




Figure 3.12: A-test scores for data sets at 100 tick intervals for 5 experiment repetitions, showing all measured ticks reporting in large difference in results.




Figure 3.13: A-test scores for data sets at 100 tick intervals for 50 experiment repetitions, showing no large differences, but majority medium difference across all time steps.




Figure 3.14: A-test scores for data sets at 100 tick intervals for 75 experiment repetitions, showing some small difference and some time steps approaching a medium difference.




Figure 3.15: A-test scores for data sets at 100 tick intervals for 100 experiment repetitions, showing all ticks reporting only a small difference, confirming statistical consistency for this number of repetitions.

3.4.2 Sensitivity

The initial hormone equation shown in Equation 3.1 contains three parameters of interest α , γ and D . The effect produced by α and γ are somewhat predictable, α producing an offset from Zero to prevent the system from reaching a permanent settling point and γ changes the rate of decay for the system, returning hormone values to the settling point more quickly or slowly after a perturbation in hormone value is experienced by a robot, depending respectively on how far or close γ is from 1.

This left the final parameter in the equation, D , the parameter influencing hormone gain/loss in the presence of other robots. This parameter was not well understood and in an attempt to gain a greater insight in to the effects of the parameter and exhaustive value sweep was performed on D for each of the test environments. The sweep covered values for D in increments of 10 between 0 and 600, running 100 experiment trials for each parameter and using an A test to give a score of difference between each parameter and the baseline parameter (selected to be 300 in each of these tests). Results are seen in Figure 3.16, 3.17, 3.18 and 3.19. The technique for this analysis is similar to that used in the consistency analysis. Data sets are compared using an A-test to identify the magnitude of distribution difference. In this case each data set (containing 100 repeats as indicated by the consistency analysis) is compared to the data set for a baseline parameter value of D , here chosen to be the middle value of 300.

From these results it is obvious that there was no single optimum value for D . Large differences in distribution to the middle value took place in low values of D in some environments and high values of D in others. This highlighted the need for a contextually aware adaptation of this value to create successful robot interaction for a general system that required no bespoke optimisation for environment changes.




Figure 3.16: A-test results for parameter D in the office environment sweeping through a range of parameter values for D between 0 and 600, incrementing in steps of 10.




Figure 3.17: A-test results for parameter D in the obstacle environment sweeping through a range of parameter values for D between 0 and 600, incrementing in steps of 10.




Figure 3.18: A-test results for parameter D in the maze environment sweeping through a range of parameter values for D between 0 and 600, incrementing in steps of 10.




Figure 3.19: A-test results for parameter D in the corridor environment sweeping through a range of parameter values for D between 0 and 600, incrementing in steps of 10.




Figure 3.20: Scalability tests in the obstacle environment showing a linear increase to the number of points mapped as the number of robots increases from 5 to 55 in steps of 5 robots.

3.4.3 Scalability

In order to prove that the composed system is indeed swarm like, tests for swarm scalability were made to see how performance changed as the number of robots in the swarm increased. In these tests the number of robots in the swarm was increased from 5 to 55 in steps of 5. The results for this are displayed in boxplots shown in Figures 3.20, 3.21, 3.22 and 3.23. These graphs show the total area mapped at the 8000th time step for swarms of robots of various sizes. Every robot in these trials ran System 2 for 100 experiment repetitions across each of the four test environments.

As would be expected, with greater swarm size comes a greater performance as more robots are able to map in parallel. In most cases, the increase to performance by adding an additional 5 robots begins to reduce as the swarm reaches a high density, unable to map more effectively in the given space. The results for the obstacle environment (Figure 3.20) are an exception to this, gaining a fairly linear rise in mapped area as the swarm size increases. This is most likely due to the large open space provided by this test environment, allowing for dispersion to prevent excessive clustering of robots.

As System 2 sees no large drop in performance in the tested swarm sizes, even with relatively large robot numbers. The system has been proven to function well at swarm-like sizes and fulfils the second criteria for true swarm-like behaviour as highlighted in Section 2.2.1.




Figure 3.21: Scalability tests in the maze environment showing a gradual increase in number of points mapped as swarm size increases, performance can be seen starting to plateau after 20 robots.




Figure 3.22: Scalability tests in the office environment showing performance increasing but with diminishing gains as more robots are added, similarly to the maze environment.




Figure 3.23: Scalability tests in the corridor environment showing performance increase initially but remaining at a similar level between 20 and 55 robots.

3.5 Chapter Summary

The experiments in this chapter offer a fair comparison between the base line and systems 1 and 2 as both methodologies are implemented on swarms of the same functional capabilities. i.e. the swarms use the same fundamental obstacle avoidance system, the same distributed communications system and the same number of robots are used in each of the compared experiments. The primary differences between these systems are the regulation of dispersion and the activation of mapping behaviour. Systems 1 and 2 regulate dispersion and the activation of the mapping state through hormone inspired methodologies, while the baseline system disperses randomly, only avoiding other robots in close proximity. The intention of this comparison was to explore initial ideas regarding how hormone adaptive systems, represented by systems 1 and 2, might outperform a brute force method, represented by the base line system.

The work described in this chapter has provided an insight into the workings of hormone inspired systems within the context of robot swarms. It has indicated that virtual hormones can be engineered to create a beneficial system for swarm motor control and that this performance can be increased by introducing additional adaptive factors to hormone equations. The results of the experiments suggested that it is possible to use adaptive values to act as stimuli within a hormone equation and as thresholds to compare hormone values against, subsequently triggering behavioural change. As hormones themselves can provide a level of adaptivity, due to their time based nature, two areas have been highlighted for further investigation in forthcoming chapters:

1. It may be worth exploring the use of hormone values acting as a stimulus to other hormones: By feeding a hormone value into another hormone equation, it may be

possible to provide an adaptive method of control closer to those seen biologically, rather than designing engineered functions to influence hormone values.

2. As hormone values can be designed to change value according to environment stimuli, it may be worth creating behavioural threshold values in the form of additional hormones. These hormones could operate across a different time scale and take into account different stimuli. Subsequently, threshold hormone values may be able to provide comparative information to allow swarm agents to change behaviours at sensible points in time, improving the performance of the swarm.

While using a virtual hormone controller to directly affect motor movements was indicated to be effective in these experiments, it is clear that the next stage in hormone control related work needs to look at larger and more complex systems. However, this work should look further than direct actuator control as the work in this chapter has highlighted that adaptive behavioural control may be where the strength of virtual hormone systems lies. The designs and experiments in the next chapter will therefore look at higher level behavioural control using hormone systems to arbitrate behaviour states. With low-level operations programmed explicitly amongst executable states. This will allow for the execution of more complex tasks, with a greater requirement for adaptation further testing virtual hormone systems without requirement for unnecessarily complex designs.

Hormone style equations may provide a method of control in situations where common implementations for control, such as PID controllers, may be difficult. In order to implement a PID controller a desired quantifiable output must be specified, the error between the current and desired output is then used to modify behaviour (Ang et al. (2005); Knospe (2006)). The PID controller will be used to reach the desired output in an acceptable manner, designed for a combination of speed, accuracy or stability. For some swarm behaviours the desired output for each agent may not be easily measurable at any given moment, as the required behaviour of swarm agents may seem counter intuitive to the goal of the task (e.g. robots sleeping to conserve energy Liu et al. (2007); Sauzé et al. (2010)). As a result, it may be difficult to provide a quantifiable, continuous measurement of desired output that is directly relevant to swarm behaviour.

Even in swarm applications where PID systems may be relevant, PID controllers may require the use of additional methods of control to be deployed for effective management of swarm behaviour. For example, in the mapping experiments addressed in this chapter, a PID controller could be implemented to obtain a pre-defined optimal separation distance for swarm agents, regulating this separation by affecting the left and right motors of the robots to guide robots away from one another. However, as also shown in this chapter, there may be different requirements for dispersion at different points in the experiment. Therefore, to achieve effective mapping with the PID system, an optimal value for separation would have to be

known at each moment of swarm exploration. Thus, for the successful use of a PID controller in this situation an additional, adaptive, means of control may need to be implemented. It may be possible to implement these adaptive means through the use of hormone systems providing environmental context to a PID controller. However, due to the aforementioned drawbacks of PID control in swarm systems, and having identified the benefit multiple hormone inspired systems might provide to behaviour state control, points 1 and 2 (previously mentioned in this summary) were deemed to be the priority for the experimental exploration featured in the following chapters.

Chapter 4

Virtual Hormones for Energy Efficient Task Allocation

4.1 Introduction

A system's parameters can often be optimised for a given task to improve performance. This optimisation will typically take into account environmental factors, thus tailoring a system for a strong performance. If a robotic system is deployed in a dynamic environment, the chosen parameters must either take a suboptimal value or be re-optimised during the task, taking into account new environmental properties. The former of these options reduces the system's capability in each setting and the latter wastes time that could be spent executing the assigned task. Additionally, both of these options are difficult to execute given that a thorough analysis of the environment would be required.

This chapter illustrates, in a series of specific environments and tasks, that it is more energy efficient to use a hormone inspired system than an optimised timer-based system to control state transitions in a foraging swarm. The work in the previous chapter showed that hormone systems could be used to directly control the motor functions of individuals amongst a swarm of robots, providing an effective method for dispersion and attraction for a group of mapping robots. In this chapter, the work diverges from the low level actuator control and instead considers the use of a hormone system to increase the performance of a foraging swarm. The hormone system will regulate the swarm on a behaviour by behaviour basis, providing conditions for the robot to shift states upon.

Rather than waiting for events or direct instructions to act as indicators for state transition, hormone systems receive information in the form of stimuli. These stimuli affect the values of virtual hormones, causing them to fluctuate over time. By having a variety of hormone values triggered by different stimuli, information about environmental aspects can be obtained by observing the relationships between each of these values.

By using a combination of hormone release, stimulated by environmental factors, and constant hormone decay, artificial hormone systems create a simple yet powerful method for controlling robots (Shen et al. (2000)). Through the appropriate selection of stimuli, hormone systems can be used to arbitrate the states of individuals in a swarm, meditating actions such as resting, searching or returning to the nest site.

Additionally, the use of such a system means that each robot in the swarm produces its own hormone without requiring a lead robot to act as an emitter, in contrast to the hormone control for the CONRO robot (Shen et al. (2000, 2004)). This decentralises the system and makes it all together more swarm-like.

Work in this chapter uses a foraging system, a well studied example in swarm robotics with a history of research on energy efficiency. Examples of this work cover: improving the efficiency of search behaviour (Schroeder et al. (2017)), exploring methods of task allocation, creating periods of inactivity when workload is low (Charbonneau & Dornhaus (2015)) and increasing the efficiency of motion, reducing congestion by allowing swarm members rest (Liu et al. (2007)). While these studies address the issue of energy efficiency, some even including methods of online adaptation, they do not apply any form hormone arbitration system in their method of control. This chapter contributes to current literature by identifying how a hormone-inspired system increases the performance of a foraging swarm or how a hormone-inspired system might make a swarm more robust to environmental change.

4.2 Hormone-Inspired Systems

Previously a number of hormone-inspired implementations have been aimed at controlling the behaviour of a single robot rather than a swarm, focusing more on providing a strong insight to the construction of a hormone system (Stradner et al. (2009)). While others have explored swarm examples for hormone behaviour arbitration, in Kuyucu et al. (2013) a system is evolved over several iterations to complete a task with notable improvements. The stimuli featured in the study took the form of virtual pheromones. The system thus required a centralised element to record pheromone values at their location, meaning that their system could not be considered completely swarm-like by definition Şahin (2004). Additionally, the experiments carried out had no real base-line comparison to show the hormone implementation as a superior system.

Other swarm-based hormone systems have exhibited chemotatic behaviour, emulating the biological diffusion of hormones in cells to organise and structure a swarm (Shen et al. (2004)). Using this system the swarm was capable of navigating and exploring an obstacle-filled environment. However, again the experiments were performed with no quantifiable base-line to compare the performance of the system against.

By examining a complex hormone system with a quantifiable measure of performance and providing a comparison optimised for each case examined, this chapter will establish a precedent for future development.

4.2.1 Comparability Of Proposed System Vs Other Hormone Inspired Systems

The system proposed in this chapter shares some similarity with other hormone systems, the hormone equations follow the typical rules for hormone messages as suggested in Shen et al. (2000) and these signals are used to influence the behaviour of other robots, similarly to the work in Kernbach et al. (2008) where a 'Hormone Driven Robot Controller' is used to create robot structures from a swarm. However, these systems are not directly comparable due to the fact that the Hormone system presented in this chapter uses the relative values of multiple hormone signals to adapt and improve the performance of the swarm in long-term tasks, while the other system coordinates robots towards known patterns and morphologies.

Systems such as those shown in Kuyucu et al. (2013) using hormone signals to make the decision to swap morphology states, are more similar to the system presented in this chapter. Although, the fact that Kuyucu et al. (2013)'s system still only uses a single hormone per robot to make these decisions, makes strong comparisons between the two systems difficult.

4.3 Controllers

Base on the definition of a swarm of robots discussed in 2.2.1, the robots used in this chapters experiments where designed to have very simple functions. This would allow the design of the robots to adhere to simple design goals set out within the field of swarm robotics (Miner (2007)) and would keep the cost of producing a swarm, if developed in hardware, to a minimum.

As a result, the robots comprising the swarm in this example had the ability to identify obstacles at short range (50cm) for avoidance purposes and were able to "see" food items within 2 meters (assuming there was line of sight between the robot and the food item). By using these sensors, the robots within the swarm were able to easily navigate towards the items once they were discovered. The robots were also equipped with a simple, two wheel, method of locomotion, changing the robots direction via skid-steering. Aside from these simple actuators, and the processors used to manage the two arbitration systems detailed below, the robots in this chapters experiments implemented no other capabilities or features and had no method of communicating directly between one another.

The swarm robot systems in these experiments were designed to perform a foraging task. Foraging has been commonly used to test swarm behaviour and, has been used in the past to test a sleep-based system for energy efficiency in a swarm of robots (discussed further in the

section 4.3.1). Within the simulated experiments conducted, food items that would be foraged were represented by black dots on the environment floor. Once discovered and "picked up" by a robot, the item would disappear and, upon the robots return to the nest, the the collection of the item would be recorded. It is worth noting that the robots in this chapters experiments were assumed to be able to only carry a single item at a time.

In order to produce a measure of energy efficiency (the value which would be used to compare the performance of the systems developed in this chapter) the approximate amount of energy consumed by each robot was measured during the task. The total energy consumed by the swarm by the completion of the experiment, in combination with the number of food items collected, formed the value of energy efficiency. The process of calculating energy efficiency is described in greater detail in Section 4.4.1.

4.3.1 System 1: Hormone Arbitration

Previous work Liu et al. (2007); Charbonneau & Dornhaus (2015) has suggested that an energy efficient system can be produced by modifying a rest time for individual swarm members upon the completion of a task. In Liu et al. (2007) this was done by modifying the duration of various counters to mediate the amount of time robots would spend 'sleeping' in the nest site and the length of time they would spend searching for food items. The length of these time periods were changed based on the number of collisions experienced by members of the swarm, along with the successes and failures of each robot. Success or failure in this case were defined by whether the robot returned to the nest with or without a food item. This work on foraging efficiency inspired the elements that would be controlled by equations 4.1, 4.2 and 4.3. These equations regulate the sleep period and conditions required for robots to return to the nest.

The artificial hormone system proposed here is constructed from several hormone inspired equations. As in the previous chapter these virtual hormones can be produced by taking a base increment (α), a decay (λ) and weighting assigned to stimuli (γ), the equations produced from these elements approximate the behaviour of hormones used for biological control.

The variables in these hormone equations work in the same manner as those in Chapter 3:

- λ defines how quickly the system reacts to a lack of stimuli; the smaller the value λ takes, the faster the system will return to the settling point without stimuli.
- α and λ combine to make the settling point of the hormone without stimuli (calculated via $\frac{\alpha}{1-\lambda}$) when α is 0, the system settles at 0.
- γ defines how quickly the hormone will deviate from the minimum settling point when stimulated.

Using this simple format, multiple hormone equations can be created, each affected by different stimuli, building together to arbitrate behaviour states. The equations produced for the experiments in this chapter are as follows:

$$\text{Avoidance Hormone: } H_A(t) = \lambda_A H_A(t-1) + \gamma_A A \quad (4.1)$$

Where A is a Boolean value detecting whether or not the robot is avoiding another robot or surface .

$$\text{Hunger Hormone: } H_h(t) = \alpha_h + \lambda_h H_h(t-1) + \gamma_h C \quad (4.2)$$

Where C is a Boolean value representing whether or not the robot successfully returned a food item to the nest site.

$$\text{Sleep Hormone: } H_S(t) = \lambda_S H_S(t-1) + \gamma_S H_A(t-1) \quad (4.3)$$

The avoidance hormone (H_A) has been designed to return swarm members to the nest when an environment is too cluttered. While in the search state the robots will move randomly to explore their environment. If an obstacle is detected during this search the H_A value is stimulated and increases slightly. Rather than specifying a search time, which could inhibit the discovery of food items, H_A is used in a more dynamic manner, with robots sent back to the nest only in the case of overcrowding. Overcrowding is detected by the relationship between H_A and the hunger hormone H_h , whenever the avoidance hormone exceeds the value of H_h the robot changes state to return to the nest. To allow some small obstacle detection prior to returning to the nest and to ensure that decay leads the system to settle in the search state, the α term in H_h is present to provide a higher settling value than in H_A .

H_h was created to build a resistance in the system against H_A . As H_h rises (depending on the success of the robot), the robot associated with the hormone becomes less likely to return to the nest. Upon each success, the resistance increases and then begins to decay slowly until the next success. The importance of this resistance is clear when you consider the implications of a successful robot; Either the robot has access to a large concentration of food or the robot has minimal competition and is clearing its immediate area of food quickly.

Considering the first of these cases, a high concentration of food may attract a large number of robots, each collecting food in close quarters. Because of this, avoidance is more likely to take place. Without accounting for food density, robots with a threshold that is too low could return to the nest before discovering food items or the swarm may be allowed to aggregate too tightly in an area of low food density with a threshold that is too large. Both of these cases would waste energy and use the members of the swarm ineffectively. By allowing a certain amount of avoidance, and allowing that amount to vary based on context, better performance

can be achieved.

In the second case, with food items close to the nest removed, the successful robot will then have to travel a longer distance to discover new food. If the robot has a single collision prior to travelling the greater distance, harder-to-find food items will not be discovered. This would leave a potential food source untapped. The swarm is thus more effective if successful robots are allowed to travel further and avoid more objects before returning to the nest.

In addition to controlling behaviour switching, the avoidance hormone acts as a stimulant for the sleep hormone (H_S). The sleep hormone controls how long the robot stays in the sleep state. In this state the robot waits in a low power mode at the nest, this minimises energy consumption when the swarm size is too large for the current work load. The amount of time spent in the low power state depends on how long H_S takes to decay below the H_h value. This means that time spent in the nest is dependent upon both robot success and the amount of avoidance performed while searching. The state transitions controlled by these hormone equations are illustrated in Figure 4.1.




Figure 4.1: State machine for foraging hormone system illustrating state names and the relative hormone ratios required for transition.

Additionally, a demonstration of the hormone value movements during the experiments can be seen illustrated in Figure 4.2. This illustration can be used to see the effect hormone values have on one another and how relative changes in hormone value will trigger behaviour state transitions. The diagram clearly shows that the presence of the avoidance hormone will increase the value of the sleep hormone. Additionally, the successful collection of an item can be seen represented by spikes in the hunger hormone.

Between approximately 800 and 1100 ticks, several item collections can be seen to gradually build the average value of the hunger hormone. This provides the robot with a slight resistance to entering the sleep state as several brief periods of avoidance behaviour are unable to trigger behaviour change by raising the avoidance hormone value above that of the hunger hormone. However, with enough time spent in the avoidance behaviour (at approximately 1200 ticks), the avoidance hormone is able to build above the hunger hormone. At which point the robot enters the sleep state, indicated on graph by the lack of collections spikes in the hunger hormone, and the robot waits for the built up sleep hormone to decay below the hunger hormone value before searching for items again.

Graph showing Hormone behaviour in a single robot.




Figure 4.2: Graph showing the hormone values of a single robot across a selected time period in single experimental trial. The sleep hormone can be seen to take a much larger value than the other two hormones. Implemented to create longer periods of sleep when the swarm is substantially cluttered. Additional behaviours such as successful collection of items or obstacle avoidance can be seen in spikes in the respective hormone values.

λ_A	γ_A	α_h	λ_h	γ_h	λ_S	γ_S
0.999	0.2	0.005	0.999	5	0.999	0.05

Table 4.1: Parameter values selected for hormone arbitration through the mathematical reasoning explained at the end of Section 4.3.1

The parameters for the hormone equations (shown in Table 4.1) were chosen empirically to create a desirable response to environmental factors.

Decay parameters were chosen by considering the amount of time across which the hormones would have to act. These time scales were chosen based on the approximate expected time of relevance for each hormone type which when included in Equation 4.4 as the value for n , in combination with an h_{sat} of 250 (the hormone saturation point) and an h_{fin} value of 1 (the lowest relevant hormone value considered in the application) appropriate decay values were produced.

$$\lambda = \sqrt[n]{\frac{H_{fin}}{H_{sat}}} \quad (4.4)$$

Parameters for stimulus weightings (γ) were the applied by considering the value for γ that would saturate the hormone given continuous stimulation. This was found by using Equation 4.5 with the relevant values λ value and S_{max} (The largest value the stimulant attached to the weighting might take). To arrive at a suitable γ value the calculated value was adjusted by considering how frequently the stimulant might be active, rather than the constant activation assumed in Equation 4.5.

$$\gamma = \frac{H_{sat}(1 - \lambda)}{S_{max}} \quad (4.5)$$

Outside of this manual tuning no optimisation was performed to increase the performance of these parameters. The parameters were also not changed across any of the trials in the experiments section of this chapter.

4.3.2 System 2: An Offline Optimised System

To form a comparison for the hormone arbitration controller, a simple timer-based system was produced. In this system, the functions performed by the hormones designed in the hormone inspired system were replaced with timers. In order to make this a fair comparison the timer lengths were optimised using an elitist genetic algorithm (GA). The GA in question was a fitness proportionate roulette wheel trained across 50 generations with a population of 30. The resultant best cases were taken as the parameters for the baseline tests. This was an attempt to simulate a system that had perfect knowledge of its environment and could choose parameter values prior to deployment, capable of producing competitive results.

The two parameters optimised were avoidance threshold (Th_A) and sleep threshold (Th_S). Th_A represents the maximum amount of time a robot would spend avoiding obstacles prior to returning to the nest site. Th_S represents the amount of time that would be spent resting in the nest upon returning. Neither of these values changed during the simulation and instead the system's performance relied upon having the best parameters for swarm size and environment type.

4.4 Experiments

The goal of the following experiments was to identify how the hormone-inspired system and the timer-based system compared. The success of the hormone-inspired system was dependent upon the rejection of the following null hypotheses:

H_0 : The optimised hormone-inspired system's performance will have no significant difference from that of an optimised timer-based system in a simple environment (rejected with a 95% confidence level) and will not provide greater energy efficiency on average.

H_1 : The optimised hormone-inspired system's performance will have no significant difference from that of an optimised timer-based system in a dynamic environment (rejected with a 95% confidence level) and will not provide greater energy efficiency on average.

To test these hypotheses two different experiments were designed. The purpose of these experiments was to discover:

- Whether System 2 could be optimised for a variety of situations and thus be capable of changing its parameters to suit an environment with high accuracy, hence it would outperform System 1.
- If given a highly dynamic environment formed from two areas with a large difference in food density, would parameters optimised for such a situation take too much of a compromise to outperform the adaptive hormone system.

It was predicted that, in the following experiments, it would be shown that it is viable to use a hormone inspired system, with parameters tuned based on simple principals, to grant a level of adaptivity to a swarm. It was also predicted that the implementation of this hormone inspired system would allow the swarm to perform at least as effectively as a swarm with parameters tuned via genetic algorithm before the start of the experiments, albeit tuned parameters that are non-adaptive.

4.4.1 Energy Efficiency

The performance of the swarms in each of the experiments was measured by energy efficiency, calculated with Equation 4.6 where E_e represents energy efficiency, E_F represents energy from total food stored, E_C represents total energy consumed by the swarm and E_A represents the total energy available in the environment from food.

$$E_e = \frac{E_F - E_C}{E_A} \quad (4.6)$$

When food items are returned to the nest they provide the swarm with 2000 units of energy. This reward value was approximately four times that of the average cost of a single robot searching for and collecting a food item in the smallest environment, under optimal conditions. With this much energy as a reward, traffic caused by the swarm and any other energy wasting difficulties arising from coordination amongst a large group should still allow for a reasonable, positive percentage value of energy efficiency to use as a comparison for the two systems tested.

The process for robot energy consumption was kept consistent between the two systems. However, the rate of consumption was based on the behaviour state of the robot at any given time step. Each robot consumed 8 units of energy per second while driving forwards, 4 while turning, 1 while stationary and an additional 2 while carrying a food item. The single unit of power consumed while stationary accounts for the power needed for basic, crucial functioning of the robot. By adding power consumption, even when robots were not moving, more value was added to having robots perform the tasks quickly; Should robots successfully clear the environment of food items before the allotted simulation time, less energy will be wasted by sleeping robots. This will provide experimental results more relevant to what would be expected in reality and means that the system cannot produce the best results by limiting foraging to a single robot.

4.4.2 Simulation

All experiments were conducted in the ARGoS simulator Pinciroli et al. (2012b) a multi-robot simulator used to simulate large robot swarms. The simulated robots used in these experiments




Figure 4.3: Graph showing the cumulative mean (red line) and confidence intervals (dashed lines) as the number of replications increases for the 20 robot hormone system used in the first set of experiments. A vertical line at $n = 21$ marks where the deviation reaches an acceptable point (less than 0.05).

where based on the marxbot Bonani et al. (2010) a miniature wheeled robot, assumed in these tests to be capable of travelling at 10cm/s and identifying food items within a 2m radius.

Consistency Analysis

The number of replicates required for consistent results was determined by performing a cumulative mean test as specified in Robinson (2004). This method was used over the A-Test previously used in Chapter 3 to identify the number of required repetitions as it required considerably fewer data sets to perform, finding an acceptable number of experimental trials much faster.

This method of testing uses the cumulative mean of a data set, along with a calculated confidence interval to give an estimate of the range in which the true mean lies. These tests indicated that amongst all of the swarm sizes and systems in the first experiments, 21 trials would be the minimum number of replicates required for the results of the experiments to be an accurate representation of the simulation responses (illustrated in Figure 4.3). For this reason 25 replicate experiments were performed in the first set of experiments.




Figure 4.4: Graph showing the cumulative mean (red line) and confidence intervals (dashed lines) as the number of replications increases for the 20 robot hormone system used in the second set of experiments. A vertical line at $n = 183$ marks where the deviation reaches an acceptable point (less than 0.05).

In the second set of experiments, involving a more complex environment, the cumulative mean test showed that a minimum of 183 repetitions would be required for reasonable consistency (illustrated in Figure 4.4). Thus 185 repetitions were selected to ensure consistency. The large increase to the number of required trials was expected due to the increase of arena complexity and environmental dynamics associated with the environment of the second experiment set.




Figure 4.5: Simulated ARGoS environment for experiment 1. The grey area represents the nest site, the black spots mark the location of food. 20 robots can be seen in the nest area awaiting the beginning of the experiment.

4.4.3 Investigating The Effect Of Swarm Density

Environment

The environment for the first set of experiments was an octagonal arena (this arena shape has previously been used in foraging research Liu et al. (2007); Lau (2012)). The arena measures 8m x 8m and the northern most 2m of arena is coloured grey to represent the nest site. Outside of the nest site 100 food items have been randomly distributed (represented by the black circles as seen in Figure 4.5).

Experimental setup

In this experiment swarms of 5, 10, 15 and 20 robots were deployed in the nest area of the foraging environment. Each trial lasted 1500 seconds (with each time step being 0.1 seconds) or until every food item was collected and returned to the nest. The energy efficiency of the swarm was measured at the termination of the simulation and used as the measure of performance. By allowing the system to terminate early upon completion of the task, the effectiveness of the swarm was challenged in addition to the efficiency. This meant that a

Swarm Size	System	
	Avoidance Counter (Th_A)	Sleep Counter (Th_S)
5	1992	40.9
10	34.9	1960
15	48.9	3672
20	35.4	4510

Table 4.2: System 2 parameters optimised via genetic algorithm across 30 generations for experiment 1. These results show that, as would be expected, the smallest group of robots require low levels of sleep with a large resistance to avoidance, allowing for continual operation. Groups larger than this required relatively low avoidance counters, allowing robots to return to the nest after brief collisions. In terms of sleep counter, times gradually increased with swarm size, reducing traffic outside of the nest to account for the increased population.

swarm capable of collecting food items quickly could also be rewarded, rather than completing the task and wasting energy waiting for the end of the simulation timer.

As mentioned in Section 4.3.2, the timer-based system was optimised for each swarm size before the experiments, these optimised parameter values are shown in Table 4.2.

Results

Observing the median performances illustrated in the box plots shown in Figure 4.6 it can be seen that the hormone-inspired system has an increased performance relative to the timer-based system in swarm sizes of 5, 10 and 15. Performing a Wilcoxon test on the data sets shows that this increase in performance is significant in the 5 and 10 robot case with p values (shown in Table 4.3) lower than 0.05, rejecting H_0 . These results also show that, in general, efficiency decreases as the swarm size increases. This was expected, as a more cluttered environment is more difficult to navigate. However, as the swarm size increased the performance of the hormone-inspired system also decreases relative to that of the timer-based system, resulting in a significantly lower performance in the 20 robot case. This suggests that the initial parameters chosen for the hormone-inspired system are more suited to swarms of a smaller size.

The significant increase in performance of the hormone system at smaller swarm sizes, compared to the timer-based system, shows that the hormone-inspired system is capable of adapting parameters to small environmental changes. The dynamics of removing food items in this case provided a distinct enough change in environment to allow the timer-based system to adapt mid-task and thus perform better. Even the 15 robot case, in which the hormone-inspired system achieves no significant difference in performance to the timer-based system, shows the benefits of a hormone arbitration method. The hormone-inspired system is capable of achieving these results having no knowledge of the environment before performing the task.




Figure 4.6: Box plots showing energy efficiency in the simple environment used for experiment 1 for both the Hormone inspired system and the timer-based system. The graph shows the results for swarm sizes of 5, 10, 15 and 20. The results show that the Hormone-inspired system is capable of obtaining competitive results versus the GA optimised system without the Hormone-inspired system having received bespoke optimisation itself.

Swarm Size	P-value
5	$P < 0.001$
10	$P < 0.001$
15	0.4034
20	$P < 0.001$

Table 4.3: Wilcoxon rank sum tests comparing the two systems for each swarm size with results for energy efficiency taken from the first, simpler, experiment environment. Significant difference can be seen at all swarm sizes with the exception of 15.




Figure 4.7: Simulated ARGoS environment for experiment 2. This environment is larger at 8mx15m and features a wall, removed at the 500 second point in the experiment to future diversify the arena properties.

In contrast to this, the significantly lower performance with a swarm size of 20 may be an issue. In the next experiment, the 20 robot swarm will be examined to identify whether a larger swarms cause a problem for the hormone system Vs the optimised system, even in a more dynamic environment.

4.4.4 Investigating The Performance Of Large Swarms In A Highly Dynamic Environment

Experimental setup

In the experiment performed in Section 4.4.3 it appeared that the swarm implementing the hormone inspired system was able to modify its behaviour according to the dynamics induced

by the consumable food sources. This allowed the swarm to achieve competitive energy efficiency versus an optimised system. After observing this behaviour the following set up was devised to further test the reaction of the hormone arbitration system to environmental dynamics. The notable changes made to the environment (shown in Figure 4.7) to test the hormone arbitration system's reaction to dynamics are as follows:

1. Zones of two different food densities were included, separated by a wall, to create a stark contrast in food availability.
2. The environment was expanded, now measuring 8x15 meters, to accommodate for the second search space.
3. The dividing wall, formed from the lower three walls of the octagon, was removed 500 seconds into the experiment to enable travel between the low food density area and the high food density area. A time of 500 seconds was chosen to enable travel between areas at a time past the point at which the hormone system would already have begun to settle, requiring reaction from the hormone system to find a new settling point.
4. To account for the larger environment and increased travel time, the maximum length of the experiment was extended to 3000 seconds.

With the listed changes implemented, both of the tested methodologies would be further challenged. The hormone-inspired system would have to quickly adjust values to obtain effective performance, compensating for time and performance lost due to initial poor values. While the timer-based system would have to produce a robust set of parameter values to perform well with greater environmental contrast.

The parameters of the timer-based system were re-optimized for the new environment. Using the same genetic algorithm as used to optimise parameters for the first experiment, the values shown in Table 4.4 were chosen for Th_A and Th_S .

It is important to note the parameters for the 5 robot swarm. In such a large environment the most efficient method for food collection was to never enter the sleep state and always persist in the environment until a food item was discovered. With a Th_A greater than 1000, robots would never return to the nest unless they had a food item and with a Th_S of 0, upon returning to the nest they would never sleep.

Results

The energy efficiency of both systems in this experiment is significantly reduced, even producing some results with a negative efficiency (i.e. more energy is used than gained from foraging). The reduction in performance was expected given the increased distance the robots were required to travel to obtain food.

Swarm Size	System	
	Avoidance Counter (Th_A)	Sleep Counter (Th_S)
5	>1000	0
10	646	6612
15	120	6609
20	101	7388

Table 4.4: System 2 parameters optimised via genetic algorithm across 30 generations for experiment 2. The second experiment shows a clearer gradient of values for avoidance counter. It is clear that robots get less resistance to collision as the swarm size increases, promoting the more frequent return to the nest of robots in more dense populations. Conversely, the sleep counter shows little difference within larger groups of robots, taking a value of 0 for the smallest. This implies that robots in the smallest group should never sleep to get the best performance. However, in groups of more than 5 robots they will sleep for approximately a fifth of the experiment time if they enter the avoidance state for long enough to return to the nest.

Swarm Size	P-value
5	$P < 0.001$
10	$P < 0.001$
15	$P < 0.001$
20	0.145

Table 4.5: Wilcoxon rank sum tests comparing the GA optimised and Hormone based systems for each swarm size with results for energy efficiency taken from the second, more dynamic, experiment environment. A significant difference can be seen in all tested swarm sizes except 20.

Wilcoxon tests between the Hormone and Timer systems of the same swarm size (illustrated in Table 4.5) show significant difference in swarm sizes of 5 to 15, but no significant difference in the 20 robot example. Observing the results displayed in Figure 4.8 the significant difference in the 5 robot example is clear, with the Timer-based system outperforming the hormone system. The relatively weak performance from the hormone system in this case is most likely due to the fact that the best performance from the optimised timer system resulted from never entering the sleep state. The hormone system would take some time to adapt towards this solution, requiring that each robot proved consistently successful at food retrieval, building the H_h to prevent robots returning home and ensuring that sleep periods are very short.

However, with the larger swarm sizes of 10 and 15 the environments become more dynamic with sleep state allocation becoming more important with a more crowded environment and the increased rate of food collection. Because of this, adaption becomes more effective versus the small swarm size where no robots are required to sleep for optimal function. As a result of this more effective adaptation, a significant increase in performance relative to the timer-based system is found. The strong performance of the hormone system at these swarm sizes shows that an online system can not just achieve the performance of an optimised system, but the sensitive adaptation provided by the hormone system can allow it to outperform an optimised case. This greater performance can be achieved as the hormone system reacts to smaller elements in an environment providing specific, context-based solutions to issues as they are encountered.

As a result of the strong performance and significant difference in the results produced by swarm sizes of 10 and 15, the null hypothesis H_1 can be rejected in these cases.

In contrast, the results produced by the swarm size of 20 in this experiment were found to be statistically similar to one another. While not outperforming the optimised system, this is an improvement from the first, simple experiment. The statistical similarity in the two systems at this swarm size shows that given a dynamic enough environment, online hormone adaptation can still produce good results without optimisation or re-calibration, even when given a large number of robots to control.

4.5 Chapter Summary

This chapter has shown that an adaptive, online hormone arbitration system can be used to increase performance (in this case energy efficiency) while foraging in environments of varying robot and food densities. In swarms of relatively small size the performance increases were significant when compared to a system with optimised but static parameter values. The observed performance of the hormone-inspired system in a simple environment depreciates as the size of swarm increases.




Figure 4.8: Box plot results showing the energy efficiency for swarm sizes of 5 to 20 in the larger, more complex, experimental environment. Even in this more complex environment, the Hormone-inspired system is capable of obtaining competitive results versus the GA optimised system without the Hormone-inspired system having received bespoke optimisation itself.

In a more complex arena, the Hormone system performed better relative to the timer system at higher swarm sizes, only struggling to create a competitive performance in the 5 robot experiments. This was partially due to the fact that the timer-based system was easy to optimise in this case, with robots never needing to sleep or return to the nest site unless to drop off food.

While the performance with larger swarm sizes in the second experiment were better than the first, the relative performance of the hormone system did decrease to match the timer-system in the 20 robot case. This indicates that in order to obtain the best performance from a swarm, an appropriate method for task division and swarm dispersion may be required. The next chapter will aim to amend this by investigating if virtual hormone systems can be used to evaluate the performance of individual swarm members, relative to others, and subsequently allocate swarm members to an area in which they would perform a task the best. This should alleviate congestion in a more productive manner than the allocation of 'Sleep States'.

Chapter 5

Virtual Hormones for Task Allocation by Self Identifying Traits

5.1 Introduction

Adaptation is commonly used as a technique for improving the performance of swarm robot systems Bahceci et al. (2003). Commonly, adaptation is performed before a task begins, tuning robot parameters to get an optimal performance for a specific problem as is the case for more genetic algorithms. However, in many situations it is not possible to know the exact requirements of a task and thus it is not possible to perform this tuning before the task begins. Systems have been proposed that attempt to transfer offline optimisation to operate during a task. Some studies have investigated migrating the concept of genetic algorithms, which are used frequently in offline optimisation, to allow swarms to adapt mid-task Haasdijk et al. (2014); Bredeche et al. (2012). This adaptation is achieved by giving each swarm member their own virtual genome. These genomes directly affect the behaviour of individual robots and by sharing genetic information with one another virtual generations are produced. By choosing appropriate fitness parameters, these genetic systems can promote successful genomes to adapt swarms to a task and obtain a better performance. Another method is to enable members of a heterogeneous swarm to chose tasks based on their abilities Jones et al. (2006). In this system, swarm members bid for tasks they are capable of performing and then work from a 'play book' to complete them. Working in this manner allows swarms to form from robots of very different types, creating what the study refers to as a 'pickup team'.

These online adaptation techniques prove successful within the context of the study's goals even with robots of mixed ability. However, the previously mentioned studies all require the members of a swarm to have an understanding of at least their own abilities. Having such an understanding is not always possible. During long term deployment factors may change: tires may wear down or robots may experience motor or actuator failure. These factors may change

the abilities of individual robots, having a negative effect on their interact with one another or the environment. Changes such as these will most likely cause reduction in performance over extended periods of time unless each robot in the swarm is capable of receiving an up to date diagnosis of their capabilities, using this to modify their behaviour to the most fitting option. Moreover, there may be situations in which there is no opportunity to inform a robot of its capabilities before a task starts. The Triangle of Life project Eiben et al. (2013) proposed a system in which robots are developed without humans in the loop, suggesting methods that would have robots in a swarm share both morphology and control systems through virtual genomes. These genomes are then shared in a reproduction-like manner. In the iterative design presented in these systems, ‘infant’ stage robots have no context for their own abilities. The combination of parent robot morphology, control system and external mutation leave the new generation of robot’s abilities ambiguous.

In the aforementioned cases, greater performance could be achieved with a method of adaptation that does not require the swarm to have any initial understanding of their own capabilities. In this chapter a method is presented that achieves this, instead of relying on an initial understanding of their abilities, robots implicitly gain information about themselves and other robots by monitoring the values of virtual hormones.

A system utilising virtual hormones was proposed in previous work, capable of arbitrating roles within a foraging swarm Wilson et al. (2018). The previous system used hormone values to select either a low-power sleep state or a searching state for each robot with the goal to conserve the overall power consumption of the swarm. The study found that an adaptive, online hormone arbitration system can be used to increase performance (energy efficiency in this case) while foraging in environments of varying swarm sizes and food densities. In swarms of relatively small size the performance increases were significant when compared to a system with optimised but static parameter values. The performance observed in the hormone-inspired system depreciated as the size of swarm increased. However, the swarms of large sizes were still capable of outperforming a simple but optimised system when introduced to an environment with more pronounced dynamic effects. The environmental change used to demonstrate this came in the form of a removable barrier which, 5000 ticks into the experiment, was removed to reveal an area of high food density.

In this chapter a new hormone-inspired system will be presented that will deal with more complex foraging examples and arbitrate the states of multiple robot types within a swarm. Arbitration will entail the decision between environment types, with each robot in the swarm using hormone values to make their environmental choice. Using a hormone responses to dictate environmental preference is not unheard of, there are natural examples of animals exhibiting exactly this. For example, desert amphibians leave spawn in pools that are intermittently filled and then dried depending on weather. Based on this environmental stimuli (i.e. water availability) hormone levels in the spawn change, accelerating or inhibiting

metamorphosis based on the need to stay or leave the pool they are currently in Denver (1999).

The system proposed in this work will not only take into account environmental stimuli, but will also use various transmitted hormone values from other swarm members. These values will then be used to gauge the capabilities of each robot individually during the task. With the information gained from the hormone values, the new system will allocate behaviour states to each robot based on how suited they are for the task.

The system presented in this Chapter will achieve a contextual awareness and environmental preference using a similar structure to the hormone examples presented in previous chapters. Equations produced from a decay, reducing the level of the hormone value over time, stimuli, a condition which when met increases the level of the hormone value, and potentially the inclusion of inhibitors, triggered by interactions in the same way as stimuli, but instead decreasing the hormone level.

The combination of decay and stimuli allow hormone values to fluctuate based on interactions, keeping a live record of the factors related to each stimuli. By using a variety of hormone values, each triggered by different stimuli, the relationships between differently affected hormone values can be examined to extrapolate information about environmental aspects. This information can then be used to educate or create preference within a swarm.

5.2 Hormone-Inspired Behaviour Arbitration System (HIBAS)

The hormone-inspired controller presented in Chapter 4 arbitrated states for a homogeneous swarm, allowing each robot to choose between sleeping at a nest site or searching the environment for food. By making this choice, the number of robots foraging was scaled by the hormone system to prevent large swarms from cluttering the environment. By reducing clutter, collisions between robots were less frequent and thus the swarm collected food in a more energy efficient manner. To build upon this work, the new system (HIBAS) removes the sleep state, instead the swarm is presented with the option to explore different environments. In addition to this the swarm is modified to contain different robot types, some more capable in one environment than others. With no prior knowledge of their capabilities individual members of the swarm are able to identify their strengths and form a preference for environment by using the hormone set shown in equations 5.1, 5.2 and 5.3.

$$H_x(t) = \lambda_1 H_x(t-1) + \gamma_1 H_{f_x} + \gamma_2 E_{stim} \quad (5.1)$$

$$H_{f_x}(t) = \lambda_2 H_{f_x}(t-1) + \gamma_3 F_x \quad (5.2)$$

Symbol	Meaning
H_x	Hormone for suiting to environment x.
H_{fx}	Food discovery hormone for environment x.
H_c	Hormone for environment/robot collisions.
λ_1	Acts as a decay affecting all environmental preference hormones (H_x), taking a value between 0 and 1.
λ_2	Acts as a decay affecting all environmental preference hormones (H_{fx}), taking a value between 0 and 1.
λ_3	Acts as a decay affecting all environmental preference hormones (H_c), taking a value between 0 and 1.
γ_1	Weighting of H_{fx} that acts as stimulus to H_x .
γ_2	Weighting of E_{stim} that acts as stimulus to H_x .
γ_3	Weighting of F_x that acts as stimulus to H_{fx}
γ_4	Weighting of C that acts as stimulus to H_c
E_{stim}	An integer variable that counts how many robots in the same state are transmitting suiting hormones that are larger than the detecting robot's.
F_x	Boolean variable that becomes true for a single time step while picking up a food item.
C	Boolean variable that becomes true if something encounters the robots obstacle avoidance sensors.
t	Current time step in experiment.

Table 5.1: Key for the symbols used in the hormone equations.

$$H_c(t) = \lambda_3 H_c(t-1) + \gamma_4 C \quad (5.3)$$

The subscript 'x' in these equations is used to denote instances where duplicate functions and variables will have to be made. In order for the system to operate, robots require one of these equations for each environmental option they are presented with. By numbering these environments and creating hormone values that relate to each environment, copies of H_x would become H_1 , H_2 and H_3 relating to environments 1, 2 and 3 respectively. Other symbols used in the hormone equations are defined in table 5.1.

H_x shown in equation 5.1 is the primary hormone for controlling environment preference. In a two environment example each robot in the swarm will have an H_1 and an H_2 value. When arriving at a nest site the robot will chose between going to environment 1 or environment 2 based on which of the two hormone values is greater. During a task every H_x value is broadcast from every robot, allowing other members of the swarm to compare hormone values. H_x values are the only values broadcast with H_{fx} and H_c being used only internally by each robot.

In a two environment foraging example (Illustrated in Figure 5.1), considering a robot with a preference for environment 1, while exploring that environment E_{stim} would keep track of how many robots are transmitting an H_1 value higher than the robot's own. E_{stim} then increases the value of every H_x value other than the hormone giving preference to

the robots current environment. In this example H_2 would be affected by E_{Stim} while H_1 would not be. This system allows robots to constantly compare their performance in their current environment and, if their performance is relatively poor given the context, start building a preference for another environment. Providing stimuli to hormones unrelated to the environment the robot is currently exploring is crucial. Without this the decay present in each hormone would slowly bring all non-stimulated hormones in the system to 0, preventing any preference for an environment from forming outside of the initial environment choice.




Figure 5.1: Two robots, A and B, operating in environment 1. Robot A has successfully discovered a food item and as a result its H_1 value is increasing. Robot B has entered the environment and is under performing relative to Robot A as indicated by its lower H_1 percentage. These robots transmit H_1 values and as Robot B receives an H_1 greater than its own H_2 is stimulated, as Robot A receives an H_1 value lower than its own no such stimulation occurs and H_2 is left to decay. The net result is an encouragement for Robot B to change preference to environment 2 and for Robot A to continue operating in the same environment.

H_{fx} is a stimuli hormone, given two environments, H_{fx} is present in the system as H_{f1} and H_{f2} , feed into H_1 and H_2 respectively. The purpose of H_{fx} is to create a stimulus for H_x which operates across a greater length of time than that of the initial stimulus trigger. This is accomplished by taking the initial impulse of the stimulus received when a robot picks up a food item (F_x) and providing a decaying value over time.

Stretching the stimuli over an additional length of time is important due to the repelling

nature of the E_{stim} variable. If the H_{fx} hormone value was to immediately increase upon picking up food, it would immediately be at its greatest value. This would mean that almost any robot that had not just picked up a food item would be encouraged to change environment preference. With the slow increase provided by H_{fx} the system is able to compare performance between robots no matter their stage in the task. This increase also better mimics the stimuli found in biology. Typically a fast acting neurological signal will trigger the production of a hormone in one organ which in turn will change the production of hormones in other organs, modifying the behaviour of the whole organism.

The hormone H_c is an element of the system kept from previous work Wilson et al. (2018). Its purpose is to monitor the frequency of collisions in the environment, returning robots to the nest should they encounter an area too cluttered with objects or other robots. Frequency of collisions is monitored with a slowly decaying hormone, stimulated by a Boolean value, C , triggered whenever the robots proximity sensors detect an entity. H_c is compared to every H_x value the robot currently stores. Should H_c exceed any of these values, the robot returns to the nest.

Figure 5.2 shows the hormone value dynamics when the swarm of robots are able to chose from three environments (one North, one South and one West). The diagram shows the robot beginning exploration in the west environment, though not performing particularly well. At just after 100 ticks of experiment time it appears that the robot swaps to explore the north environment (the environment this particular robot is designed for) where a good performance is then seen with multiple items returned to the nest (indicated by the arcing increases in hormone value). After 500 ticks these items become scarce, preventing success and reducing the hormone value for the environment currently being explored. It is foreseen that when there are minimal items available for collection, the hormone system will not be able to accurately allocate robots to their most suitable environment. However, this is not an issue, as a lack of items would indicate the completion of a foraging task, thus no longer requiring environment categorisation.


Graph showing Hormone behaviour in a single robot.

Figure 5.2: A graph displaying the dynamics in the hormone values used in an experiment in which a swarm of three different robot types had an option of three varied environments to forage in. Curved increases in hormone value seen in this graph typically represent the collection of an item, while more abrupt increases will be caused by other robots transmitting hormone values.

Wheel Type	Environment Floor Type	
	Wood Speed (cm/s)	Grass Speed (cm/s)
Wood Wheel	30.9	18.8
Grass Wheel	24.6	21.1

Table 5.2: This Table lists the speeds recorded when each of the wheels were tested in hardware examples. The figures listed represent the average speed recorded across 10 repetitions.

5.3 Creating a Heterogeneous Swarm

In order to test a system in simulation with robots capable of self classifying, robots of different capabilities had to be identified in hardware that could provide a realistic reference point for parameters in the simulated experiments. To keep the system simple, an existing swarm formed from the psi swarm robot Hilder et al. (2016) developed by the York Robotics Laboratory was altered to allow the attachment of different wheel types to their drive train. A disparity in robot capability was then created by designing different wheel types that could be 3D printed and easily attached, creating a heterogeneous swarm from groups of robots with the same fundamental construction. Once each of the designed wheel types were printed, they were tested in trials of 10. The average speeds when moving in each environment were recorded, allowing loss of traction and instability in these terrains to affect these speed values (shown in table 5.2).

5.3.1 Wheel Type: Wood Environment

The first wheel type (shown in figure 5.3) was a simple design, the only constraints being that the wheels would have to: allow a robot to travel quickly on a wooden surface and the robot should at least be capable of entering and exiting the grass environment.

The first factor to consider in designing the wheel was the diameter. By increasing the wheel diameter from the 31mm of the standard robot to 60mm, the robot would gain an additional 14mm clearance between the bottom of the robot and the ground (initially 6mm, now 20mm) and have a largely increased wheel circumference. This additional circumference allowed the robots to travel faster and, with the added height, they were able to effectively transition from a wooden floor to the deep grass whereas previously this was impossible.

The next consideration was the wheel width, given that the wheels were being designed for a smooth surface, there was no real benefit to having wide wheels, as such a thickness of only 3mm was chosen. These thin wheels reduced the amount of force required to turn them by minimising weight, allowing them to accelerate more quickly.




Figure 5.3: Wheel designed to specialise in the environment with a wooden floor. The wheel has a large diameter, creating a higher top speed as well as a lightweight, thin and hollow design to allow for quick acceleration. The wheel affixes to the robot via bolts attached through the three holes in the middle of the design. The wheel was designed in Autodesk Inventor Professional 2018.

5.3.2 Wheel Type: Grass Environment

The second wheel (shown in figure 5.4) was a more complex design as it had to perform well in the grass environment. To achieve this, a spoke-like design was produced. These spokes gave the robot additional traction in soft grounded environments, digging into the surface and catching on imperfections in the ground to propel the robot forwards. At 14mm this wheel was also much wider than the wheel designed for the wooden environment. This additional width made the robot much more stable when travelling through the rough grass environment and, with more area in contact with the ground, made it less likely for wheels to fall into divots in the environment, causing momentary wheel slip.

5.3.3 Wheel Type: Conceptual

This wheel was not designed in hardware but instead was created as a theoretical competitor to the first two wheels. This wheel was designed to travel at a constant speed in any environment. This speed was lower than the slowest speed in either the grass or wooden floored environments but would still travel at the same speed in areas of very difficult terrain where the other two robot types would be much slower. This was an attempt to simulate a robot with either large tracks, very wide wheels or even a robot with hovering capabilities.




Figure 5.4: Wheel designed to specialise in the environment with a grass floor. The spoke-like outer rim of the wheel provided traction for the wheel, preventing slippage when travelling over uneven surfaces. The substantially wider wheel also provided stability, preventing the robots from moving off course when encountering small bumps. The wheel affixes to the robot via bolts attached through the three holes in the middle of the design. The wheel was designed in Autodesk Inventor Professional 2018.

Parameter Symbol	λ_1	λ_2	λ_3	γ_1	γ_2	γ_3	γ_4
Parameter Value	0.999	0.999	0.995	0.005	0.1	50	0.2

Table 5.3: Table listing the parameters used in all software experiments. Parameter symbol meanings are described in Table 5.1

5.4 Experiments

To test the proposed HIBAS two experiments were designed. The goals of these experiments were to identify the categorising capabilities of the HIBAS and the performance increases such a system may give. In each of the experiments the swarms are given a foraging challenge. Their task is to discover, pick up and return food items to a nest area. Foraging is a well studied task in swarm robotics and has previously been found to provide a strong experimental testbed for coordinated systems and acts as a convenient abstraction for real-world applications Lu et al. (2016). By choosing foraging as the task for this experiment it is possible to build upon previous work in which the effect of hormone systems on foraging tasks were examined Wilson et al. (2018).

Success in these experiments will be measured by the percentage of correct categorisations made by the systems and by the rate at which food is collected.

5.4.1 Parameters

All of the following experiments were conducted with the parameters listed in Table 5.3. Decay parameters were chosen by considering the amount of time across which the hormones would have to act, in the case of the suiting hormones (H_x) item collection should be relevant across a large time scale. This time scale was arbitrarily chosen as 5000 ticks in this experiment, which when included in Equation 5.4 in combination with an h_{sat} of 100 (the hormone saturation point) and an h_{fin} value of 1 (the lowest relevant hormone value) a decay value of 0.999 is produced for λ_1 and λ_2 . A similar calculation was conducted for the collision hormone though instead the relevant time scale for collisions was chosen at 1000 ticks to produce a λ_2 value of 0.995.

$$\lambda = \sqrt[n]{\frac{H_{fin}}{H_{sat}}} \quad (5.4)$$

$$\gamma = \frac{H_{sat}(1 - \lambda)}{S_{max}} \quad (5.5)$$

Once the decay values were chosen, the parameters for stimulus weightings (γ) were assigned by using Equation 5.5. This equation takes into account the maximum hormone value that could be present H_{sat} , the decay rate and the maximum value a stimulus might take (S_{max}). A *gamma* value is then produced that will only allow the hormone values to saturate if a constant input stimulus is given at maximum magnitude (implying the worst of best case scenario for the performance of the hormone value). An exception to this method of calculation is found in *gamma*₃, as this weighting attributes food collection it will be activated, at most, once per search. As a result of this *gamma*₃ immediately adds a large quantity of stimulus to the relevant hormone value and if said hormone is not able to decay quickly enough and saturates, it simply means that there is an abundant food source very close to the nest site of which all robots in a swarm should be taking advantage of.

5.4.2 Simulation

All experiments were conducted in the ARGoS simulator Pinciroli et al. (2012b) a multi-robot simulator used to simulate large robot swarms. As previously mentioned the robots used in these tests were assumed move at the speeds shown in table 5.2 based on the simulated wheel type and terrain. Additionally, it was assumed that each of the robots was equipped with a food sensor, allowing them to identify food items within a 2m radius.

Each experiment was set up to run for 1000 seconds, each simulation time step was 0.1 seconds with samples recorded for every 10 seconds of simulated time.

The number of repeat trials required for consistent results were determined via cumulative

mean tests. These tests and the number of replicates required are detailed in Section 5.6.1.

5.4.3 Comparison System

To provide baseline data in these experiments a random arbitration system was produced. This system performs the exact same tasks as the HIBAS with the exception of the two following changes:

Random Arbitration

Rather than using a series of hormones to decide which environment should be explored by each robot, each robot picks randomly giving each environment an equal weighting.

Collision Hormone Threshold

With no H_x values to compare H_c against the system instead uses a flat rate of 10 as the threshold value. If this value is exceeded before a robot finds a food item, the robot returns to the nest and picks another environment to explore at random. This threshold value indicates that the robot has been colliding with a robot or obstacle for a notable amount of time and will ensure robots do not get stuck in a single environment.

5.4.4 Experiment 1 - Swarm Preference Between Two Environments with Static Features

The swarm in the experiment is made up of 7 robots with the wheels specialising in wooden floors and 7 robots with the wheels specialising in grass flooring. The environment for this experiment (shown in figure 5.5) measures 8mx20m, split into three parts. The two larger areas are both 8mx9m containing 50 food items, each of these areas has a different floor type. The third section is a strip down the middle acting as a nest site measuring 8mx2m.

5.4.5 Experiment 2 - Swarm Preference Between Two Environments with Mid Experiment Terrain Swap

This experiment was designed to test the robustness of the categorisation technique. Using the exact same setup as the first experiment, this test had only one difference: when the simulation reached 100 seconds the environment floor types switch. This sudden change should test the swarms ability to categorise once already acclimatised to the environment. This change could represent a landslide or other catastrophe, clearing one side of a task environment but making the other more difficult to travel in. The percentage of correct categorisations from




Figure 5.5: Screenshot of simulation environment used in experiment 1. Black dots represent food items, white ground represents wooden flooring, green represents grass and the grey area in the centre represents the nest site. Robots of different types (represented by red and black wheels either side of the robots) can be seen moving to and from respective environments, some carrying food items (represented by black rings above them).

this experiment should be expected to suddenly drop at the 100 second point. A successful system will then steadily increase back to the same or greater categorisation percentage than that of the switch point as the system re-adapts.

5.4.6 Experiment 3 - Swarm Preference Between Three Environments With Static Features

The third experiment is a more challenging test of the system. The arena is much larger (see Figure 5.6), introducing a third floor type for robots to explore with the same dimensions (8mx9m) as the environments in the first experiments. The new environment also included an additional 50 food items, making for a total of 150 in the whole arena. The nest area is also expanded in this arena, measuring 8mx8m to give an equal perimeter to each of the three environments. This experiment also sees the addition of a third robot type bringing the swarm composition to: 5 robots with grass specialising wheels, 5 robots with wood specialising wheels and 5 robots with wheels specialising in the difficult terrain (red flooring). In the new environment, the robots that specialised in the grass and wooden floored environments moved at 11.1cm/s and 8.8cm/s respectively, both 10cm/s slower than in the grass environment to account for additional difficulty.

The purpose of this new environment, while having measurements from hardware, is to give an example of a hazardous area in which two thirds of the swarm are not able to viably operate in. With such a large disparity in ability, successful categorisation of robots will benefit the overall performance of the swarm as items are foraged faster and more efficiently.




Figure 5.6: Screenshot of simulation environment used in experiment 3. Black dots represent food items, white ground represents wooden flooring, green represents grass, red represents very rough terrain and the grey area represents the nest site. Robots of varying type (indicated by the presence of red, green and black wheels) can be seen moving to and from respective environments, some carrying food items (represented by black rings above them).

5.5 Results

5.5.1 Experiment 1 - Swarm Preference Between Two Environments With Static Features

The results from the first experiment are shown in Figure 5.7. From this graph it can be seen that in experiment 1 the HIBAS outperforms random arbitration in terms of ability to categorise. For the first few samples it appears as though the hormone system performs identically to the random system. To confirm this Wilcoxon tests were performed, comparing the categorisation percentage datasets recorded for both systems at each time step. These tests showed that the 600th time step was the first sample with a significant difference between each systems results, giving a p value of less than 0.05. This marked the point at which the HIBAS and random system diverge. This initial starting period is to be expected; it will take time for the HIBAS to begin adapting to the new environment. From the 600th time step onwards it can be seen that the random arbitration remains with a correct categorisation percentage of roughly 45% while the the hormone-inspired arbitration increases gradually, peaking at just over 75%. Showing that the HIBAS gives a large improvement to categorisation over random allocation.

After this peak, the HIBAS starts to decrease in its ability to categorise environment, falling gradually to just below 50% and then fluctuating near the performance of the random system. This can be explained by the reduction in food items in the environment. As the source of primary stimuli reduces, the hormone system has no reward for item discovery and is therefore unable to accurately categorise. Once this source of stimuli is fully depleted, the system will behave essentially the same as a system arbitrating at random. This drop in performance is of no concern as, when the task nears completion, there is little to no need for the system to categorise successfully.

5.5.2 Experiment 2 - Swarm Preference Between Two Environments with Mid Experiment Terrain Swap

In the second experiment, the average performance between random system and the HIBAS is less disparate than the first (shown in Figure 5.8). By swapping the arena floor types just as the system starts to acclimatise, the hormone values must be re-evaluated by the swarm. This re-evaluation can be seen between the 1000th time step, where the switch occurs, to just after the 5000th. During this time period, robots reallocate themselves as their hormone values decay and it becomes apparent that their performance is lacking in their current environment. After the 5000th time step, in a similar manner to the first experiment, there is not enough food left in the environment for the robots to appropriately categorise themselves and as a result, the percentage of correct categorisation begins to tend towards random selection.




Figure 5.7: Graph showing the mean results across 150 trials in experiment 1. The performance of both the random and hormone-inspired system is shown in terms of correct categorisation and food items foraged. The small difference between food collection rate verses the large disparity in allocation type shows that, even given a small difference in collection capability, categorisation is still possible using the hormone system in this environment.




Figure 5.8: Graph showing the mean results across 150 trials in experiment 2. The performance of both the random and hormone-inspired system is shown in terms of correct categorisation and food items foraged. The large flip in correct categorisation percentage seen in the hormone system illustrates the point at which the environment types swapped. The hormone system is subsequently able to recover from this, eventually outperforming the random system once more.

These first two experiments shared some commonality in that the rate of food collection was marginal between system types. Performing Wilcoxon tests on the food collection data at every time step showed that, for almost all of the time steps past the 150th, there was a significant difference in the food collection data. However, performing an effect magnitude test using the A-test Alden et al. (2013) of the 100 time steps sampled for each experiment, only 22 datasets from the first experiment and 0 datasets from the second had a significantly large difference (an A-test score of over 0.75) from the random systems datasets. This lack of difference in food collection rate is due to the minimal difference in speed between the two robot types. This actually speaks to the benefit of the HIBAS as it was capable of assisting robots in their choice of environment even with a small difference in robot ability.

5.5.3 Experiment 3 - Swarm Preference Between Three Environments With Static Features

The third experiment shows what the system is capable of achieving in a more complex system and how the HIBAS can give a large increase to performance when there are larger disparities in robot ability.

Even given the added complexity in this final experiment (results shown in Figure 5.9), the HIBAS behaves similarly to the previous two experiments. The categorisation percentage takes an initial dip and then begins to diverge from the percentage of the random categorisation. However, in this experiment the percentage of robots correctly categorised fluctuates at around 40% which is much lower than the previous tests. With an additional robot type and environment choice this is still a good result as the HIBAS still outperforms the random system by upwards of 10% once the system has adapted.

The third experiment highlights two other key features of the hormone-system. First, due to the increased number of food items in the three environment experiment, not all of the food is foraged. As a result of this, the categorisation percentage does not taper off by the end of the experiment. Second, due to the increased difference in robot capability, for the first time in these experiments there is a clear difference in food collection. This is confirmed by running A-tests for each of the 100 time steps sampled showing a significantly large difference between the food collection in the two systems across all results from the 1500th time step onwards.

The results for the third experiment show that, given a large enough difference in capabilities, the HIBAS can provide a large improvement to foraging collection by correctly categorising robot ability to environment.

5.6 Analysis

Surplus to the results confirming the performance of the systems presented in this chapter, this section conducts additional analysis to verify the validity of the system. Starting with a consistency analysis, this section identifies the number of replicate results required for results that accurately represent the performance of the system. The section then looks into the scalability of the system in the environments tested in the previous experiments with the intention of identifying whether or not the system can be considered swarm-like.

5.6.1 Consistency

The number of replicates required for consistent results were determined by performing cumulative mean tests as specified in Robinson (2004). By using the cumulative mean of a data set, along with a calculated confidence interval, an estimate can be produced for a range




Figure 5.9: Graph showing the mean results across 180 trials in experiment 3. The performance of both the random and hormone-inspired system is shown in terms of correct categorisation and food items foraged. With the introduction of the red environment there is a greater penalty to assigning robots to search an arena area they are not suited to. As a result, the effect of correct categorisation can be seen more clearly in the rate at which food is collected by the hormone system versus that of the random.

in which the true mean lies. By taking cumulative mean tests across multiple time steps it was indicated that 150 trials was required for the two terrain experiments (graph shown in Figure 5.10) and 180 trials was required for the three terrain experiments (illustrated in Figure 5.11) in order to accurately represent the simulation responses.




Figure 5.10: 2 robot types - Cumulative mean and confidence intervals for experiment with two terrain types and time step requiring the largest number of replicates. The vertical line at $n = 150$ marks where the deviation reaches an acceptable value (less than 0.05).




Figure 5.11: 3 robot types - Cumulative mean and confidence intervals for experiment with three terrain types and time step requiring the largest number of replicates. The vertical line at $n = 180$ marks where the deviation reaches an acceptable value (less than 0.05).

5.6.2 Scalability

To introduce a key element of difficulty to the system presented here, with more robots in the swarm, the more difficult it will be to categorise them effectively. The system will have to cope with elements such as cluttering and high levels of communication. Cluttering will come into play as the swarm increases the environment remains unchanged, creating a high robot density area in which collisions will be frequent, potentially having a greater effect on collection than that of the wheel types and, as a result, it will be very difficult for the swarm to identify robot suitability to environment. The high level of communication may also effect the system, with transmitted hormones potentially over stimulating neighbouring robots. Due to the line of sight nature in the communication systems the robots in the swarm are equipped with, this should be less of a factor in the reduction of performance, with closer robots blocking the transmitted hormone values to others as they forage.

Data for the scalability tests was taken from the 500th time step, giving the system the opportunity for multiple food retrievals but without operation being diminished by a lack of food items in the environment. Observing the results of the scalability tests shown in Figures 5.12 and 5.13. The success of the system in the two environment is clear, showing that a categorisation median of above 50% (the performance expected of a randomly allocating system) was possible up until a swarm size of 42. At which point performance levels out to effectively the same as a random system, with it being very difficult to distinguish between wheel types for the reasons previously mentioned. The systems ability to obtain a performance




Figure 5.12: Correct categorisation percentages recorded 500 seconds into the experiment, at a variety of swarm sizes in an experiment containing two robot types and two terrain types to chose from. The horizontal line through the graph displays the expected median percentage performance for a system categorising at random between two options.

increase at a group size of 42 confirms that the system can indeed function at swarm like sizes within the two environment context.

The experiment containing three robot types and environments did not fair as robustly when it came to increasing swarm size. With results for 18 and 20 robots being the only data sets with a mean above 33% (the expected correct categorisation percentage for a system with three choices, decided at random). However, it should be taken into account that, while marked at 33% due to the size of the nest, typical categorisation percentage for a random system is most likely below this number. This is confirmed by the previous results illustrated in figure x in which the random system settles to a correct categorisation percentage of approximately 27%. If 27% is taken as the boundary for a poor median result from the data sets, the three robot system performs similarly to the two robot, dropping below the expected random performance threshold at a swarm size of 42.




Figure 5.13: Correct categorisation percentages recorded 500 seconds into the experiment, at a variety of swarm sizes in an experiment containing three robot types and three terrain types to choose from. Statistical expected median percentage is marked with a dark solid line, the expected median percentage based off of previous results is marked with a dark dashed line.

5.7 Chapter Summary

The work presented in this Chapter has shown that by using a hormone-inspired behaviour arbitration system a heterogeneous swarm of robots can categorise their abilities based on their performance in a selection of environments. In all experiments it has been shown that given stimulus availability the HIBAS was able to outperform the random system in terms of percentage of correct categorisation. It is also clear from the presented results that, given a simple choice between two environments, the hormone-system is capable of categorising successfully with even a small difference in robot traits.

By observing the collection of food in the environment with three terrain types and an additional robot type, it is clear that:

1. The HIBAS can increase likelihood of correct categorisation when presented with more complex choices.
2. The categorisation provided by the hormone-inspired system can be beneficial to the performance of a task, so long as there is a large enough change in robot capability.

These results found in this chapter indicate that the HIBAS is a system worth exploring further in terms of its ability to categorise hormone ability, though it does also highlight a weakness to the system. Realistically robots within the swarm will encounter a range of difficulties not limited to the type of terrain they are navigating. While a strong categorisation percentage may be achievable in these cases, it has been shown that strong categorisation does not always equate to strong task performance. This was seen in the set of experiments with only two environment types to chose from. In said experiment the difference between the robot wheel type was relatively small regarding their performance across each terrain and it was seen that the rate of collection was only marginally better on average than the random allocating system. These difficulties could be overcome by introducing additional levels of hormone adaptation assisting the swarm with not only preference but more advance behaviour switching and explicit control elements. To an effect the research conducted within this thesis has provided a strong foundation for this, having presented and examined multiple levels of hormone control. These layers have existed as direct control over motors, behavioural control and, most recently, behavioural preference. The next chapter within this thesis will attempt to combine each of these hormone control types controlling a swarm at multiple levels of control at once. Through this combination of hormone control types it will be identified whether or not the implementation of a hormone amalgamation to improve the energy efficiency of item collection is possible.

Chapter 6

A Multi-Hormone System for Arbitrating Traits and Behaviours in Dynamic Environments

6.1 Introduction

To this point previous chapters have only tested swarm systems with relatively simple experimental scenarios. While these chapters have shown that virtual hormone systems can be engineered to arbitrate and adapt swarms of robots amongst a small set of behaviours, it is yet to be shown how hormone systems could be used when a large array of behaviours and task types are available to a swarm. Evidence of virtual hormones being used to control such systems in simulation would prove the viability of virtual hormone control in non-abstracted tasks and create an argument for their implementation in physical systems.

Having already explored several applications for hormone inspired systems in previous chapters, virtual hormone systems have been shown to effectively regulate behaviours and preferences, respectively selecting appropriate states in dynamic environments and allocating robots to environments based on their performance across different terrains. This final chapter of experiments will attempt to combine these applications to create an energy efficient foraging swarm regulated by numerous, simultaneously functioning hormones. This will show that virtual hormone systems can be used to effectively adapt and regulate large, complex systems with diverse behaviours.

The hormones comprising the amalgamation operate at different levels of a behavioural hierarchy (illustrated in Figure 6.1), controlling preference, behavioural control and actuator control. Combining systems acting at these different levels of behaviour allows for the swarm to be controlled by hormones at every stage of operation, truly testing the combined systems capabilities and compatibility. This, alongside the fact that more than three times the number

of individual hormone types previously studied have been used in these experiments means that the number of hormones used in this amalgamation can be considered numerous.




Figure 6.1: Behavioural hierarchy for the behaviours investigated within this Thesis.

To begin the experiments, Section 6.2 first investigates virtual hormone driven motor control as a method to improve energy efficiency in the foraging swarm. This will focus on the need for adaptive motor speeds and their implementation, building upon the work described in Chapter 3.

Section 6.3 explores the compatibility between this new system and one governing sleep, similar to the system designed in in Chapter 4 will be investigated. The potential energy efficiency benefits of combining a sleep system and a virtual hormone framework are examined.

In Section 6.4, the swarm will be diversified, using the heterogeneous wheel types designed in Chapter 5, and a system capable of self analysis for task reallocation is combined with the previously established hormone speed and sleep regulation. Thus creating a system with 6 or more simultaneously acting virtual hormones in each member of the swarm, depending on the number of environments available to the swarm.

By testing the combination of these systems, a complex hormone system capable of arbitrating different elements of behaviour, will be shown to be effective for live adaptation. The implementation of this complex virtual hormone system will be effective for live adaptation and produce significant improvements to energy efficiency in foraging examples over individual hormone systems.

Finally, Section 6.6 gives a number of conclusions of the work and suggests future areas of investigation.

6.2 HIBAS Implementation for Control of a Foraging System with Deviating Motor Speeds

Hormone Inspired behavioural arbitration systems (HIBAS) have been studied using energy efficiency as the target output, as seen in previous chapters. However, the speed at which robots move and the efficiency of their movement, vital to energy efficiency, have not been investigated. When simulating the energy consumption of robots it is typically assumed that robots in the swarm are either moving at a specific speed, stationary or consuming a fixed quantity of energy in a given behaviour state Wilson et al. (2018); Liu et al. (2007); Lee & Ahn (2011); Pang et al. (2017). The start of this chapter will investigate the viability of virtual hormone implementation to directly control and adapt wheel speeds to achieve improved energy efficiency when foraging. A ‘demand’ concept will be present in the task that allows the user to specify, prior to or during use, the number of items to be gathered in a given time period. The purpose of this is to add an additional complexity for the swarm to overcome through adaptation.

6.2.1 Energy Characteristics of Psi Swarm Robot Hardware

To obtain realistic results from the simulated experiments, data was taken from the PSI swarm robot platform (Hilder et al. (2016)) to obtain a power model similar to that produced in Bonani et al. (2010) for the MarXbot. To construct a power model, power consumption was measured using a Keysight N6705B power analyser (Key (2016)). Results for power consumption as speed increases were recorded through 10 repetitions and a quartic trend line was fit to the mean of these results, this is illustrated in Figure 6.2. The resultant equation for power consumption with speed as the input was:

$$1.05 - 7.76 \times 10^{-3}s + 2.2 \times 10^{-3}s^2 - 8.89 \times 10^{-5}s^3 + 1.14 \times 10^{-6}s^4 \quad (6.1)$$

Where P is power consumption per second (Watts) and s is the current speed of the robot (cm/s).

When implementing this equation in the robot swarm simulation, the offset of 1.05 was reduced to 0.05, as it was assumed that most of the offset was due to the base consumption of energy used by robot peripherals. The offset of 0.05 was left to ensure a negative power was never experienced during experiments. The equation was also scaled for the appropriate time frame, ensuring that the correct amount of power per wheel was collected per experiment tick. This equation was then used at each time step to calculate the current energy consumption based on the speed of each individual robot. Energy consumption could then be used to feed into the value of energy efficiency that would be used to measure the fitness of the systems tested in the experiments presented in this chapter’s experiments.




Figure 6.2: Graph displaying the results of the power consumption of a Psi Swarm robot increasing motor speed gradually. 10 repetitions were taken for these results and a trend line has been fit to the mean of these results and is shown in red. The equation forming the trend line is shown in Equation 6.1.

After implementing Equation 6.1 in the simulation, the analyses of energy efficiencies at different speeds were conducted. In these tests, 20 robots foraged in a simple environment for 500 simulated seconds or until 100 food items were gathered. The average final energy efficiency (food item per unit of energy consumed) from 50 trials at speeds ranging from 1 to 50 cm/s were then plotted (illustrated in Figure 6.3). Taking the peak value of energy efficiency for a given speed, a value was chosen to act as a baseline for the following experiments.

6.2.2 Hormone Interaction With Motor Speed

To produce a hormone equation that controlled motor speed in a direct manner and at appropriate speeds given context, it was decided that the two primary influencing factors should be the item demand from the user and the evidence of negative performance.

The presence of frequent collisions and the decay present from failing to achieve task goals have been demonstrated as good indicators of negatives performance in the previous examples of hormone systems. These features were therefore used as the first step in the implementation of the new hormone system. The decay would reduce the hormone, and subsequently the speed, to an efficient settling point. Collisions would also reduce the hormone, thus inhibiting the speed of poorly performing robots and limiting their impact on energy consumption.




Figure 6.3: Graph showing average food gathered per energy unit consumed in a swarm of 20 foraging robots across 50 trials. The speed showing the greatest food collection per energy consumed is marked at 35cm/s.

Demand

‘Demand’, as a new feature to the virtual hormone system, required the development of a novel formula accounting for: a target number of items to be collected (to be specified before deployment), the allotted time to collect said items, the current collection rate throughout the experiment. Following this, Equation 6.2 was created:

$$D(t) = \frac{I_T}{t_T} - \frac{I_c + 1}{t} \quad (6.2)$$

In this equation $D(t)$ represents the demand function, I_T is the total number of items desired by the end of the allotted time period, t_T is the end time for the allotted period, I_c is the current number of stored items and t is the current time step. Decentralisation is required to remain ‘swarm like’ during the experiment, therefore the ‘demand value’ is only accessible to individual robots in the nest. The value is updated as they leave and used as their stimuli throughout their next period of exploration.

Equation 6.2 models the demand value to fluctuate as items were collected without incurring an exponential increase near the end of the experiment should the swarm only be a few items away from the target collection. By setting the demand as the difference between the required average rate of collection and the current rate of collection, the hormone value and speed could increase with repeated failure to meet target collection rates. This meant that speed would only slightly deviate from the optimal speed of travel. Gradual incrementation in this manner prevented an inefficient burst of speed late in the experiment to compensate for a lack of items collected.

With a function for demand in place, the two Hormone equations were produced (Return Hormone and Speed Hormone, shown respectively in Equations 6.3 and 6.4) to regulate the speeds and behaviours of each robot in the swarm. The hormones produced in these experiments were designed in the same format as Chapters 3, 4 and 5 with λ representing decay and γ representing the coefficient of stimuli.

Return Hormone

$$H_r = \lambda_r H_r + \gamma_r C \quad (6.3)$$

Where t is the current time step H_r is the return hormone, λ_r is the decay for the system and γ_r is the stimuli weighting.

The return hormone has a single stimulus, C , for collision detection. Although it does not regulate speed, it does feed into the speed hormone. The primary function of the Return Hormone is to identify the frequency of collisions detected by a robot, between walls or other robots. This information can then be used to decide if an individual robot should return to the nest having been unsuccessful, typically by exceeding either a fixed or similarly adaptive threshold. At this stage the threshold for returning was set to 50, with any value of H_r exceeding that resulting in a given robot changing behaviour state and travelling back to the nest site.

Speed Hormone

$$H_s = \lambda_s H_s + \gamma_{s1} D(t) - \gamma_{s2} H_r \quad (6.4)$$

Where $H_s(t)$ is the Speed Hormone, λ_s is the decay rate for the system, γ_{s1} is the weighting for the stimuli and γ_{s2} is the weighting for the inhibitor.

The speed hormone had two influencing factors. A stimulus, $D(t)$ (Demand), and an inhibitor, H_r . With these features in place, higher demand would result in faster activity, consuming more energy but reducing the item demand. Conversely, the system would slow down robots in poor positions or in areas densely populated by other members of the swarm, consuming

less energy while in a compromised position. It is worth noting that H_r was used in this case rather than C in order to smooth the response to collisions, rather than experiencing a sudden, large value inhibiting the system upon encountering a collision, H_r allows for the reduction to H_s to be smooth and gradual. This avoids the sudden loss of mobility in what could potentially be a one off collision.

While the speed of a robot does increase with the Speed Hormone, it doesn't have true direct control over the motor speed as has been seen in studies such as Stradner et al. (2009). Instead, the Speed Hormone system allows the robot to operate at the optimal travelling speed for energy efficiency. To avoid deviation from this speed at low hormone levels, the speed hormone has no effect on speed until it exceeds the value of 10. Values below 10 in speed hormone would have very minimal effect on the actual speed of the robot while still reducing energy efficiency by deviating from the optimal speed. After the value of 10, the speed hormone effects the speed with the relationship shown in Equation 6.5, providing potential speeds ranging between 35, for H_s values below 10, and 50 when H_s is fully saturated.

$$S = 33.33 + \frac{H_s}{6} \quad (6.5)$$

Parameters

Parameter values for the hormone equations (shown in Table 6.1) were selected empirically using the context of the experiments to decide on appropriate time scales for decay, these time scales were then converted to decay values using Equation 6.6, taking values for H_{sat} (the numerical value at which the virtual hormone will saturate) and H_{fin} (the smallest value deemed relevant to the hormone system) as 100 and 1 respectively. The period of decay chosen for the sleep hormone was based on the amount of time it would take for an ideally operating robot to locate and retrieve two food items. i.e., the time it would take to reach the centre of available items and return twice, travelling in a straight line while operating at optimal speed. This meant that under ideal operation stimuli from the previous collection would still be present when returning for the second time, allowing the hormone value to build. The period for decay for the return hormone was calculated for only a single full collection and the collisions in a previous search period should have minimal bearing on that of the next.

$$\lambda = \sqrt[n]{\frac{H_{fin}}{H_{sat}}} \quad (6.6)$$

Stimuli coefficients were subsequently chosen to provide adequate response when interacting at expected minimum and maximum values of decay and rate of collision.

Figure 6.4 shows the hormone value dynamics of the speed regulating hormone system. As previously mentioned the return hormone, indicated by the green line, inhibits the speed

hormone. The relationship between these two hormones can be seen in the graph. With severe reductions to the speed hormone value at the points in time where the return hormone is present. It is also worth noting, between 300 and 400 ticks, despite the lack of return hormone value, the speed hormone does not increase. This is indicative of the fact that item demand will have been met by the swarm for the given period. However, the speed hormone is seen to increase at approximately 450 ticks, indicating that the swarm is again behind schedule in item collection. As a result, the individual robot being monitored increases its speed to address this.

Graph showing Hormone behaviour in a single robot.




Figure 6.4: A graph displaying the dynamics in the hormone values used in an experiment in which a swarm with hormone regulated speeds conducted a foraging behaviour. The graph displays the hormone values of an individual robot within this swarm across 500 experimental ticks.

λ_r	γ_r	λ_s	γ_{s1}	γ_{s2}
0.9977	5	0.999	9	0.01

Table 6.1: Parameter values for the Return and Speed Hormones. Where $H_s(t)$ is the Speed Hormone, λ_s is the decay rate for the system, γ_{s1} is the weighting for the stimuli and γ_{s2} is the weighting for the inhibitor.

6.2.3 Comparison Systems

In order to test how effective the designed hormone systems were, two additional systems were produced for comparison. The first had no adaptive element, keeping all robots at optimal speed (35 cm/s) while foraging. This system was not influenced by ‘demand’ and should highlight the point at which speed adaptation is required to obtain remaining items required in the collection. In order to keep environmental awareness consistent across the three systems, the return hormone was implemented across all systems, allowing swarm members to return to the nest site should they encounter too many collisions.

Engineered Adaptive Comparison

The second comparison system featured an on-line adaptation method similar to reinforcement learning. This engineered adaptation was driven by the same function for demand as featured in the virtual hormone system. This style of online engineered adaptation has been used in the past to modify swarm traits, finding optimal partition lengths in Buchanan et al. (2016) modifying travel distances based on success and failure of swarm individuals.

The adaptive system, designed for speed control, stepped the robot motor speeds up or down depending on the value of demand upon returning to the nest site. Positive demand values would increase speed, and negative values would decrease it. As with the hormone system, this would allow speed to be increased or decreased (and hence increase or decrease energy expenditure) in relation to collection requirements.

The increments and decrements made by the engineered system were influenced by demand, providing a variable adaptation to the system. A base change of 1 was applied based on the sign of the demand in addition to a change proportionate to the value of demand itself, increased by a coefficient of 20 to make suitable changes to the speed value. These values were tuned via iterative selection to produce strong rates of collection and energy efficiency across a wide variety of task demands.

The base change was used so that the swarm can catch up to the required collection rate even when demand is small. If this change was not implemented, increments based solely on demand would be too small to have a perceivable effect on robot speed. The same effect could not be achieved by increasing the coefficient of demand because the system could react too quickly to large disparities in current collection rate versus required rate, overcompensating

by a large margin.

6.2.4 Analysis Of Systems Highlighting The Need For Adaptation

After designing these systems, preliminary tests were conducted demonstrating why adaptation is required for the foraging task. This section will elaborate on the environment in which the systems were tested, detail the key features of the simulations and discuss the results produced from the experiments.

Environments

The three systems discussed in this paper were tested in two environments. The first is a square environment measuring 15×15 m. The first 2 meters of the environment were assigned as the nest area, highlighted in grey as illustrated in Figure 6.5. This environment provided an arena for simple operation, identifying whether the system, under only the pressure of the specified demand could operate effectively.




Figure 6.5: Screenshot of first simulated environment used in initial experimentation with the new speed regulating systems. Food items are shown as black circles in the white environment, puck robots can be seen waiting in the nest area (light grey).




Figure 6.6: Screenshot of second simulated environment used in initial experimentation with the new speed regulating systems. Food items are shown as black circles in the white environment, puck robots can be seen waiting in the nest area (light grey). Obstacles creating corridors are illustrated in dark grey.

The second environment (illustrated in Figure 6.6) instead measured 20x10 metres though retained a similar nest layout to the first. Four funnelled corridors were included in this environment to act as obstacles. These increase swarm density during exploration and provides additional difficulty to the tested systems, akin to that of a group of robots attempting to complete tasks in industrial settings such as mines, power plants or drainage systems, where space could be limited. This congestion will not only limit the success of the robots by slowing them down, but short range collision sensors will be triggered more frequently, meaning that the return hormone will potentially instruct robots to return home too early. This will heavily test the adaptability of the system, giving the combination of hormone systems a greater challenge, making the probability of one system disrupting the other in a negative fashion more likely.

Simulation

The experiments were performed in the ARGoS simulator (Pinciroli et al. (2012b)) a multi-robot simulator used to simulate large robot swarms. It was assumed that each of the robots

was equipped with a food sensor, allowing them to identify food items within a 2m radius.

Each test was executed for 500 simulated seconds (each simulation time step lasting 0.1 seconds) or until the target number of food items were collected.

The number of replicates required for consistent results were determined by performing cumulative mean tests as specified in Robinson (2004). This test indicated that the minimum number of trials required for consistency was 36. Therefore, 36 was the lowest number of replicates used when testing these systems.

6.2.5 Results

The data collected from the experiments conducted was compiled into box plots, showing the performance of each system in terms of energy efficiency and collection rate. These results are illustrated in Figure 6.7 for environment 1 and Figure 6.8 for environment 2.

Environment 1

Visual inspection of the first environment (Figure 6.7) shows that the static speed system has a fairly consistent level of food collected per energy unit used as the demand increases. This is expected due to the lack of change in speed, though the lowest target number for item collection does see a drop in energy efficiency when compared with the rest of the collection rates. This is because not all of the robots in the swarm will have returned to the nest by the time the experiment terminates having reached the target number of items. This will result in unnecessary energy consumption from the robots unable to return food items within the short period of the experiment.

The downside of this consistent energy consumption is the inability to reach greater item target numbers. This drawback can be seen in the discolouration of the box plots starting at 100 food items required and saturating to red, indicating a collection of less than 70% of the required items, by 130 required items.

Disregarding the lack of success in large item demand experiments, the results from the static speed system provide a strong baseline for energy efficiency. Giving a clear target for the other two more intelligent systems to aim for.

When inspecting the results of the two adaptive system it is immediately obvious that target collections are met more consistently with the demand function introduced to the system, with discolouration starting at 120 in the engineered system and 130 in the hormone system. In the engineered system the collection rate drops to approximately 80% by the 150 item goal while the hormone system still manages to collect upwards of 90%.

In terms of energy efficiency the engineered adaptive system follows a similar initial trend to the none adaptive system. The similarity is maintained until an item target of 50, at




Figure 6.7: Results for the three systems tested in environment 1. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).

which point the engineered system becomes increasingly less efficient. Table 6.2 supports this, showing that there is no significant difference in the data sets of the Engineered and static systems until 70 target items. At this point the systems diverge as the engineered system consumes more energy.

These results also show that the hormone system managed to outperform both systems in regard to energy efficiency. With a significant difference versus the engineered adaptation and increased median result at every collection target excluding 10, the hormone system results can be seen arcing over those of the engineered system after starting at a similar point. Similarly, when compared to the static system, the hormone system shows significant increases to the food collected per energy used in all cases but targets of 10, 120 and 130 items. The similarity in energy efficiency of the hormone and speed systems at item targets of 120 and 130 can be explained by the speed increase of the hormone system in cases of very high item demand, actually reaching collection targets while the static system misses them by a large margin.

The efficiency of the hormone system over the static and engineered systems was explained by three factors:

Sensitivity: The hormone system is sensitive to collisions and capable of not only returning robots to the nest due to collisions, but also reducing speed due to the prolonged influence of collisions.

Dispersion: Rather than consistent speeds, or speeds of specific increments, the speeds of the hormone driven robots fluctuate during the search. This leads to not only more efficient speeds, but also more heavily dispersed robots, as a by product of diverse speeds amongst the swarm. This in turn will lead to less traffic and more energy efficient item collection.

Gradual Variability : Due to the fact the hormone system fluctuates over time, speed can build across the length of a search rather than having to react immediately at the nest in a manner that is potentially exaggerated or understated. As was the case in the engineered system, in which relatively large changes in speed had to be taken upon returning to the nest, potentially stepping over the best value of speed for the next exploration.

Environment 2

The results for the second environment, the increased length of environment and introduction of corridors, predictably show a notable decrease in percentage of target collection completed. The static system started to fail collection targets at 50 items and the engineered adaptive

System Type		Engineered Vs Static	Hormone Vs Static	Hormone Vs Engineered
Item Number	Target			
10		0.8550	0.0330	0.0053
20		0.1648	p < 0.0001	p < 0.0001
30		0.1800	p < 0.0001	p < 0.0001
40		0.2626	p < 0.0001	p < 0.0001
50		0.0906	p < 0.0001	p < 0.0001
60		0.8227	p < 0.0001	p < 0.0001
70		0.0068	p < 0.0001	p < 0.0001
80		0.0262	p < 0.0001	p < 0.0001
90		p < 0.0001	p < 0.0001	p < 0.0001
100		p < 0.0001	p < 0.0001	p < 0.0001
110		p < 0.0001	p < 0.0001	p < 0.0001
120		p < 0.0001	0.0199	p < 0.0001
130		p < 0.0001	0.3984	p < 0.0001
140		p < 0.0001	p < 0.0001	p < 0.0001
150		p < 0.0001	p < 0.0001	p < 0.0001

Table 6.2: Environment 1: Wilcoxon rank sum tests comparing the three systems for the tested item collection targets between 10-150 in terms of energy efficiency. Significant differences (indicated by a p value of < 0.05) are highlighted in **bold**.

system starting to fail at 70. Compared with these, the change to collection rate in the hormone system is substantially less reduced. The results show the hormone system falling to a 70% collection rate at the 130 item target mark, showing a considerable increase in collection performance versus the two comparison systems.

In terms of energy efficiency there is again an expected drop in performance, when compared to the first environment, across all experiments due to the larger, more cluttered arena.

Analysing the systems tested in this environment, there is very little statistical similarity. Table 6.3 shows that almost all of the data sets at each item target number, with the exception of the first 5 item targets of the engineered versus static system, are all significantly different. The data produced from this environment does however follow very similar patterns those of the first environment. The static system maintains a consistent energy efficiency, though dipping slightly in the case of the smallest collection target. The Engineered system, while improving collection, does little to benefit energy consumption and lessens as target numbers increase. The hormone system, while exceeding the two comparison systems in both collection and energy efficiency, as it did in the first environment, does so in a much more exaggerated manner in the second environment.

6.2.6 Section Summary

The results from these experiments show that with a system taking variable demand for a number of actions complete/items collected adaptive systems can be implemented to




Figure 6.8: Box plot results for the three systems tested in environment 2. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).

System Type		Engineered Vs Static	Hormone Vs Static	Hormone Vs Engineered
Item Number	Target			
10		0.2482	p < 0.0001	p < 0.0001
20		0.6918	p < 0.0001	p < 0.0001
30		0.3432	p < 0.0001	p < 0.0001
40		0.1010	p < 0.0001	p < 0.0001
50		0.0817	p < 0.0001	p < 0.0001
60		0.0020	p < 0.0001	p < 0.0001
70		0.0002	p < 0.0001	p < 0.0001
80		p < 0.0001	p < 0.0001	p < 0.0001
90		p < 0.0001	p < 0.0001	p < 0.0001
100		p < 0.0001	p < 0.0001	p < 0.0001
110		p < 0.0001	p < 0.0001	p < 0.0001
120		p < 0.0001	p < 0.0001	p < 0.0001
130		p < 0.0001	p < 0.0001	p < 0.0001
140		p < 0.0001	p < 0.0001	p < 0.0001
150		p < 0.0001	p < 0.0001	p < 0.0001

Table 6.3: Environment 2: Wilcoxon rank sum tests comparing the three systems for the tested item collection targets between 10-150 in terms of energy efficiency. Significant differences (indicated by a p value of < 0.05) are highlighted in **bold**.

achieve requested goals effectively. The results have shown that, while in some cases this adaptation can contribute to additionally expended energy, adaptation can be implemented in a manner that improved energy efficiency in low demand tasks and sacrifices very little in high demand tasks. This was proven through the successful implementation of the HIBAS in two environments of varying complexity, and the significant improvements the hormone inspired system subsequently provided versus both the adaptive and static systems.

This has shown that the a hormone system for direct motor speed control can viably improve operation in a foraging context. The next section investigates the practicality of combining the hormone system presented in this section with a hormone sleep system, based on the controller discussed in Chapter 4.

6.3 Introduction Of The Sleep Hormone To A Foraging Swarm

The foundations of this introduced sleep hormone system are very similar to those presented in Chapter 4, following the same behaviour states as shown in Figure 6.9. The hunger hormone detailed in previous chapters was given an identical structure. However, due to the slight change in context to the foraging system, the stimuli to the sleep hormone (now represented by $H_{\sigma}(t)$ to avoid confusion with the speed hormone system) in the system was edited from

α_σ	λ_σ	$\gamma_{\sigma 1}$	$\gamma_{\sigma 2}$	α_h	λ_h	γ_h
0.01	0.999	0.01	0.06	0.0.015	0.999	10

Table 6.4: Parameter values for the hunger hormone and new sleep Hormone. Parameter descriptions can be seen alongside Equations 6.7, 6.8 & 6.9.

the original equation:

$$\text{Sleep Hormone (Chapter 4): } H_\sigma(t) = \lambda_\sigma H_\sigma(t-1) + \gamma_\sigma H_A(t-1) \quad (6.7)$$

To include both an α value and an inhibitor in the form of $\gamma_{SI}d$ (where $d(t)$ is the function of demand presented earlier in this chapter) resulting in the new equation:

$$\text{New Sleep Hormone: } H_\sigma(t) = \alpha_\sigma + \lambda_\sigma H_\sigma(t-1) + \gamma_{\sigma 1} H_A(t-1) - \gamma_{\sigma 2} d_t \quad (6.8)$$

The introduction of an α value offsets the settling point of the hormone. This allowed for the implementation of the demand based inhibitor ($\gamma_{AI}d_t$) and ensured that the hormone could fluctuate below the settling point without producing a negative value. The demand inhibitor itself created a larger decrease to the sleep hormone under high demand circumstances, assisting the decay already present in the hormone and reducing sleep times when the swarm's rate of collection was inadequate.

Meanwhile H_h was kept in the same format, using the equation:

$$\text{Hunger Hormone: } H_h(t) = \alpha_h + \lambda_h H_h(t-1) + \gamma_h C \quad (6.9)$$

Where C is a Boolean value representing whether the robot successfully returned a food item to the nest site or not.

The parameters used for the hunger and sleep hormones were calculated in a similar manner as Section 6.2.2, using the approximate time scale across which the hormones were expected to operate and thereafter tuning stimuli for the fitting reaction. The parameter values selected for the coming experiments are displayed in Table 6.4.

6.3.1 Preliminary Tests For Sleep Hormone In A Demand Lead Foraging Task

The initial tests conducted on the new sleep hormone system used the same environments as the previous section and operated until a time limit of 500 seconds or until the target number of items was reached. A cumulative mean test indicated that a minimum of 14 trials were required. To ensure certainty, 20 trials were run.




Figure 6.9: State machine for foraging hormone system illustrating state names and the relative hormone ratios required for transition as first displayed in Chapter 4. H_S is now replaced by H_σ .

Environment 1

Observing the results for the first environment (illustrated in Figure 6.10) the results appear very different to those of the previous three systems. The energy efficiency starts low, peaking momentarily and, after a dip to median performance, increases as the number of target items does. This pattern leads to an increased energy efficiency at all item targets compared to the previously tested static speed system and considerably better efficiency performance at item targets greater than 120 for the other two adaptive systems.

The initial spike in performance from this pattern is explained by the removal of poorly positioned robots at deployment. Those robots starting off in large groups will enter the sleep state either immediately or very soon after exploration. This initial state selection is then diluted as robots make more passes between the nest and the food area, seen as the Food Collected Per Energy used (FPE) reduces to a similar level as the non-adaptive system seen in Section 6.2.5. The gradual increase to FPE thereafter is due to the sleeping of poorly performing robots across greater periods of time, while robots with better positioning within the arena are able to collect food items more effectively.

While this system sees several increases to performance in terms of energy efficiency, it sacrifices this for poor performance in terms of item collection, with collection starting to drop at item targets of 90, lower than even the static system in the previous section. This is expected as the system actively impedes collection speed, with the sleep state removing




Figure 6.10: Box plot results for the hormone inspired sleep system tested in environment 1. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).

swarm members for brief periods of time.

Though the collection percentage is lower in the sleep system than in the systems previously examined, this does show that it may be beneficial from the perspective of energy efficiency to combine the speed and sleep systems. With the intention of reducing the decrease to FPE seen in the adaptive speed systems as item target increases and using the speed system to compensate for the poor collection performance seen at targets above 90.

Environment 2

The benefit of this enhanced hormone system is further proven in the second environment. Following a similar pattern to the first environment, the energy efficiency increases with the target number (illustrated in Figure 6.11). In this environment, the sleep system is able to outperform the static and engineered system in terms of energy efficiency across all item target values. In addition to this, while not able to compete at lower item targets, after 80 items the sleep system largely outperforms the hormone speed system.




Figure 6.11: Box plot results for the hormone inspired sleep system tested ed in in environment 2. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).

This increase to energy efficiency is due to the sleep system regulating the number of robots present in the corridors at any given moment, retaining poorly performing robots until demand is high and as a result increasing the productivity of the foraging swarm.

Again these results, while producing good values for energy efficiency, sacrifice collection rate. With collection similar to the static system, failing past 50 items.

6.3.2 Combining the Sleep Hormone With The Speed Deviating System

In an attempt to combine the benefits of both hormone systems and to identify the viability of combining existing hormone systems, the speed hormone was added to the already established sleep system. The parameter values established in prior testing were again used for the combined system. The speed hormone acted explicitly on motor speeds during exploration and the sleep hormone system regulated higher level behaviours.

The performance of this new combined system is illustrated in Figure 6.12. The first obvious improvement to the system can be seen in the results from the first environment, this set of




Figure 6.12: Results for the combined hormone sleep and speed system tested in both environments. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).

data achieves the highest average collection rate of any system at a required collection of 150, obtaining an average of 92.3% of the needed items.

In addition to this the combined system achieves a significantly greater energy efficiency versus the sleep system at all item targets between 50 and 110 in the first environment and all item targets before 100 in environment 2 (p values for these tests can be found in Table 6.5). At higher value item targets the energy efficiency still crosses over, though the exceptional item collection rate more than compensates for this.

Relative to the adaptive speed system the combined hormone system obtains very similar results in the first environment at target item values below 70. Though there are large improvements to the energy efficiency at item targets larger than this. This increase to performance is mirrored in environment 2, though with a consistent increase at all item target values.

The substantial improvement in performance is proposed to be the mutually beneficial actions of the separate systems. It allows the system to avoid the circumstance in which positioned poorly robots in a high demand context might travel at high speeds that a cause large drain to power for poor returns.

It is clear from these results that these systems work better in combination than separately. This shows a strong symbiosis of already established hormone control, verifying the viability of combining hormone systems.

Item Target Number	Hormone Speed Vs Hormone Combination (Environment 1)	(Environment 2)	Hormone Sleep Vs Hormone Combination (Environment 1)	(Environment 2)
10	0.9680	0.0047	0.2315	0.019
20	0.8830	0.0040	0.1653	0.0056
30	0.5290	p < 0.0001	0.5831	p < 0.0001
40	0.6017	p < 0.0001	0.0810	p < 0.0001
50	0.5290	p < 0.0001	p < 0.0001	p < 0.0001
60	0.0809	p < 0.0001	p < 0.0001	p < 0.0001
70	0.0283	p < 0.0001	p < 0.0001	p < 0.0001
80	0.0047	p < 0.0001	p < 0.0001	p < 0.0001
90	0.0675	p < 0.0001	p < 0.0001	p < 0.0001
100	0.0024	p < 0.0001	p < 0.0001	0.0430
110	0.0910	p < 0.0001	0.0024	0.0002
120	0.0763	p < 0.0001	0.9042	p < 0.0001
130	0.1081	p < 0.0001	0.0211	p < 0.0001
140	0.0227	p < 0.0001	p < 0.0001	p < 0.0001
150	0.2648	p < 0.0001	p < 0.0001	p < 0.0001

Table 6.5: Wilcoxon rank sum tests comparing the combined hormone system with both the speed hormone system and the sleep hormone system, in both of the previously established environments. Tests were conducted for item collection targets between 10-150 in terms of energy efficiency. Significant differences (indicated by a p value of < 0.05) are highlighted in bold.

6.4 Introduction Of Environment Selection Hormones With Sleep and Speed Regulating Hormone Systems

With the viability of a larger hormone system confirmed, the next step taken was to combine the hormone combination system presented in the last section with yet more hormone arbitration. This was an important step because, while it has been shown that hormone systems can interact to produce satisfactory results, the current combination of hormone systems experience minimal detrimental interactions. In terms of behavioural arbitration, the speed and sleep hormone systems do not interfere with one another.

This section presents the amalgamation of the speed hormone, sleep hormone and a hormone system capable of monitoring the emergent success of the swarm under different conditions, implementing the designs shown in Chapter 5. The monitoring of success and ensuing environmental preference, was driven by the speed at which items could be collected from the environments. Therefore, it is essential to investigate if the preference system will still be capable of categorising robots within a heterogeneous swarm effectively with an adaptive speed mechanism in place.

As such, the environment for these tests required a diverse terrain and multiple directional options alongside heterogeneity amongst the swarm.

Terrain Type	Wood Suited Wheels	Grass Suited Wheels
Grass	0.6	0.7
Wood	1	0.8

Table 6.6: Speed Coefficients for Heterogeneous Robot Wheels on Different Terrains. These values were used to simulate different capabilities amongst a heterogeneous swarm of robots. This created a requirement to assign robots to appropriate environments in order to obtain efficient performance from the swarm.

6.4.1 Environmental Setup

The new environment used for testing in this scenario was identical to that used in Chapter 5 Section 5.4.4, featuring two different terrains designed to challenge robots with two specific wheel types.

In order to incorporate heterogeneity into the swarm, while still using the energy characteristics presented in Section 6.2.1 to measure energy efficiency at different speeds, each wheel type was given a speed coefficient for respective terrains. These coefficients (displayed in Table 6.6) inhibited the speed properties of the wheels based on the ground a given robot was travelling on, these values are shown in Table 6.6. While not as realistic as the data used for wheel speeds in previous environment preference experiments, this allowed for the testing the combination of systems without extensive testing of robotic hardware.

6.4.2 Effect Of Demand On Environment Selection When The Speed Hormone Is Combined With Environmental Preference Hormone

Before fully combining the systems, the speed hormone was added to the Environmental Preference System. The performance of the selection system was then measured in the same manner as Chapter 5, looking at the proportion of robots active in the environments they were best suited to as a percentage.

In order to incorporate the speed hormone to the directional preference hormone system, demand functions identical to that previously produced in Section 6.2.2 were created for both the north and south environments, taking only items collected in the respective environment into account when producing demand. Depending on the environmental preference when returning to the nest site, robots within the swarm would then update their demand stimuli with the corresponding demand value.

The full results of these tests are illustrated in Figure 6.13. Minimal differences were found in median categorisation across the range of item targets. Further, these were not found to differ from median categorisation found when the speed hormone was not included in the system. As the speed hormone did not appear to have a negative effect on the environmental preference hormones, it was deemed reasonable to further add the sleep hormone to the




Figure 6.13: Effect of item target values driving different demands in the speed hormone on the percentage of robots taking preference to their optimal environment. The categorisation system running with no speed hormone present is marked as 'no adapt'. Results show a fairly consistent percentage of categorisation relative to item target, indicating that the speed hormone has minimal effect on the categorisation abilities of the combined hormone system.

system.

With minimal negative interaction between the speed regulating and environmental preference hormones, it was deemed reasonable to continue with the implementation of the combined hormone system with the introduction of the hormone driven sleep system.

6.4.3 System combining Sleep, Speed and Preference Hormones.

To observe the performance of this new system, the various combinations of hormone systems were tested in combination in the new multi-terrain environment. First the preference system was tested on its own, the results from this are illustrated in Figure 6.14. These results would act as a baseline to the additional systems as results from this new environment, with new




Figure 6.14: Box plot results for the hormone preference system tested in the environment containing two difference terrain types. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).

task complexities, would be incomparable with data from previous experiments.

When adding the sleep hormone to the system (results illustrated in Figure 6.15), results for energy efficiency are consistently raised past the first item targets of 5, as shown by median results increasing by approximately 25% for item targets past 70. While there is a large improvement in terms of energy efficiency, adding the sleep hormone only results in a slight increase to collection rate, with the cut off point for collection becoming poor (below 90% of the target item collection) shifted from 80 to 90.

When the speed hormone is added to the system in the absence of the sleep hormone energy efficiency suffers considerably. This is seen with the consistent drop in efficiency results illustrated in Figure 6.16 when compared to the baseline results. However, this drop in efficiency is traded for a substantial improvement to collection rate, moving the cut off point for poor collection to 120 target items. These results are expected with the additional speed fluctuation and if the results from previous hormone combinations hold consistent, the




Figure 6.15: Box plot results for the hormone preference system, combined with the sleep hormone system tested in the environment containing two difference terrain types. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).




Figure 6.16: Box plot results for the hormone preference system, combined with the speed hormone system tested in the environment containing two difference terrain types. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).

addition of the sleep system to the preference/speed hormone regulation should amend the poor energy efficiency while maintaining the item collection rate.

Combining all three systems provides the best result in terms of item collection, maintaining adequate collection until the 130 target item trial (illustrated in Figure 6.17). Simultaneously, the system that combines all three hormone types is capable of accomplishing competitive values for energy efficiency. These values show improvements across all item targets for the standard preference and combined speed hormone results. The three hormone system only marginally under performs in energy efficiency versus the combined preference and sleep system, although it shows much greater item collection percentages. These results suggest that the fully combined system as the strongest of the system permutations when considering both item collection and energy efficiency.

These results show the strength of these combined hormone systems and as a result the




Figure 6.17: Box plot results for the hormone preference system combined with both the sleep and speed hormone system, tested in the environment containing two difference terrain types. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).

feasibility of large, connected, hormone systems. In order to further test these systems, the next section investigates the robustness of combined hormone systems, presenting the system with additional complexities within the given task.

6.4.4 Scalability of Final Amalgamated Hormone System

To introduce a key element of difficulty to the system presented in this chapter a scalability test was conducted. With more robots in the swarm, the more difficult it will be to move effectively within the environment without slowing due to clutter. Along side this, with greater swarm density robots may receive over-stimulation from the transmitted hormones of the increased number of robots or there may be too much competition for food items, with multiple robots travelling to the same item simultaneously. With these additional negative features present it will be difficult for robots to form accurate preferences to terrain due to the fact that these negative features may have a greater effect on performance than the speed variance provided by the different wheel types.

The scalability tests were conducted by increasing the number of robots in each simulation by 6 for each set of trials, testing swarm sizes ranging between 12 and 60. In each experiment the target number of items was set to 100 and in every test all of these items were retrieved. The experiments conducted terminated after 500 simulated seconds or if the target number of items was reached. The item target of 100 was chosen due to the variability in performance at said target number across each of the previously tested systems. This indicated that this number of items is an area of interest, providing substantial challenge to some systems while still an achievable goal to others.

The results of the scalability test can be seen illustrated in Figure 6.18. It can be seen that energy efficiency decreases linearly with the increase in members of the swarm. This was expected with the increased difficulty to the task as, while the amalgamated system is able to augment performance with a given swarm size, additional or unneeded robots will still create detriment to performance. Through the linear nature of this performance degradation, a user can select a swarm size which is suitable for a given task, trading off energy efficiency for the speed at which items should be gathered.




Figure 6.18: Results for energy efficiency as the number of robots in the swarm increases from 12 to 60. The introduction of additional robots can be seen to decrease energy efficiency in an almost linear fashion. This is to be expected when increasing swarm size but confining operation to a confined environment. With more robots present traffic when collecting items will slowly build, making it more difficult for the swarm to collect items efficiently.




Figure 6.19: Test environment for hormone amalgamation containing two areas of different terrain, highlighted in white and green as well as funnel like obstacles designed to increase swarm density for a section of both terrains.

6.4.5 Reintroduction Of Obstacles For Amalgamated System

As previous experiments had introduced a second environment containing an obstacle in the form of corridors, it was deemed appropriate to reintroduce these obstacles to test the amalgamated system and be certain in its capabilities compared to the other hormone systems presented in this chapter. The environment, with these obstacles introduced to both available terrains, can be seen illustrated in Figure 6.19.

As was found with the introduction of previous obstacles, performance in terms of energy efficiency and collection rate was reduced (results shown in Figure 6.20).

The amalgamated system seems to deal quite well with the addition of obstacles. The global median of the results taken from the environment with the obstacles Vs the clear environment shows a change of no more than 23%. The robustness of the amalgamated system to the introduction can be seen when comparing the to the drop in energy efficiency to other systems that have been introduced to an obstacle filled environment. The speed regulated system tested in section 6.2 showed a drop of 41.5% and the sleep regulated system tested in section 6.3 showed a 36.7%.




Figure 6.20: Results from the hormone amalgamated system in an environment with two directions to forage in, each with different terrain types and both containing obstacles. Target number for items collected ranged from 10 to 150 items of food. Percentage of the items requested versus those collected by the end of the simulation is indicated by colour (Green 100% and Red <70%).

In addition to this the drop in collection rate sits between the two initially tested systems, experiencing a 41.6% drop in target collection success (target collection success referring to the last item target the system was capable of collecting above 70% of the required items) versus the 44.4% drop seen in the sleep system and the 33.3% drop seen in the speed system. Showing that collection rates were not dropping by a large margin to accommodate for the success of the amalgamations energy efficiency robustness.

The relative robustness of the amalgamated system is due to the fact that there are multiple arena options for the robots to explore. Subsequently they are able to avoid areas of high density, instead performing tasks in a different environment, still acting productively. This shows that the system is able to work well to reallocate robots when presented with a challenge and furthers the argument for the benefit of multiple hormone adaptations acting at once.

6.5 Robustness Analysis - Modelling Wear Over Time To Test System Robustness

As robots explore environments, they experience degradation of their actuators due to exposure to environmental factors. Over time performance can be effected by dirt clogging gear boxes, heat putting strain on motors or tires breaking down, among other faults. Testing systems susceptibility to faults is not a new concept. However, most previous work look only at the extreme cases of faults in locomoting robots i.e. complete motor or sensor failure (Tarapore et al. (2015); Winfield & Nembrini (2006); Ferrell (1994)). The robustness tests executed on the combined hormone system explored in this section seek to identify what effect the degradation of robots locomotive systems has over time.

The model the deterioration of the robots in the simulation, a coefficient (M_c) was added to the speed at the wheels. The coefficient would take a value between 1 and 0, reducing the speed at which the robots could travel at, while still consuming the same amount of energy as they would had the speed not been decreased. This multiplier was defined by the following equation:

$$\text{Motor Degradation Coefficient: } M_c(t) = M_c(t - 1) - \frac{s}{\zeta} \quad (6.10)$$

Where s is the current speed and ζ is a weighting attached to the motor speed. ζ was introduced to modify the severity of the deterioration. This variable would be modified from trial to trial in the robustness tests, decreasing and increasing the rate of decay in speed with the respective rise and fall of ζ .

6.5.1 Wheel Wear Results

Using the Spartan package Alden et al. (2014) an effect magnitude test was performed across a range of ζ values, producing A-test scores for the various data sets versus a system where the motor coefficient had no effect. ζ was modified across a range that resulted in an expected reduction to speed of 0%-100% after travelling 1.75Km (the expected distance travelled if maintaining optimum speed consistently across the entire experiment). In this test the item target was fixed at 120 as this had previously been shown as a value that required adaptation and could be improved by each of the independent systems operating constructively.

The results for this robustness analysis are illustrated in Figure 6.21. As can be seen, the system is capable of adapting and obtaining results with little difference to that of a system in perfect working conditions. Through the several adaptations available to the system, resting robots that are poorly performing, increasing the speed when collection rates do not match demand and allocating environments based on success or density. Due to this the system only




Figure 6.21: Graph showing the A-Test scores of the combined hormone system from trial to trial as Zeta increases in effect, degrading the motor system more quickly over time. The Zeta value can be seen starting to have a greater effect on the system as it approaches 100%.

encountered substantially different data sets when the system was presented with failure of motor speed on a large scale i.e. greater than 90%.

6.6 Chapter Summary

This Chapter has explored the viability of numerous simultaneously functioning hormone inspired systems. To address this, a speed controller for a foraging swarm was designed using a hormone inspired system and proven to be effective for energy efficient item collection at a number of different item targets. This system was then combined with a previously developed sleep system. The combination of these two systems addressed issues found amongst each of the individual systems, creating large increases to performance with minimal drawbacks. Based on this success a third hormone system was introduced, allowing members of a heterogeneous swarm to form a preference for environment, based on how successful individual robots assessed themselves to be in a given terrain. This new system tested with the speed adapting virtual hormone, identified as the system that would cause the most issue when attempting to categorise robots, was still able to effectively categorise robots, with limited change as demand increased.

Finally, the combination of all of the hormone systems was tested. While not producing the best energy efficiency of the tested systems, the amalgamated hormone system produced the best combination of collection rate and energy efficiency for the environment the system was tested in. Considering the total performance of the system should definitely take into account both energy efficiency and item collection, as the values they represent show task effectiveness and task completion respectively.

This system was also found to be very robust, showing an expected reaction to increases in swarm density and scalability. Additionally, the system was found to be resistant to wear over time. With little change to the baseline system found as the rate of motor degradation was modified across a large spectrum.

The results from the work in this chapter have shown that a complex system controlled almost entirely by virtual hormones can be an effective adaptation system within a swarm robotic context.

Chapter 7

Conclusions and Future Work

This chapter compiles the findings and contributions of this thesis, presenting the conclusions and contributions made in each chapter. Content from the thesis will be summarised within both the context of these chapters and the thesis as a whole, readdressing the General Hypothesis made in the introduction.

7.1 Thesis Summary

The breakdown of each chapter within this thesis is as follows:

7.1.1 Chapter 2 - Background

This chapter reviewed a wide range of articles relevant to the understanding of later chapters. In particular, this chapter provided background specifying robot swarm features presenting arguments for definitions along side the positives and negatives of their implementation. The chapter then visits the topic of biomimicry, explaining its relevance and previous engineering work that it has been utilised to improve. Of particular importance this section introduces virtual hormones for the first time, highlighting instances of the first cases uses of virtual hormone system for robotic control. Finally the chapter reviews methods of robotic adaptation, highlighting the fact that typical adaptation is performed offline and how this may be unsuitable for the tasks swarm robots are best suited for.

7.1.2 Chapter 3 - Virtual Hormones For Explicit Control

As the first experimental chapter, the work introduced here is preliminary and identifies the initial feasibility of re purposing virtual hormones to improve performance in within a swarm context. The findings here showed that virtual hormones could be used to improve the mapping potential of a swarm. This was achieved through the addition of the adaptive traits

inherent to hormone systems, allowing for appropriate dispersion of the swarm in a variety of environment types. Additionally this work identified that, in order to further test virtual hormone systems effectively, investigating different levels of control will more productive than creating bespoke hormone control systems for actuators.

7.1.3 Chapter 4 - Virtual Hormones For Energy Efficient Task Allocation

Continuing from the previous chapter, this work now moves to behavioural control. In this chapter a previously tested system introducing a sleep state to a swarm of foraging robots was augmented with a virtual hormone system. The introduction of this system allowed for fluctuation in sleep time, based on the environmental interaction of individual robots. This hormone system proved more effective than a similar swarm with a general sleep time optimised via genetic algorithm. The success of this systems control of behaviour states suggested that it may be interesting to investigate an even higher level of abstracted control, prompting the design and examination of the system introducing the concept of preference to a foraging swarm in the next chapter.

7.1.4 Chapter 5 - Virtual Hormones For Task Allocation By Self Identifying Traits

This chapter features a more complex and diversified heterogeneous swarm. This intention of this chapter's investigation was to identify whether a freshly deployed swarm with no explicit knowledge of the capabilities of individual robots could use hormones to rank one another on their performance across different terrain types. The robots in the swarm would then use this information to collect items in the environment they were best suited to exploring. This was found to be more effective than a system randomly selecting robots to an environment even with a relatively small difference in robot capability. Having tested a variety of different levels of control across the experimental chapters covered so far, it seemed that with these promising results the final set of experiments should research the amalgamation of the control systems designed so far.

7.1.5 Chapter 6 - A Multi-Hormone System For Arbitrating Traits and Behaviours in Dynamic Environments

This chapter investigated the amalgamation of all of the levels of hormone control previously analysed within the thesis, combining actuator control, behavioural control and behavioural preference. Creating an amalgamated system first required that a new actuator control system was developed, relevant to a foraging system. For this reason the chapter first investigates the potential of a speed regulating hormone which optimises the energy consumption of the

swarm based on demand by fluctuating the speed appropriately. Once this system was found effective the chapter investigates the combination of the sleep hormone system and preference hormone system to ensure that the adaptive traits they provide a swarm could be implemented without negative disruption. The implementation of these systems was found effective, with results showing that each system was typically capable of introducing benefit to the swarms performance in terms of either collection rate or energy efficiency. These positives were found to be introduced with minimal negative trade offs. Showing that, in the context of the tested systems, hormone amalgamation is indeed viable, the next steps being that a multi-hormone system could be found useful implemented in less abstracted tasks as a solution to real world problems.

7.2 Concluding Remarks

Throughout this thesis it has been regarded that swarms of robots excel in dynamic and dangerous environments. The destruction or removal of individual robots is not a problem due to the relative low cost and ability to complete tasks as a group even with many members of the swarm expended. However, for these swarms to perform consistently well they must adapt and remain versatile long after deployment. The initial tuning of parameters for task execution or the initially established behaviour cycle may need to adapt to changes in the environment as tasks develop. Virtual hormone systems as explored in this thesis offer a solution to this, adapting over time from deployment as they interact with the environment and encounter problems. This thesis began with the following general hypothesis:

Hypothesis: *A swarm robotic system can obtain a greater efficiency or effectiveness against a comparison technique through the implementation of a hormone inspired system. Hormone inspired systems will help agents within the swarm adapt over time, without prior knowledge of the environment properties. Adaptation provided by the hormone systems will regulate either robot features or behaviour states.*

In each of the systems produced, including the amalgamation of the systems adaptability has been consistently demonstrated with hormone system competing against systems equipped with random selections, optimised parameters and even competing engineered systems. By matching the performance of or out performing said systems, it has been shown that virtual hormones can be implemented to adapt a system over time without prior knowledge of a task or the use of computational optimisation techniques. In addition to this these hormone systems have not only regulated robot features and behaviour states, but have done so concurrently with the introduction of the amalgamated hormone system.

Virtual hormone systems have been shown as an alternative to optimisation both when optimisation is not possible due to a lack of environmental knowledge or when dynamics

within an environment make producing single parameter values very difficult. Additionally virtual hormone systems can act as an alternative to reinforcement learning (as seen in the comparison to the engineered speed system in Chapter 6) as the decay element within the hormone equations means that values can fluctuate based on time rather than only based on interactions, returning to baseline values for normal operation if no stimuli are encountered across enough time. This also means that behaviours and arbitration are less likely to reach a state of non-recovery as adaptations, as explored in this thesis, modify only a known selection of states or actions rather than the robotic behaviour at a fundamental level.

7.3 Further Work

Having rigorously tested Virtual Hormone systems in a simulated setting the next steps in testing will be to test the systems using a physically implemented system, rather than further decreasing the reality gap. By using genuine robots to perform similar abstracted tasks such as those executed in this thesis (i.e. foraging), the demonstrated capabilities of the physical systems will provide evidence to suggest that swarms, equipped with complex hormone systems, would be capable of functioning well in real world applications that require on-line adaptation. While these new tests would be valuable to all of the experiments conducted throughout this thesis, real world tests would prove particularly valuable to the preference selection system presented in Chapters 5 and 6. These tests would provide an opportunity to study the viability of selecting preference in real environments and if, in reality, performance would be consistently different enough in environments of alternate terrains to appropriately categorise heterogeneous swarm members. Additionally, it may be worth exploring the potential of physically modifying the morphology of swarm agents during a task via virtual hormone. This would allow systems to use hormone values to not only select appropriate robots for a task, but modify swarm members to create a swarm more suited or specialised to a current terrain or assignment. Such a system should also be able to revert to the initial morphological state, keeping the versatility that is required for the effective application of a swarm system.

Following these abstracted, though real, experiments, the element of abstraction should be removed, using the swarm to execute tasks that would provide an actual real world service. These tasks will not need to be significantly more complex than the abstracted tasks previously tested, but will potentially require bespoke hardware to perform activities such as item sorting, area searching or environmental monitoring. These relatively simple tasks scaled up to more consequential activities, could see virtual hormone equipped swarms involved in applications involving disaster relief work. Robot swarms using virtual hormone systems could be tasked with investigating vast areas of volatile and changing environments associated with disaster aftermath, securing survivors or sustaining resources. Equally, hormone systems could be

implemented to enable searching for valuable minerals or suitable areas for habitation on foreign planets with hostile and erratic weather, potentially making a large impact to the future of extra-planetary industry and exploration. The implementation of hormone amalgamations in these roles would grant large portions of autonomy to the agents performing the tasks. These agents would be able to regulate themselves and one another through the use of shared hormone values, providing a frame of reference to the tasks they are executing. This shared information maintains the distribution required for a swarm to function robustly, as robots are not issuing commands to one another in a centralised manner. Instead, swarm agents offer one another contextual information which individual robots can then use to make decisions in a decentralised manner. These hormone amalgamations would require little to no influence from human users once deployed to robustly enact a task, while still leaving room for user input should the desired output from the swarm change.

The future of swarm robotics could see vast numbers of versatile robots supporting societies by taking advantage of parallel operations. Their efficiency in performing tasks will result in not just establishing industries or cities, but also helping to maintain them. The rapid and task efficient methods that swarm technology affords, allows for quick construction or repair projects on infrastructure such as roads or railways, creating minimal disturbances. This sort of technology will not be viable for implementation until the capabilities of swarm systems are acknowledged. This will not happen without proof of robot swarms robustness and adaptability to presented tasks, something which this thesis provides evidence for, furthering the argument for the use of swarm robotics as a solution to real problems.

Bibliography

- (2016), N6700 Modular Power System Family, Keysight Technologies.
- Alden, K., Read, M., Andrews, P. S., Timmis, J. & Coles, M. (2014), 'Applying spartan to understand parameter uncertainty in simulations.', *R Journal*, **6**(2).
- Alden, K., Read, M., Timmis, J., Andrews, P. S., Veiga-Fernandes, H. & Coles, M. (2013), 'Spartan: a comprehensive tool for understanding uncertainty in simulations of biological systems', *PLoS computational biology*, **9**(2), p. e1002916.
- Ang, K. H., Chong, G. & Li, Y. (2005), 'Pid control system analysis, design, and technology', *IEEE transactions on control systems technology*, **13**(4), pp. 559–576.
- Bahceci, E., Soysal, O. & Sahin, E. (2003), 'A review: Pattern formation and adaptation in multi-robot systems', *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-03-43*.
- Beni, G. (2004), 'From swarm intelligence to swarm robotics', in 'International Workshop on Swarm Robotics', pp. 1–9, Springer.
- Benyus, J. M. (1997), 'Biomimicry: Innovation inspired by nature', .
- Bjerknes, J. D. & Winfield, A. F. (2013), 'On fault tolerance and scalability of swarm robotic systems', in 'Distributed autonomous robotic systems', pp. 431–444, Springer.
- Bonani, M., Longchamp, V., Magnenat, S., Rétornaz, P., Burnier, D., Roulet, G., Vaussard, F., Bleuler, H. & Mondada, F. (2010), 'The marxbot, a miniature mobile robot opening new perspectives for the collective-robotic research', in 'Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on', pp. 4187–4193, IEEE.
- Brambilla, M., Ferrante, E., Birattari, M. & Dorigo, M. (2013), 'Swarm robotics: a review from the swarm engineering perspective', *Swarm Intelligence*, **7**(1), pp. 1–41.
- Bredeche, N., Haasdijk, E. & Eiben, A. (2009), 'On-line, on-board evolution of robot controllers', in 'International Conference on Artificial Evolution (Evolution Artificielle)', pp. 110–121, Springer.
- Bredeche, N., Montanier, J.-M., Liu, W. & Winfield, A. F. (2012), 'Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents', *Mathematical and Computer Modelling of Dynamical Systems*, **18**(1), pp. 101–129.
- Buchanan, E., Pomfret, A. & Timmis, J. (2016), 'Dynamic task partitioning for foraging robot swarms', in 'International Conference on Swarm Intelligence', pp. 113–124, Springer.

- Campo, A. & Dorigo, M. (2007), 'Efficient multi-foraging in swarm robotics', in 'European Conference on Artificial Life', pp. 696–705, Springer.
- Cannon, W. B. (1929), 'Bodily changes in pain, hunger, fear and rage.', .
- Charbonneau, D. & Dornhaus, A. (2015), 'When doing nothing is something. how task allocation strategies compromise between flexibility, efficiency, and inactive agents', *Journal of Bioeconomics*, **17**(3), pp. 217–242.
- Christensen, A. L., OGrady, R. & Dorigo, M. (2009), 'From fireflies to fault-tolerant swarms of robots', *IEEE Transactions on Evolutionary Computation*, **13**(4), pp. 754–766, ISSN 1089-778X.
- Christensen, A. L., O'Grady, R. & Dorigo, M. (2008), 'Swarmorph-script: a language for arbitrary morphology generation in self-assembling robots', .
- Clark, J., Goldman, D., Lin, P.-C., Lynch, G., Chen, T., Komsuoglu, H., Full, R. J. & Koditschek, D. E. (2007), 'Design of a bio-inspired dynamical vertical climbing robot.', in 'Robotics: Science and Systems', volume 1.
- Daler, L., Lecoer, J., Hählen, P. B. & Floreano, D. (2013), 'A flying robot with adaptive morphology for multi-modal locomotion', in 'Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on', pp. 1361–1366, IEEE.
- Davis, L. (1991), 'Handbook of genetic algorithms', .
- Denver, R. J. (1999), 'Evolution of the corticotropin-releasing hormone signaling system and its role in stress-induced phenotypic plasticity', *Annals of the New York Academy of Sciences*, **897**(1), pp. 46–53.
- Deveza, R., Thiel, D., Russell, A. & Mackay-Sim, A. (1994), 'Odor sensing for robot guidance', *The International Journal of Robotics Research*, **13**(3), pp. 232–239.
- Duarte, M., Costa, V., Gomes, J., Rodrigues, T., Silva, F., Oliveira, S. M. & Christensen, A. L. (2016), 'Evolution of collective behaviors for a real swarm of aquatic surface robots', *PLoS one*, **11**(3), p. e0151834.
- Ducatelle, F., Di Caro, G. A., Förster, A., Bonani, M., Dorigo, M., Magnenat, S., Mondada, F., O'Grady, R., Pinciroli, C., Réturnaz, P. et al. (2014), 'Cooperative navigation in robotic swarms', *Swarm Intelligence*, **8**(1), pp. 1–33.
- Ducatelle, F., Di Caro, G. A., Pinciroli, C. & Gambardella, L. M. (2011), 'Self-organized cooperation between robotic swarms', *Swarm Intelligence*, **5**(2), p. 73.
- Eberhart, R. & Kennedy, J. (1995), 'Particle swarm optimization, proceeding of iee international conference on neural network', *Perth, Australia*, pp. 1942–1948.

- Eiben, A., Bredeche, N., Hoogendoorn, M., Stradner, J., Timmis, J., Tyrrell, A., Winfield, A. et al. (2013), 'The triangle of life: Evolving robots in real-time and real-space', *Advances in artificial life, ECAL*, **2013**, p. 1056.
- Ferrell, C. (1994), 'Failure recognition and fault tolerance of an autonomous robot', *Adaptive behavior*, **2**(4), pp. 375–398.
- Fras, J. & Althoefer, K. (2018), 'Soft pneumatic prosthetic hand', in 'Annual Conference Towards Autonomous Robotic Systems', pp. 112–120, Springer.
- Fujisawa, R., Dobata, S., Kubota, D., Imamura, H. & Matsuno, F. (2008), 'Dependency by concentration of pheromone trail for multiple robots', in 'International Conference on Ant Colony Optimization and Swarm Intelligence', pp. 283–290, Springer.
- Garnier, S., Tache, F., Combe, M., Grimal, A. & Theraulaz, G. (2007), 'Alice in pheromone land: An experimental setup for the study of ant-like robots', in 'Swarm Intelligence Symposium, 2007. SIS 2007. IEEE', pp. 37–44, Citeseer.
- Gordon, D. M. & Mehdiabadi, N. J. (1999), 'Encounter rate and task allocation in harvester ants', *Behavioral Ecology and Sociobiology*, **45**(5), pp. 370–377.
- Haasdijk, E., Bredeche, N. & Eiben, A. (2014), 'Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics', *PloS one*, **9**(6), p. e98466.
- HARANO, K.-I., Sasaki, M., Nagao, T. & Sasaki, K. (2008), 'Dopamine influences locomotor activity in honeybee queens: implications for a behavioural change after mating', *Physiological entomology*, **33**(4), pp. 395–399.
- Hauert, S., Winkler, L., Zufferey, J.-C. & Floreano, D. (2008), 'Ant-based swarming with positionless micro air vehicles for communication relay', *Swarm Intelligence*, **2**(2-4), pp. 167–188.
- Hecker, J. P. (2015), 'Evolving efficient foraging behavior in biologically-inspired robot swarms', .
- Hilder, J., Horsfield, A., Millard, A. G. & Timmis, J. (2016), 'The psi swarm: A low-cost robotics platform and its use in an education setting', in 'Conference Towards Autonomous Robotic Systems', pp. 158–164, Springer.
- Ioannidis, K., Sirakoulis, G. C. & Andreadis, I. (2011), 'Cellular ants: A method to create collision free trajectories for a cooperative robot team', *Robotics and Autonomous Systems*, **59**(2), pp. 113–127.
- Ismail, A. R., Bjercknes, J. D., Timmis, J. & Winfield, A. (2015), 'An artificial immune system for self-healing in swarm robotic systems', in Lones, M., Tyrrell, A., Smith, S. & Fogel, G., eds.,

- 'Information Processing in Cells and Tissues', pp. 61–74, Springer International Publishing, Cham, ISBN 978-3-319-23108-2.
- Jakobi, N., Husbands, P. & Harvey, I. (1995), 'Noise and the reality gap: The use of simulation in evolutionary robotics', in 'European Conference on Artificial Life', pp. 704–720, Springer.
- Jin, Y., Guo, H. & Meng, Y. (2009), 'Robustness analysis and failure recovery of a bio-inspired self-organizing multi-robot system', in '2009 Third IEEE International Conference on Self-Adaptive and Self-Organizing Systems', pp. 154–164, IEEE.
- Jones, E. G., Browning, B., Dias, M. B., Argall, B., Veloso, M. & Stentz, A. (2006), 'Dynamically formed heterogeneous robot teams performing tightly-coordinated tasks', in 'Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on', pp. 570–575, IEEE.
- Karaboga, D. & Akay, B. (2009), 'A comparative study of artificial bee colony algorithm', *Applied mathematics and computation*, **214**(1), pp. 108–132.
- Kernbach, S. & Kernbach, O. (2011), 'Collective energy homeostasis in a large-scale microrobotic swarm', *Robotics and Autonomous Systems*, **59**(12), pp. 1090–1101.
- Kernbach, S., Meister, E., Schlachter, F., Jebens, K., Szymanski, M., Liedke, J., Laneri, D., Winkler, L., Schmickl, T., Thenius, R. et al. (2008), 'Symbiotic robot organisms: Replicator and symbion projects', in 'Proceedings of the 8th workshop on performance metrics for intelligent systems', pp. 62–69, ACM.
- Kim, Y.-S., Jung, G.-P., Kim, H., Cho, K.-J. & Chu, C.-N. (2013), 'Wheel transformer: A miniaturized terrain adaptive robot with passively transformed wheels', in 'Robotics and Automation (ICRA), 2013 IEEE International Conference on', pp. 5625–5630, IEEE.
- Knospe, C. (2006), 'Pid control', *IEEE Control Systems Magazine*, **26**(1), pp. 30–31.
- Krieger, M. J., Billeter, J.-B. & Keller, L. (2000), 'Ant-like task allocation and recruitment in cooperative robots', *Nature*, **406**(6799), pp. 992–995.
- Kuindersma, S., Deits, R., Fallon, M., Valenzuela, A., Dai, H., Permenter, F., Koolen, T., Marion, P. & Tedrake, R. (2016), 'Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot', *Autonomous Robots*, **40**(3), pp. 429–455.
- Kuyucu, T., Tanev, I. & Shimohara, K. (2013), 'Hormone-inspired behaviour switching for the control of collective robotic organisms', *Robotics*, **2**(3), pp. 165–184.
- Lau, H. (2012), *Error Detection in Swarm Robotics: A Focus on Adaptivity to Dynamic Environments*, Ph.D. thesis, University of York.

- Lee, J.-H. & Ahn, C. W. (2011), 'Improving energy efficiency in cooperative foraging swarm robots using behavioral model', in '2011 Sixth International Conference on Bio-Inspired Computing: Theories and Applications', pp. 39–44, IEEE.
- Lehman, J., Clune, J., Misevic, D., Adami, C., Beaulieu, J., Bentley, P. J., Bernard, S., Belson, G., Bryson, D. M., Cheney, N. et al. (2018), 'The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities', *arXiv preprint arXiv:1803.03453*.
- Lerman, K. & Galstyan, A. (2002), 'Mathematical model of foraging in a group of robots: Effect of interference', *Autonomous Robots*, **13**(2), pp. 127–141.
- Levi, P. & Kernbach, S. (2010), *Symbiotic multi-robot organisms: reliability, adaptability, evolution*, volume 7, Springer Science & Business Media.
- Liu, W., FT, A. W., Sa, J., Chen, J. & Dou, L. (2007), 'Towards energy optimization: Emergent task allocation in a swarm of foraging robots', *Adaptive behavior*, **15**(3), pp. 289–305.
- Liu, W., Winfield, A., Sa, J., Chen, J. & Dou, L. (2006), 'Strategies for energy optimisation in a swarm of foraging robots', in 'International Workshop on Swarm Robotics', pp. 14–26, Springer.
- Lu, Q., Hecker, J. P. & Moses, M. E. (2016), 'The mpfa: A multiple-place foraging algorithm for biologically-inspired robot swarms', in 'Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on', pp. 3815–3821, IEEE.
- McLurkin, J. & Yamins, D. (2005), 'Dynamic task assignment in robot swarms.', in 'Robotics: Science and Systems', volume 8, Citeseer.
- Miner, D. (2007), 'Swarm robotics algorithms: A survey', *Report, MAPLE lab, University of Maryland*.
- Mintchev, S., Donati, E., Marrazza, S. & Stefanini, C. (2014), 'Mechatronic design of a miniature underwater robot for swarm operations', in '2014 IEEE International Conference on Robotics and Automation (ICRA)', pp. 2938–2943, ISSN 1050-4729.
- Nam, C. & Shell, D. A. (2018), 'Bundling policies for sequential stochastic tasks in multi-robot systems', in 'Distributed Autonomous Robotic Systems', pp. 237–252, Springer.
- Navarro Oiza, I. & Matía Espada, F. (2013), 'An introduction to swarm robotics', *ISRN Robotics*, **2013**, pp. 1–10.
- Neal, M. & Timmis, J. (2003), 'Timidity: A useful mechanism for robot control?', *Informatica*, **27**(2), pp. 197–204.

- Neal, M. & Timmis, J. (2005), 'Once more unto the breach: Towards artificial homeostasis', in 'Recent Developments in Biologically Inspired Computing', pp. 340–366, Igi Global.
- O'Keefe, J., Tarapore, D., Millard, A. G. & Timmis, J. (2017), 'Towards fault diagnosis in robot swarms: An online behaviour characterisation approach', in Gao, Y., Fallah, S., Jin, Y. & Lekakou, C., eds., 'Towards Autonomous Robotic Systems', pp. 393–407, Springer International Publishing, Cham, ISBN 978-3-319-64107-2.
- Palmieri, N., Yang, X.-S., De Rango, F. & Marano, S. (2017), 'Comparison of bio-inspired algorithms applied to the coordination of mobile robots considering the energy consumption', *Neural Computing and Applications*, pp. 1–24.
- Pang, B., Zhang, C., Song, Y. & Wang, H. (2017), 'Self-organized task allocation in swarm robotics foraging based on dynamical response threshold approach', in '2017 18th International Conference on Advanced Robotics (ICAR)', pp. 256–261, IEEE.
- Pinciroli, C., Trianni, V., O'Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., Birattari, M., Gambardella, L. M. & Dorigo, M. (2012a), 'ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems', *Swarm Intelligence*, **6**(4), pp. 271–295.
- Pinciroli, C., Trianni, V., O'Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., Birattari, M., Gambardella, L. M. & Dorigo, M. (2012b), 'ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems', *Swarm Intelligence*, **6**(4), pp. 271–295.
- Raibert, M., Blankespoor, K., Nelson, G. & Playter, R. (2008), 'Bigdog, the rough-terrain quadruped robot', *IFAC Proceedings Volumes*, **41**(2), pp. 10822–10825.
- Ribeiro, P., Miyazawa, A., Li, W., Cavalcanti, A. & Timmis, J. (2017), 'Modelling and verification of timed robotic controllers', in Polikarpova, N. & Schneider, S., eds., 'Integrated Formal Methods', pp. 18–33, Springer International Publishing, Cham, ISBN 978-3-319-66845-1.
- Riddiford, L. M. & Truman, J. W. (1993), 'Hormone receptors and the regulation of insect metamorphosis', *American Zoologist*, **33**(3), pp. 340–347.
- Robinson, S. (2004), *Simulation: the practice of model development and use*, Wiley Chichester.
- Şahin, E. (2004), 'Swarm robotics: From sources of inspiration to domains of application', in 'International workshop on swarm robotics', pp. 10–20, Springer.
- Şahin, E. (2005), 'Swarm robotics: From sources of inspiration to domains of application', in Şahin, E. & Spears, W. M., eds., 'Swarm Robotics', pp. 10–20, Springer Berlin Heidelberg, Berlin, Heidelberg, ISBN 978-3-540-30552-1.

- Sauzé, C., Neal, M. et al. (2010), 'A neuro-endocrine inspired approach to long term energy autonomy in sailing robots', *Proceedings of TAROS (Towards Autonomous Robotic Systems)*, pp. 255–262.
- Schmickl, T. & Crailsheim, K. (2006), 'Trophallaxis among swarm-robots: A biologically inspired strategy for swarm robotics', in 'Biomedical Robotics and Biomechanics, 2006. BioRob 2006. The First IEEE/RAS-EMBS International Conference on', pp. 377–382, IEEE.
- Schmickl, T., Thenius, R., Moslinger, C., Timmis, J., Tyrrell, A., Read, M., Hilder, J., Halloy, J., Campo, A., Stefanini, C., Manfredi, L., Orofino, S., Kernbach, S., Dipper, T. & Sutanty, D. (2011), 'Cocoro – the self-aware underwater swarm', in '2011 Fifth IEEE Conference on Self-Adaptive and Self-Organizing Systems Workshops', pp. 120–126.
- Schroeder, A., Ramakrishnan, S., Kumar, M. & Trease, B. (2017), 'Efficient spatial coverage by a robot swarm based on an ant foraging model and the lévy distribution', *Swarm Intelligence*, **11**(1), pp. 39–69.
- Shen, W., Will, P., Galstyan, A. & Chuong, C. (2004), 'Hormone-inspired self-organization and distributed control of robotic swarms', *Autonomous Robots*, **17**(1), pp. 93–105.
- Shen, W.-M., Lu, Y. & Will, P. (2000), 'Hormone-based control for self-reconfigurable robots', in 'Proceedings of the fourth international conference on Autonomous agents', pp. 1–8, ACM.
- Sims, K. (1994), 'Evolving 3d morphology and behavior by competition', *Artificial life*, **1**(4), pp. 353–372.
- Stradner, J., Hamann, H., Schmickl, T. & Crailsheim, K. (2009), 'Analysis and implementation of an artificial homeostatic hormone system: A first case study in robotic hardware', in 'Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on', pp. 595–600, IEEE.
- Tarapore, D., Lima, P. U., Carneiro, J. & Christensen, A. L. (2015), 'To err is robotic, to tolerate immunological: fault detection in multirobot systems', *Bioinspiration & biomimetics*, **10**(1), p. 016014.
- Tarapore, D., Timmis, J. & Christensen, A. L. (2019), 'Fault detection in a swarm of physical robots based on behavioral outlier detection', *IEEE Transactions on Robotics*, **35**(6), pp. 1516–1522.
- Timmis, J., Murray, L. & Neal, M. (2010), 'A neural-endocrine architecture for foraging in swarm robotic systems', in 'Nature Inspired Cooperative Strategies for Optimization (NICSO 2010)', pp. 319–330, Springer.

- Trivedi, D., Rahn, C. D., Kier, W. M. & Walker, I. D. (2008), 'Soft robotics: Biological inspiration, state of the art, and future research', *Applied bionics and biomechanics*, **5**(3), pp. 99–117.
- Vargas, P., Moioli, R., de Castro, L. N., Timmis, J., Neal, M. & Von Zuben, F. J. (2005), 'Artificial homeostatic system: a novel approach', in 'European Conference on Artificial Life', pp. 754–764, Springer.
- Watanabe, S. & Kuczaj, S. (2012), *Emotions of animals and humans: Comparative perspectives*, Springer Science & Business Media.
- Wilson, J., Timmis, J. & Tyrrell, A. (2018), 'A hormone arbitration system for energy efficient foraging in robot swarms', in 'Annual Conference Towards Autonomous Robotic Systems', pp. 305–316, Springer.
- Winfield, A. F. & Nembrini, J. (2006), 'Safety in numbers: Fault tolerance in robot swarms', *International Journal on Modelling Identification and Control*, **1**(ARTICLE), pp. 30–37.
- Yang, X.-S. (2009), 'Firefly algorithms for multimodal optimization', in 'International symposium on stochastic algorithms', pp. 169–178, Springer.
- Zhao, K. (2013), 'Multiple-agent task allocation algorithm utilizing ant colony optimization.', *JNW*, **8**(11), pp. 2599–2606.
- Zhao, N., Luo, Y., Deng, H. & Shen, Y. (2017), 'The deformable quad-rotor: Design, kinematics and dynamics characterization, and flight performance validation', in 'Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on', pp. 2391–2396, IEEE.