A Neuro-Endocrine Inspired Approach to Power Management in Sailing Robots

Colin Sauzé

December 2010

This thesis is submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy of The University of Wales.

DECLARATION

This work has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Signed ..... (candidate)
Date .....

STATEMENT 1

This thesis is the result of my own investigations, except where otherwise stated.
Other sources are acknowledged by footnotes giving explicit references. A bibliography is appended.

Signed ..... (candidate)
Date .....

STATEMENT 2

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

Signed ..... (candidate)
Date .....

Abstract

Sailing robots represent a new and interesting approach to oceanographic monitoring, with advantages over traditional data buoys, survey ships and satellites. They can transport themselves to and from target areas and can potentially remain on site for months, with little or no risk to human operators. However, they are likely to be restricted to low bandwidth, high latency and high cost communication systems and cannot rely upon human operators to keep them out of danger. Ideally they should autonomously avoid danger and manage their power consumption.

Sailing robots represent one of the few currently available robot platforms, which can theoretically sustain themselves entirely from energy obtained from their environment. Currently the most viable strategy to achieve this, is to obtain electrical power from photo-voltaic solar panels and locomotive power from the wind using a sail. Compared with other outdoor robotic systems they are also relatively cheap, can be used relatively safely and without major legal restrictions. Given these properties they present an ideal opportunity to study long term autonomy in robotics.

A major challenge for sailing robots is that electrical power budgets are extremely limited. It would therefore seem logical to attempt to vary the behaviour of the robot in response to changes over time in the amount of electrical power available to the robot. A potential strategy to achieve this is to borrow ideas from biology. Biological systems have demonstrated themselves to be highly capable in coping with varying energy levels. Key to this is the endocrine system, which through the secretion of chemical messengers known as hormones is able to maintain a near consistent internal state with regards to a series of parameters, including blood sugar levels.

This work developed an artificial endocrine system and used it to modulate the behaviour of an artificial neural network. The neural network was responsible for steering and sail setting. The artificial endocrine system produced a hormone in proportion to the battery level, this promoted and suppressed the neural network. A series of experiments were undertaken using two small sailing robots and a series of simulations. By modulating the neural network it was possible to reduce the power consumption in simulation by almost 13 times and in the real robot by two times, without any noticeable effects upon sailing performance. Further increases to the modulation level would bring power consumption to near zero, but came at the expense of any useful sailing ability. In additional simulation experiments, a hormone was linked to sun light levels. This helped to reinforce power saving behaviours and was able to facilitate indefinite operations in a simulation which would have otherwise been an impossible power budget. Through this, it has been shown that an artificial neuro-endocrine controller is a potentially useful in modulating the power consumption of sailing robots. In comparison to traditional approaches of creating distinctive power management modes, it provides a near continuous method for modulating power consumption.

Acknowledgements

There are many people who without their contribution, this thesis would not have been completed or even started in the first place. There are too many to name in this acknowledgement, but here are some of the key people who I wish to thank.

Firstly, I would like to thank my grandfather who sadly died a few weeks before the initial submission of the thesis. It was due to him taking me sailing as a child that I developed an interest in sailing. Without this interest, I would not have undertaken this work.

I would also like to give thanks to the rest of my family and particularly my long suffering fiancee Claire who has become a PhD widow during the past few months of writing up and who has also been the first person to proof read each chapter for me.

To everyone in the Department of Computer Science at Aberystwyth University. In particular to my supervisors, Mark Neal and Fred Labrosse. To Ian Izzet, Barry Thomas, Tom Blanchard and Rokas Žmuidzinavičius for their help developing in developing hardware and software for the sailing robots. To Alan Woodland and Richard Shipman for their help in deploying and chasing after the robots. To everyone else who has offered the use of their boat or driven chase boats.

To everyone in C57 during this time who have provided support, advice and interesting discussions during the last few years.

To everyone involved with the Microtransat, World Robotic Sailing Championships and Sailbot competitions for providing fun and friendly competitions, interesting discussions, inspiration for robot designs and support and encouragement of this work.

To the residents near Llyn-Yr-Orefa for not thinking it was too strange to see somebody stood by the lake all day in the middle of winter with a laptop, while looking worryingly at two small boats in the middle of the lake.

Finally I would like to thank Aberystwyth University for funding this work through all the part time jobs they gave me , the staff fee waiver program and the Aberystwyth Postgraduate Research Studentship.

Contents

List of Algorithms

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background and Motivations

1.1.1 Sailing Robots

Sailing robots offer the opportunity to provide a flexible and low cost ocean monitoring platform which can augment or even replace current systems. Effective power management will allow sailing robots to perform additional tasks, without increasing the total power budget. This allows the same set of tasks to be carried out with a smaller power budget. Either strategy represents a potential method to reduce the overall cost of building or operating the robot and improve its usefulness. This thesis presents a biologically inspired power management system based on an artificial analogue of the neural and endocrine systems.

At present ocean monitoring is typically performed by data buoys moored to the seabed, drifting data buoys, sensors attached to commercial shipping, dedicated survey vessels and satellites. These ocean monitoring systems suffer from a number of shortcomings. Moored data buoys require a cable (sometimes thousands of metres long) to link a weight on the seabed with the buoy. Once deployed they can autonomously relay data via a satellite link, but must be visited regularly for servicing. Drifting buoys typically transmit data continuously, but usually end their lives when they are washed ashore[1]. Sensors attached to commercial shipping only cover shipping channels and often avoid more remote, dangerous and potentially scientifically interesting areas. While dedicated survey ships can enter these areas, they do so by putting their crews in potential danger and at considerable financial cost. Satellite systems offer the benefit of covering wide areas, but costs can exceed £100 million[2], data resolution is limited and subsurface measurements are not possible. Sailing robots could

[1] `http://www.adp.noaa.gov/faq.html` accessed 04/02/2011

[2] `http://www.eumetsat.int/Home/Main/Satellites/Jason-2/MissionOverview/index.htm?l=en` accessed 01/02/2011

offer a platform which can remain at sea for several months at a time either holding station at a fixed position, sailing a pre-determined route or drifting with currents for some time. At the end of their mission they can autonomously return to a port or rendezvous point, eliminating much of the deployment cost associated with data buoys. It is expected that sailing robots can be constructed at a relatively low cost of around £20,000 (2005 prices) each [90]. At these rates it becomes practical to deploy a small fleet of robots into a potentially dangerous area instead of sending a survey ship, this will eliminate risk to the crew of the survey ship and reduce overall costs.

It is envisaged that a sailing robot would only send and receive data occasionally through a satellite or short wave radio link. Therefore, it cannot rely upon receiving instructions from human operators and must be able to take decisions that will enable its continued survival. Although locomotive power will be obtained from the wind, electrical power must be provided to move actuators and run on-board sensors, computers and communications equipment. Electrical power will be obtained either from on-board power sources such as a large battery, fuel cell or generator or by obtaining energy from the environment through solar panels, wind power or wave power. Given current technology, the most viable solution appears to be the use of photo-voltaic solar panels and batteries.

When compared with robots using motorised propulsion (e.g. wheeled vehicles, aeroplanes, helicopters or propeller driven boats) sailing robots face the favourable advantage of having their locomotive power provided by the wind. This makes it possible to obtain all the required power from the robot's environment. However, it still leaves a large challenge. Smaller robots present a particular challenge as the relationship between total deck area, which can be covered in solar panels and power consumption is not linear, leaving small boats with extremely limited power budgets.

Sailing robots also provide a useful test bed for developing power management systems applicable to other types of robot as well. They are particularly suited because they can be operated in reasonable safety, with few legal problems and for long periods of time. They also are one of the few platforms which could operate for prolonged time periods with current energy generation and storage technology. For example, a current generation wheeled or airborne robot would be considered exceptional if it could sustain 24 hours of continuous movement. Current legal restrictions do not allow wheeled or airborne robots to share road/airspace with other users given the high probability of accidents. However, no such legal restriction appears to exist at sea and the speeds and relative vehicle sizes involved dramatically reduce the chance of any damage to other vessels. Separated water space is easily achievable by operating in lakes with no other vessels. Given these advantages sailing robots can also be left unattended for extended periods of time without too many

difficulties.

Since 2004 the Department of Computer Science at Aberystwyth University has been involved in developing autonomous sailing robots. These are intended with the triple purpose of performing oceanographic observations autonomously, researching control systems techniques aimed at improving longevity of robots and racing against boats built by other organisations. At least one racing competition has taken place every year since 2006 and in 2010 Aberystwyth University made the first (ultimately unsuccessful) attempt to race a sailing robot across the Atlantic ocean.

1.1.2 Biologically Inspired Control

Autonomous mobile robots operating in harsh environments, such as the open sea or outer-space, must be able to operate in a sustainable manner with regards to power management. They must either carry sufficient energy on-board in the form of batteries or fuel or they must obtain it from their environment. They must be able to make appropriate decisions regarding when to use additional energy to avoid danger and perform their mission or when to attempt to gather more energy or conserve what remains. It is also desirable for them to be able to take advantage of unforeseen opportunities presented by excess energy availability to improve their performance or undertake extra data gathering.

Traditionally its has been common for robot engineers to leave many of these decisions with human operators, to ignore them all together and build in large margins of error. Such approaches suffer from the obvious drawbacks associated with increasing weight, power consumption, financial cost and a reduction of autonomy, ultimately these reduce the effectiveness and flexibility of such robots. One approach to resolving these issues is to draw inspiration from nature. Biological systems have proven themselves to be highly resilient and adaptable to harsh conditions. One particular property that could be of benefit to robotics is homeostasis, the ability to maintain a constant internal state in the face of extreme variations in external environment. Through homeostasis [31] many key internal variables such as blood sugar level, body temperature, salt levels, calcium levels and blood pressure are regulated.

In mammals three systems play a key role in maintaining homeostasis, these are the neural, endocrine and immune systems. These interact heavily and are capable of reinforcing or regulating each other's behaviour. Each system covers different functions and is responsible for different timescales. Consider, for example the reactions which take place when you accidentally touch a sharp object. The near immediate (sub second response time) reaction is provided by the neural system and might include moving suddenly away from the sharp object and feelings of pain. This will then be followed by a burst of adrenaline from the endocrine system. Adrenaline, a chemical messenger known as a hormone, spreads through

the blood and triggers behavioural changes in certain target cells. Its effects will increase the heart rate and blood flow, speeding reaction times, increasing the uptake of glucose into the blood to provide a burst of energy and temporarily reduce any pain. This will last anything from a few seconds to a few minutes depending on the severity of the problem. Finally, the immune system will be responsible for fighting any infection resulting from bacteria that were living on the sharp object (assuming it pierced the skin), this process can take several days. If presented with a new and unseen infection the immune system may have to learn how to fight it and will save this knowledge for future encounters with this particular bacterium. It is believed that these systems underwent a common evolution and that they can essentially be viewed as a large super-system [43, 101, 17].

Computer Scientists have created artificial models of all three systems [72, 110, 25, 7, 42, 35, 92] and have applied them to a number of problems including robotics [85, 122, 135, 136, 134, 103]. However, to date this work has focused on simulation and laboratory robots and nobody has demonstrated them as a practical method of keeping a real robot working in harsh conditions operational. This thesis investigates if an artificial neuro-endocrine controller can be useful in managing power consumption of a sailing robot over relatively long time periods. Despite its obvious importance in biology, the immune system's function is not being considered for a number of reasons. Firstly the complexity of many current artificial immune systems algorithms do not lend themselves to implementations on low power embedded computers. Many species also function without immune systems or with only very limited immune systems and given that the immune system covers the longest time scales of any part of the neural-endocrine-immune system it was felt more important to focus on the shorter time scales of the neural and endocrine systems. Even achieving a functional neuro-endocrine system operating in a real robot over a multi-hour to multi-day timescale was considered to be sufficiently challenging and useful.

1.2 Problem Summary and Background

Most existing power management strategies are limited to switching between a few distinctive predefined modes and cannot gradually adjust over the long term. As sailing robots must operate within highly restricted power budgets and in particularly dynamic and unpredictable environments, their power consumption requirements are likely to vary considerably but must be kept within acceptable limits. Previous work [8, 7, 85, 84, 92] on artificial homeostasis and artificial endocrine controllers offer a potential solution to this problem. However, to date this work has been limited to simulated or laboratory robots.

Moioli, Vargas, Von Zuben and Husbands (2008) [85, 84] implemented a power manage-

ment system in which their robot would seek energy when the battery level was low but they did not modify its behaviour in any other way to reduce power consumption. Such a strategy would be difficult to apply to a sailing robot as a presumably solar powered sailing robot will have little need (or chance of success if it did) to go off in search of energy sources. Instead it is able to recharge in parallel to normal operation although some benefit could be obtained by changing course to either reduce shadows on the solar panels or to cause the boat to tilt towards the sun. A more appropriate strategy might be to modify the behaviour of the robot in order to reduce power consumption when battery levels are low.

Arkin (1992,1993) [8, 7] demonstrated a system which modified a route planning algorithm to take riskier but shorter paths when fuel levels were low. Such an approach could be applied to sailing robots operating in coastal waters where riskier paths would be those which sail closer to the coast in order to take a more direct route. However, this approach will not make any difference on the open sea where there are few obstacles to avoid and paths to re-plan.

Neal and Timmis (2003) [92] demonstrated that an artificial endocrine controller, which secretes artificial hormones that could modulate the behaviour of an artificial neural network. Their approach modulated a neural network in response to distances from sonar sensors, giving a panic like behaviour when the robot was enclosed with obstacles. Although this technique was not applied to power management it could be reapplied to modulate the network in response to battery levels not sonar distances.

What is needed for sailing robots (and arguably any other robot which is capable of recharging in parallel with normal operation) is a method, or series of methods, which can modify its behaviour to dramatically reduce power consumption while still managing to perform at least some basic operations to continue the mission. This will most likely come at the expense of reduced performance, or temporary suspension of the mission, or parts of the mission (for example turning off scientific sensors or reducing or eliminating communications). Conversely when ample power is available this may also present the opportunity to improve performance beyond normal levels or to perform extra tasks. When applied specifically to sailing robots, sailing performance could be adjusted through varying the duty cycle and/or magnitude of rudder and sail actuator movements, the frequency at which sensors are switched on and sampled or the frequency and extent of communications with the robot's operators. With regards to actuator movements, this might mean that during periods of low power availability, fewer actuator movements may mean that the desired course is not held particularly well or that sail settings are sub-optimal but that considerable power savings can be made.

1.2.1 Applying Neuro-Endocrine Architectures to Power Management in Sailing Robots.

A key observation is that through a homeostatic process controlled by the endocrine system, biological systems are able to maintain a stable internal state with respect to available energy levels. It was envisaged that a potential control system for a sailing robot could consist of an artificial neural network for controlling both the rudder and sail(s) with respect to the compass heading and wind direction. These could be modulated through an artificial endocrine controller which is able to suppress or promote the movement of actuators. The endocrine controller could respond to changes in available energy from a battery and/or solar panels, this would create a situation analogous to insulin regulating the uptake of glucose. The artificial insulin could be used to increase power consumption as energy levels increase. This should create a homeostatic mechanism with respect to maintaining the battery level within a zone of viability, despite fluctuations in power consumption and generation. The achievement of an artificial homeostatic state will help to ensure that such a robot is energy autonomous and capable of long term operation without being resupplied with energy from a human operator (e.g. somebody plugging it in, changing batteries or refuelling it). Although we should be somewhat hesitant to use the term "capable of indefinite operation", as operations in extreme latitudes during winter time will not be practical due to a lack of sun light and the eventual degradation of batteries and solar panels will not allow for truly indefinite operations. Such a system can also create the potential to exploit the opportunities presented by the availability of excess power. In the proposed architecture any such excess will increase the magnitude of actuator movement, bringing it closer to (if it was previously suppressed) or even over (if it was not previously suppressed) its "normal" magnitude and thus creating an improvement in performance.

There is a possibility that this will actually be counter productive. By decreasing actuator movement we may decrease the accuracy with which a particular course can be held and this may result in oscillation of the robot attempting to get back on course. This may in turn destroy any power saving advantage achieved through the initial decrease in actuator movements. There are a large number of dynamic factors in real world environments which could cause such problems, simply identifying all these factors let alone quantifying them would be a massive undertaking. Therefore it cannot be realistically established through simulation or shown analytically whether such problems will prevent a neuro-endocrine system from effectively controlling power consumption. Instead actual experiments with a real robot need to be performed. Such experiments have not been performed in any of the existing literature. Existing approaches have been limited to path re-planning or using action selection to enable

food seeking behaviours and even these are restricted to laboratory or simulation environments. These robot experiments were carried out using a pair of small sailing robots known as MOOPs (Miniature Ocean Observation Platform), these are only 72cm long and weigh only 4kg. Their small size makes them easy to deploy and recover, reducing the overhead incurred in operating a sailing robot to perform these experiments.

1.3 Hypothesis and Research Question

The key hypothesis for this thesis is:

> "That an artificial neural and endocrine system can manage power consumption in a sailing robot operating in a real environment through the modulation of actuator movement. This will result in an improvement in power management when compared to an unmodulated neural network."

Which leads to the research question:

> "To what extent will a control system based upon an abstracted artificial neuro-endocrine control system, be able to manage power consumption in a sailing robot through the modulation of actuator movement and what effects will this have upon sailing performance, in comparison to an unmodulated neural network controller?"

Answering this research question will provide insights into how applicable artificial neuro-endocrine controllers are in the real world context of a sailing robot, when compared with a traditional unmodulated neural network. If successful, this will ultimately provide a mechanism to control power consumption in sailing robots and thus improve their longevity and/or reduce power budgets and financial cost. Assuming that a neuro-endocrine controller can reduce power consumption of a sailing robot by reducing actuator movement there are further questions as to the extent to which this can be carried out before the robot is no longer able to sail correctly. A trade off between power consumption and sailing performance is likely to exist and there is likely to be a point at which sailing the desired course (or even an alternative course) can no longer be achieved. This point may not be fixed either and is likely to be dependent upon the weather and sea conditions in which the robot is operating. Conversely when ample power is available the neuro-endocrine controller may be able to improve sailing performance by increasing actuator movements, thus exploiting opportunities which arise almost as an emergent property rather than through explicitly defined behaviours. This creates a mechanism by which the poor performance, during a period when little energy is

available, can be compensated by improved performance during a subsequent period when energy is plentiful.

1.4 Aims and Objectives

In order to answer this research question a series of key aims and objectives can be derived, these are:

1. To develop a set of artificial neural networks which are capable of controlling a sailing robot.

2. To develop an artificial endocrine controller which can modulate those neural networks in order to manage power consumption.

3. To develop a metric or series of metrics to quantify the success of an algorithm with respect to power consumption and performance in a sailing robot.

4. To develop a suitable simulator to simulate both the internal and external environment of the robot, this is in order to validate the feasibility of the algorithms before they are deployed on a real robot.

5. To develop a suitable robot that is physically capable of operating for reasonably long periods of time in a real world environment across a wide variety of weather conditions.

6. To construct a series of experiments to test the feasibility of the neuro-endocrine algorithms in simulation.

7. To repeat the same experiments upon the real robot while it is actually operating in a real outdoor environment.

Experimental work for this thesis was performed in both simulation and on a real robot. Simulation experiments typically lasted in excess of 24 hours while robot experiments (due only to physical limitations of the robot) lasted between two and four hours. An artificial neural network was used to control the positions of rudder and sail actuators. An artificial endocrine controller produced artificial hormones, which modulated the neural network. This modulation changes the behaviour of the neural network, by modifying the magnitude of actuator movement.

1.5 Results Overview

The first experiment fixed a hormone at pre-determined levels for the duration of each run. It was performed in both simulation and on the real robot. The robot was set to sail a pre-defined course and the amount of energy used was recorded. This process was then repeated with varying hormone concentrations to discover the extent to which power consumption could be controlled and at what point the boat would fail to complete the course. By reducing the magnitude of actuator movements, statistically significant reductions in power consumption were achieved. Increasing the level of modulation initially appeared to cause little reduction in sailing performance, but when further increases in modulation occurred sailing performance dropped significantly.

A second experiment allowed the hormone level to vary in response to the remaining battery capacity. This caused gradually increasing levels of modulation during the course of the experiment. It was possible to extend the battery lifespan by nearly 13 times in simulation (compared to an unmodulated neural network) and nearly double in the real robot.

Additional simulation experiments took place which added a photo-voltaic solar panel to the simulated robot. The first of these used only the battery hormone of the previous experiment and set the robot the task of completing a 100km course, without completely discharging the battery. The simulation was run considering the sun light levels in June, September and December at a latitude of 52 degrees north. Only in the June simulation was the robot able to complete the course. An additional hormone was then added, which represented the amount of sun light available. This was able to reinforce the modulatory behaviour of the battery hormone and lead to indefinite operation in which a healthy battery level was maintained. Although, this was achieved by suppressing the control system to the point that it was no longer to able to sail in the correct direction for much of the day.

1.6 Key Contributions

A number of key novel contributions have been made by this thesis:

- The development of an artificial neuro-endocrine controller to steer a sailing robot and manage its power.
- A series of metrics and a methodology for testing the power consumption of a sailing robot.
- Refinements to existing artificial neuro-endocrine control algorithms.

- Showed it's possible to control power consumption in a sailing robot, through the modulation of the magnitude of actuator movement.
- Performed the most extensive investigation into autonomous power consumption in a sailing robot to date.

This work has the potential to enable sailing robots (or other autonomous systems) to adapt their behaviour to cope with unpredictable supplies and demands for energy. It allows for small and gradual changes to be made to the robot's behaviour, although large and sudden changes are also possible. This could assist sailing robots in operating autonomously for long periods of time, to cope with the degradation of batteries and solar panels and to enable them to operate with narrow margins of error within their power budgets. Ultimately this could increase the length of missions achievable by a sailing robot and reduce the need to over-engineer their power systems, reducing financial costs.

1.7 Chapter Summaries

- Chapter 2 reviews existing work in robotic control systems, biologically inspired control architectures and autonomous power management techniques. Key techniques are evaluated for their viability in a sailing robot control system.
- Chapter 3 reviews the development of sailing robots and focuses upon power management strategies in sailing robots.
- Chapter 4 outlines the development of sailing robots at Aberystwyth University, the lessons learnt during this process and how they contributed to the identification of the problems this thesis aims to solve.
- Chapter 5 describes the mathematics behind artificial neuro-endocrine controllers and makes refinements to existing architectures. It also describes the potential roles of hormones in an artificial endocrine controller for a sailing robot.
- Chapter 6 presents results of a series of experiments simulating a sailing robot. It also discusses a set of metrics for evaluating the performance of power management strategies upon sailing performance. In these an artificial neuro-endocrine controller is applied to power management in a simulated robot.
- Chapter 7 presents the results of experiments from a real sailing robot. These are repeats of some of the experiments performed in simulation.

- Chapter 8 presents a discussion of the results and compares the relative performance of the simulator and robot results.
- Chapter 9 presents the conclusions, discusses the implications of this work and makes suggestions for future work.
- Appendix A provides a glossary of sailing and biology terminology.
- Appendix B shows the distributed motor controller circuit described in section 4.2.1.2.
- Appendix C contains the details of the component layout and dimensions of a MOOP sailing robot.
- Appendix D presents raw data from the experiments in chapter 6 and 7.
- Appendix E contains the data used to estimate the power consumption of the actuators on a MOOP sailing robot.
- Appendix F contains a list of papers published during this work.

Chapter 2

Robotic Control Strategies

2.1 Introduction

This chapter provides an overview of existing methods for controlling robots. It covers the evolution of early cybernetic techniques for robot control systems and the development of deliberative, reactive and hybrid robot control architectures. A review of biologically inspired techniques and an introduction to the biology behind them is also covered. This includes a review of artificial homeostasis (maintaining a stable internal state), artificial neural networks, artificial endocrine controllers and artificial immune systems, with an emphasis on their applicability for performing power management. Finally a review of approaches to power management in robotic systems and sensor networks is covered.

2.2 Cybernetic Approaches

The term Cybernetics was coined by Wiener (1948) [145], it derives from the Greek word kybernētēs which means steersman, governor or pilot. Much of the inspiration behind cybernetics had come during the second world war and in particular Wiener's own work on RADAR directed anti-aircraft guns. He was able to identify the importance of feedback systems in both artificial control systems such as RADAR directed gunfire as well as realising their importance in biological and physical systems. Nehmzow (2003) [94] summed up cybernetics as:

> "The goal of cybernetics is to define the control function f and all control parameters (such as delays, lags, input and output signals etc) in such a manner that the system will respond appropriately to sensory stimuli: intelligence is the minimisation of an error function".

Cybernetics drew inspiration from many areas including biology, physics, electronics, economics and political science. This section presents an overview of some of the early attempts by cyberneticists to produce robotic systems and artificial brains.

2.2.1 Grey Walter's Tortoise

Many cybernetic concepts were applied by Grey Walter when he constructed his robots Elsie and Elmer [57, 142], known as the tortoises. These were three wheeled robots which were driven by an electric motor, could steer themselves and could sense light sources using a photo electric cell. By modern standards they were incredibly primitive lacking digital computers and operating only using analogue electronics and valves. Yet, they were able to exhibit many behaviours including obstacle avoidance and object seeking. A light on top of a charging station allowed them to navigate back to their charging station when their batteries ran low. This could be seen as a primitive version of what would later be known as artificial homeostasis in which the robot is attempting to stabilise its internal state (the battery level) through modification of its overall behaviour (returning to the charging station).

2.2.2 Ashby's Homeostat

Ashby (1952) [11] explored the idea of creating a form of artificial homeostasis with a machine he called the homeostat. It consisted of four units, each of which had a pivoting magnet on top of it. Each unit outputs an electrical current in proportion to the distance of the magnet from its central position. The units were configured so that they each had three inputs which were connected to the outputs of the other units. Each of these inputs was wired to an electromagnetic coil which affected the position of the pivoting magnet. However, before reaching the coils a commutator and potentiometer determined the polarity and the amount of current which would reach the coils. These acted as user configurable parameters. Depending upon the setting of the commutators and potentiometers the magnets would either settle to a stable position, oscillate in a stable pattern or enter a 'runaway' state in which the movements of the magnets would increase with each oscillation. At this point the system was not what Ashby called 'ultrastable'. He gave the example of an ultrastable system as an aircraft autopilot which when the controls between the left and right ailerons were reversed would continue to fly the plane in level flight or the example of how humans can adapt to damaged neurons following an accident and eventually regain the ability to move. To give the homeostat ultrastability a series of electro-mechanical stepping switches (often found in automatic mechanical phone exchanges) were connected to the potentiometers and commutators. Each of these had 25 unique positions giving 390,625 unique combinations

across all the inputs in all four units. An additional electromagnetic coil changed the position of the stepping switches when the output current of the unit exceeded a certain threshold. This caused the combinations of the commutator and potentiometers to change when the output current exceeded this threshold. This effectively regulated the whole system and prevented it from entering the runaway state. At this point Ashby claimed to have developed the homeostat into a form of artificial ultrastability. He admitted that even ultrastable systems have limits beyond which they cannot adapt. Despite being built from what are, by modern standards simplistic components and only being of limited immediate practical use the homeostat demonstrated many of the foundations of an artificial homeostatic system.

2.2.3 Summary of Cybernetics

These early cybernetic architectures provided some of the first attempts to produce biologically inspired control systems for robotic, electronic and mechanical systems. Many of these ideas have been resurrected since the 1980s to form the basis of modern reactive and biologically inspired control systems. Some of the key concepts introduced by cybernetics are those of viability and stability within artificial systems and that a system must exist within certain environmental limits in order to operate successfully. Systems must internally regulate themselves to maintain this viable state. Such principles will be key to an autonomous sailing robot which must maintain a number of key parameters (e.g. battery level) within viable limits.

2.3 Deliberative Architectures

The growing popularity of highly serial digital computers during the 1960s lead to a paradigm shift in thinking away from biologically inspired cybernetic models towards symbolic models and deliberative behaviours. These focused on building explicit models of the world and decomposing the world into a set of objects. The robot will contain in its memory a model of how it perceives the world. This model must either be based on pre-supplied data or from data received from the robot's sensors. The process which developed is commonly referred to as "sense, model, plan, act" in which the robot will sense its environment, build a model of that environment, plan its actions and then act upon them. The process is then repeated so that new data is read from the sensors and the model is updated to reflect any changes.

One of the first attempts to apply such techniques to a real robot was between 1966 and 1972 at the Stanford Research Institute using their robot SHAKEY [97]. SHAKEY was driven by two drive wheels and a caster wheel and was approximately as tall as a typical person. It

featured a laser range finder, television camera, bump sensor, and a radio link to a powerful (for the time) mini computer and a less powerful onboard computer. The remote computer performed planning tasks while the onboard computer was responsible for running the robot. SHAKEY used data obtained from its sensors and an internal model of the laboratory to navigate around. The remote computer planned paths and this data was then downloaded to the robot which would make the required movements and re-sense its environment. Should it detect discrepancies between the environment and its model then the model would be updated and a new plan generated. Due to the speed of the computers being used SHAKEY was limited to very slow movements and its environment was incredibly simplistic consisting of walls and large blocky obstacles. This work was continued with the Stanford Cart [88], a remotely controlled mobile robot with a stereo television camera onboard which resembled a go-kart. A remote computer was able to take control of the cart and navigate around obstacles in a room, having built models from data received from the cameras and planned a route around them.

By the 1980s there was a realisation that deliberative systems suffered from difficulties when operating in real environments. Although robots such as SHAKEY were able to happily navigate around flat laboratories where the only objects were walls and square obstacles over 1 metre tall they ran into difficulties in real world environments where objects could not be classified so easily. It was realised that creating sufficiently detailed models of the world inside the robot was an incredibly complex problem.

Hanard (1990) [52] discussed the difficulties of modelling the real world and identified the "Symbol Grounding Problem" in which to determine what something is you must first have prior knowledge of it. Within a constrained environment, such as a laboratory, it may be quite possible to give the robot a sufficiently accurate description of everything it is likely to encounter. This is not the case in less constrained environments such as outdoors, underwater or on another planet.

2.4 Reactive Architectures

The limitations of deliberative architectures sparked an interest in much more simplistic reactive systems. These made no attempt to build models of the world but instead responded directly to whatever they sensed in their environment. This dispensed with the need for slow and complex planning routines and the reliance on vast amounts of computing power to perform these planning tasks. Instead robots would react almost immediately to the world around them. The underlying philosophy of reactive robotics is often summarised by a quote from Brooks (1991) [24]: "the world is its own best model". Brooks compared his

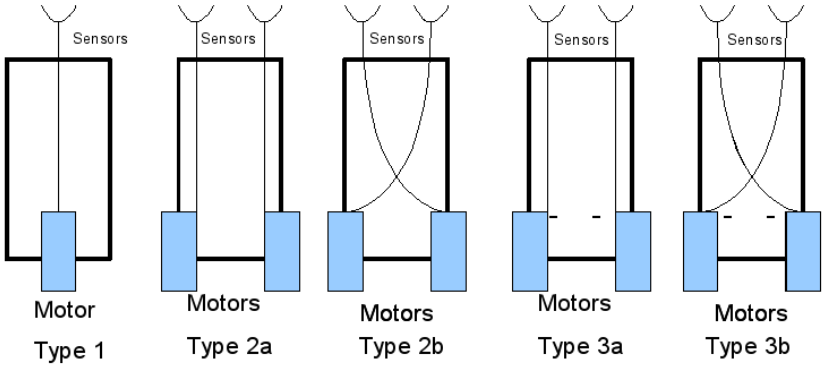

Figure 2.1: A diagram showing type 1,2 and 3 Braitenberg Vehicles.

reactive approaches with trying to mimic the behaviour of insects while previous deliberative attempts had focused (and in his opinion failed) to implement more human like behaviour. Brooks defined four key points to his approach to robotics. He called these "Situatedness", "Embodiment", "Intelligence" and "Emergence". He defined situatedness as being situated in the real world. Embodiment as the robot having a body through which it experiences the world and who's actions will effect the world around them. Intelligence as the robots being observed to be intelligent with respect to how the robot interacts with the world. Finally he defined emergence as the result of a collection of activities giving rise to a particular behaviour and that it can be difficult to attribute these to a single feature of the robot.

2.4.1 Braitenberg Vehicles

Braintenberg (1984) [19], sets out a series of abstract designs for simple vehicles which could exhibit seemingly complex behaviours without the need for any planning or data processing. Although this was essentially the approach taken by Grey Walter in the 1950s [142] it regained popularity in the reactive movement of 1980s and 90s.

Examples of some of Braitenbergs' vehicles are shown in figure 2.1. The simplest of these vehicles (type 1) used a single sensor (Braitenberg deliberately didn't specify the type of sensors to be used) and a motor. There are many possible sensor types which could be used, common examples include a light sensor or temperature sensor. The motor speed will be directly proportional to the output from the sensor, so as the sensor readings increase so will

the motor speed and the vehicle will move towards whatever stimulates the sensor (although it is unable to turn). As the vehicle gets closer to the stimuli it will get faster and faster. More complex variants include two motors and two sensors. Each motor drives one of two drive wheels and one sensor is on the left hand side and the other on the right hand side. In a type 2a vehicle the sensors are connected to the motor on the same side (left sensor to left motor, right sensor to right motor). This causes the vehicle to drive towards the stimulus when it is directly ahead but if it is to one side then the vehicle will turn away. The type 2b vehicle reverses the connection so that the left sensor controls the right motor and vice-versa. This too will drive towards any stimulus which is directly ahead but when the reading from one sensor is higher than the other it will turn towards it. The type 3 vehicles change the connection from the sensor to the motor from an excitory one to an inhibitory one. In these greater sensor readings will result in slower motor speeds, so as the vehicle nears the stimuli it will slow down. The type 3a vehicle will drive towards the source of the stimuli and then come to rest near it while the type 3b will approach it and then turn away.

Braitenberg goes on to define many other types of vehicle along similar lines. Most of these could easily be constructed with basic analogue electronics or very simple digital computers. He shows that despite using seemingly simple techniques very complex behaviours can emerge from them.

2.4.2 Subsumption Architecture

Brooks (1986) [23] devised a system for combining multiple reactive behaviours known as subsumption architecture. Subsumption is a layered approach with each behaviour that the robot can exhibit occupying its own layer. In an example system in Brooks (1986) [23], the lowest level behaviour was a collision avoidance behaviour which uses a distance sensor to detect obstacles and will cause the robot to turn away from them. The next layer was responsible for wandering around the world while the top layer is responsible for exploring. Each layer has full access to all the robot's sensors and actuators but can inhibit a higher layer from having access to the sensors or actuators. The subsumption architecture uses a bottom up approach in which the lowest level layers are implemented first. Only once these have been fully developed, tested and debugged are the subsequent layers developed. This breaks down the task of implementing a robot control system into small and (hopefully) manageable stages.

Subsumption architecture offers the hope of building robots which are capable of dealing with real world problems without the need to resort to centralised control or building world models. In the years following the initial development of the subsumption architecture, Brooks' group at MIT produced several promising robots [27] which showed a great deal

of success in dealing with real environments, at least when compared with the previous generation of deliberative systems. Subsumption architecture can also suffer from scalability problems when many different behaviours are introduced. As there is no central process co-ordinating behaviours they are free to compete with each other and the end result may be counter productive. Such interactions can be difficult to predict in advance and can result in unforeseen behaviour of the robot during real operations. In many areas of robotics (e.g. domestic, industrial or military), safety concerns might prohibit the use of such systems if there is the potential that such actions could kill or injure humans or even if they could damage the robot.

2.4.3 Motor Schemas

Arkin (1989) [9] demonstrated a motor schema architecture in which a series of independent agents each compute a vector for their desired movement of the robot. The eventual movement of the robot is determined by the summation of all the vectors. Each schema focuses on only a single simple behaviour such as staying on the road, avoiding fixed obstacles or navigating towards a given point. When a schema does not need to take part (for example a collision avoidance schema which currently has no obstacles to avoid) then it simply sets the magnitude of its vector to zero. This method has a resemblance to Brooks' subsumption architecture but gives no way for one schema to completely disable another.

2.4.4 Conclusions on Reactive Architectures

Despite the advances in reactive architectures many applications show some level of planning can still be highly beneficial, consider for example a maze solving robot. If we attempt to solve a maze purely by moving around it and sensing the walls then we will probably eventually be able to exit the maze, but given a map of that maze and the ability to localise our position within it then the route to leave the maze can be pre-calculated and the exit achieved immediately. Of course, if the maze changes as we navigate through it (most mazes probably don't change but real environments like a collapsed building might) then our route needs to be re-planned. There are also applications which still lend themselves to deliberative architectures. For example, manufacturing robots which operate in highly constrained environments, could be best suited to deliberative architectures. However these robots do not meet the definition of an agent which is situated as defined by Brooks and it could be argued they are not robots.

2.5 Hybrid Architectures

Shiffrin and Schneider (1977) [120] demonstrated through psychological experiments that the human brain operates in both a deliberative and reactive manner. Given this and the shortcomings of both deliberative and reactive control architectures it might seem logical to follow the same approach in robotics. Taking this concept as a basis, Arkin (1989) [10] devised a hybrid robot control architecture known as AuRA (Autonomous Robot Architecture) which used an A* route planner combined with a series of reactive motor schemas.

Firby (1989) [41] and Gat (1998) [49] proposed three layer hybrid architectures, which involved a real time reactive feedback control layer, a deliberative planning layer and an intermediate sequencing layer which linked the two together. Their architectures did not arise from any biological observations but rather from the practicalities of engineering robots to solve real world problems and the realisation that properties of both reactive and deliberative systems were required and needed to be combined effectively.

Hybrid architectures can potentially draw from the strengths of both reactive and deliberative approaches. They can still incorporate higher level deliberative planning which may be useful to obtain an overall longer term view of the robot's situation while avoiding the need to build detailed models of the world and leaving the second by second (or even minute by minute) needs to a lower level reactive system.

2.5.1 The Legacy of the Reactive versus Deliberative Debate

The debate between reactive and deliberative architectures was often described as resembling a religious conflict, with many on each side believing that their position was absolutely correct and that the other architecture could never work in the real world. Improvements to both sensors and computers during the 1990s and 2000s allowed robots to be deployed into increasingly dynamic environments and began to put these claims to the test. A series of robotic competitions provided realistic scenarios and inspired research into solving real world problems. The AAAI mobile robot competition began in 1992 and initially consisted of obstacle avoidance and navigation challenges. Balch and Yanco (2002) [13] describe the evolution of the competition until 2001. They describe how robots with purely reactive, purely deliberative and hybrid architectures all managed to complete this task. By 1995 all robots were completing the navigation tasks successfully so an additional task of picking up tennis balls (including some balls which could move themselves around) was added in 1996. As the competitors continued to advance, additional tasks were added. In 1997 the robots were tasked with serving Hors d'Oeuvres to spectators, in 1998 with delivering items around an office and in 2000 a robot rescue event was added.

In 2004 and 2005 the DARPA Grand Challenge[1] challenged teams to drive autonomous cars across a section of desert in Nevada. After this was successfully completed in 2005 it was followed up with the Urban Challenge in 2007 in which the robot cars had to navigate city streets and simulated traffic. Of the 6 teams which completed the 2007 Challenge [40, 83, 29, 131, 138, 132] all used some form of deliberative architecture. Advances in computer technology had now ensured that real time deliberative architectures operating in complex real world environments were now a possibility. Although at present they still require a car full of computers, a laser scanner and high resolution cameras to process a real world traffic situation deliberatively. It is not (yet) practical to do this in small, light weight, lower power and low cost robots. Currently this restricts such approaches from low cost personal robots or from small and low power budget airborne systems, underwater robots or autonomous surface craft including sailing robots.

2.6 Biologically Inspired Approaches

Biologically inspired robotics covers both bio-mimetic approaches to locomotion and actuation (for example humanoid robots or insect like robots) and approaches to the underlying control system even if the robot itself does not physically resemble a biological entity. It has long been understood that many of the problems which need to be solved in robotics have already been solved by biology. Copying, or at the very least, taking inspiration from biology has the potential to solve many robotics problems. Although in many cases creating a full copy of the biological solution is beyond the limits of current technology. In many ways biologically inspired robotics actually pre-dates the field of robotics itself, as early science fiction robots [32] were often humanoids with many human like features. Attempts at actually producing biologically inspired systems date back to some of the cybernetic systems of the 1940s and 50s such as Grey Walter's tortoises[142] (discussed in section 2.2), Ashby's Homeostat [11], early work on Neural Networks [110, 72] also have roots in the reactive systems of the 1980s [19, 23, 27]. Biologically inspired algorithms have been applied to many different areas including multi robot co-operation[66], evolutionary learning strategies[98] and artificial homeostasis. This section reviews attempts to build biologically inspired robot controllers based upon artificial analogues of the neural, endocrine and immune systems. It also briefly introduces the biology behind each of these, but is merely intended to equip a reader with little or not prior knowledge of these areas with enough information to understand the key biologically concepts.

[1] `http://archive.darpa.mil/grandchallenge/overview.asp` accessed 15/04/2011

2.6.1 Neural Systems

This section discusses neural systems, it covers both biological neural systems and attempts to produce artificial analogues of them. It reviews the various approaches to constructing artificial neural networks, presents the mathematics behind them, provides an overview of their use of neural networks in robotics and discusses their limitations.

2.6.1.1 Neural signalling in Biology

The neural system is typically associated with the roles of sensing, coordinating motor functions, performing memory and learning. It is generally considered to consist of two key parts: the Central Nervous System (CNS) and Peripheral Nervous System (PNS). The CNS is composed of the brain, (and in vertebrates) the retina and spinal cord. While the PNS consists of the sensory and motor neurons and the nerves which connect them to the CNS. The neural system operates through two key signalling systems, electrical impulses and chemical neurotransmitters. It is capable of coordinating sub-second responses over distances of several metres (in the case of large mammals such as whales or elephants).

The nervous system is made up of billions of interconnected cells known as neurons. A diagram of a typical neuron is shown in figure 2.2. Electrical input signals are received into the nucleus via the dendrites and are aggregated together in the cell body. If the total strength of these signals exceeds a threshold at the axon hillock then the neuron will "fire" and send an electrical impulse along the axon to the axon terminals. At the end of each axon terminal is a synapse which forms a link to the dendrites of another neuron. The synapse is shown in detail in figure 2.3. A small gap (100 nanometres), known as a synaptic cleft, is left between the synapse and the dendrite of the next neuron.

In addition to the electrical impulses jumping across the synaptic cleft, chemical neurotransmitters which effect the target dendrite, may also be released from the synaptic vesicles in the axon terminal. These then bind to receptors on the dendrite, given that they only need to cross the synaptic cleft these are fast acting, having an effect upon the target neuron within milliseconds. Neurotransmitters can have either an inhibitory or excitory effect upon the target neuron and these effects can be long lasting. Common examples include endorphins, often associated with responses to pain, epinephrine (adrenaline) associated with the "fight or flight" response (Epinephrine is also a hormone and effects other organs too, this will be discussed in section 2.6.2).

A third method of neural signalling has also been discovered. This is that neurons are capable of modulating neighbouring neurons [48], to which they have no direct connection, through the release of modulatory gases such as nitric oxide. These slowly diffuse in three

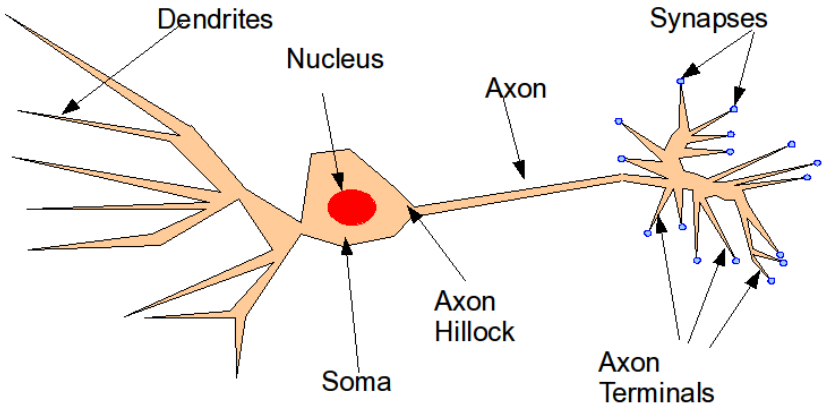

Figure 2.2: A illustration of a neuron.

dimensions over short distances through the neural tissue, modulating the behaviour of neighbouring neurons.

2.6.1.2 Artificial Neural Networks

Artificial Neural Networks are a series of techniques which attempt to form an artificial analogue of neurons found within the brain. As discussed in the previous section, in biological neurons, a series of branching fibres known as dendrites receive electrical impulses from other neurons or cells which perform sensing (e.g. retinal cells). These dendrites act as the inputs to the neuron and eventually link into the cell body or soma, which is responsible for integrating and processing these inputs. The soma will then determine whether or not the neuron should "fire" and release an electrical impulse through the axon. The axon will, in turn be linked to the dendrites of other cells.

The development of Artificial Neural Networks began with work by McCulloch and Pitts (1943) [72]. They devised the mathematics behind a simplified artificial model of a neuron. In their model a neuron had one or more binary inputs (the dendrites) and a single binary output (the axon). They took the sum of the inputs and then applied them through a thresholding unit. If the sum of the inputs exceeded the threshold then the output of the neuron would be a one if not it would be zero. They envisaged cases in which multiple artificial neurons could be connected together and that this would give the neurons the ability to perform calculations in much the same way that a collection of logic gates form the basis of a digital computer.

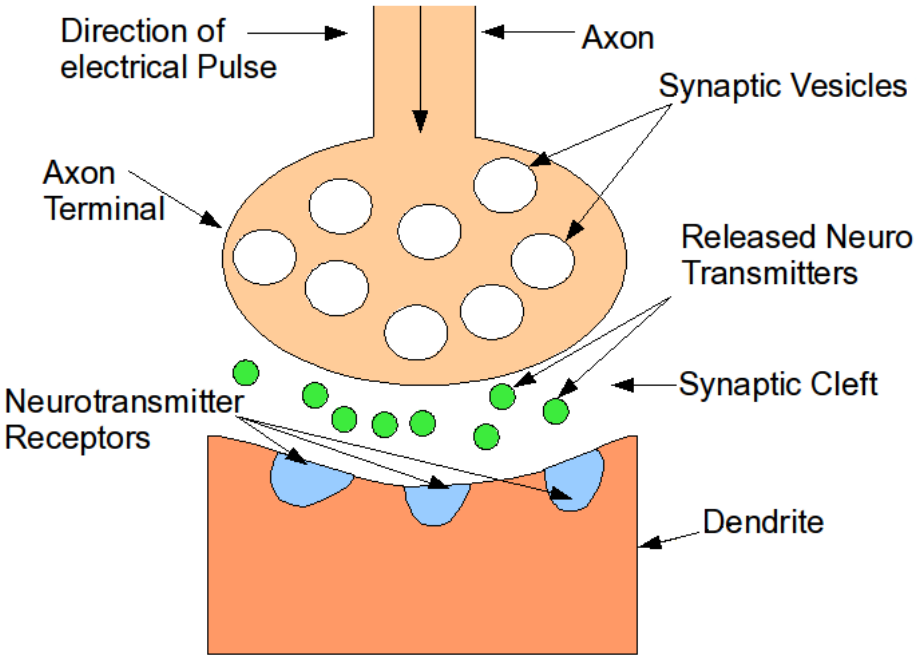

Figure 2.3: An illustration of a synapse.

While this concept underpins the basics of a neural network the restriction to using binary inputs reduces the network's flexibility. An example of a McCulloch Pitts neuron is shown in figure 2.4.

Further development of neural networks continued during the 1950s inspired by the goal of many cyberneticists and computer scientists of the time to create an artificial brain. The artificial neuron was further refined by Rosenblatt (1958) [110]. In this model multiple neurons are joined together in a layered architecture, known as a Perceptron. In each neuron inputs were encoded as continuous values instead of binary ones (as had been used by McCulloch Pitts). Each input $(x_1, x_2.....x_n)$ value was also multiplied by a weight $(w_1, w_2.....w_n)$. The sum of all the inputs multiplied by their weights is then taken. This is expressed mathematically in formula 2.1.

$$\sum_{i=0}^{n} w_i x_i \tag{2.1}$$

This sum is then run through a thresholding function as in the McColluch Pitts neuron. By adjusting the values of the weights and the connections within the network the behaviour of the network can be modified. An example of a Rosenblatt style neuron is shown in figure 2.5. Perceptrons are organised in two layers, an input layer and an output layer. The input layer consists of one or more neurons which receive input from some external stimuli. These

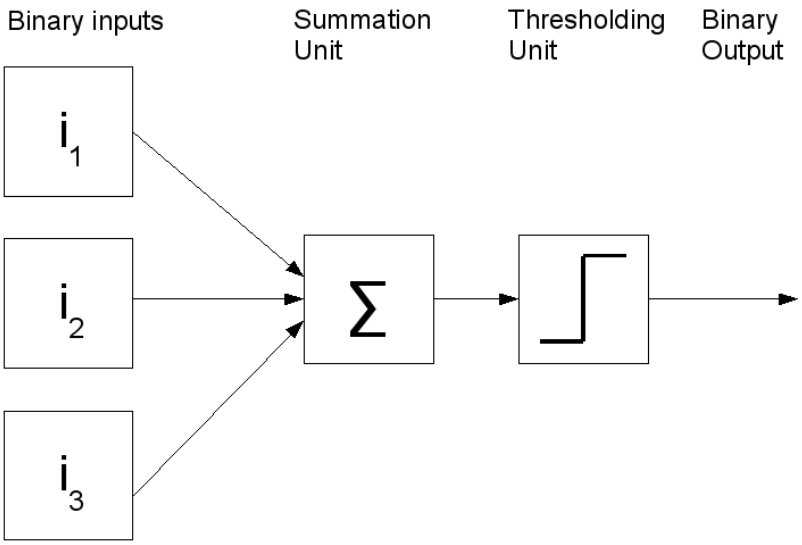

Figure 2.4: A McCulloh Pitts Neuron.

are then connected to several more neurons in the output layer. A commonly cited example [102, 70, 106] of a perceptron is in performing character recognition. In this example a grid of pixels forming a kind of artificial retina are used as an input area and each pixel is connected to a single input neuron. These are then connected to some output neurons, which will each represent a possible character which the system could recognise. Figure 2.6 shows an example of such a multi-layer perceptron which can distinguish between the letters 'T', 'U' and 'V' and is recognising the input of a letter 'T'. Such networks are often referred to as "feed forward" networks as each layer feeds forward into the next and no looping back to previous layers is allowed.

2.6.1.3 Learning in Neural Networks

Inspired by the idea from Hebb (1949) [54] that neurons adjust their behaviour in order to encode learning or that:

> "When an axon of cell A is near enough to excite cell B and repeatedly or persistently takes place in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased."

This lead to the idea that through adjusting the weights of an artificial neuron it could be made to learn to perform a given function. Rosenblatt's original perceptron learning algorithm focuses on training a single neuron to solve a classification problem (such as classifying written characters) with multiple inputs and two output classes. In this algorithm a set of

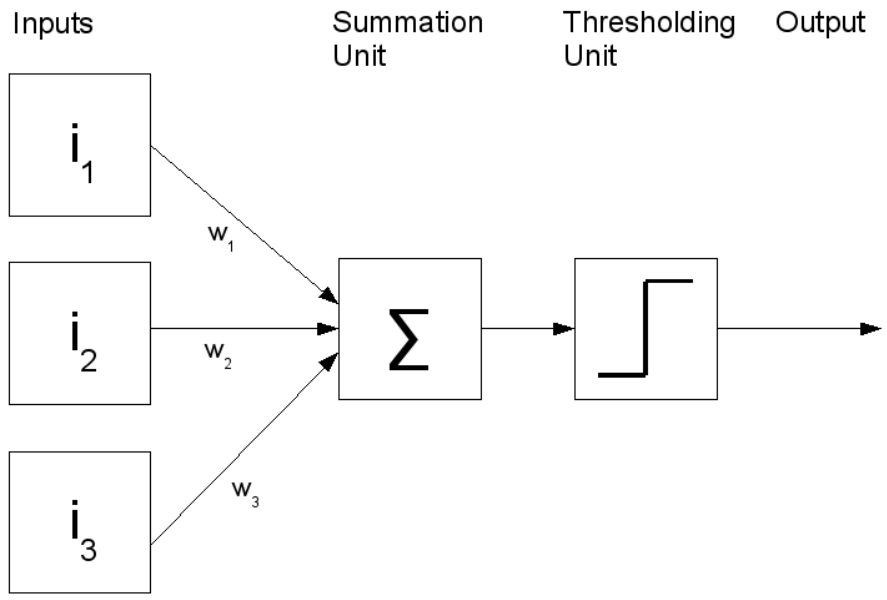

Figure 2.5: A Rosenblatt neuron.

training data is supplied which links inputs to the correct output classification. The algorithm operates by starting with random values for each weight. The output of the network is then compared with the training data and the difference is calculated. This value is then used to adjust the weights. The full algorithm is shown in algorithm 2.1.

Rosenblatt showed in his perceptron convergence theorem that this algorithm would always converge on a correct solution. Although, this depended upon the function to be learnt being linearly separable. This means that the output space can be divided between the two output classes using a straight line, an example is shown in figure 2.7.

After much initial excitement about perceptrons many researchers began to understand that they also faced a number of limitations. The perceptron learning algorithm could only be applied to a single neuron and was restricted to linearly separable functions. In 1969 Minsky and Papert [81] showed that even trying to solve a simple exclusive OR function was not possible using this method due to the linear separability problem.

It was not until the 1980s that methods capable of training an entire network and overcoming the linear separability problem (and thus learning XOR) were devised. Perhaps the best known of these is the back-propagation method devised by Rumelhart, Hinton and Williams in 1986 [112]. This method makes use of a gradient descent to train the entire neural network. In order to do this the threshold activation function had to be replaced with a function that was continuous, non-linear and differentiable. Therefore, the threshold activation was replaced with a sigmoid activation function (shown in equation 2.2), this also

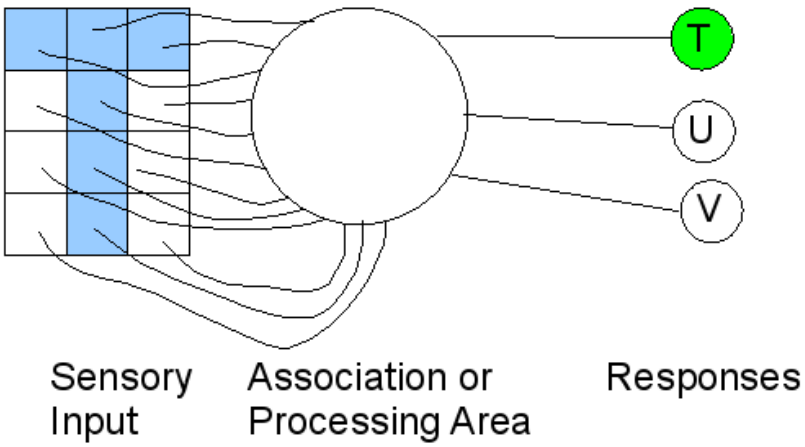

Figure 2.6: An example of a Rosenblatt style multi-layer perceptron being used to recognise some text on the left. The output of the network is shown on the right, where it has correctly identified the letter shown to it.

limits the output of each neuron to between 0 and 1.

$$y = \frac{1}{(1 + e^{-x})} \tag{2.2}$$

Typically back-propagation is applied to networks with 3 or more layers. Each input node is then connected to every node in the next layer, known as the hidden layer and each of these will be connected to every node in the output layer. This is known as a fully connected network and is common practice, although is not required for back-propagation. In networks designed to solve more complex problems there might be additional hidden layers or additional nodes in the hidden layer. Figure 2.8 shows an example of such a network.

Like the perceptron learning algorithm, back-propagation begins by randomly initialising all weights in the network and a set of training data is used to teach the network. The output of the network is then calculated with the first item of the training data. A set of error values are then calculated by subtracting the value of each output from its expected value in the training data. This error is then propagated back through the network starting at the output layer and working back through the hidden layer(s). This allows for weights in the hidden layer(s) to be modified based upon the errors calculated on the output layer. The process will be repeated in a loop until the sum of the output layer errors is less than some threshold value. Through the use of three or more layers the back-propagation algorithm is able to train a neural network to solve problems which do not exhibit linear separability including the XOR problem. Despite this, back-propagation is not perfect, it can take a long time to converge on the correct solution and the total error typically falls at an inverse exponential

Algorithm 2.1 The Perceptron learning algorithm.

```
//loop until error is less than an error threshold
while err < training_threshold
    //loop through each training example
    for each trainingexample e
       //loop through each input/weight
       for each input i
           //set it to the training input
           x[i] = training_inputs[e][i]
           //calculate output
           sum = sum + x[i]*w[i]
       //perform threshold function
       output = Threshold( sum )
       //work out difference between output and training data
       err = training_outputs[e] - output
       //adjust weights accordingly
       for each weight i
           w[i] = w[i] * err * learning_rate
```

rate with the largest changes occurring in the first few iterations and the final few percent of the error reduction taking hundreds or thousands of iterations. The back-propagation algorithm can also become stuck in local minima depending upon the function being trained. This can be at least partially overcome through the introduction of a momentum term which is multiplied by any change to the weight values. The back-propagation is also restricted to only operating on feed-forward networks in which data is fed from the input layer to the hidden layer and onto the output layer. No loops (e.g. linking a layer back to a previous layer) are allowed. Some implementations of back-propagation also include a bias node on the hidden and output layer. These have variable weights, but a fixed input value of one. This forces a non-zero value into the network, as zero values can cause the learning process to become caught in local minima.

Even in the face of these limitations the combination of multi layer perceptrons and back-propagation do allow neural networks to be trained to solve many classification problems. They are considered to be tolerant of noise [113]. Although they lack the ability to offer any way of explaining how they arrive at their results. Perceptrons can be applied to a very wide range of problems, typical applications include classification or pattern recognition systems, forecasting and prediction systems and digital signal processing.

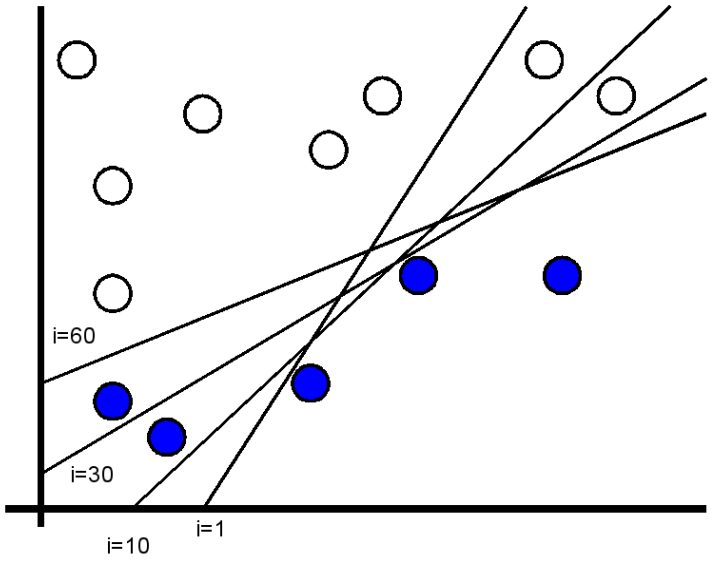

Figure 2.7: An example of linear separability. The lines represent the classification made by the perceptron learning algorithm, as the number of iterations increases the separation converges on the correct solution.

2.6.1.4 Neural Networks in Robotics

During the late 1980s and 1990s, neural networks began to gain some level of popularity in robotic control, sensing and vision systems. This section outlines some of those implementations and discusses the limitations they encountered.

Nguyen and Widrow (1989) [96] presented a theoretical problem of reversing a delivery truck to a load bay. This is a highly complex task with a non-linear solution, which they did not believe could easily be implemented through a rule based system. They applied back-propagation learning to a simulated truck and were able to create a solution which could backup the truck even when it was initially jack knifed. Although only in simulation this work suggested that neural networks could be used to produce controllers for a wide range of non-linear control problems including the control of robot vehicles.

Modada and Floreano (1995) [87] combined a neural network and genetic algorithm to produce a Braitenberg [19] style controller for a small (55mm diameter) Khepera mobile robot. They replaced back-propagation learning with a genetic algorithm. This allowed for online learning and instead of needing to specify explicit training data patterns they only needed to supply a fitness function which represented a specification of the desired global behaviour of the robot. Through this process they were able to have the robot solve small mazes using 8 infrared proximity sensors to detect walls. The fitness function penalised

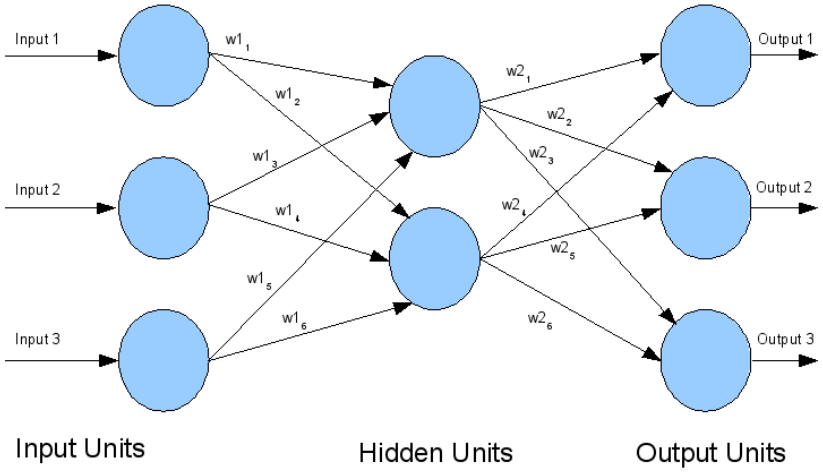

Figure 2.8: A diagram showing a typical multi-layer perceptron. On the left are three input units, in the middle two hidden units and on the right 3 outputs.

collisions and tried to keep the robot moving. This approach caused initial generations of their controller to perform seemingly random movements, but as the fitness improved across generations collision avoiding behaviour emerged.

The NAVLAB project at Carnegie Mellon university aimed to build a fully autonomous car capable of driving on normal roads. A neural network based vision system known as ALVINN (Autonomous Land Vehicle In a Neural Network) [103] was developed. ALVINN used a 960 pixel (30x32) camera (and later laser scanner) generated image connected to a 4 layer perceptron. 30 output neurons determined the position to place the steering wheel in. The left most neuron turned the steering wheel sharply left, the central neuron would centre and the rightmost would produce a sharp right turn. The network was trained using back-propagation with visual data taken from a road facing camera during human controlled drives. This training data consisted of a variety of road types and bends in the road. ALVINN essentially repeated the human driver's response to bends in the road. It demonstrated the robustness of neural networks to create a generalised solution which was capable of finding the edge of the road and keeping the vehicle on the road across a wide variety of road types and in a wide variety of lighting conditions.

Thrun et al (1998) [133] developed a system to interpret sonar data using a neural network as part of work to build a mobile robot to navigate indoor environments. They aimed to establish an occupancy value for a given point on a map grid, this determined the probability

that the part of the map in question contained an obstacle. Training data was created by driving the robot through a known environment and manually assigning occupancy values. They noted that neural networks provided the advantage of easily adapting to new circumstances and cite the example of the walls at the 1994 AAAI robot competition being much smoother than the walls the robot was originally developed to detect. Despite this difference the network could be quickly retrained to recognise the smoother walls and detect them correctly. They also noted that a neural network provided a useful way to evaluate multiple sensor readings together in comparison to traditional approaches which evaluated each sensor individually.

Howell, Wood and Koksal (2003) [58] developed a neural network system for classifying sounds received by a passive sonar system. Their system operated by analysing fast Fourier transforms of sound recordings and was able to distinguish between various marine animals, seismic activity, boats and underwater volcanoes. Although not embedded in a robotic system, they note that the system would be capable of operating in an intelligent sensor network which would only begin reporting when particular species were detected. This work demonstrated the capabilities of neural networks to classify complex sensor data and reduce it into a few simple outputs.

These examples all demonstrate the ability of neural networks to be applied to solve either sensor classification problems or to act as robot controllers. They also highlight the difficulties involved with trying to obtain suitable training data for a robot and some solutions to produce online learning systems which can overcome or at least reduce this problem.

2.6.1.5 Limitations of Neural Networks

While the examples in the previous section demonstrate the robustness of neural networks in learning generalised solutions and their suitability to classifying of data, neural networks also suffer from a number of limitations. These limitations fall into four key categories, firstly neural networks have difficulty analysing information which varies over time. Secondly, they are ideally suited to highly parallelised computer architectures while most digital computers are highly serialised. Thirdly, they offer no way of providing an explanation for how they arrive at their output. Finally, there are a number of problems associated with the back-propagation learning algorithm creating solutions which are either too specific to the training data or too general.

Feed-forward neural networks are not able to easily deal with temporal information. They are capable of producing a set of outputs for a given set of inputs, but are not capable of considering past values of those inputs and outputs. For example, if a neural network were controlling a robot actuator, which is attempting to move the actuator to a target position.

A feed-forward multi-layer perceptron could be trained to determine the amount of time the actuator must be turned on in order to reach the target, given its current distance from the target. This simple proportional control strategy may be suitable for many situations, but may suffer from problems with overshoot, undershoot or oscillations. Control engineers frequently solve this problem [64] by adding an integral controller which examines the rate of change (in the difference between target and current position) over time. However, adding this feature to a pure feed-forward neural network is not easily achieved, without relying on code outside of the neural network to perform this calculation. Brooks (1991) [26] noted how neural network implementations are often augmented with extra code to help them overcome this limitation.

Ideally each neuron in a layer of a neural network could be processed in parallel with all others in the same layer. This would lend itself to a computer which contains many parallel units, with each only requiring enough complexity to receive some inputs, multiply them by a weight, perform summation and thresholding and finally to communicate these to the next layer. However, digital computers have emerged along the lines of the Von Neumann architecture which are able to process individual instructions rapidly and in series. Arbib (1989) [6] discusses how computer hardware has over time gained limited parallel processing features and how he hopes that future computers will lend themselves better to parallel processing. Arguably at the time of writing Arbib's vision for the future still remains only partially fulfilled. Modern (low cost) computers typically only have a few generalised processing cores rather than tens or hundreds of simple parallel units that would lend themselves to processing a neural network.

A neural network also offers very little in the way of an explanation for how it reached a particular decision. It must instead effectively be viewed as a black box, to which data goes in and out but who's internal workings remain a mystery. This can cause problems when software does not perform as desired and the developer needs to understand why a given output is being produced. It can also be problematic when there is a need to explain to the user why a particular decision has been arrived at.

A commonly observed problem [102] of back-propagation does not lend itself to online learning, as the learning algorithm can be slow to execute and requires new training data to be generated through some form of feedback process. As previously discussed, to overcome these limitations many neural networks in robotics are often augmented with other artificial intelligence strategies. For example applying genetic algorithms as an alternative online learning mechanism. Another difficulty with back-propagation is that it is difficult to discover the point at which training should stop. This is normally achieved by using a threshold value, which the difference between the network output and the training data must be below. If

this value is set to be too small then the network may be over trained (or over fitted) for the training data, this reduces the network's ability to generalise its solutions which reduces performance when used on anything but the training data. This also highlights the need for sufficiently accurate training data, which represents the real situations the network may encounter. In robotics it is often tempting to generate training data through simulation as it can be difficult to obtain enough training examples and an even distribution across the entire range of potential inputs. However, simulation data is likely to differ from real world data and may not provide a sufficiently accurate or diverse fit to real world situations.

Despite these limitations, neural networks still remain a powerful mechanism for providing interpreting sensor data and providing actuator control in a robot. Given their limitations they may not always be suitable to form the entirety of a robot control system, but are still a contender to make up key parts of it.

2.6.2 The Endocrine System

This section discusses the endocrine system, it covers background of biological endocrine systems and attempts to produce artificial analogues of them. As with the previous section on the neural system, it is only intended to serve as a brief overview of how biological endocrine systems function. An outline of how the endocrine system functions, its role in homeostasis and interaction with the neural and immune systems is presented. A number of artificial endocrine inspired algorithms are discussed and the mathematics behind some approaches, which are relied upon in later chapters are presented.

2.6.2.1 Biological Endocrine Systems

Endocrine signalling involves the secretion of chemical messengers known as hormones from glands around the body into the bloodstream. These are then carried rapidly around the body (typically within a few seconds) and reach virtually all cells in the body. They will only interact with certain target cells, this interaction takes place either (depending on the type of hormone) by bind with a receptor on the cell surface or by entering the cell body and interacting with the behaviour of the cell's internal workings. In contrast with neural signalling, the endocrine system can be seen as a global broadcast system. It is capable of sending messages to target cells distributed across the body rather than to a those nearby the signal originator. A diagram illustrating hormone secretion and binding to a surface receptor are shown in figure 2.9.

Hormones were first discovered in the early 1900s [14] and are, in comparison to many other areas of biology (for example the immune system), relatively well understood. How-

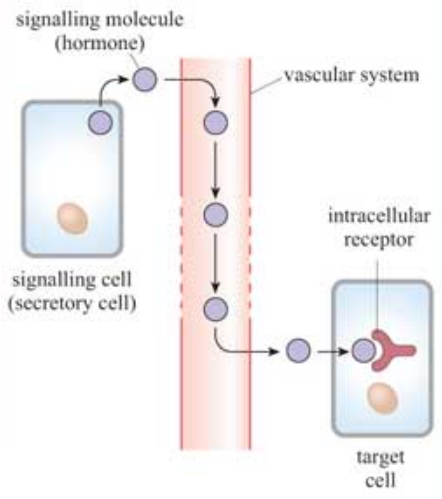

Figure 2.9: An illustration of endocrine signalling, showing a hormone travelling through the bloodstream to a receptor on a target cell. Image from Open University Learning Space [137]

ever, much is still not understood about their operation and studies into hormone functions and interactions often focus on solving medical problems rather than creating a complete understanding [34].

There are several different classes of hormone with differing chemical structures. Peptide hormones are relatively large, bind with receptors on the surface of target cells and tend to trigger reactions quickly (between seconds and minutes). Upon binding with a target cell they trigger the release of "second messenger" chemicals which perform intra-cellular messaging and cause behavioural changes within the cell. Steroid hormones are derived from cholesterol molecules and are physically much smaller than peptides. They can enter the body of the cell and interact directly with the internal apparatus of the cell. This can interfere with the synthesis of proteins within the cell causing a change in behaviour of that cell. This is a much slower process than relying on a second messenger chemical and these hormones tend to act on much longer time scales (hours to days) than peptide hormones. A number of hormones are neither steroids or peptides, notable examples include epinephrine (adrenaline) and thyroxine. Despite not being peptides or steroids, epinephrine behaves in a manner similar to a peptide, while thyroxine behaves similarly to a steroid.

In vertebrates (including humans) several key glands are responsible for secreting hormones. These include the hypothalamus, pituitary gland, gonads (testes in males or ovaries in females), adrenal gland, pancreas, thymus, thyroid gland, and pineal gland. Figure 2.10 shows the location of these glands (except the hypothalamus) in the human body. Of these glands the hypothalamus is unique as it does not release hormones directly into the blood stream but instead releases them into the hypothalmic-pituitary portal circulation which

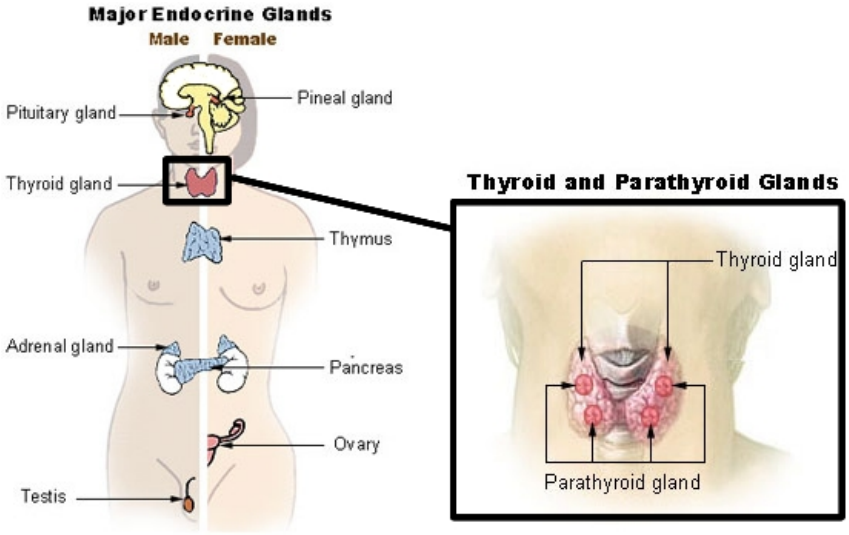

Figure 2.10: A diagram of the major endocrine glands in humans. The hypothalamus is the yellow portion directly above the Pituitary gland. Image courtesy of the National Cancer Institute `http://training.seer.cancer.gov/anatomy/endocrine/glands/` accessed 15/04/2011

connects to the neighbouring pituitary gland. Hypothalmic hormones then act upon the pituitary gland, which in turn secretes hormones into the blood stream where they can affect other cells. Table 2.1 contains an example list of some human hormones, the glands which produce them and the class of the hormone.

Gland Name	Hormones Produced	Hormone Action	Class of Hormone
Pineal	Melatonin	Regulation of sleep patterns	Other
Testes	Testosterone	Male reproductive system	Steroid
Ovaries	Progesterone, Estradiol	Female reproductive system	Steroid
Adrenal	Cortisol	Stress Response	Steroid
	Aldosterone	Salt Conservation	Steroid
	Epinephrine (adrenaline)	Stress Response, Fight or Flight Response	Other
Thyroid	Thyroxine (T4)	Regulates metabolic processes	Other
Pituitary	Adrencortical-Stimulating Hormone (ACTH)	Controls production of other hormones in the adrenal gland	Peptide
	Growth Hormone (GH)	Regulation of growth	Peptide
	Luteinizing Hormone (LH)	Stimulates ovulation, estrogen secretion	Peptide
	Prolactin	Growth of mammary glands and milk production during lactation.	Peptide
	Thyroid Stimulating Hormone (TSH)	Stimulates thyroid hormone synthesis	Peptide
Pancreas	Glucagon, Insulin	Regulation of glucose levels	Other
Thymus	Thymosin	Regulation of immune functions	Other
Parathyroid	Parathyroid Hormone	Regulation of calcium	Peptide
Hypothalamus	Corticotropin Releasing Hormone (CRH)	Stimulates production of ACTH in the pituitary gland	Peptide
	Tyrotropin Releasing Hormone (TRH)	Stimulates TSH production in the pituitary gland.	Peptide

Table 2.1: An example list of the glands and hormones of the human endocrine system and their function.

2.6.2.2 Hormones and Homeostasis

Homeostasis is the term first coined by Cannon (1927,1929) [30, 31] for the process by which the body maintains itself at an optimal equilibrium which is conducive to sustain life. The term literally means the "same state" and refers to how the biological systems manage to maintain an internal state of equilibrium despite massive changes in its external environment. This is achieved through a series of negative feedback systems. Common examples of parameters regulated by homeostasis include: blood pressure, blood glucose levels, body temperature, salt levels and calcium levels. Many of these are regulated through the interaction of hormones and the external environment. The hormones which control this regulation rarely act alone and are often part of larger processes known as hormone cascades. A hormone cascade typically involves a loop of several hormones with each hormone having a role in controlling the production of another. Take, for example the CRH, ACTH and Cortisol cascade (described in detail on page 43 of Crapo (1985) [34] and as shown in figure 2.11). The Hypothalamus releases CRH which is secreted into the hypothalamic-pituitary portal circulation and triggers the pituitary gland to secrete ACTH into the bloodstream. ACTH then stimulates the adrenal glands to secrete Cortisol. One effect of Cortisol is to trigger the liver to convert excess glucose (a sugar which can be immediately used by the cells in the body for energy) into glycogen to store energy for long term use instead of consuming it immediately. Increased levels of cortisol also trigger a negative feedback mechanism to the hypothalamus and pituitary glands which stops the production of CRH and ACTH, which in turn stop the further production of cortisol. Page 159 of Bentley (1982) [16] notes how this process also sees an amplification in which only 0.1 μg of CRH initially being secreted leads to 1.0 μg of ACTH being secreted and this is followed by 40 μg of cortisol and results in 5600 μg of glycogen. In addition to large multi hormone cascades, hormones also operate in pairs that work together to maintain a stable homeostatic state. For example, glucagon and insulin operate together to maintain a constant blood sugar level, as discussed on page 30 of Hardie (1994) [51].

2.6.3 Artificial Endocrine Systems

This section reviews a number of attempts to produce artificial homeostatic systems and algorithms inspired by the endocrine system. As discussed in section 2.6.2.2, homeostasis is the biological process in which various parameters are maintained within an equilibrium and regulated within set limits through a series of negative feedback mechanisms. Artificial homeostasis attempts to produce a stable internal state within an artificial system such as a robot. A number of examples of artificial homeostasis have emerged since the advent

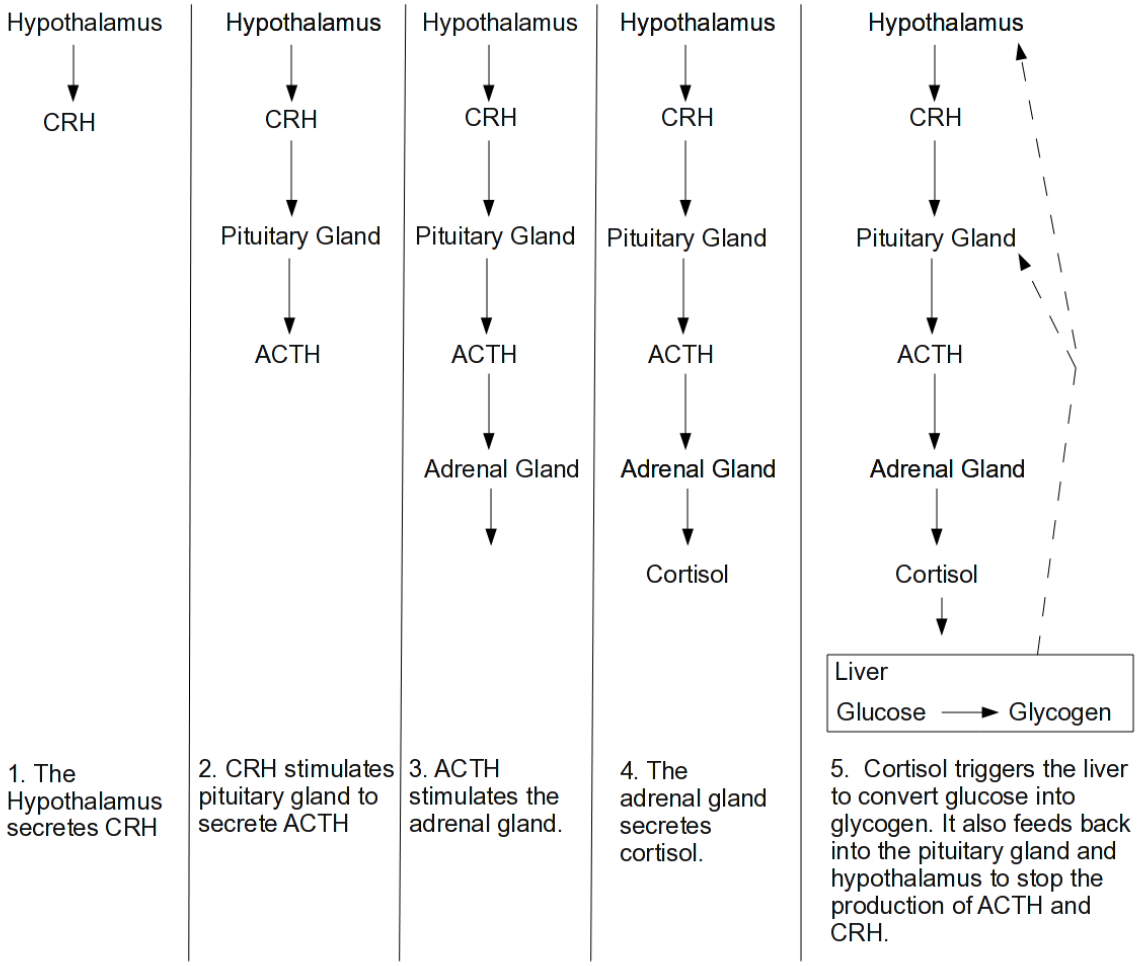

Figure 2.11: A diagram showing the stages of the CRH, ACTH, Cortisol hormone cascade.

of reactive robotic architectures in the mid 1980s, some of the early cybernetic systems discussed in section 2.2 could also be considered examples of artificial homeostasis. The concept of artificial homeostasis particularly lends itself to power management as the process of trying to keep a robots battery charged can be likened to a human body attempting to maintain a constant blood glucose level. Additional applications could include regulating the temperature of electronic components or trying to maintain equal duty cycles of two (or more) actuators which can be used to perform the same task. The concept could also be applied to achieving a new stable state in the event of component failure. When considering potentially useful behaviours for a robot which must operate autonomously for long time periods (such as a sailing robot) the concept of artificial homeostasis is particularly appealing.

2.6.3.1 Arkin's Homeostatic Control

Arkin (1992,1993) [8, 7] noted that truly autonomous robots would need to perform self-monitoring and manage their own internal functions with little or no interference from their operators. He observed that people can survive for a limited time when sleeping or comatose despite no high level planning and that much of these life sustaining processes take place below the level of the central nervous system in the endocrine system. He suggested that internal state factors such as available fuel reserves and internal temperature should be taken into consideration when planning a robot's activities. He devised a system which extended the AuRA motor schema approach [10, 9] described in section 2.5 and added receptors for artificial hormones to the motor schemas. A hormone would be be released in response to low levels of fuel in the robot and the receptors would influence the calculation of the vector which drove the robot's motors and determined the direction of travel. The hormone would suppress the variation in the planned path taking the robot on more direct (and thus requiring less fuel) but riskier routes that come closer to obstacles.

He also defined two forms of "stress" relating to heating of components. Local stress as heat which must be redistributed around the robot in order to maintain the operation of a particular subsystem. Global stress is heat which must be transferred from within the robot to the surrounding environment. In order to mitigate this stress the robot releases hormones which cause the robot to slow down in order to reduce the temperature of its motor.

This approach is able to modulate the behaviour of a robot in response to its internal state without any explicit symbolic reasoning about that state. Such an approach lends itself well to integration within a reactive control system and helps to minimise the amount of processing power required to implement it. Although Arkin limited his study to simulations, such an approach offers potential for real robots. However, it maybe questionable as to how much the difference in distance there is between a riskier shorter path and a less risky longer one and whether or not this is the best way to reduce power consumption. Unless the cost of steering were significantly higher than the cost of moving forwards (as it might be in a sailing robot where forward movement is essentially free but steering is not). There is also a danger in incorporating a path planning system where changing fuel levels could constantly trigger new plans to be formulated and constant change to occur. Some kind of dampening mechanism maybe required to ensure that a particular choice of path is kept for a given period of time.

2.6.3.2 Endocrine Inspired Action Selection

Brooks (1991) [25] identified that hormones in biological systems set gains or bias outputs of systems, he also took on the concept of using thresholds to perform action selection from Maes (1989) [69]. He implemented an artificial endocrine system using a walking robot, in which hormones were released by releasers. Hormones were able to influence behaviours which operated in a subsumption architecture. Each behaviour was implemented as an augmented finite state machine. These could take on one of three types, regular behaviours which cannot be disabled, haltable behaviours which can be disabled and inhibitable behaviours who's output can be inhibited but which continue to run and maintain their internal state. An activation threshold was specified for each behaviour. Behaviours were able to excite releasers in order to trigger the production of a particular hormone and thus activate or deactivate other behaviours. Some behaviours also used the hormone levels to influence parameters (such as the speed of a walking robot). Through the use of this architecture, Brooks was able to demonstrate control over several behaviours within his walking robots. He suggested it as an approach to increase the complexity of behaviour based systems as the hormone system creates processes which can regulate a number of behaviours without resorting to traditional deliberative systems to perform this oversight.

Yamamoto (1993) [147] observed that robots using the subsumption architecture could become deadlocked and repetitive when two or more layers competed with each other. He suggested that a hormone system could observe longer term changes and break this deadlock by altering the priority of some layers or by completely enabling or disabling them. He also suggested that a hormone system could change an overall emotional behaviour of the robot. He demonstrated this concept using a small vacuum cleaner robot known as SOZZY, which was able to detect obstacles through bump and proximity sensors and could detect its charging station via an infrared beacon detector. Four emotions were devised for the robot: joy, desperation, fatigue and sadness. The joy emotion would become dominant when the robot was able "eat" dust and to see its beacon, once dominant the robot would continue to vacuum dust until another behaviour took over. The desperation behaviour occurred when the motor constantly stalled due to the robot hitting an obstacle, at this point it would stop trying to perform obstacle avoidance and act somewhat aggressively to try and free itself. The fatigue behaviour became active when the battery was low and triggered the robot to return to its charging station. The sadness emotion occurred when the robot lost sight of the beacon. To overcome this the robot would give up on attempting to navigate or vacuum and instead wander around, if it detected a human (by detecting their body heat) it would edge up to them in order to "ask" for help. A stabiliser process prevented excessive buildup of any hormone which might have stopped a behaviour switch from occurring, this could be

likened to the process of hormone decay in biological systems.

2.6.3.3 Autonomous and Emotional Agents

McFarland (1995) and McFardland and Spier (1997) [73, 74] present the concept of robot activity cycles which split a robot's time between three key tasks: working, finding fuel and refuelling. When working the robot will be consuming fuel but will be performing its intended purpose and creating a reward (described as "money") for its "employer" (owner/operator). When finding fuel the robot will be searching around its environment for a fuel source such as a charging station and will be consuming fuel and not earning money. When refuelling the robot will be increasing its fuel levels but not earning money. The robot must remain in a viable state in which it does not run out of fuel and manages to achieve some work. Therefore, certain limits for the operation of the robot can be defined and the robot must remain within these limits in the same way that a biological system must remain within certain homeostatic viability limits. The robot will have to adjust its behaviour from time to time in order to remain within these limits. For example, switching from a working state to a fuel seeking and refuelling state. McFarland and Spier [74] also discuss the possibility of opportunistic refuelling occurring when the robot is not low on fuel but happens to come across some fuel during the course of working.

This work took place in a highly artificial and simulated environment in which the robot would randomly encounter fuel throughout its environment. When considering applying such techniques to sailing robots (which are most likely to gather their electrical power from photo-voltaic solar panels) the process of refuelling will be somewhat different. Firstly the robot is able to refuel while still performing its work and so the need to seek fuel is reduced. Secondly their ability to search for sunlight in order to obtain energy is limited given the slow speed of a sailing robot, although changes of course could be used to avoid shadows being cast by the sail upon the solar panels. Finally in the event of battery levels becoming low it may become necessary to stop the robot's work and refuel. Opportunities to perform extra work or improve sailing efficiency may also present themselves when more energy than was expected becomes available.

A hormone based mechanism was also adopted by Cañamero (1997) [33] in her work on emotional agents. In this each agent, known as an Abbott had a series of physiological variables including adrenaline, blood pressure, blood sugar, heart rate and pain. These (combined with external stimuli) acted as both stimuli and as variables to be modified from the Abbott's set of behaviours. For example, the eating behaviour was stimulated by the presence of food while the withdraw behaviour was stimulated by the presence of pain. Cañamero and Garcia (1997,2004) [33, 12] extended this work and implemented a system on

lab based robots. Their robots performed tasks based upon the robot's current motivations which resulted from its emotional state. Motivations for the robot were hunger (a lack of energy) and being cold, a winner takes all action selection mechanism would choose a behaviour based on the strongest motivation. This behaviour would in turn attempt to correct the original cause of the motivation by turning on a heater or attempting to find energy. They noted that the robot must be maintained within a zone of viability in order to continue operation, but that simply staying within the zone of viability is not optimal. Instead they stated that the robot should aim to achieve a quality of life by trying to maintain a position well within the viability zone rather than oscillating around its boundaries. They split the viable zone into two sub zones: an ideal zone in which the robot should aim to always be and a critical zone with lethal boundaries. This was calculated through life span and overall comfort which determined how the long the robot could be expected to survive and how comfortable it was with its current situation. A hormonal system was implemented to try and maintain the robot within its ideal zone by heightening the robot's response to a lack of food or heat.

As with McFarland and Spier's work, this took place in a highly artificial environment and questions can be asked about Cañamero and Garcia's observations regarding viability zones. In biological systems there is a (reasonably) clear cut off point that when crossed by the organism will result in death, for example, not eating for a prolonged period of time. When a certain threshold is crossed the organisms key systems shut down, it dies, cannot be revived and will begin to decay. Most robots can cope with prolonged periods of being powered down without energy without suffering irreparable damage. If we consider a robot which is powered by photo-voltaic solar panels and batteries it can have its batteries completely discharge, loose all of its onboard computing ability and still recover, providing the solar panels can recharge the batteries. Such a situation may not be ideal in a dynamic environment where the powered down robot has no ability to avoid dangerous situations (such as a collision with an obstacle) but is not instantly fatal to the robot either. Therefore although a control system should try and keep the robot within its viability zone it is not totally impossible that it might be permissible to allow the robot to essentially enter its death zone and there could be situations when this is actually preferable to an extended period in the critical zone.

Gadanho and Hallam (1997,1998,2001) [45, 46, 47] used an artificial hormone system to encode the emotions of a robot. They associated one hormone with each of the emotions the robot could experience. A dominant emotion would be selected based on the emotion with the highest intensity which also exceeded a selection threshold. Hormones could build up rapidly but only decay slowly, creating a mechanism which could trigger the persistence of a dominant emotion for some time. Their work linked the emotion system with a reinforcement

learning algorithm, this used the dominant emotion to help reinforce particular behaviours. They also used the dominant emotion to inhibit frequent state changes (which are particularly computationally expensive in reinforcement learning). This concept of emotional persistence could be useful in a number of robotic scenarios to ensure that a given behaviour has time to complete an action that it has begun or that behaviour switches or state transitions do not occur too frequently.

2.6.3.4 Sleeping Robots

The hormone melatonin, produced by the pineal gland is often associated with the diurnal sleep cycles in mammals [68]. This triggers a number of behavioural changes at night including slowing the metabolic rate and triggering feelings of tiredness. It would seem sensible for a robot operating for long periods of time to have some notion of sleeping, this would be particularly important in a solar powered robot as it is dependent on a diurnal solar cycle to perform its task and for its survival.

Mirolli and Parisi (2003) [82] devised a simulated environment populated by artificial organisms. These organisms must seek food during daylight in order to survive as they cannot see and therefore cannot move at night. The first type of organism has a light sensor and can sense when there is daylight. However, when placed into an environment where some organisms live in caves they cannot determine if there is daylight or not and so never leave their caves to search for food. A second type have internal biological clocks and are able to cope with this problem but are totally dependent on their internal clock being set correctly for the environment (in which day and night are both 12 hours long). A third type of organism combines both an internal biological clock and a light sensor, so it is able to adapt to changes in the length of day while still managing to cope with the cave problem.

Parisi (2004) [99] elaborated on this with work on internal robotics, the concept that robots interact with both the external environment around them and their own internal environment. He presented examples of this to include managing sleep, hunger/thirst and pain.

Rocks and Barnes (2004) [108] adopted an artificial circadian rhythm based on the entrainment of an internal oscillator from external sun light cues in order to maintain an internal biological clock. This system of entrainment would allow the robot to adapt to changes in the length of day caused by long distance travel or seasonal changes. The use of an internal biological clock also allows the robot to maintain its sense of day and night even if temporarily it loses sight of any daylight as in Mirolli and Parisi's [82] cave example. They applied this to a simulation of a Mars exploration rover and used it to help schedule tasks around both daily and seasonal solar cycles.

2.6.3.5 Digital Hormones in Modular Robotics

Shen, Choung and Will (2002,2004) [118, 119] took an endocrine inspired approach to communications in modular robotics which they called "digital hormones". Their robots are cube shaped and can connect to up to four other modules, each contains onboard power, computing, actuation, orientation sensing and the ability to communicate with its neighbours. These robots can join together to perform various types of locomotion including walking, moving like a snake or forming into a circle and rolling like a wheel. To facilitate this they required a decentralised and scalable communications architecture. They took on a mechanism inspired by the propagation of hormones through the body, each message has no addressing and will propagate from one unit to its neighbours. To prevent routing loops messages have finite lifespans. This architecture allows for a totally decentralised control system which can coordinate a variable and unknown number of robot modules together to perform locomotion. This is perhaps the only example in this literature review of a robot architecture with an actual physiology where (like in a biological system) the propagation path and speed of the hormone can have serious effects upon the robot's overall performance. However, as most robots do not have much, if any flexibility to their physiology in most non-modular scenarios this will be of less importance.

2.6.3.6 GasNets

Neuro-scientists have observed that the brain does not operate purely through electrical synaptic connections but is also reliant on neuro-modulation where neurons emit modulatory gases such as nitric oxide [48] which affect the behaviour of their neighbouring neurons. These gases diffuse slowly and only over relatively short distances. Husbands et al. (1998) [60] devised an artificial neural network known as a GasNet based upon this observation. In a GasNet the network is represented as two dimensional topology and connections were evolved through a genetic algorithm. Neurons are then able to modulate neighbouring neurons through the release of the modulatory gases, which slowly diffuse through the network. This makes the topology of the network very important to its behaviour and also allows neighbouring, but otherwise unconnected neurons to modulate each others behaviour. Experiments were conducted on a visually guided robot and showed that they could be more resilient to noise than traditional genetic algorithms or neural networks. It has been suggested that this is due to the interactions between the "electrical" connections of the neural network's synapses and the gaseous modulators.

2.6.3.7 Artificial Neuro-Endocrine Systems

Based on the observation that homeostasis is the product of the interaction of the nervous, endocrine and (in some cases) immune systems Neal, Timmis and Mendao (2003,2004,2007) [92, 76, 78] devised an artificial neuro-endocrine architecture which integrated together an artificial neural network and an artificial endocrine system. Unlike the Gasnet approach discussed in the previous section their approach was inspired entirely by hormonal modulation of electrical neural networks and was not dependent on network topology. Their system consisted of one or more artificial neural networks and one or more artificial glands which would secrete hormones. These hormones would modulate the neural networks through the modification of their weights. In biological systems hormones only bind with cells that poses target receptors for them, to mimic this in an artificial system they assigned each neural network a sensitivity level to each hormone. The level of modulation produced by a hormone was proportional to the neural networks sensitivity to it. In biological systems there are variable levels of matching between the target cell's receptors and the hormone. To mimic this property they introduced a second property called the "receptor matching distance" which measures the distance between each target cell and each hormone. Hormones are released from the gland into the "bloodstream" through a variable rate release mechanism and then decay according to a geometric decay function.

Neal and Timmis (2003) [92] tested an implementation of this artificial neuro-endocrine controller on a two wheeled Pioneer 2DX robot. This robot is driven by two electric motors and has 16 sonar range finders placed around it for detecting obstacles. They implemented an artificial neural network with an input from each sonar and an output to each motor. This network was trained to wander around an arena and turn away from any obstacles it encountered. A single gland was also connected to the sonar sensors and as the ranges from the sonars got smaller it would excite all the weights in the neural network. As the values from the sonars got larger the gland would inhibit the weights. This gave a behaviour which could be likened to adrenaline as the robot would speed up and rotate quickly giving the appearance of panicking when it became trapped in a corner or surrounded by obstacles. Once it became clear of obstacles the hormone would begin to decay and after a short time it would calm down and return to its normal behaviour. They compared using no hormone, a constant fixed level of hormone and a variable hormone regime. The variable hormone regime would typically be the fastest to negotiate obstacles and "run away" from them.

The mathematics behind their controller is shown here as it is relied upon in later chapters. They took the traditional formula for calculating the output of a perceptron, which takes the sum of one or more inputs which have each been multiplied by a weight:

$$\sum_{i=0}^{n} w_i x_i \tag{2.3}$$

Where w is the weight assigned to a connection, x is the input value and i is the input in question.

They then modified this formula to include modulation from a set of artificial hormones, the modified formula is shown below in equation 2.4.

$$\sum_{i=0}^{nx} w_i \cdot x_i \cdot \prod_{j=0}^{ng} C_j \cdot S_{ij} \cdot M_{ij} \tag{2.4}$$

These hormones are secreted by a set of artificial glands g. Each hormone specifies a concentration value C, this represents the quantity of the hormone that is available and is analogous to the concentration of the hormone in the bloodstream. Each neural network also has a sensitivity level S to each hormone, this determines the size of the effect the hormone has upon the neuron in question. In biological systems there are variable levels of matching between the target cell's receptors and the hormone. In the artificial system this is simulated with the variable M who's value is calculated in equation 2.5 (shown below).

$$M = \frac{1}{1 + dis(i,j)} \tag{2.5}$$

The rate of hormone release is governed by the formula shown below in equation 2.6. Where $\alpha_g$is the rate at which hormone is released for a specific gland (g), $x_i$is the input to the gland and n is the number of inputs to the gland.

$$r_g = \alpha_g \sum_{i=0}^{nx} x_i \tag{2.6}$$

In biological systems the concentration of a hormone in the blood stream will gradually decay over time. Of course, this is assuming that no additional hormone is secreted by the gland which produces it. This decay is due to hormone molecules binding with receptors upon target cells. However, in an artificial system no such physical process is present to remove the hormone from the bloodstream. Therefore, the blood steam hormone concentration must be artificially decayed over time. Neal and Timmis followed a simple geometric decay function shown in equation 2.7.

$$C(t+1)_g = C(t)_g \cdot \beta \tag{2.7}$$

Where C is the hormone concentration, g specifies the gland in question, $C(t)_g$ specifies

the bloodstream hormone concentration for a given point in time t and β specifies a constant decay rate for the hormone.

Henley (2006) [55] applied a neuro-endocrine controller to a legged robot attempting to negotiate its way over rough terrain. He used hormonal modulation to vary the gait and stride length of the legs to overcome difficult obstacles. By using a slow decaying hormone a memory effect was established and the robot would continue to operate with an adjusted stride and gait for sometime after it encountered the initial obstacle. Typically rough terrain would not be encountered for a single step but would extend for some distance and this allowed the robot to maintain the correct mode to cope with it and increased the overall speed at which it could be crossed.

Mendao (2006,2007) [77, 78] adopted a similar architecture and used a camera to have a robot seek either black or white objects. He used two hormones to decide which colour to seek and two modulated neural networks which detected and drove the robot towards coloured objects. By varying the hormone levels he would completely inhibit one of the neural networks and excite the other and thus control which colour the robot was seeking. Initially he fixed hormone levels to a sinusoidal oscillator which would gradually change the robot from seeking black to white and vice-versa. He found that his system would often loose momentum and stabilise upon an intermediate state in which neither colour was sought. He identified three factors causing this: the symmetrical nature of the environment, the lack of a topology through which hormones travel and the lack of pools in which hormones build up before they are secreted. He decided to implement a system of pools that would store hormones before they were released. In this model hormones are produced when a certain stimulus occurs, but stored in the pool until a threshold value is reached. Upon reaching this threshold the pool is emptied, releasing a surge of hormone to the receptors on the neural networks and triggering a sudden behavioural change. The hormone then decays at a geometric rate. This was sufficient to overcome the stagnation of the behaviour switching and cause the behaviour of the robot to oscillate between seeking the black and white objects.

Walker (2007) [141] used a set of pooling hormones to perform behaviour action selection on a pair of robots. Three behaviours were developed and switching between them was governed by a "winner takes all" approach in which the behaviour that had the highest hormone concentration would be selected. All hormones built up in pools in a similar fashion to those used by Mendao, only when the pool completely filled and released its contents could it trigger a behavioural change. This helped to ensure that task switching did not occur too frequently. The approach was extended to also coordinate behaviours across two co-operative robots.

Further work by Timmis, Thorniley and Neal (2009) [136] applied Hebbian learning to

adjust sensitivity levels to a hormone. This was applied to a robot attempting to navigate an environment with the help of a neural network for wandering and another for avoiding obstacles. Proximity and bump sensors provided stimulation to glands which were used to trigger the switch between wandering and avoiding. Through the use of Hebbian learning the robot was able to adjust hormone sensitivities so that it activated the avoidance neural network quickly enough to avoid collisions. Timmis, Murray and Neal (2010) [135] scaled up the work to cover a swarm of robots attempting to collect rubbish from their environment and keep their batteries charged. The robots implemented a number of behaviours including obstacle avoidance, rubbish seeking, power seeking, bin seeking and wandering all coordinated through a neural endocrine controller and using a set of hormones to perform behaviour switching. This goes at least some way to demonstrating the scalability of the neuro-endocrine approach to managing large sets of behaviours. They also noted the importance of being able to adapt across varying timescales from reactions lasting a few seconds to those lasting months. In biological systems the neural system is often seen as performing the shortest term sub-second response, the endocrine system multi-second to multi-month responses and the immune system for multi-minute to lifelong responses. It is this varied timescale adaptation that helps to create a robust ability to adapt to various problems that biological systems will experience during their lifetimes.

2.6.3.8 Hybrid Neuro-Endocrine and GasNet Controllers

Moioli, Vargas, Von Zuben and Husbands (2008) [85, 84] extended Neal, Timmis and Mendao's concept of an artificial neuro-endocrine controller by replacing the neural network element with a GasNet controller. Their GasNet also differed from those in previous work by Husbands et al. (1998) [60] as this was a Non-Spatial GasNet in which the network topology was not important and gases were free to diffuse across the entire network and modulate any neuron within it. They demonstrated switching between several tasks in the robot using an artificial endocrine system to control the behaviour switching. They also demonstrated an artificial homeostasis through the maintenance of a constant battery level, when the robot experienced low battery readings it would create a "desire to recharge" through the release of a hormone which would trigger the robot to go and search for energy (in this case a light source) until it was sufficiently recharged.

2.6.3.9 Summary of Artificial Endocrine and Homeostatic Systems

This section has reviewed the biological background to the endocrine system and existing work on artificial endocrine and homeostatic systems. These are capable of promoting and

suppressing other systems including artificial neural networks. A number of key advantages for a robotic system have been identified. The first and most general property is that robot behaviour can be adapted to match its environment, for example behaviour can be adjusted to reduce power consumption when battery levels are low. Secondly robots may take advantage of unforeseen opportunities which they may encounter in their environment, for example finding themselves with more power than expected. Third, that artificial endocrine and homeostatic systems can maintain a robot within a zone of viability, in this zone a number of parameters which are vital for continued operation of the robot are maintained within acceptable limits. Fourth, that it is possible to synchronise the behaviour of a robot with cyclically varying parameters in its external environment such as diurnal and annual solar cycles or tides. This allows the robot to adjust its behaviour with some prior knowledge of future events and allows this in a totally sub symbolic manner, without explicit reasoning regarding them. Fifth, they allow for decentralised communication across a number of subsystems within one robot and/or across multiple robots. Sixth, they allow for a variety of behaviours across differing timescales. Finally they allow for competing demands, each encoded as separate hormones, to act upon systems within a robot.

All of these properties are of potential use to an autonomous sailing robot which will benefit from modifying its behaviour to adapt to changing circumstances. A sailing robot could potentially use an artificial neuro-endocrine controller, with the neural network performing actuator control on a sub second basis and the endocrine controller provider longer term management of power and robot's overall health. The robot could modify its behaviour to manage power consumption and maintain its battery levels within a zone of viability. This behaviour could also be adjusted on a daily or seasonal basis in conjunction with knowledge of the expected power from a photo-voltaic solar panel. Should the robot encounter a day with higher than average sunlight levels then it maybe able to exploit this opportunity to improve its sailing performance or to switch on some additional equipment (for example perform optional communications or additional ocean monitoring tasks). An artificial endocrine controller could be used to achieve a decentralised communications mechanism between multiple processes or computers within the robot. This could potentially create a mechanism to coordinate tasks between redundant equipment in order to choose which piece of redundant equipment was used at any time and to switch between them in the event of failure or (under normal conditions) to see that each used an even amount of time to reduce wear. Finally adaptation across different timescales could be achieved with several hormones, each varying their release rate on a different timescale. By restricting the quantity of a hormone to only being able to change gradually over time, a form of noise filtering can also be created. Although, caution must be applied not to over restrict the rate of change as genuine state

changes may also be lost in the process.

2.6.4 Immune Systems

The immune system in mammals is often viewed as being two distinct components, the innate and the adaptive immune system. The innate system produces immediate and non-specific responses to infection, while the adaptive system learns how to fight specific infections and retains this knowledge for future use. Many organisms, such as invertebrates and insects do not have an adaptive immune system and rely solely on an innate immune system. Although often viewed as separate systems, the adaptive and innate system are highly interdependent and rely on a variety of signalling mechanisms to suppress and promote the other's activities.

The immune response to a pathogen (infection) essentially operates in two stages. The first stage is the identification of the pathogen and the second stage is its destruction and removal. Both the innate and adaptive immune system offer mechanisms which contribute to the identification, killing and removal of the pathogen. There is some debate as to how the immune system actually identifies pathogens. One theory is that it is able to distinguish foreign objects from those which are part of the body, a theory known as "self from non-self" [100]. A refinement on this, known as danger theory [71] states that the immune system will attack entities which do "damage" rather than those which are classed as being foreign. These theories are not completely mutually exclusive and the debate is far from settled. When selecting ideas to base a robot control system on, both theories have potential merits. This also brings up the question as to whether or not biologically inspired robotics should draw upon ideas in biology which are not completely confirmed, fully understood or which have even been found to be incorrect. It could be argued that if a particular algorithm performs well then we should make use of it, regardless of whether or not the underlying biology is correct. After all, we are not seeking to produce a perfect artificial simulation of biology (although other areas of computer science maybe interested in doing this for medical purposes) but merely to borrow key concepts and ideas which could be useful in controlling a robot.

2.6.4.1 The Innate Immune System

The innate immune system is believed to have evolved before the adaptive system and many simple life forms (e.g. plants and fungi) only have innate systems. The first layer of the innate immune system is often considered to be the physical barriers of the body, these include the skin, lining of lungs and stomach acid. It is these barriers that prevent most pathogens from entering or having any effect upon the body.

Before attempting to destroy a pathogen, the innate immune system must first identify it. Phagocytic (able to kill pathogens) white blood cells circulate around the bloodstream. One type of phagocytic cell known as macrophages, operate by recognising common pathogen characteristics and attempt to engulf (phagocytosis) the pathogen to kill it. One of the key mechanisms to achieve this is by detecting Pathogen Associated Molecular Patterns (PAMPs), molecular patterns commonly associated with pathogens but which are not present on any of the body's own cells [89]. These bind to Toll Like Receptors (TLRs) upon the surface of immune cells (including macrophages). Phagocytic cells also recruit other immune cells through a variety of inter-cellular signalling mechanisms. Key to these signalling mechanisms are cytokines, protein based messengers which are used throughout the immune system to trigger behavioural changes in target cells which have appropriate receptors. In many ways they are not dissimilar to hormones, although they can act upon the cell releasing them (autocrine action), they can diffuse through extra cellular fluid to act upon neighbouring cells (paracrine action) or they can even be absorbed into the bloodstream and transported to other parts of the body (endocrine action). As with hormones these messengers operate through a cascade mechanism, in which one messenger triggers the release of another before any action is taken.

Another type of phagocyte is the neutrophil, these are recruited by macrophages through the release of chemical messengers such as cytokines and interleukins. Neutrophil's take around 6 hours to begin operation, while macrophages are immediately on hand to deal with infection [89] and only recruit neutrophil's when they are unable to deal with an infection. Some cytokines are also toxic to many pathogens and help to trigger the inflammatory response. This response also increases blood flow to the affected area by dilating the capillaries. Which makes their walls more porous, allowing more white blood cells to reach the area. This triggers pain signals to the neural system and increases the temperature, which can help kill pathogens. Macrophages which have engulfed a pathogen, present receptors on their surface which help to recruit components of the adaptive immune system.

2.6.4.2 The Adaptive Immune System

The adaptive immune system operates through two key types of specialised white blood cells known as lymphocytes. These move around the body in the blood and and lymph glands. Unlike macrophages in the innate immune system, they are specific to particular pathogens and must learn to deal with any new pathogen they encounter. There are two main classes of lymphocytes known as B and T-cells. As with much of the immune system, the operation of these cells is not fully understood. Many of the studies relating to them come from HIV and AIDS research where the goal is find out why HIV/AIDS causes them to malfunction.

As a result their "normal" operation is not as heavily researched and not fully understood.

B-cells produce proteins known as antibodies which bind to proteins on the surface of a specific pathogen, these surface proteins are known as antigens. The antibody has two roles, firstly it can inhibit some processes within the pathogen which prevents it from dividing and producing more pathogen cells and secondly it highlights or tags (opsonization) the pathogen for destruction by macrophages. In order to reach the point at which a B-cell can produce antibodies that bind to a specific antigen, the B-cells must first undergo a process which tunes this affinity. This process is known as clonal selection and see's B-cells effectively evolve to deal with a specific pathogen. This occurs through mutations of the genetic makeup of B-cells during the process of cloning through cell division. B-cells which improve their binding affinity (ability to bind) will tend to proliferate faster.

As concentrations of B-cells are relatively low, when a B-cell binds to an antigen it begins to divide and clone itself to increase the numbers of B-cells available to deal with the infection. Some of the B-cells with the highest binding affinities will be left over at the end of an infection and will be retained should the immune system have to deal with that particular pathogen again in future. This is the basis of immune memory and retains long-lived immunity to particular pathogens.

There are various types of T-cell, including: killer, helper and regulatory. Killer T-cells attack viruses which have parasitically invaded a host cell (and are now beyond the reach of antibodies). They do this by binding to the infected cell and injecting proteins which trigger the host cell to commit suicide (apoptosis). Helper T-cells aid macrophages to kill cells they have phagocytosed but cannot kill and help B-cells to produce antibodies. Regulatory T-cells help to end the immune response and prevent the immune system from misidentifying the body as a target for attack. As with B-cells, T-cells are antigen specific and undergo a clonal selection process to produce sufficient numbers to fight an infection and some of these will eventually contribute to immune memory functions. Due to the relatively low concentrations of adaptive lymphocyte cells they can take some period of time (a few days) to activate and often only do so in response to signals from the innate immune system.

2.6.4.3 Neural, Endocrine and Immune Interaction

The neural, endocrine and immune systems do no exist in isolation but interact heavily with each other and can almost be viewed as a single super system [17]. They frequently share the same chemical messengers, for example, epinephrine is both a hormone and a neurotransmitter. The hormones of the thymus are responsible for the regulation of the production of immune cells. While cytokines play a signalling role in both the immune and endocrine system. There is evidence to suggest that these systems actually co-evolved [43]

and that they are highly interdependent. There is also evidence that these systems not only work together but that at times they can also work in competition [101], for example the neuro-endocrine system can limit certain immune reactions which are so powerful they will actually kill the host. The neural, endocrine and immune system each contribute to homeostasis on a different time scale, the neural system being responsible for the sub-second to multi-second scale, the endocrine system the multi-second to multi-day scale and the immune system the multi-minute to lifetime scale.

2.6.4.4 Artificial Immune Systems

The immune system has been identified as having many properties which could be useful for artificial control systems being applied to robotics. The concept of producing an artificial analogue of the immune system began with Forrest, Perelson, Allen and Cherukuri (1994) [42] who decided to try and apply the principle of distinguishing self from non-self in order to combat computer viruses. De Castro and Von Zuben (2002) [35] created a clonal selection algorithm based on how B and T cells are selected and cloned in response to an infection. This creates a random set of antibodies and evolves them to recognise a set of patterns. Those which evolve the strongest affinity to the pattern are cloned and saved for future use. This creates a form of immune memory and allows for the detection of a repetition of a previously observed scenario. Within robotics this algorithm has been applied to robot learning and fault detection and recovery. Singh and Nair (2005)[122] used the clonal selection algorithm to help a robot learn and recover from fault conditions. In their scenario two robots travelled in concentric circles. The outer robot was able to help the inner robot in the event that it became misaligned and lost its track. The inner robot used the clonal selection algorithm to record the circumstances in which it had become misaligned in order to avoid repeating this mistake in future. In this Singh and Nair identified three key stages to error handling, the first is error processing which identifies an error has occurred, the second is fault recovery which recovers the robot to a non-fault state and the third is fault treatment which prevents the conditions causing the fault from being repeated.

Neal and Timmis (2005) [93] proposed that homeostasis is the result of interactions between the neural, endocrine and immune systems and that in order to achieve artificial homeostasis all three must be combined. They outline artificial versions of each system and mechanisms for them to interact. Their neural and endocrine systems incorporate the architecture seen in their previous work [92] discussed in section 2.6.3.7. However, in this work they incorporated a mechanism for adding new neurons into the neural network. The adaptive immune system recognises these as being "non-self" and attempts to kill them but is regulated by the artificial endocrine system which produces a hormone that suppresses this

killing process when the performance of the network improves. Thus the system is able to reinforce any new successful behaviours which may emerge.

Neal et al. (2006) [91] devised an artificial innate immune system for responding to damage in a robot. In the innate immune system as series of Toll-Like Receptors (TLRs) respond to certain Pathogen Associated Molecular Patterns (PAMPs) caused by a certain pathogen. The TLRs will bind to the pathogen and trigger the activation of other immune system reactions such as the inflammatory response or triggering feelings of sleepiness. Within the robotic system the TLRs each respond to a single set of sensors such as the temperature of a particular motor. The PAMPs represent abnormal conditions for that sensor. The TLRs are combined to form a set of vectors which are used to generate a self organising map which is able to represent the overall state of the robot. An artificial neuro-endocrine controller will then, depending on the state of the self organising map, release hormones to activate remedial behaviours, this closes the loop between the detection and remedial action stages. In the example of an overheating motor this could either be to reduce the duty cycle of that motor or turn on a fan to cool it down. This architecture has not been implemented in a real system but offers the potential to represent the entire health state of a robot through a single self organising map. It could be of particular use in larger more complex robotic systems with many proprioceptive sensors where obtaining an overall view of the robot's health status would otherwise be difficult.

Humza et al. (2009) [59] implemented a robotic organism made up of a series of interconnected modules. Each module contains its own battery but is also connected to a shared power bus which interconnects it to other modules. An artificial immune system was responsible for maintaining both the health of individual components and the entire organism, it was able to monitor typical battery levels from each robot module and obtain a picture of normal homeostatic behaviour of the organism. This allows the generation of a single "health measure" for the entire swarm. Faults within the organism as whole can be detected as changes in the health measure occur. This could potentially be used to take remedial action to maintain homeostasis within the organism as a whole.

Whitbrook, Aickelin and Garibaldi (2007) [143] developed an approach based on the Idiotypic Immune Network Theory [107], which states that antibodies recognise each other and suppress and promote each other's production. They combined the Idiotypic network with a reinforcement learning algorithm and compared it with a standard reinforcement learning algorithm in solving a series of robot navigation problems. In these a robot was required to navigate its way through a maze containing a series of obstacles towards a goal point. The idiotypic controller showed a distinct improvement over the pure reinforcement learning, in particular it was noted that the idiotypic network helped to overcome repeated

error patterns and generate new behaviours. In further work, Whitbrook, Aickelin and Garibaldi (2010) [144] developed a two time scale approach to robot navigation. This involved a short term learning (STL) strategy based upon an idiotypic controller and a long term learning (LTL) strategy based upon a genetic algorithm. The genetic algorithm evolved a set of behaviours in simulation and these were then used to seed the initial behaviour of the artificial immune system. They used an e-puck robot [86], a small (7.5cm diameter) desktop robot to navigate through a set of obstacles to a target point both in real life and simulation. Bump sensors and infrared proximity sensors were used to detect obstacles and an RGB camera to detect the target. They specified antigens to the AIS as detecting or colliding with obstacles and gaining or loosing sight of the target. These were responded to with antibodies which activated wandering, moving forward, turning, reverse turning or target tracking behaviours. The LTL system was able to generate a set of initial antibodies within 10 minutes for a static environment and 25 minutes for a dynamic one, while the STL system was able to provide online adaptation. This two time scale strategy could be compared with the three timescale strategy proposed by Neal and Timmis [93], in both cases there is the ability of the adaptation at each timescale to affect the other timescales. Whitbrook, Aickelin and Garibaldi's strategy also offers the advantage of an offline pre-learning phase in simulation, this creates an adequately bootstrapped system to begin with rather than leaving the robot to learn basic operations in a real world environment. However, this adds the requirement of an adequately realistic simulation to perform the bootstrapping with. The "real" environment in this situation was also quite simplistic, consisting of a series of rectangular and circular blocks placed on a flat surface. It is not entirely clear if these results would scale to more complex worlds such as those encountered by a sailing robot.

In this section we have reviewed several attempts to control robots with artificial immune systems. Artificial immune systems provide a potential mechanism to create behaviours which adapt to changing circumstances over long time periods. Within the context of sailing robots this could be applied to coping with the failure or degradation of components. Additionally, as proposed by Neal and Timmis [93] an artificial immune system could be combined with an artificial neuro-endocrine system to cover multiple timescales. This could see a neural network providing second by second control of actuator movements, an artificial endocrine system providing longer term responses to fluctuations in power level or to component degradation, while an artificial immune system could provide fault tolerance and recovery functions in response to failed or severely degraded components.

2.7 Approaches to Power management in robotics

A number of the biologically inspired methods discussed in the previous section can be applied to solving power management problems in robotics. However, a number of non-biologically inspired approaches to autonomous power management have also been taken. This section provides a brief overview of some of these methods and the results they have achieved.

Mei et al. (2005) [75] analysed the power consumption of a Pioneer 3DX robot. They found that only 50% of the power consumption was used by the motors and that the rest was used by the onboard computers and sensors. Therefore they focused on reducing the power consumption of the computers rather than the motors. Through the use of dynamic power management techniques that are often used in portable computers and mobile phones [15], CPU voltages and clock frequencies were altered according to demands on the processor. They also used a real-time scheduling system to switch off unused components and guarantee that they would be switched back on again before they were required. One of the components controlled by this mechanism was the network card, switching this off for prolonged time periods will cause any established TCP (and many other stateful protocols) connection to timeout. This creates a series of tradeoffs between power consumption and the quality of communications available. This approach might also be made less relevant in robots where actuators consume considerably more power or where computers are more efficient. As the Pioneer 3DX uses a full PC motherboard with computing power that might be considered excessive for many potential applications these savings might also be achieved simply through using more power efficient computers. However, in applications where bursts of intense computation are required they may have potential use.

Sun and Reif (2003) [130] proposed a path planning algorithm which would plan the most energy efficient path through an open environment (assuming no roads, pathways etc) based upon a digital elevation model (DEM) of that environment. This model was able to take into account the energy lost through overcoming gravity and surface friction and would plan routes intended for a car sized robot. This approach is reliant on having an accurate DEM of the area and it does not take into account variations in terrain type such as mud, loose gravel or water run off. DEMs are often derived from relatively low resolution satellite data and will consider the terrain height of forests to be at the tree tops and may therefore may not give totally accurate models of the ground shape, they might also assume that a large vehicle can drive across the top of a forest. However, presented with data of sufficient quality such an approach should be able to allow a robot to select the most energy efficient route. The data could also be restricted to only include valid roads and paths.

Ngo and Schioler (2006) [95] produced a series of robots which were capable of sharing

energy from one to another through the transfer of electrical energy via terminals located on the outside of the robots. Their approach is broadly inspired by the mechanism of food distribution in ant or bee colonies and is also somewhat similar to the approach taken by Humza et al. [59] in the previous section. This approach allows for energy to be redistributed across a group of robots and for robots who's batteries are too flat to reach a recharging station to be "rescued" by others that are. They also suggest a heterogeneous approach in which a few robots amongst a group are equipped with energy harvesting abilities such as solar cells or wind turbines and these are responsible for distributing power to the rest of the group.

Ray et al. (2007) [105, 67] constructed a polar exploration rover intended to operate for at least 2 weeks while driving across Antarctica, although to date this has only been tested in Greenland. Their robot was cube shaped and powered by 5 photo-voltaic solar panels (one on each exposed surface). The intended operation was during the Antarctic summer when near perpetual daylight is available. A maximum power point tracking system optimises the voltage and current output of the solar panel in order to give an optimal power output. The effects of constant variations in the solar panel angle while the robot is moving create additional challenges for optimising the maximum power point tracking. The robot aims to keep its batteries at approximately 80% charged and with a zero net current flow through the battery. The robot is able to vary speed in order to help maintain this and has a series of predefined power consumption modes. These include a stationary mode to be used while the robot is stopped and performing scientific measurements, an emergency mode which is used to navigate without range finders or cameras and a safe mode which is used when the wind is too strong for the robot to drive. In early tests the power control mechanism is not able to vary the speed of the robot to manage power consumption, although it can request the robot stops to conserve power although this was planned for future use.

Shillcutt (2000) [121] proposed a solar synchronous vehicle for polar exploration. This circumnavigates near the pole during summer time, by choosing a latitude that is sufficiently close to the pole and maintaining constant motion the vehicle is able to track the sun and essentially always experience a mid day sun. Shillcutt targeted an antarctic terrestrial rover and found that it would be required to avoid shadows cast by terrain in order to continue a solar synchronous journey. She therefore created a path planning system to calculate shadows and predefine a path (and the allowable deviation from that path) to drive which would ensure continuous solar coverage. She also devised an algorithm to allow a robot which had entered the shadows to find its way back to the closest spot in sunlight in order to recharge. Such an approach could also be applied on other planets or to airborne robots. Although it is primarily suitable for summer time polar operation such techniques could be

applied to give robots operating at lower latitudes knowledge of when and where they can expect to obtain solar energy. Potentially this could be of use to any solar powered robot attempting to optimise its charging regardless of location although a reasonably accurate map of local terrain is required in order to estimate shadows.

Klesh and Kabamba (2009) [65] present a design for a solar powered unmanned aerial vehicle (UAV) intended for perpetual operation. They devised a path planning and optimisation system which considered the aircraft kinematics, energy losses and energy gains of the aircraft. This bears many similarities to Shillcutt's terrestrial path solar planning in optimising the path according to solar constraints. What neither of these models are able to consider is the effects of the weather and cloud cover upon the performance of solar panels. Both instead assume clear skies, which while simpler is likely to overestimate the available energy. However, determining future availability of solar energy requires linking to weather forecasts and accuracy will still be severely limited. Both mechanisms although not biologically inspired could be compared with homeostatic mechanisms in that they must maintain the system within the boundaries of a certain steady state otherwise they risk sub optimal performance, damage to the robot or even the total loss of the robot.

2.8 Power Management in Sensor Networks

Sensor network systems are increasingly being used for environmental monitoring applications. These require that sensors be distributed across an environment and communicate their data via wireless networks. Such sensors are often placed in difficult to access locations away from mains electricity and must rely solely on batteries and/or their environment for any energy needs. Therefore these sensor networks face many of the same problems faced by robots operating in isolated environments away from human operators. They must tradeoff between maximising the number of readings that are taken and transmitted and minimising power consumption in order to maintain regular transmission intervals. Much in the same way a robot might optimise between minimising power consumption and moving in the correct direction and at the desired speed. However, unlike robots they are unlikely to face any danger from being completely switched off. Depending on the application requirements they may also not have a minimum transmission interval. While a robot will have a series of minimal requirements for motor speeds or actuator duty cycles in order to maintain movement.

Kansal (2004) [63] developed a series of small sensor nodes known as Helimotes. These had a small 60 milliamp (peak) solar panel and 1800 milliamp hour battery both operating at nominal voltage of 3.3 volts. By varying the duty cycle of the sensor sampling and applying

dynamic voltage scaling to the microprocessor he was able to completely control the power consumption and achieve what he described as "eternally sustainable operation".

Jiang, Polastr and Culler (2005) [61] developed a similar sensor node with a dual power storage system consisting of a super capacitor and rechargeable battery. Their system was able to switch between the power sources depending on energy availability. This allowed them to maximise battery lifetime by reducing the number of charge cycles the battery experienced. They used a solar cell to maintain the charge of the battery and super capacitor. Like Kansal they also adjusted the duty cycle of transmissions and readings to manage overall power consumption.

The work from the sensor networks field suggests similar approaches might be attempted in robotics although as previously discussed there are limits beyond which robots cannot alter power consumption. To date there does not appear to be any attempt to implement any kind of pre-emptive behaviour which predicts diurnal or seasonal trends in solar power to adjust behaviour to give more equal transmission rates, instead batteries are simply used to buffer out some level of variation.

2.9 Chapter Summary

In this chapter a number of robot control architectures have been discussed. We have explored deliberative approaches, reactive approaches, hybrid deliberative-reactive, some of the early cybernetic approaches, neural networks, biologically inspired control systems and approaches to power management in robotics. Within the context of biologically inspired approaches, we have focused on the process of homeostasis, the maintenance of a stable internal state within an organism. This is in part achieved through the neural, endocrine and immune systems, with each system covering a separate timescale. The neural system operates on a sub-second to multi-second scale, the endocrine system on a multi-second to multi-day scale and the immune system on a multi-minute to lifetime scale. We have reviewed strategies to produce artificial analogues for each of these systems and attempts to produce integrated artificial neuro-endocrine and neuro-endocrine-immune systems.

These integrated strategies show promising potential to produce a control system capable of ensuring the continued long term operation of an autonomous robot. In particular using a neuro-endocrine controller, the sub second control of the robot could be achieved with a set of neural networks. Within the context of a sailing robot these would be responsible for setting the sail with respect to the wind direction and setting the rudder with respect to the current and target heading. These neural networks could then be modulated by a series of artificial hormones generated by an artificial endocrine controller. The hormones would be able to

modulate neural networks to reduce or increase the level of actuator use and thus control the amount of power being used. The hormones could also be used to coordinate between redundant motor controllers or redundant actuators to select which motor controller/actuator to use. Hormones could be produced in response to battery levels, sunlight levels or motor controller/actuator temperature. This would allow the robot to maintain a stable state with respect to battery level and actuator temperature and to also coordinate its behaviour with solar cycles. It will also create a control system which is capable of gradual adjustment of its behaviour in response to changes in the environment. This could be exploited over a long timescale to gradually adjust behaviour to slowly degrading components, for example as battery capacity reduces over its lifetime or as actuators wear and increase their power consumption.

By creating a control system which allows a sailing robot to autonomously manage its power, we can improve its overall autonomy and life expectancy. Ultimately, this may create the opportunity to reduce the level of over-engineering required and fit smaller batteries or remove backup power systems reducing the overall cost and complexity of the robot. To date, the deployment of artificial neural endocrine controllers into real world robotic systems has been very limited and it remains to be seen if such techniques will actually scale to be of genuine use in a real situation, such as managing sailing robot power consumption.

Chapter 3

Review of Sailing Robots

3.1 Introduction

Sailing robots are a form of Autonomous Surface Craft (ASC) capable of propelling themselves from the wind and autonomously making decisions to adjust their sails and rudders in order to reach a desired geographical position. A typical sailing robot includes at least one sail and one rudder that can be controlled through some form of actuator, a compass to sense its direction, a wind sensor to sense wind direction and a GPS to sense its geographic position. Some more sophisticated designs might include some form of tilt sensor, water speed sensors, range finding devices such as RADARs , AIS (Automatic Identification Systems - a transponder system for ships), and proprioceptive sensors such as battery state, internal humidity or actuator temperature sensors.

A number of applications are envisaged for sailing robots, these include sensing various ocean surface parameters such as water temperature, salinity, dissolved carbon dioxide levels, chlorophyll levels and turbidity. More ambitious applications include towing sonar arrays, deploying hydrophones to listen for marine life, surveillance of other vessels or of coastal features. At the time of writing there are only about 20 ongoing developments of sailing robots (including companies, universities, schools and private individuals). All of these can essentially be regarded as prototypes or research projects, although the Harbor Wing boat[1] appears to be close to commercialisation. Despite this apparent infancy of development the concept of sailing robots has been in existence since at least 1969 when Ernest Schleiben devised SKAMP (Station Keeping and Mobile Platform) [124]. The concept seems to have again been described in the 1980s with the proposal of the "Zero Handed Transatlantic Race" [36] and again in the late 1990s when Abril, Calvo and Salmon [1] developed a fuzzy logic controller for a modified 1 metre long radio controlled boat.

[1] `http://www.harborwingtech.com/` accessed 8/3/2011

Since 2005 the popularity of sailing robots has grown considerably and this appears to have mainly been driven through the creation of several sailing robot competitions and the publicity they have generated. In 2005 the Microtransat Challenge[2] [21] was created by Yves Briere and Mark Neal with the aim of crossing the Atlantic ocean with a sailing robot and to generally stimulate the development of autonomous sailing robots. Initially boats were restricted to being from academic organisations, under 2 metres in length, 40kg in weight and 60,000 euros in cost, eventually these restrictions were lifted to broaden the range of entrants and the only restriction to remain was that the boat was under 4 metres in length. Two warm-up events were held in 2006 and 2007. The first of these was in Toulouse, France in June 2006 and second in Aberystwyth, Wales in September 2007. The original intention was to cross the Atlantic from Portugal in 2008. Unfortunately the Portuguese authorities objected and the race was cancelled, it was then rescheduled to start from Ireland in 2009 until all but one team pulled out causing another cancellation. The first competition finally started from Ireland in September 2010. Only one team (from Aberystwyth University) competed. A second transatlantic competition is planned for September 2011.

A spin off competition known as the World Robotic Sailing Championships (WRSC)[3] began in 2008, this aimed to race sailing robots in a variety of long and short races over the course of several days (which is effectively what the 2006 and 2007 Microtransat events had done) with a focus on autonomy. The first WRSC was held in Austria in May 2008, the second in Portugal in July 2009 and the third in Canada in June 2010. At almost the same time the Microtransat competition began the Sailbot competition was created in Canada [4] with the (slightly) less ambitious target of racing autonomous and semi-autonomous boats no longer than 2 metres in length and constructed by teams consisting mostly of undergraduate students. Sailbot has been repeated every year since (except 2007) and in 2010 was hosted alongside the WRSC by Queens University in Kingston, Ontario, Canada.

The rest of this chapter examines the details of many of the sailing robots that have been developed and in particular examines their control systems and approaches to power management.

3.2 The Atlantis Project and Harbor Wing

The Atlantis Project at Stanford University between 1998 and 2001 reused an R-19 catamaran and replaced the mast/sail assembly with a wing sail. Unlike all the other boats described

[2] `http://www.microtransat.org` accessed 8/3/2011
[3] `http://www.roboticsailing.org` accessed 8/3/2011
[4] `http://www.qmast.ca/competition/index.htm` accessed 8/3/2011

here the wing sail was controlled by a small rudder attached to the back of the sail, by adjusting the position of the rudder the wind would cause the entire wing to rotate. This allows a relatively small, low power and simple actuator to control the large wing sail and does not place excessive strain on that actuator when compared to the strains encountered by an actuator which directly controls the wing position (as is used in all other wing sail boats discussed in this chapter). This system is similar to the Walker Wing Sail which was developed in the 1980s [140] and was used to equip pleasure boats with wing sails.

Elkaim (2004) [37] states that sensors are sampled at 100hz, the computer response rate is 5hz and that the boat was capable of sailing a course within 0.3 metres of accuracy through its use of differential GPS techniques. Since 2006 further development has been undertaken by Harbor Wing technologies and their "HWT X-1" boat, a modified 30 ft (9.1 metre) long Stiletto catamaran [38] with a 10.7 metre high wing sail. This is by far the largest boat reviewed in this chapter and the only one built after 1997 that was not built by a team entering into the Microtransat or Sailbot competitions. The total power consumption of the electronics is stated as being an average of 120 Watts [38] with actuator movements requiring several kilowatts, this places the X-1 in a very different league to any other boat reviewed here as many of the others have average power budgets of under 10 watts. Despite this large power budget, with such a large hull there should be no problem placing enough solar panels to generate 120 watts. The control system consists of a set of PI controllers which attempt to reduce both heading error and cross track errors (the distance from the line between the last waypoint and the next waypoint), this system is claimed to be capable of sailing a course to sub-metre accuracy [37].

The use of a multihull helps to improve stability in calm waters and doubles the theoretical maximum hull speed, however multihulls are notorious for being more stable upside down and are incredibly difficult to self right so this design might not be best suited to heavy seas where there is a greater probability of capsizing. Despite this risk the use of a catamaran provides extra space in which to place equipment and allows the wing sail to be mounted between the hulls allowing a larger wing than would be possible in a monohull. Utilising a rudder mechanism to maintain the wing sail position with respect to the wind should consume dramatically less power than a direct drive mechanism. It may be possible to extend similar techniques to the rudder through the use of wind based self steering gear [139].

This boat represents a design with the potential to sail for long periods without assistance, its size allows space for a large scientific payload and/or for electricity generating equipment. The size also presents a challenge with regards to collision avoidance as a boat of this size could easily inflict lethal damage upon a manned vessel should a collision occur, this may render it impractical for coastal use without further development of autonomous collision

avoidance systems or constant (but perhaps remote) observation by human operators. Although theoretically capable of long duration missions, there is no evidence in the literature that, to date, any of the Harbor Wing robots have performed missions lasting in excess of 24 hours.

3.3 IBoat

Figure 3.1: IBoat at the 2006 Microtransat on Lake Saint Nicholas de la Grave in France.

Since 2005 at least 8 academic institutions have begun developing sailing robots, many of them with the aim of competing in the Microtransat. The first of these was "IBoat" developed since 2005 by Briere et al. [20] at Institut Supérieur de l'Aéronautique et de l'Escape (ISAE) . This boat uses a custom made 2.4 metre long hull and weighs only 40kg, it originally featured both a mainsail and jib which were attached along their foot (bottom of the sail) to a boom which could be positioned by an actuator underneath it. Figure 3.1 shows "IBoat" in 2006 when this configuration was in use, two holes can be seen near the top of the sail, these were cut to reduce sail area against high winds. This sail configuration was later abandoned in favour of a single main sail configuration as shown in the photograph of IBoat in 2009 in figure 3.2. The control system of IBoat was originally implemented as a simple state machine

Figure 3.2: IBoat in 2009, the boat has been upgraded with solar panels, an ultrasonic wind sensor and a smaller sail.

and then converted to a neural network which was trained to copy the state machine [20]. Later a fuzzy logic control system was developed [22]. Average power consumption is stated as being 7 Watts [20], with lead acid batteries offering 14 amp hours (the voltage is not stated) and solar panels capable of generating only 20 watts peak. Briere (2007) [20] admits that these solar panels are insufficient for sustained operation and that a solar panel with 90 watts peak is proposed for a future version of the boat and that this should provide an average of 10 watts which would be sufficient for sustained operation. No mention is made of any autonomous controls over power consumption.

3.4 Robbe Atlantis and ASV Roboat

Since 2006 the Österreichische Gesellschaft für innovative Computerwissenschaften (Austrian Association for Innovative Computer Science or INNOC) have built two sailing robots. The first named "Robbe Atlantis" was adapted from a 1.38 metre long model sailing boat [126]. The second named "ASV Roboat" is based on a 3.75m long Learling dinghy [128]. Both used a traditional Bermudan rig and a Mini-ITX computer running Linux. Robbe Atlantis used only rotary wind sensors while the ASV Roboat introduced an ultrasonic wind sensor

Figure 3.3: ASV Roboat at the 2008 World Robotic Sailing Championships in Breitenbrunn, Austria .

as well as solar panels and later a methanol fuel cell. The control system of both boats is based around a multi-layered fuzzy logic system [129] which can deal with multiple levels of routing ranging from local obstacle avoidance and course following to more strategic and longer term decisions such as weather routing. The power system can deliver 285 Watt (peak) of solar power and an additional 65 watts from the fuel cell which is treated as a backup system should sufficient sunlight not be available. Stelzer and Jafarmadar (2009) [128] claim that the fuel cell is capable of running the robot for up to 4 weeks and that average power consumption is in the range of 30 watts. As they are currently developing a new boat capable of observing marine mammals which will require additional power to run scientific equipment they hope to save power through efficiency improvements to the electrical components, he control algorithms and through the use of a balanced rig design. A balanced rig will reduce the power demands on the sail actuator in comparison to the current design. Despite the identification of this need no details of changes to the control algorithms or electrical systems have yet been disclosed.

Figure 3.4: FASt during the 2008 World Robotic Sailing Championships in Breitenbrunn, Austria.

3.5 FASt

In 2007 the Faculdade de Engenharia da Universidade do Porto (Engineering department of the University of Porto or FEUP) began working on a sailing robot known as "FASt" (FEUP Autonomous Sailboat) intended for transatlantic racing and performing oceanographic missions [3, 2]. Their design was based upon an Open 60 hull, but scaled down to be only 2.5m long and weigh only 50kg. Computer control is based around an FPGA system running the Linux operating system. The choice of an FPGA allowed for reconfigurable hardware, which matched the exact specifications required by the team. The wind sensor is a rotary system but is able to isolate the electronics from immersion in water through the use of a magnetic position encoder [3] with the magnet placed on the wind vane and the electronics encased inside epoxy resin. Power is provided by a 45 Watt (peak) photo-voltaic solar panel and 2x 95 Watt-Hour Lithium Ion batteries [3], 50% of the power budget is used by computers and sensors (possibly in part due to using an FPGA) but further optimisation is possible. This system appears to present a potentially viable platform for continuous ocean sailing.

3.6 The University of Lübeck's MicroMagic boats

In 2009 the Department of Computer Science at the University of Lübeck began building small (53cm long) sailing robots [28] based around the Graupner MicroMagic radio controlled boat. These boats are typically raced by radio control enthusiasts on inland waters. However, the team from Lübeck adapted them to sail autonomously with the addition of a PDA, microcontroller, GPS, compass, tilt sensor and wind sensor. These boats are highly agile and required a fast control system with a response rate of over 20hz. In order to sail accurately a high frequency GPS was added [4], this enabled the possibility of receiving differential GPS corrections over the Internet in order to achieve an accuracy of less than 0.1m (compared to a typical L1 consumer grade GPS accuracy of 5m). Low power and relatively long range communications (approximately 1km) are provided by a class 1 bluetooth transceiver. For the 2010 WRSC and Sailbot competitions a larger 1 metre boat was constructed [5], this is based on the IOM (International One Metre) radio control boat hull. A lithium polymer battery with a capacity of 2.1 Amp hours supplies power at 7.4 volts and this is claimed to be sufficient for 3 hours of sailing [28]. Such small and high performance hulls require very high frequency actuator movements to maintain accurate sailing and it would be difficult if not impossible to place enough solar panels on the boat to generate enough solar power to sail indefinitely. Despite this such boats provide an ideal low cost and simple test bed for developing control algorithms for sailing robots.

3.7 Florida Atlantic University

In 2007 a team from Florida Atlantic University designed and constructed a 4.2m long sailing robot with a single wing sail [114], it was primarily intended to perform oceanographic monitoring and included several oceanographic sensors. Power efficiency does not yet seem to be a major concern with the average power consumption quoted at between 50 and 100 watts. Rynne (2009) [114] also discusses the idea of generating electrical power from either a wind turbine or a submerged water turbine. He suggests that in light winds the excess drag created would create a heavy penalty for sailing performance but in high winds when the boat is approaching its maximum hull speed the trade off might be more reasonable. He calculates that a 20% efficient turbine on a boat moving at 3.9 knots could generate approximately 500W. This should be sufficient for even the most inefficient control system and plenty of oceanographic equipment. Unfortunately it would only be feasible when high wind strengths could be guaranteed. The long term reliability of such a system is also questionable and Rynne does suggest solar power would be more feasible.

Figure 3.5: Luce Canon at the 2010 Sailbot and WRSC Competition in Kingston, Ontario, Canada.

3.8 The Boat's of the United States Naval Academy

Since 2008 the United States Naval Academy has been building boats initially aimed at the sailbot competition [80, 79] but with a wider aim of building ocean going boats capable of crossing the Atlantic or performing long term missions. Given the nature of the US Naval Academy much of the focus to date has been around hull and rig design rather than the design of the control systems (although fully autonomous control systems have been developed). Their second boat known as "Luce Canon" had two separate power systems to supply various voltages required by different components. Two 6 volt 1100mA C cell NiMH battery supplied the main control system and the rudder, while a 22 volt Lithium Polymer battery ran the sail winch. Their latest boat "Gill the Boat" [79] still uses two separate batteries but capacity has been increased to 16.8 amp hours at 6 volts for the processor battery and 18 amp hours at 6 volts for the actuators. A 150 milliamp solar panel can recharge the actuator battery but no comment is made as to whether or not this is sufficient to sustain the batteries indefinitely and with no solar panel powering the processor and electronics their separate battery will eventually become discharged.

Figure 3.6: Breizh Spirit at the 2009 WRSC in Portugal.

3.9 Breizh Spirit

École Nationale Supérieure de Techniques avancées Bretange (ENSTA Bretagne) in Brest, France have built a 1.2m long sailing robot called "Breizh Spirit" [123]. Despite its small size it is intended to undertake a transatlantic crossing and has already made several relatively long sea voyages[5]. The control system is implemented with a PIC microcontroller, wind sensing is provided by a CV7 ultrasonic sensor, radio control servo motors control the rudder and sail positions, a 5hz EB-85A GPS and HMC6352 compass provide navigation. Power is supplied by lithium polymer batteries and it is noted that solar panels will be added in future. To further save power and improve simplicity the team have attempted to infer wind direction information from only knowledge of sail setting and boat movement and orientation [146]. This allows for the removal of the wind sensor reducing overall power consumption, weight, complexity and removing a potential point of failure. Given the need to place the wind sensor outside the hull it is perhaps the most exposed of the sensors commonly attached to sailing robots and attempting to remove it would seem sensible when attempting to improve the robustness of a sailing robot. At only 1.2m long Breizh Spirit appears to make a reasonable

[5] `http://media.ensta-bretagne.fr/robotics/index.php/Travers%C3%A9e_de_la_rade_le_14_Septembre_2009` accessed 11/03/2011

compromise between the speed and agility of a larger boat and the simplicity, low cost and ease of transport of a smaller one.

3.10 Avalon

In 2008 a student team from the Swiss Federal Institute of Technology Zürich (Eldgenossische Technische Hochschule Zürich) began building a 4 metre long sailing robot known as Avalon with the main aim of entering the 2009 Microtransat [50]. At a cost of approximately 209,000 sFr (approximately £120,000) it is probably more expensive than all other boats built for the Microtransat, WRSC and Sailbot combined and is probably the most expensive boat reviewed in this chapter with the exception of the Harbor Wing X-1. Avalon featured a balanced rig design with an L shaped mast and fabric sail, this essentially combined the design of a wing sail in that there were no ropes (sheets) controlling the sail and the sail was kept taut and had no opportunity to collapse (luff) while still retaining features of a traditional sail as it can be raised and lowered. A 200 watt Maxon motor controlled the sail position. Rudder control was provided by a second 150 watt Maxon motor. Almost the entire deck surface is covered in photo-voltaic panels capable of supplying a peak of 360 watts, these are connected to 70 Lithium Manganese batteries with a capacity of 600 watt hours. A methanol fuel cell provides a backup source of power should the solar panel supply be insufficient. A 500mhz Linux based computer controlled the system with a Gumstix Single Board Computer providing a backup should the main system fail.

Frey (2009) [44] designed and implemented an automatic power control system which was able to control the interval at which the boat moved its actuators, read sensors and sent/received satellite communications. He identified the main source of power consumption to be the actuators and although exact figures were not available found the average power consumption (including actuator movements, communications equipment, sensors and computers) to be around 40 watts. He proposed that changing between duty cycles of 0%, 25%, 50%, 75% and 100% would allow a reasonable level of control over power consumption and decided to implement this using one hour blocks so for example a 75% duty cycle would leave the system turned off for one hour in four. A finer grain approach might have improved performance but possibly at the cost of slightly higher power consumption, should there be any penalty to pay when waiting for devices to startup before they become usable. Unfortunately Frey did not manage to test his software on Avalon itself but instead had to resort to simulation. During the 2009 WRSC "Avalon" suffered problems when making large turns which caused excessive stress to be placed on the Maxon motor controlling the sail. This resulted in an overheating motor and damage to the gearbox. This was further compounded

Figure 3.7: Avalon at the 2009 WRSC. Photo courtesy of ETHZ Students Sail Autonomously Team (`http://www.gallery.ethz.ch/ssa/` accessed 15/04/2011)

by damage to a GPS antenna and problems with magnetic interference to the compass.

3.11 Queens University's Mostly Autonomous Sailboat Team

The Mostly Autonomous Sailboat Team (MAST) is a student team from Queens University in Kingston, Ontario, Canada. Although the team has primarily focused on the Sailbot competition, their long term aim has been to build a Microtransat entry. To date all their boats have been 2 metres long and built primarily for high speed sailing in calm waters. Their second boat "North Star" struggled to cope with the wind and sea conditions in both the 2007 Microtransat and 2008 WRSC.

Queens University are currently designing a 4 metre long transatlantic boat. Blair (2010) [18] undertook an analysis of the power requirements for this boat and produced a design for an appropriate power system. She calculated that the system would need an average of 19 Watts and suggested three different power consumption profiles: "full power", "reduced power" and "power save". In "full power" no attempt to save power is made, in "reduced

power" the frequency of actuator movement was reduced and in "power save" the GPS and wind sensor are turned off, the sails are placed in a mid-reach position (half out) and steering is achieved against the compass only. Threshold remaining battery levels of 1050 Watt Hours and 150 Watt Hours were selected for the activation of "reduced power" and "power save" modes respectively. It was estimated that without any charging input the boat could sail for 5 days in "reduced power" and 70 days in "power save" modes.

There are several potential problems with this approach, firstly that when crossing the Atlantic conditions are likely to remain similar for hours or possibly days. A well balanced boat should not need to adjust sail or rudder positions regularly, so the boat could probably spend most of the voyage in the "reduced power" mode without any noticeable loss of performance. Secondly, a gradual reduction of actuator duty cycle in response to reduced battery level might improve power efficiency and eliminate the need for explicit power saving modes.

3.12 Discussion and Conclusions

This chapter has reviewed the present and past attempts at building sailing robots and shown that they are still at an embryonic stage in their development. Nobody has yet been able to successfully deploy one for a long term mission lasting more than 48 hours. Despite this, the field is fast moving and since work on this thesis began in 2005 the total number of people working on sailing robots globally has increased substantially. As much of the focus to date has been on getting basic short course sailing to work, little attention has yet been paid to longer term issues including power management.

Of the literature reviewed in this chapter, only the work by Stelzer and Jafarmadar (2009) [128], Blair (2010) [18] and Frey (2009) [44] have detailed the power generation and average power consumption of their robots. Unfortunately none of them have been able to validate their approach through any prolonged real world testing and they have all been limited to short experiments, theoretical calculations and simulation. All three boats generate power through a photo voltaic solar panel and store excess power in batteries and all three come close to covering the entire available deck area in solar panels. ASV Roboat has a peak solar output of 285W and a 30W average power consumption [127], Avalon has a peak solar power output of 360 Watts and an average of 40W [50, 44]. The boat described by Blair is currently only a design and does not exist for real. It is proposed to have 200W of solar panel on a 4 metre long hull and to use 17.5W on average. Table 3.1 compares the expected power output of the solar panels on the equinox at a latitude of 50 and 25 degrees (the range in which the Microtransat race would mainly take place). These figures assume that the solar panel is place on a flat surface, that there is no cloud and does not take into account degradation of

Team/Boat	Queens University		ETHZ - Avalon		INNOC - ASV Roboat	
Latitude	25 Degrees	50 Degrees	25 Degrees	50 Degrees	25 Degrees	50 Degrees
Average Power Consumption	17.5W		40W		30W	
Peak solar output	200W		360W		285W	
Daily solar output	4.9 mega joules or 1386 watt hours	3.5 mega joules or 985 watt hours	8.9 mega joules or 2500 watt hours	6.38 mega joules or 1700 watt hours	7.1 mega joules or 1972 watt hours	5.1 mega joules or 1417 watt hours
Average solar output	56.7 watts	40.5 watts	104 watts	73 watts	82.2 watts	59 watts
Error Margin	324%	231%	261%	183%	274%	196%

Table 3.1: A comparison of the solar power generation capabilities and power consumption of the robot boats Avalon, ASV Roboat and a theoretical boat by Queens University. Power generation is compared at a latitude of 50 and 25 degrees at the equinox. The error margin states the difference between the stated power budget and the average solar output, values over 100% mean that the robot is generating more than its stated power budget.

the panel due to a build up of salt or biofouling. Given these issues it would not be unrealistic to expect the solar panel power output to be cut by 50% or more due to these problems. If this is the case then there would not be sufficient power to run Avalon or ASV Roboat and barely enough to run the Queens boat at 50 degrees latitude on the equinox, let alone during winter. These figures also do not include any power budget for running environmental sensing equipment, which would be required to operate any kind of oceanographic survey mission. This goes to highlight the difficulties involved in sustaining a small sailing robot. Both Blair and Frey have devised mechanisms to modify their robot's behaviour in order to reduce power consumption, but these are through incredibly coarse changes and it would not be desirable to see the robot spending excessive amounts of time using the power saving mode. It would seem more sensible to produce some kind of mechanism which produces a gradual change in behaviour and can create proportional responses to low battery states rather than one which jumps from doing nothing to an incredibly conservative strategy.

Current power management strategies also seem to only be concerned with feedback from the current battery level rather than attempting to estimate future energy availability by considering how much sun light will be falling on solar panels in the near future. There also exists a tendency by many (such as the INNOC or ETHZ Microtransat teams) to simply

"over-engineer" the power system using excessively large batteries, solar panels and backing them up with fuel cell systems. While this approach is likely to guarantee that the boat always has sufficient electrical power it can increases cost, weight and the complexity of electrical systems.

A system based upon an artificial neuro-endocrine controller (as discussed in the previous chapter) could be capable of gradual adjustment of the robot's behaviour to manage power consumption and maintain the battery charge within a stable state. This could reduce the need for over engineering and help the robot to cope with a wide range of conditions. It also presents a mechanism by which the robot could adjust its behaviour to match future expected sunlight or tidal conditions rather than simply reacting to current or past conditions.

Chapter 4

Design and Evaluation of Sailing Robots

This chapter provides a background to the development of sailing robot hardware and software at Aberystwyth University since 2004. It covers a design overview of each boat, presents the experiments performed with each and provides a discussion of the implications of these and how they influenced the research aims of this thesis. Of these boats the MOOP (Miniature Ocean Observation Platform) is the most recent design and has been designed from the ground up rather than reusing existing hulls or hull designs as many of the other boats have. It is these boats which have formed the basis of the work covered in subsequent chapters.

The physical design and construction of these robots has not been my sole work and has been mostly undertaken by Dr. Mark Neal and Mr. Barry Thomas of the Computer Science Department at Aberystwyth University. The only exception to this was BeagleB (described in section 4.1.3) who's construction was outsourced to Robosoft. All software written for these robots has been my work, with two exceptions. The tacking algorithm described in algorithm 4.2 was initially devised by Dr. Mark Neal but has since undergone further modification to handle some scenarios he did not envisage when designing it. He also wrote the software UART used for PIC to Gumstix communication (described in section 4.1.5).

4.1 Boat Designs

4.1.1 AROO 2004/2005

AROO (Autonomous Robot for Ocean Observation) was constructed in late 2004 by Dr. Mark Neal and was used for my Masters thesis [115], although not directly part of this research, it formed the basis for many later design decisions and control system architectures. AROO consisted of a 1.5m long ABS plastic hull based on an off the shelf design for a remote control racing boat. The sail was an aluminium wing sail, built from a single aluminium

Figure 4.1: A photograph of AROO.

sheet folded in half to form an aerofoil. A wing sail was chosen over a traditional fabric sail in order to improve reliability and to reduce problems associated with wear and tear to the sail itself and to any ropes controlling it. A potentiometer on top of the sail acted as a wind sensor and a magnetic compass provided heading information. No GPS receiver was used in the control system, although a handheld Garmin eTrex was placed inside the hull during experiments to log the robot's position. A DC motor rotated the sail and a servo controlled the rudder position. The control system was split across two computers, a Basic Stamp microcontroller controlled all the actuators and sensors this provided an interface via a serial port to a higher level system. This higher level system was implemented on a PDA running Linux. Full hardware specifications are shown in table 4.1, a schematic is shown in figure 4.2 and a photograph is shown in figure 4.1.

A number of lessons were learnt in the construction and operation of AROO. Firstly, the compass was not tilt compensated and this introduced significant errors into compass readings resulting in unwanted oscillations of the boat. Secondly, the communications between the Basic Stamp Microcontroller and Linux PDA were unreliable and slow, limiting the control system's response rate to about 1hz. The lack of a GPS also limited navigation to simply following a fixed compass heading. A simple experiment was devised to sail a course of 45 degrees for 3 minutes and then to adjust the course to 225 degrees and sail it indefinitely. This returned the boat to a point a few metres down wind of its approximate starting location. AROO demonstrated the basic feasibility of operating a sailing robot autonomously and provided the basis for future work.

Name	AROO (Autonomous Robot for Ocean Observation) also known as "Number 1"
Hull	1.5m long ABS "hard chine" plastic hull
Weight	12kg
Power Sources	12 Volt, 4.2 Amp Hour Lead Acid Battery
Computers	Linux PDA (Psion 5mx later replaced with an HP Jornada 720) and a Basic Stamp Microcontroller
Sail	Aluminium Wing Sail 1.25m x 0.18m. Rotated with a DC motor (geared down to 2rpm) with potentiometer feedback.
Rudder	Single rudder, RC servo controller.
Compass	Devantech CMPS03
GPS	None
Wind Sensor	Continuous rotation potentiometer mounted on top of the sail.
Control System	High level control system ran on the PDA, communicated via serial to the Basic Stamp which ran all low level operations. Approximate response rate of 1hz.
Cost	Around £500 (2004), many components were scrap and therefore free.

Table 4.1: AROO's specifications

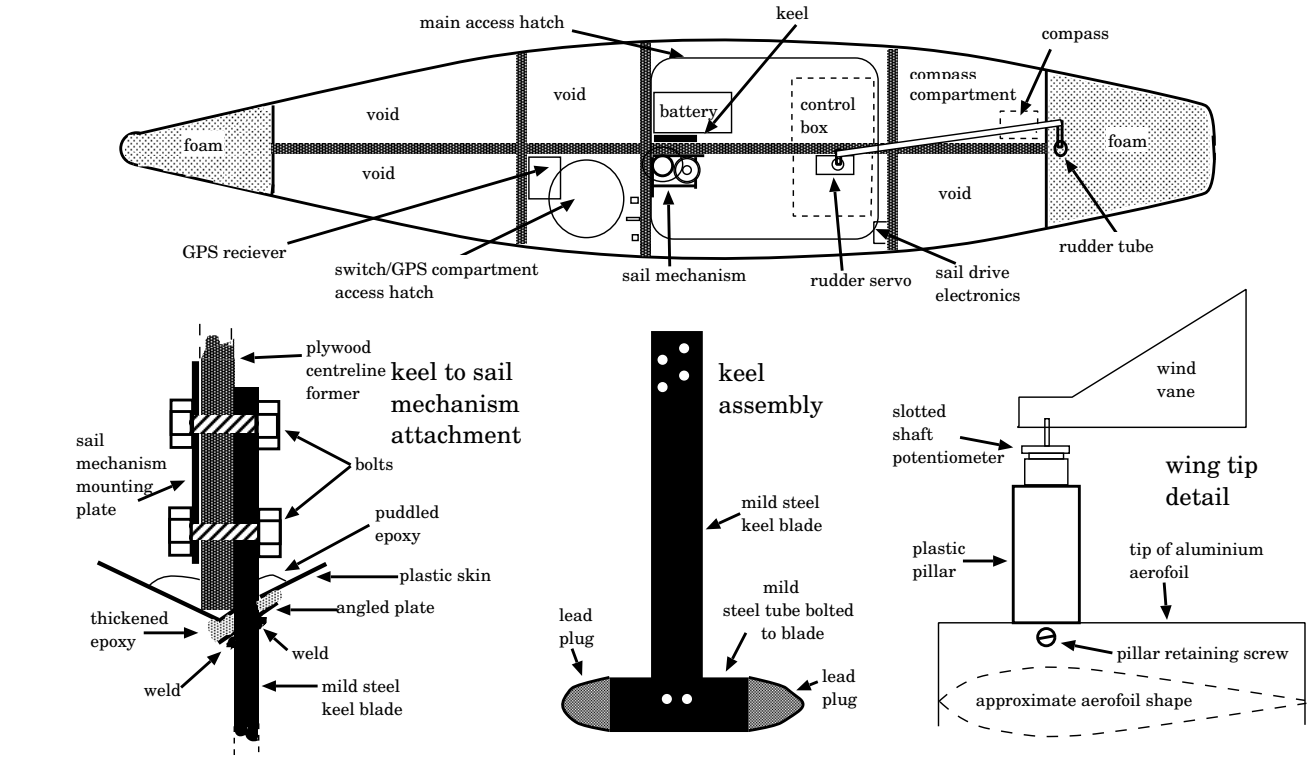

Figure 4.2: A diagram of AROOs hull and internal layout. From Neal (2006) [90].

4.1.2 ARC 2006-2008

Figure 4.3: ARC sailing at Saint Nicolas de la Grave lake in Southern France.

Name	ARC (Autonomous Robotic sailing Craft)
Hull	1.5m long hard chine plywood hull
Weight	15kg
Power	20 AA 1.2 volt 2300mAH Nickle Metal Hydride Batteries.
Computers	Gumstix Connex (200mhz)
Sail	Twin acrylic wing sails, controlled by two stepper motors.
Rudder	Twin rudders controlled by a single stepper motor.
Compass	Devantech CMPS03 on a horizontally swinging (1 DOF) gimble
GPS	Psion 12 channel GPS
Wind Sensor	Continuous rotation potentiometer mounted on a pole at the stern
Other Sensors	Power transistor temperature sensors for each actuator. Sail position feedback potentiometers.
Control System	Entire control system running on a Gumstix, low level control via GPIO14 IO extenders.
Cost	Around £300 (2006), many components were scrap or taken from AROO and therefore free.

Table 4.2: ARC's specifications

ARC (Autonomous Robotic sailing Craft) was developed in 2006 and aimed to address many of the shortcomings of AROO. Its hull was similar in length but wider, this was intended to make the boat more stable and easier to control. There was a general aim to create redundancy in all systems, with twin wing sails, twin rudders (although only a single

rudder actuator) and redundant motor controllers. The introduction of a Gumstix Single Board computer offered significant computing power combined with wireless networking and relatively high speed I/O systems (at least in comparison to AROO's), while the introduction of a GPS offered proper navigational abilities. A summary of ARC's hardware specifications is shown in table 4.2.

4.1.3 BeagleB 2007 - present

BeagleB was constructed for Aberystwyth University by the French robotics company Robosoft in late 2006 and early 2007. Robosoft were responsible for the construction of the hull, wing sail and identifying the compass, GPS, wind sensor, actuators and motor controller but did not supply any computing systems. BeagleB was intended to take the lessons learned from AROO and ARC and produce a large scale boat which would be capable of remaining at sea for long periods of time. It was built using the hull of a Mini-J sailing dinghy but equipped with a carbon fibre wing sail. It is equipped with solar panels and lead acid batteries which should be sufficient to theoretically enable it to remain at sea continuously. Full specifications of BeagleB are shown in table 4.3. BeagleB's rudder and sail are operated through LA12 linear actuators, which provide position feedback through an analogue signal line. Initially motor speed and direction control was provided with a Devantech MD22[1] motor controller and position feedback was read through a GPIO14 I/O extender. Both of these were controlled by a Gumstix single board computer and communicated via an I2C bus. Initially this approach appeared to work well but later the I2C bus was identified as a potential cause of several crashes. After not using the boat for approximately a year between the autumn of 2007 and 2008 the I2C communications became dramatically less reliable for an unknown reason and it was decided to replace the system with something more reliable. Around the same time bad experiences with I2C on ARC had also suggested this was wrong technology to be using. The motor controller was replaced with four relays (two for each actuator) controlled by a PIC microcontroller and a 9 bit (8 data bits and 1 bit to select which motor to use) parallel bus between the Gumstix and the PIC. The desired actuator position was placed onto the data bus by the Gumstix and the PIC would then control the movement of the actuator to that position. This removed the need for the Gumstix to perform real time processing of actuator positions and was more effective as the PIC could respond to position feedback much faster since virtually no communications delays were incurred in reading the analogue position feedback line (unlike reading these via I2C from the Gumstix). This system was also far more resilient to electrical noise which was suspected of corrupting data on the

[1] `http://www.robot-electronics.co.uk/htm/md22tech.htm` accessed 20/3/2011

Name	BeagleB
Hull	3.65m long, built from the hull of a "Mini J" dinghy built for disabled sailors.
Weight	280kg
Power	2x 45 Watts (peak) photo-voltaic solar panels, Four 12 volt, 60 amp hour lead acid batteries.
Computers	Gumstix Connex (200mhz)
Sail	Single Carbon Fibre wing sail 3m x 1m, Linak LA12 linear actuator.
Rudder	Single rudder, Linak LA12 linear actuator.
Compass	Furuno PG500 compass.
GPS	Furuno GPS
Wind Sensor	Furuno Rowind ultrasonic sensor.
Other Sensors	Analogue position feedback sensors in the linear actuators.
Control System	Initially entire system running on Gumstix. MD22 I2C motor controller, motor feedback via I2C on a GPIO14 I/O extender, GPS/Compass/Wind sensor connected via serial ports. Later revision of the control system moved all motor control to a PIC 18F4550. Approximately 1hz response rate.
Cost	£30,000 (2007)

Table 4.3: BeagleB's specifications

I2C bus. Additionally it allowed for actuator movement to occur concurrently with other activity and for a new actuator position to be specified before the previous movement had even completed.

Figure 4.4: BeagleB sailing off the coast of Matoshinos, Portugal at the World Robotic Sailing Championships in July 2009.

4.1.3.1 Tacking Algorithms

Algorithm 4.1 Heading difference algorithm, used by the tacking algorithm.

```
function get_hdg_diff(int heading1,int heading2)
    int result
    result = heading1 - heading2
    if result<-180 then
        result = 360 + result
        return result
    if result>180 then
        result = 0 - (360-result)
    return result
```

The control systems of AROO and ARC had no method for tacking upwind. A simple tacking algorithm was devised which would pick an arbitrary tack and sail it until the course became directly sailable. This is shown in algorithm 4.1 and 4.2. The downside of sailing an upwind course in only one tack is that the boat will travel a long way from its original course. When operating in a confined space this creates a risk of accidentally sailing into the shoreline.

Algorithm 4.2 Tacking algorithm

```
function check_tacking(int relwind,int heading,int des_hdg)
  HOW_CLOSE=55
  truewind = relwind + heading
  if truewind > 360 then truewind -= 360
//work out if the course is sailable using the heading difference algorithm
  if abs(get_hdg_diff(truewind,des_hdg)) < HOW_CLOSE then
     SAILABLE = 0
  else
     SAILABLE = 1
     PORT_TACK = 0
     STBD_TACK = 0
//if it wasn't previously sailable, workout which tack to sail on
  if SAILABLE == 0 AND PORT_TACK ==0 AND STBD_TACK == 0 then
     tmp_truewind = truewind
     tmp_des_hdg = des_hdg
     if des_hdg < HOW_CLOSE then
        tmp_des_hdg = tmp_des_hdg + 180
        tmp_truewind = tmp_truewind + 180
        if tmp_truewind > 360 then
           tmp_truewind =  tmp_truewind - 360
     if des_hdg > (360 - HOW_CLOSE) then
        tmp_deshdg = tmp_deshdg - 180
        tmp_truewind = tmp_truewind - 180
        if tmp_truewind < 0 then
           tmp_truewind += 360
     if tmp_des_hdg > tmp_truewind then
          PORT_TACK = 1
     else
        STBD_TACK = 1
  if SAILABLE == 0 then
     if PORT_TACK == 1 then
        des_hdg = truewind + HOW_CLOSE
     if STBD_TACK == 1 then
        des_hdg = truewind - HOW_CLOSE
  if des_hdg > 359 then des_hdg = des_hdg - 360
  if des_hdg < 0 then des_hdg = des_hdg + 360
  return des_hdg
```

4.1.3.2 Power Budgets

Component	Power Consumption	Duty Cycle	Average Power Consumption
Computer - Gumstix Verdex	200mA, 5V = 1W	100%	1W
Wind Sensor - Furuno Rowind [2]	50mA,12V =0.6W	100%	0.6W
Compass - Furuno PG-500 [3]	2W	100%	2W
GPS - Furuno GP-320B [4]	105mA,12V = 1.26W	100%	1.26W
Satellite Modem - Iridium 9601 [5]	0.35W (standby power, transmission cycle only lasts 80ms, effect on overall consumption is negligible)	100%	0.35W
Sail Actuator - Linak LA12	60W	1%	0.6W
Rudder Actuator - Linak LA12	60W	1%	0.6W
Total			6.41W

Table 4.4: A liberal power budget for BeagleB assuming that all sensors and the Gumstix are left switched on.

Component	Power Consumption	Duty Cycle	Average Power Consumption
Computer - Gumstix Verdex	200mA, 5V = 1W	10%	0.1W
Wind Sensor - Furuno Rowind	50mA,12V =0.6W	10%	0.06W
Compass - Furuno PG-500	2W	10%	0.2W
GPS - Furuno GP-320B	105mA,12V = 1.26W	1%	0.0126W
Satellite Modem - Iridium 9601	0.35W (standby power, transmission cycle only lasts 80ms, effect on overall consumption is negligible)	2%	0.007W
Sail Actuator - Linak LA12	60W	1%	0.6W
Rudder Actuator - Linak LA12	60W	1%	0.6W
Total			1.5656W

Table 4.5: A conservative power budget for BeagleB assuming that the Gumstix, Rowind and Compass can operate at a 10% duty cycle, that the GPS can operate at a 1% duty cycle and the Iridium modem at 2%. It is also assumed that each device will draw no current when switched off.

Tables 4.4 and 4.5 present two potential power budget scenarios for BeagleB. Table 4.4 shows a liberal power budget in which the Gumstix, all the sensors and the Iridium modem are left switched on at all times and that the actuators are moving 1% of the time. This brings a

total of 6.41 Watts on average. A more conservative power budget is shown in table 4.5, in this scenario it is assumed that the Gumstix, Rowind and Compass operate on a 10% duty cycle, the Iridium modem on a 2% duty cycle and the GPS on a 1% cycle. In this scenario only 1.5656 watts are used on average.

An estimation of solar panel power output can be calculated by considering the power output of a solar panel to be equal to the sine of the angle of the sun's elevation and that a cloudy day produces half the energy of a sunny day. On the summer solstice at 52 degrees north (the latitude of Aberystwyth) this will result in an average power output of between 16.44 and 32.88 watts. On the winter solstice it will produce an average power output of between 2.37 and 4.74 watts. This shows that the liberal power budget is realistically sustainable in the summer months and the conservative one is just sustainable during winter. It should also be noted that many of the intended applications for a sailing robot will involve running additional equipment and that spare power needs to be available for these. Therefore there is a desire to create a control system which provides as much spare power for applications as possible. Additional power can always be used to power extra scientific monitoring equipment or more frequent communications.

4.1.4 Pinta 2008 - 2010

Name	Pinta
Hull	2.95m long hull, built from a Topper Taz sailing Dinghy
Weight	Approximately 150kg
Power	16x 12 volt, 7 amp hour lead acid batteries. 120 watts peak photo-voltaic solar panels.
Computers	Gumstix Verdex Pro (400mhz) and later a PIC 18LF4550 microcontroller
Sail	Traditional Bermudan rig main sail, approximately 2 m^2 sail area. DC motor controlling main sheet wound around a drum.
Rudder	Simrad TP22 [6]tiller pilot to control the rudder. Interfaced buttons via I/O lines on the Gumstix.
Compass	Furuno PG500 compass, later replaced with a lower cost Devantech CMPS09
GPS	Telit GM862, replaced with Garmin eTrex and finally with a SiRF3 module.
Wind Sensor	Furuno Rowind ultrasonic wind sensor.
Other Sensors	Hall effect sensor to detect rotations of drum. Later replaced with a multi-turn potentiometer on the drum.
Control System	Entire control system running on a Gumstix, motor interface via Gumstix GPIO lines. Response rate approx 1hz. Moving to same control board as BeagleB featuring a PIC and Gumstix combination.
Lighting	Tri-Colour (Port/Red, Starboard/Green, Front/White) navigation light on the deck above the bow. Later replaced with lower power consumption Aqua Signal 32000 All Round White LED.
Cost	Around £3000 (2008)

Table 4.6: Pinta's specifications

Pinta, was built with the sole aim of competing in the Microtransat Challenge (discussed in section 3.1), a transatlantic race of autonomous sailing robots. Although BeagleB was probably capable of a transatlantic crossing it was considered to be too expensive to risk. Therefore, a cheaper boat was required that would closely replicate BeagleB's performance. Pinta, named after one of the ships which accompanied Columbus on his first voyage to the Americas. Pinta, was constructed from a Topper Taz[7] child's sailing dinghy. Unlike all the other robots developed at Aberystwyth University it used a fabric sail instead of making use of a rigid wing sail, this decision was taken simply to ease construction and reduce costs. However, the existing sail was considered too large and its sprit rig design too fragile for transatlantic sailing. The sail was trimmed to form something more resembling a Bermudan

[7] `http://www.toppersailboats.com/taz.aspx` accessed 20/3/2011

Figure 4.5: A photograph showing Pinta with its original sprit rig.

rig, this reduced the total sail area to about 2 m^2. Photographs of Pinta before and after this modification are shown in figures 4.5 and 4.6. This sail was later replaced with a custom made sail of similar size and shape.

Pinta used the same Gumstix single board computer, Furuno PG500 compass and Furuno Rowind wind sensor that had all been used in BeagleB. The GPS was initially provided through a Telit GM862 combined GSM modem and GPS, which allowed for telemetry to be sent via GSM while simultaneously reading the GPS. However, this device suffered reliability problems and eventually stopped working when salt water leaked into the electronics compartment, it was initially replaced with a Garmin eTrex handheld receiver and finally a SiRF3 GPS module, the GSM functionality was lost in this process. Steering was achieved with a Simrad TP-22 tiller pilot which was modified so that the control buttons were triggered through I/O lines on the Gumstix and the state LEDs could be read from the Gumstix. The tiller pilot has three modes, the first does not actively move the tiller but allows the user to control it manually using a left and right button, the second holds the current heading and the third takes an input via NMEA commands sent over a serial line which can be used to specify a target heading to follow. This first version of Pinta's control system used the manual mode with a proportional controller to get the boat on course within 20 degrees and would then switch to the heading hold mode removing the need for the Gumstix to make any decisions about course holding. This was planned to allow the Gumstix to be shutdown much of the time to reduce power consumption. However, this was never implemented (and was not appropriate for short course sailing) and the control system continued to monitor the heading and if the difference between the target heading and current heading varied by

Figure 4.6: A photograph showing Pinta with its cut down sail and Bermudan rig.

more than 20 degrees then the Gumstix would retake control over the tiller pilot until the heading error was again reduced to less than 20 degrees. The sail was controlled with a rotating drum which wound up the main sheet, a magnet and hall effect sensor on the drum were used to count rotations. In total 5 rotations were required to wind up the entire main sheet and therefore 5 different sail positions could be easily achieved. Although this may not be optimal it was found to be more than sufficient to allow Pinta to sail and complete test courses in all conditions in which it was tested (Beaufort force 1-5). The full specifications of Pinta are shown in table 4.6.

4.1.4.1 Power Budget

Component	Power Consumption	Duty Cycle	Average Power Consumption
Gumstix Verdex Pro XL6P	200mA,5V = 1W	100%	1W
Rowind [8]	50mA,12V =0.6W	100%	0.6W
GPS (EM-408 SiRF3 [9])	44mA, 5V = 0.22W	100%	0.22W
Compass (Deventech CMPS09 [10])	25mA, 5V = 0.125W	100%	0.125W
Iridium 9601 Modem [11]	70mA, 5V = 0.35W	100%	0.35W (Transmit/receive burst lasts only 80ms and consumes 2A. At 1 message per hour this does not noticeably alter average power consumption)
PIC18LF4550 [12]	25mA,5V = 0.125W	100%	0.125W
Sail Motor	60W (measured moving under normal load)	1% (sufficient for 6 full movements per hour)	0.6W
Tiller Pilot [13]	500mA,12V = 6W (average figure according to data sheet)	N/A	6W
Navigation Light (Aqua Signal 32000 All Round White LED [14])	1.5W	50%	0.75W
Total			9.175W
Solar Panels	120W Peak		Approximately 25W (optimal) or 10W (poor weather) at 52N in mid September.

Table 4.7: Pinta's power budget (September 2010 version of Pinta) showing the typical consumption for each device, its duty cycle and the resulting average. 3.3 volts is supplied through a linear voltage regulator which regulates 5 volts to 3.3. This is assumed to waste excess energy as heat, therefore all 3.3 volt devices are calculated as if they ran at 5 volts. Energy losses from the sail actuator's relays, DC-DC converters which convert 12 to 5 volts and solar charge controllers are not included.

As the 90W Peak of solar panels on BeagleB were shown in section 4.1.3.2 to only produce just enough power to keep the robot running, Pinta was over engineered and given 120W peak of solar panels. At 52 degrees North these should be capable of producing an average of between 13.5 and 27 watts during September (the intended departure time for the Microtransat). Relatively little attention was paid to intelligent power management and all sensors, the Gumstix and Iridium modem were left switched on 100% of the time. Pinta's power budget (describing Pinta as it was equipped in September 2010) is shown in table 4.7. This shows the average power consumption as being 9.175W. However, the values shown in table 4.7 are mainly based on data sheet figures not actual measurements. Also these do not include power inefficiencies of voltage regulators, DC-DC converters, solar charge controllers and the consumption of relays used to switch the sail motor and the navigation light.

4.1.4.2 Microtransat 2010 Modifications

Pinta underwent further modifications to prepare for the 2010 Microtransat Challenge. The solar panels were installed and the computer system replaced with a new Gumstix and a PIC microcontroller for delegating low level commands to. This computer system is nearly identical to the system being used in BeagleB and the MOOPs (see section 4.1.5). Sail position feedback was changed from a hall effect sensor and magnet to a multi-turn potentiometer on the base of the sail winch, this gave more accurate sail position feedback and returned absolute positions rather than relative ones. The tiller pilot control system was changed so that it was always in automatic mode and the desired heading was adjusted through a series of button presses, this eliminated the need for position feedback readings from the tiller pilot. To adjust the target heading the left or right button had to be pressed a series of times, one short press for each one degree change and one long press for each ten degree change. A delay of one second was required between each button press, so entering large heading changes was a time consuming process. Occasionally button presses didn't seem to register, to ensure the correct heading was entered the control system would resend the target heading every 90 seconds if the boat was more than 15 degrees off course. Sometimes it could take a few minutes for the boat to settle down on the right course after reaching a waypoint. However, given that sailing across the Atlantic would involve sailing the same course for days at a time this was not perceived to be a major problem.

The compass was replaced with a lower cost Devantech CMPS09, although in final testing the day before departure this was showing an occasional fault where a heading of zero was sometimes returned. Rather than risk an intermittent fault, which was believed to be due to loose wires which were no longer easily accessible or fixable the decision was taken to use the GPS heading instead. This had the side effect of giving inaccurate readings when the boat

moved slower than about 1km/h or going backwards. However, the GPS returns the actual course made good instead of the direction the boat is pointing. This provides the control system with information about the progress the boat is actually making rather than the direction the hull is pointing in. When sailing up wind these can differ and it is actually the course made good that the control system needs to evaluate to decide if the boat is travelling in the correct direction. This risk was considered to favourable to using intermittently faulty data from the compass. A series of short test runs before the departure showed that this still produced a viable control system which was able to sail a triangular course.

An Iridium 9601 Short Burst Data modem was added for satellite communications in addition to an independently powered SPOT satellite tracking device. The Iridium modem provided a telemetry message every hour which included the boat's position, target waypoint number, distance to waypoint, wind direction and speed, air temperature, battery state and information about whether or not the waypoint was directly sailable or required tacking. The SPOT provided position reports every 6 hours via a simplex transmission to a satellite.

4.1.5 MOOP 2008 - Present

The development of the MOOPs (Miniature Ocean Observation Platforms) began during the summer of 2008 and aimed to develop an alternative Microtransat contender that was small, simple, light weight, easy to transport, easy to deploy and that could be produced in large numbers. The rationale was that several (perhaps 4 or 5) MOOPs would cost about the same as a larger boat (e.g. Pinta) but there would be a higher probability that at least one of them would cross the Atlantic. It was decided in early 2009 that they would also be ideal candidates for performing research into autonomous power management for this thesis. Previously ARC and BeagleB had been intended to fill this role, however the MOOPs were seen as preferable because they offered the opportunity of multiple boats and avoided the operational difficulties of using BeagleB.

The MOOPs are based around a 72cm long glass fibre hull loosely based on a traditional long keel boat design. This design uses a solid keel which is filled with lead for ballast and forms an integral part of the hull, running about three quarters of the hull length. The rudder is directly behind the keel and this helps to prevent any weed or other floating debris from becoming caught on the rudder. Figure 4.7 shows a photograph of a MOOP sailing and figure 4.9 shows a photograph of two MOOPs out of the water and shows the layout of the keel and rudder. Additional details of the MOOP's dimensions and electronics layout are shown in appendix C. At the time of writing 8 MOOPs have been completed with another 3 under construction, so far no two have been identical and each has been built with slightly different goals in mind. All of the MOOPs have been equipped with a solid state Honeywell HMC6343

Figure 4.7: MOOP1 being sailed on a lake near Aberystwyth during initial testing and development work.

tilt compensated compass, a SiRF3 GPS and a PIC18F4550 microcontroller. Some have also had a Gumstix Verdex Pro for additional computing power. A variety of wind sensors have been trialled including a custom built ultrasonic sensor, a continuous rotation potentiometer, a Furuno Rowind, no wind sensor (and attempting to derive wind direction based on sailing performance) and most recently an Austria Microsystems AS5040 rotary magnetic encoder. The magnetic encoder uses a magnet connected to a rotating shaft from the vane, this is then separated from the encoder by a thin layer of epoxy resin which ensures no water can reach the electronics. The encoder transmits the magnet position using a pulse width modulated signal at a frequency of approximately 1kHz.

Both single wing and twin wing varieties have been built and work has begun (although at time of writing is far from completion) to steer entirely using the sails, thus removing the need for a rudder. The sail actuator is a Futaba S3306 servo, this provides a rotation range of approximately 200 degrees and can move through this range in under one second. The rudder servo is a Futaba Micro 2BB MG and is connected to the rudder through a magnetic linkage, this removes the need for a hole in the hull and removes a potential source of leaks. Should the rudder be knocked out of position by wave then it will force the magnets out of alignment rather than placing force on the servo, after sufficient servo movement the magnets eventually realign and the mechanism will recover.

Both servos (although particularly the sail servo) suffer from problems of poor accuracy

Position	Sail Angle	Rudder Angle
-5	70	140
-4	55	148
-3	35	156
-2	25	164
-1	10	172
0	0	180
1	350	188
2	335	196
3	325	204
4	305	212
5	290	220

Table 4.8: The angles of MOOP sail and rudder shown against their position numbers. Sail and rudder angles are with respect to the angle of the bow (front) end of the sail and rudder.

repeatability. During development it was found that up to 10 degrees difference in position could be experienced when requesting the servo to move to the same target position. This was found to be in part dependent upon the originating position and the load on the servo. As no additional feedback sensing was available there was no simple way to overcome this limitation. This lead to only 11 unique positions for each servo being used across a travel range of approximately 180 degrees on the sail servo and 80 degrees on the rudder servo, this typically guaranteed that at least some level of movement would occur between any two positions. The approximate angles of these positions are shown in table 4.8.

Some later variants of the MOOP featured a modified version of this servo which was able to continuously rotate through 360 degrees and used a rotary magnetic encoder for position feedback. This was to allow "non-standard" sail positions (the equivalent of allowing the boom out more than 90 degrees from its centralised position on a normal sailing boat) on twin wing boats. This also gave far greater control over the exact positioning of the servo and allowed the control system access to accurate position feedback information.

4.1.5.1 MOOP0 and MOOP1

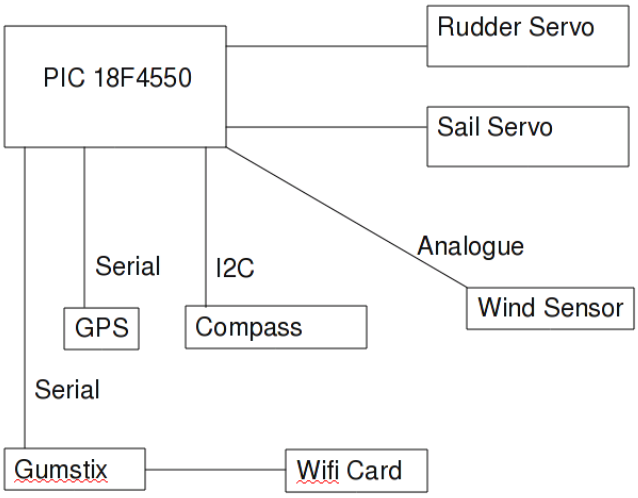

Figure 4.8: A block diagram showing the hardware connections within the MOOP boats.

From March 2009 one MOOP was dedicated towards this thesis with a second being obtained in June 2009. These were the first and second MOOPs built and were named MOOP0 and MOOP1. These are nearly identical in their specifications except that the hulls are slightly different in shape. MOOP0's hull was built around a polystyrene mold while MOOP1's hull was built from a mold and has a much smoother look and feel. MOOP0's deck is split into two parts forward and aft of the mast while MOOP1's deck is a single piece which improves access to the electronics but requires the entire sail assembly to be removed with the deck. The rudder assemblies also differ slightly, the servo on MOOP0 is mounted directly above rudder and the join between the rudder and the hull is on a vertical axis. While MOOP1 has the servo pointing backwards at approximately 45 degrees, this causes the rudder movement properties to be slightly different and gives MOOP0 a more effective rudder but the system in MOOP1 was easier to construct.

The rudders of both MOOPs were extended to be approximately double their original area. This was after observing that the MOOPs struggled to manoeuvre when sailing either in high winds (high winds for a MOOP being Beaufort force 3 and above) or when a tack failed, causing the boat to loose its momentum and steering ability. The different rudder layouts can be seen in figure 4.9. Both boats were equipped with a polystyrene and glass fibre wing sail with a rotary wind sensor on top, initially these were continuous rotation potentiometers and later they were replaced by rotary magnetic encoders which did not suffer from a dead spot as the potentiometers do at the point when they wrap around. Both were equipped with

Figure 4.9: MOOP1 (left) and MOOP0 (right) out of the water showing the keel and rudder designs.

Gumstix single board computers to run the main control system and PIC microcontrollers to provide low level control, the two being linked by a serial port. MOOP0 used an older Gumstix Connex with a compact flash card providing wireless network connectivity while MOOP1 used a newer Gumstix Verdrex Pro with a miniature wireless access point connected via ethernet providing its wireless network connectivity. Functionally these two Gumstix are identical, however the external access point means that MOOP1 requires more electrical power and that the physical distribution of electronics is very different. Electrical power is provided through 18 AA batteries in three packs of six batteries wired together in parallel providing a total of 8.1 amp hours at a nominal voltage of 7.2 volts. Both boats were also equipped with two ultra bright LEDs pressed up against the deck for locating them in darkness or fog. It had been found during development that it was quite easy to loose sight of a boat after sunset and if it was beyond wifi range it became nearly impossible to establish its location. One LED was located at the stern and another at the bow, each boat used a different colour combination (MOOP0 green at the stern and white at the bow, MOOP1 white at the bow and blue at the stern), this made it possible to identify the boat and determine which direction it was pointing. MOOP1 was also equipped with a white LED at the stern which flashed every time a command was received by the PIC from the Gumstix, this created a continuous flashing when the control system was running and meant that a crashed system could be identified at a glance. MOOP0 was not equipped with this

Name	MOOP (Miniature Ocean Observation Platform)
Hull	72cm long, glass fibre hull
Weight	4kg
Power	18 AA 1.2 volt 2700mAH Nickel Metal Hydride Batteries, connected in parallel packs of 6.
Computers	PIC18F4550[15] and Gumstix Verdex Pro or Gumstix Connex [16]
Sail	Single Polystyrene/glass fibre wing sail. Controlled by a Futaba S3306 RC servo.
Rudder	Single rudder controller by an RC servo. Magnetic linkage between servo and rudder to prevent potential source of hull leaks.
Compass	Honeywell HMC6343 [17]
GPS	EM408 SiRF3 GPS module[18]
Wind Sensor	Continuous rotation potentiometer or Austria Microsystems AS5040 continuous magnetic encoder on top of sail.
Communications	Wifi via either a D-Link DWL-G730AP [19] Travel Wifi Access Point or Prism 2.5 Compact Flash Wifi Card
Control System	Gumstix and PIC18F4550 combination. Communication via serial. Response rate of 10-12hz.
Other	LEDs located at the bow and stern for locating the boat in the dark, determining which boat it was (MOOP0 white/green, MOOP1 white/blue) and which way direction it was pointing.
Cost	Around £500 (2009) per boat - excluding labour.

Table 4.9: MOOP's specifications

due to difficulties in wiring it due to the physical layout of components in the boat. The full specifications of MOOP0 and 1 are shown in table 4.9 and a block diagram of the system is shown in figure 4.8. Appendix C shows detailed annotated photographs of the internals of MOOP1 and size measurements for the hull.

4.2 Experiments

This section details preliminary experiments undertaken with ARC, BeagleB, Pinta and the MOOPs. These experiments represent an effort to produce a working and stable control system for sailing robots in order to carry out research into long term power management.

4.2.1 ARC Experiments

4.2.1.1 Microtransat 2006

The only major test of ARC was undertaken during the first Microtransat Challenge on Lake Saint Nicholas de le Grave near Toulouse in France (44.0843 degrees North , 1.033 degrees East) on June 9th 2006. Unfortunately a gear driving the sails broke during on shore testing and it was not possible to move the sails. Instead the sails were fixed in position and the control system was setup to sail a fixed compass heading. A beam reach (across the wind) course was chosen as this is typically the most stable and easiest to control point of sail. Two runs were undertaken and GPS plots of these are shown in figures 4.10 and 4.11. The resulting sailing appeared to be very stable and did not see the oscillations that AROO had suffered from.

To test if this stability was a result of the control system or the hull design another run was conducted without any control over the rudder. A GPS plot of it is shown in figure 4.12. At one point during this run a motor boat passed by at high speed and generated a large wash which spun ARC around 180 degrees, despite this it spun back around and proceeded to sail on its original heading despite there being no active control system or intervention from the chase boat crew. This suggests that once the sails and initial course are set correctly that ARC could sail between waypoints with little or no actuator movement and therefore an efficient control system would suppress all (or almost all) actuator movement between waypoints.

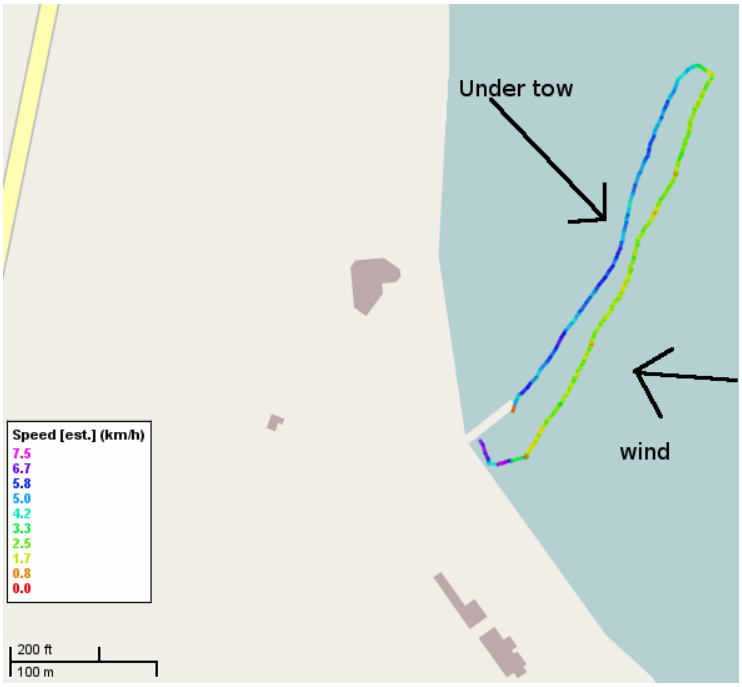

Figure 4.10: A GPS plot showing the first test run of ARC on Lake Saint Nicholas de la Grave in France in June 2006. The sail positions were fixed for this so the boat only controlled the rudder and attempted to sail a fixed course. The return leg (blue line) is under tow from a motor boat. Map courtesy of gpsvisualizer.com and OpenStreetMap.

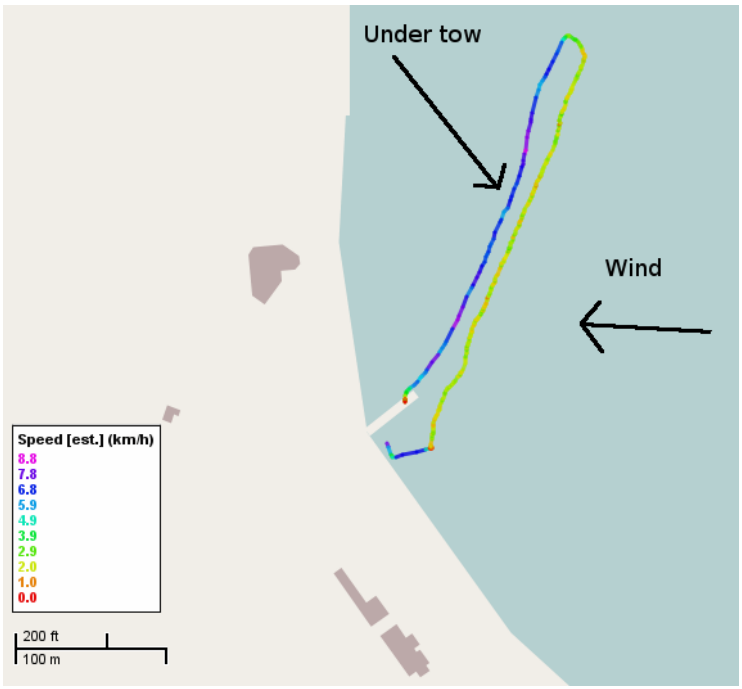

Figure 4.11: A GPS plot showing the second test run of ARC on Lake Saint Nicholas de la Grave in France in June 2006. The sail positions were fixed for this so the boat only controlled the rudder and attempted to sail a fixed course. The return leg (blue line) is under tow from a motor boat. Map courtesy of gpsvisualizer.com and OpenStreetMap.

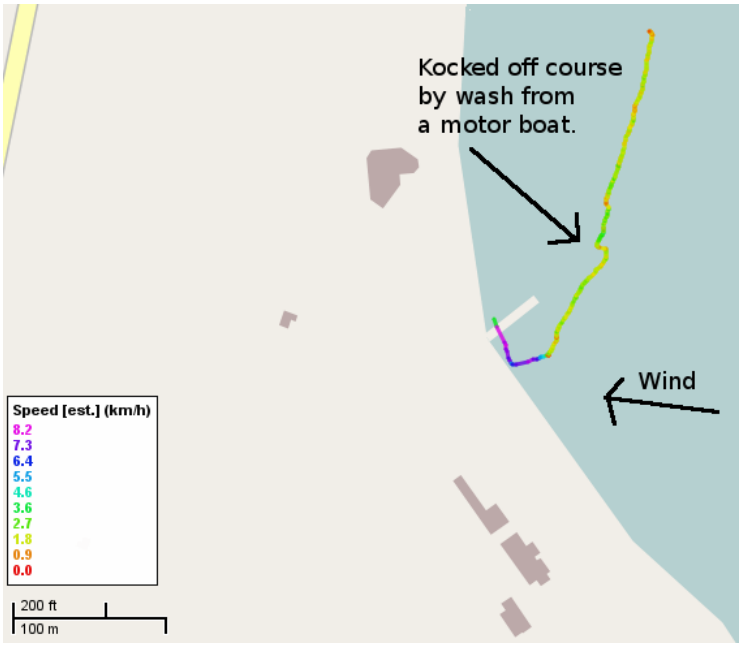

Figure 4.12: A map showing the third test run of ARC at Lake Saint Nicholas de la Grave in France in June 2006. For this run there was no autonomous control of the boat with the aim being to test its passive stability. Approximately halfway along the course a motor boat passed creating a wash which spun ARC around 180 degrees, but it was able to correct this without intervention. The only side effect being that the boat ended up further down wind.

4.2.1.2 Redundant Control Architecture

ARC was built with the intention that it would be used to develop an artificial neuro-endocrine controller similar to those developed by Mendao [76] and Neal and Timmis [92]. One of the possibilities for such a controller was envisaged to be the creation of a redundant control system. Such a system would have the ability to switch between different actuators or motor controllers in order to overcome actuator failures or overheating components. This could be viewed as an artificial stress response to actuator failure. It was envisaged that an artificial neuro-endocrine controller would release a hormone in response to actuator failures and that this would signal behavioural changes such as, changing to a different motor controller or reducing the actuator's duty cycle.

Each actuator in ARC was driven by a motor control board with a series of power transistors. These had a tendency to overheat and to create redundancy, two sets of transistors controlled each actuator. Each power transistor was bolted to a heat sink, a temperature sensor was also attached to the heat sink. A photograph of this is shown in figure 4.13. In the software system the artificial endocrine system was used to select which motor controller was to be used. The temperature of each motor controller was monitored and two hormone's

Algorithm 4.3 The controller selection algorithm.

```
/*get temperatures in degrees Celsius
consider 55 degrees the maximum allowed
as this is the temperature of the sensor the other side of the
heat sink the real power transistor temperature is much higher*/
hormone1 = 55 - get_temp1()
hormone2 = 55 - get_temp2()
if hormone1 < 0 then hormone1 = 0 //ensure values are 0-55
if hormone2 < 0 then hormone2 = 0
if hormone1 > 55 then hormone1 = 55
if hormone2 > 55 then hormone2 = 55
//generate a random number between 0 and 110
random = rand(110)
//if the random number is less than hormone1 then use controller 1
if random<hormone1 then
    controller = 1
//if it is between hormone1 and hormone1+hormone2 then use controller 2
else if random >= hormone1 AND random < hormone1 + hormone2
    controller = 2
//if the random number is more than controller1+2 then use neither
else
   controller = 0 //don't use either controller
```

produced in inverse proportion to the motor controller temperature. These hormones can be used to suppress the control of an individual motor controller. A random number generator is used to determine which motor controller to use, but as that motor controller's temperature increases its probability of being chosen falls. When the temperature sensors register over 55 degrees Celsius the motor controller will be completely out of use. 55 degrees may seem like a low threshold, but given that the temperature sensor is not directly attached to the power transistors but only to the same (rather large) heat sink, when the sensor registers 55 degrees the power transistors are actually experiencing much higher temperatures. The pseudo code of this algorithm is shown in algorithm 4.3.

ARC also has an additional layer of redundancy in its control surfaces. By having two sails and one rudder it is possible to steer with the sails, this creates several useful scenarios. First, the sails can be used to assist in steering and reduce "weather helm" (where the rudder must constantly be held off centre in strong winds). Second, in the event of the rudder failing the sails can become the sole source of steering. The hormone architecture previously discussed could be extended to influence the use of another actuator to assist an action. For example if the rudder controllers are both overheating then the hormones which are suppressing them could also activate another behaviour which begins to use the sails for steering. This work

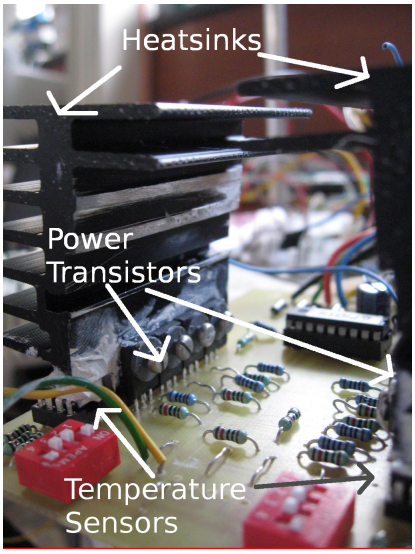

Figure 4.13: A photograph of one of the motor controller's from ARC.

is discussed further in Sauzé and Neal (2008) [116]. A copy of the circuit diagram for the redundant motor controllers is also included in appendix B.

Unfortunately the hardware of ARC did not allow these plans to mature into a working control system. The I2C bus which connected all the motor controllers, temperature sensors, position feedback potentiometers and the compass was prone to picking up interference from the electromagnetic field produced when the motors turned. This in turn caused the I2C device drivers in the Gumstix's Linux kernel to crash, typically requiring the Gumstix to be rebooted after only a few movements of the motors. Additional electrical problems with the motor controllers caused several power transistors to short circuit and the turn rate of the stepper motors for the sails was considered to be at the low end of the acceptable range. For these reason's ARC was abandoned and replaced with the MOOPs.

4.2.2 BeagleB Experiments

4.2.2.1 The 2007 Microtransat

BeagleB's first long and fully autonomous sail was during the second Microtransat Challenge in Aberystwyth in September 2007. This competition set an open event in which each competitor was free to sail a course of their choice for a maximum of 24 hours with the aim of demonstrating the capabilities of their boat. A large triangular course of approximately 17km was setup, a map and GPS plot of this course can be seen in figure 4.14. At the start the wind was roughly a force 2 on the Beaufort scale (7-11 km/h) and blowing from the south. After sailing approximately 2km west the wind shifted from a southerly to a

westerly triggering the boat to begin attempting to tack towards the south west. As can be seen in figure 4.14 and 4.15 that the boat was not able to maintain a south westerly heading. This is believed to be due to two factors: a bug (discovered later) in the control system which caused heading calculations to be wrong under certain conditions and that very little averaging of the wind direction was performed so small changes in wind direction could cause the system to alternate between attempting to tack up wind and sailing directly to the waypoint. After sailing (roughly) south west for approximately 2km the chase boat noticed a lack of movement from the actuators and that the boat was not sailing in what appeared to be a sensible fashion, it was concluded that the computer may have crashed and it was manually rebooted. The decision was then taken to abandon attempting to reach waypoint 1 and instead sail for waypoint 2. The robot then sailed without incident towards waypoint 2 but overshot the waypoint slightly and ended up about 1km north of it, this is believed to have been due to the action of the tide. At this point (approximately 1am) the wind had entirely dropped and the robot circled for several hours unable to correctly determine the wind direction. Around 6am the first signs of daylight appeared and a gentle westerly breeze began to blow, at this point the robot was able to determine the wind direction correctly and sailed onto waypoints 2 and 3.

This experiment showed that BeagleB still had a few bugs in the system but otherwise was capable of sailing for 19 hours and longer. During this experiment the crew of the chase boat noted that they rarely noticed the rudder and sail moving and that the boat seemed very stable once on course. This suggests that, as with ARC that when sailing between waypoints some distance apart that the control system can be suppressed and only infrequent checks of the compass and GPS are required.

Figure 4.15 shows that there is a consistent gap between the target heading and actual heading. This trend is repeated in figure 4.17 in the next section, despite this being with a significantly different control system. The most likely explanation for this is that it is due to the deadband around the target heading. Once inside the deadband the control system makes no further attempt to adjust the rudder to reduce the difference between the actual and target heading. Another possibility is that the accuracy and repeatability of the rudder actuator is sufficiently small that when the target heading nears the desired heading the rudder is sufficiently centred to make no difference. As the controller is only a proportional controller no further attempt is made to reduce the error level, if an integral controller were added to this then this error should be overcome.

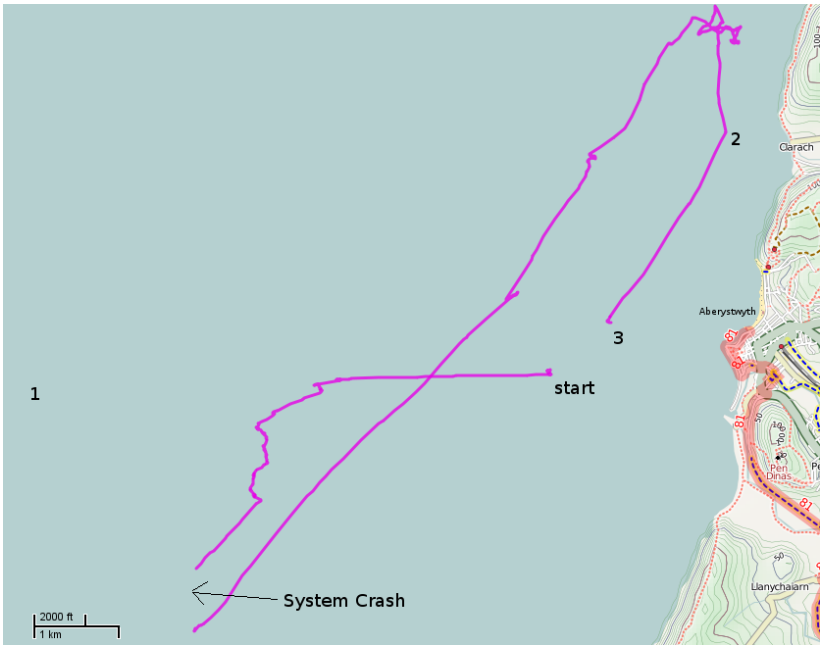

Figure 4.14: A GPS trace of the BeagleB from the September 2007 during the Microtransat Challenge in Aberystwyth. At the start point the wind was blowing from the South but it then shifted to the South West and the boat can be seen adjusting its course to tack south west around 2km west of the start point. After sailing on this course for approximately another 2km the system crashed for a reason that has never been fully identified and it was decided to head for waypoint 2 instead. As waypoint 2 was approached the tide was at its strongest and was pulling the boat north, at the same time the wind dropped and this is the cause of the circling to the north of waypoint 2. When the tide eased and the wind picked up the boat was able to reach waypoint 2 and finally waypoint 3. Map courtesy of gpsvisualizer.com and OpenStreetMap.

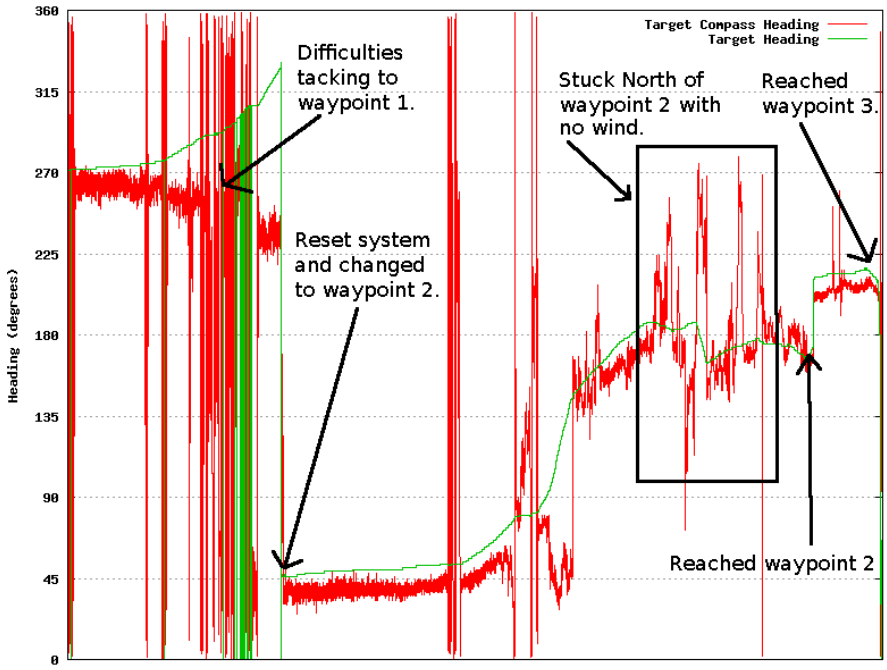

Figure 4.15: The compass heading of BeagleB during the 2007 Microtransat (see figure 4.14 for corresponding GPS trace).

4.2.2.2 2009 World Robotic Sailing Championship

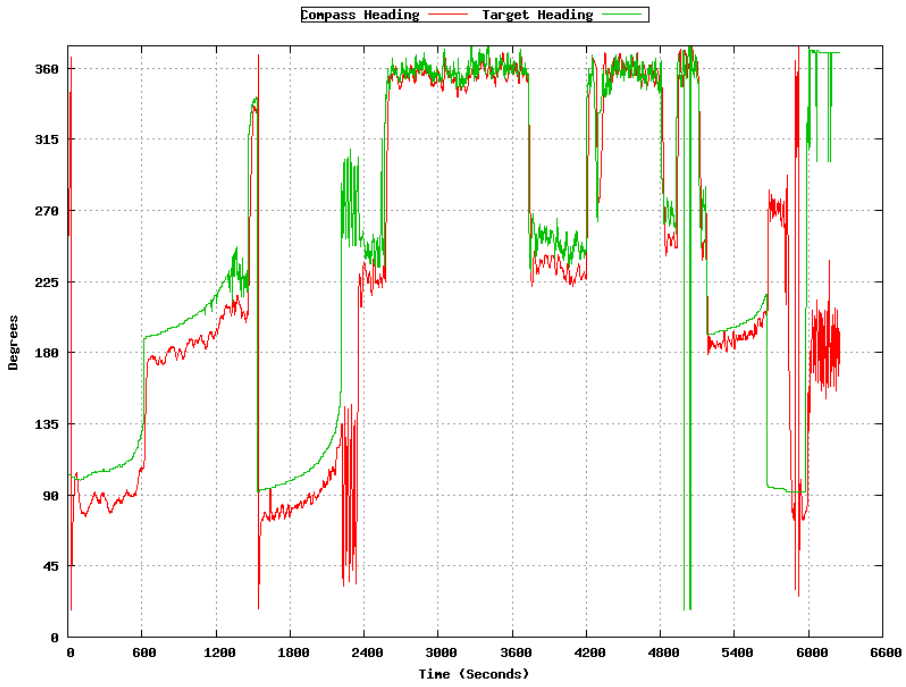

Figure 4.17: Compass headings of BeagleB during the 2009 World Robotic Sailing Championships (see figure 4.16 for corresponding GPS trace). Note that headings between 0 and 15 degrees have been graphed as being 360-375 to make this graph easier to view.

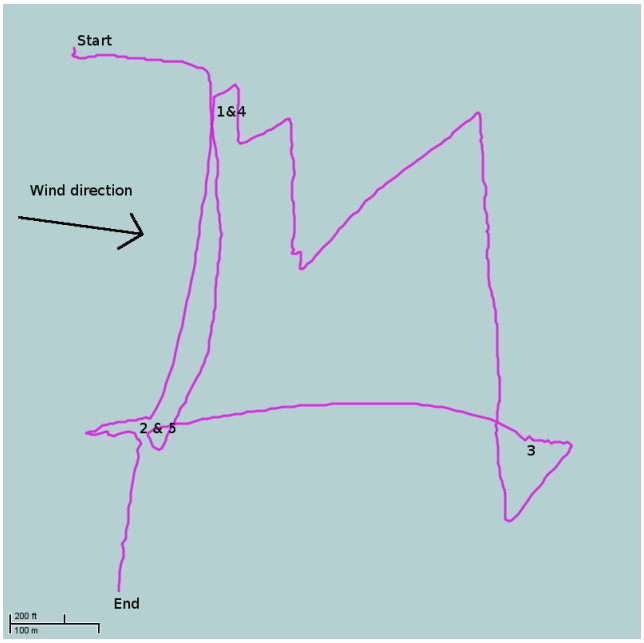

Figure 4.16: A GPS trace of BeagleB at the World Robotic Sailing Championships in Matosinhos, Portugal in July 2009. Map courtesy of gpsvisualizer.com.

BeagleB was entered into the 2009 World Robotic Sailing Championship (WRSC) held in Matosinhos, Portugal in July 2009. It had not been sailed since the 2007 Microtransat, although an attempt was made in October 2008 when the problems with I2C discussed in section 4.1.3 began to appear. At this point the decision was taken to redesign the control system and move to a PIC microcontroller and relays instead of an I2C based motor controller and I/O extender (for reading motor feedback). Much of the time during the WRSC was spent performing this re-engineering and by the final day the system was working in a credible manner. A triangular course about 1km in total length was set and BeagleB was able to sail 1 and 1/3rd laps of this course before it was aborted due to time constraints. The only problem encountered during this process was a bug prevented turns over 127 degrees. As the GPS plot in figure 4.16 shows, there was a problem rounding waypoints 2, 3 and 5, eventually the chase boat had to assist. The tacking between waypoint 3 and waypoint 4 is also of interest, this was using the same algorithm previously used (algorithm 4.2) which was supposed to sail the entire course in just 1 tack. However what happened was that a small wind shift caused the boat to believe it could sail directly to the waypoint when it could not and the result was several tacks each smaller than the last. This could actually be regarded as a useful property as it reduces the deviation distance from the direct course between waypoint 3 and 4. Ideally this tacking algorithm requires further work so that it will remain within a fixed corridor in a similar fashion to the approach taken by Stelzer and Pröll (2008) [129].

4.2.3 Development of the MOOP Control System

The MOOP control system was able to inherit much of the software developed for BegaleB and Pinta. A low level layer was written for the PIC microcontroller to interface with the compass, GPS, servos and wind sensor and allowed interaction with the Gumstix using a set of simple commands sent over a serial line. The higher level control system on the Gumstix was responsible for initiating all sensor readings and all actuator movements. It was discovered early in the development process that the MOOPs require a much higher rudder response rate than any of the previous boats. While rudder response rates of 1hz or less had previously been acceptable it was found that MOOPs could at times require rates as high as 10 to 12hz. Sail response rates needed to be no higher than 0.25hz. This created some engineering challenges and required the serial commands to the PIC be kept as short as possible (as transmission times were responsible for much of the lag in the system) and that duplicate commands not be sent (e.g. don't send "set servo to position x" twice in a row).

The servos both suffered from poor position repeatability and eventually it was decided that only 11 unique positions could be achieved. While this was perfectly sufficient for normal sailing and maintaining a target heading it did not allow small adjustments of the sail position which could have been suitable for implementing a sail steering system similar to that described for ARC in section 4.2.1.2. The lack of a sail steering mechanism increased the difficultly in sail upwind. It also made tacking harder, especially when the first attempt at tacking had failed and the boat had lost its momentum.

Testing of the MOOPs took place mainly on Llyn-Yr-Oerfa, approximately 12 miles east of Aberystwyth (52.4 degrees north, 3.87 degrees west), this provided a reasonably safe location without any marine traffic and where under most conditions a robot in difficulty will eventually wash up on an easily accessibly shoreline. However, the lake's small size caused very short wavelengths which tended to cause the MOOPs to constantly rock, slowing their progress especially when attempting to sail upwind. The tacking algorithm previously described in algorithm 4.2 was reused, while it was found that although the boat's could sail as close as 35 degrees to the wind the control system had trouble holding them on this heading and that the boat would frequently tack after sailing upwind for only a short period of time. To combat this the software was modified so that the boat would only attempt to sail 55 degrees to the wind. This allowed a sufficient margin of error for it to point too high up wind and not accidentally tack but reduced the overall velocity made good (VMG). Sailing downwind also tended to result in frequent jibes and the control system would struggle to sail towards the target heading. To prevent this a downwind tacking strategy was also implemented using the same technique as the upwind tacking algorithm and keeping at least 55 degrees from sailing straight down wind. Given these limitations, it was suggested that for

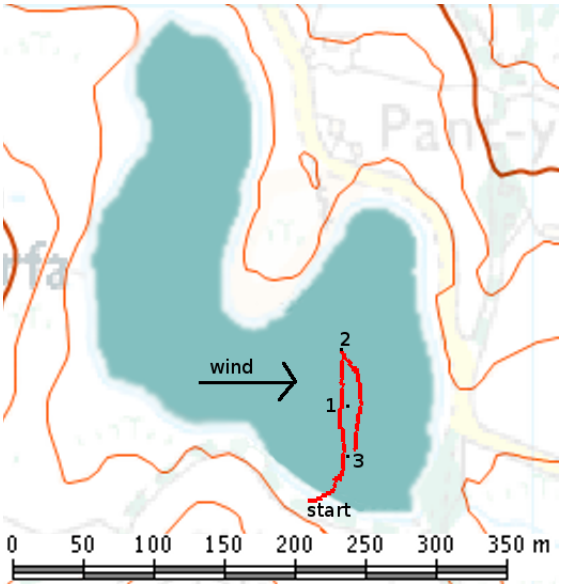

Figure 4.18: A GPS plot showing an example MOOP course. © Crown Copyright/database right 2009. An Ordnance Survey/(Datacentre) supplied service.

experiments requiring the boat to remain sailing for several hours in the constrained space of a lake, that it would be easier simply to sail back and forth across the wind on a beam reach rather than to attempt a triangular course (as is typically used in sailing boat races).

A GPS plot of an example run from the final version of the control system, is shown in figure 4.18 and a graph of the compass headings and target headings is shown in figure 4.19. These show that the MOOP was able to sail a beam reach course and return to the starting point without any major difficulties.

Figure 4.19: A graph showing the headings of a MOOP during the example course (the corresponding GPS trace can be seen in figure 4.18).

4.2.4 Pinta Experiments

4.2.4.1 2008 World Robotic Championships

Pinta was taken to the first World Robotic Sailing Championships (WRSC) held on Lake Neusiedl, Breitenbrunn, Austria between May 19th and 24th 2008. The first successful sail was achieved on the May 21st. May 20th and 21st had been set aside for a long range 48 hour endurance event in which the boats were supposed to sail a large triangular course. Pinta successfully sailed across the lake to the second waypoint but encountered problems turning towards the 3rd waypoint. At this point the chase boat crew noticed that the wind sensor had rotated out of its original alignment. This caused incorrect readings to be taken and every time the sail moved the sensor rotated further from its original position. Therefore the decision was taken to abandon the race and tow Pinta back to the start point. A bracket was then attached to the wind sensor to prevent rotation. This still provided the first test of the tiller pilot system and showed it worked well enough to maintain a straight course for several kilometres, figure 4.20 shows a GPS plot of the run and figure 4.21 shows the compass

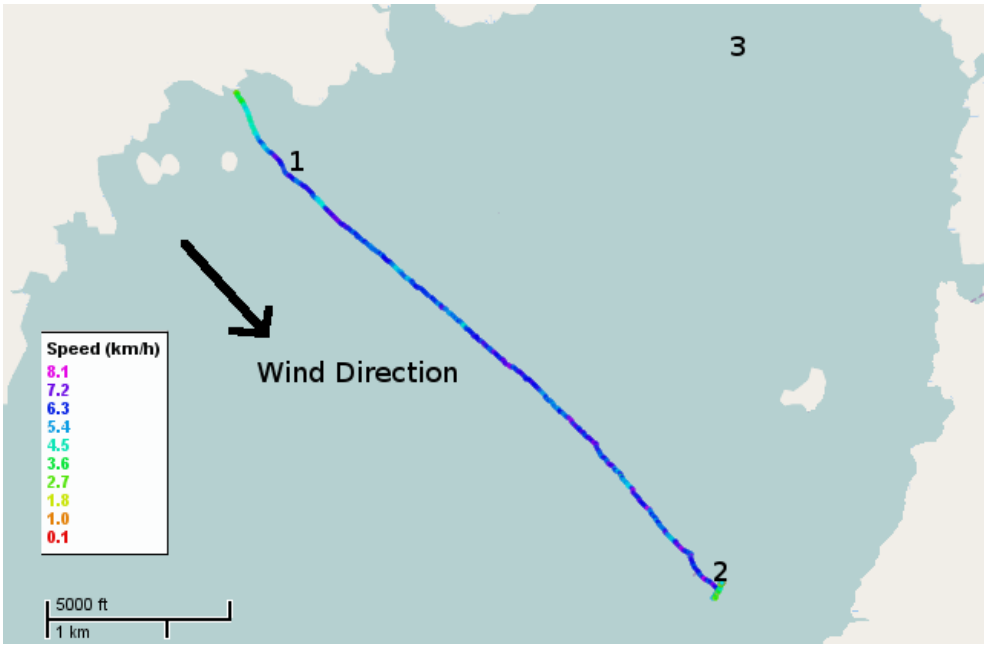

Figure 4.20: A map of Pinta's first test on May 21st 2008. Map courtesy of gpsvisualizer.com and OpenStreetMap.

heading during this run . The straighter portions between 2250 and 2500, 2800 and 3300, 3800 and 4050 and 4400 and 4800 seconds are where the tiller pilot was in control, the other portions with higher levels of oscillation are where the Gumstix was controlling the steering. The large variations in heading after 5400 seconds are due to Pinta missing waypoint 2 and attempting to turn around to reach it, because of the twisting wind sensor it was unable to reach the waypoint.

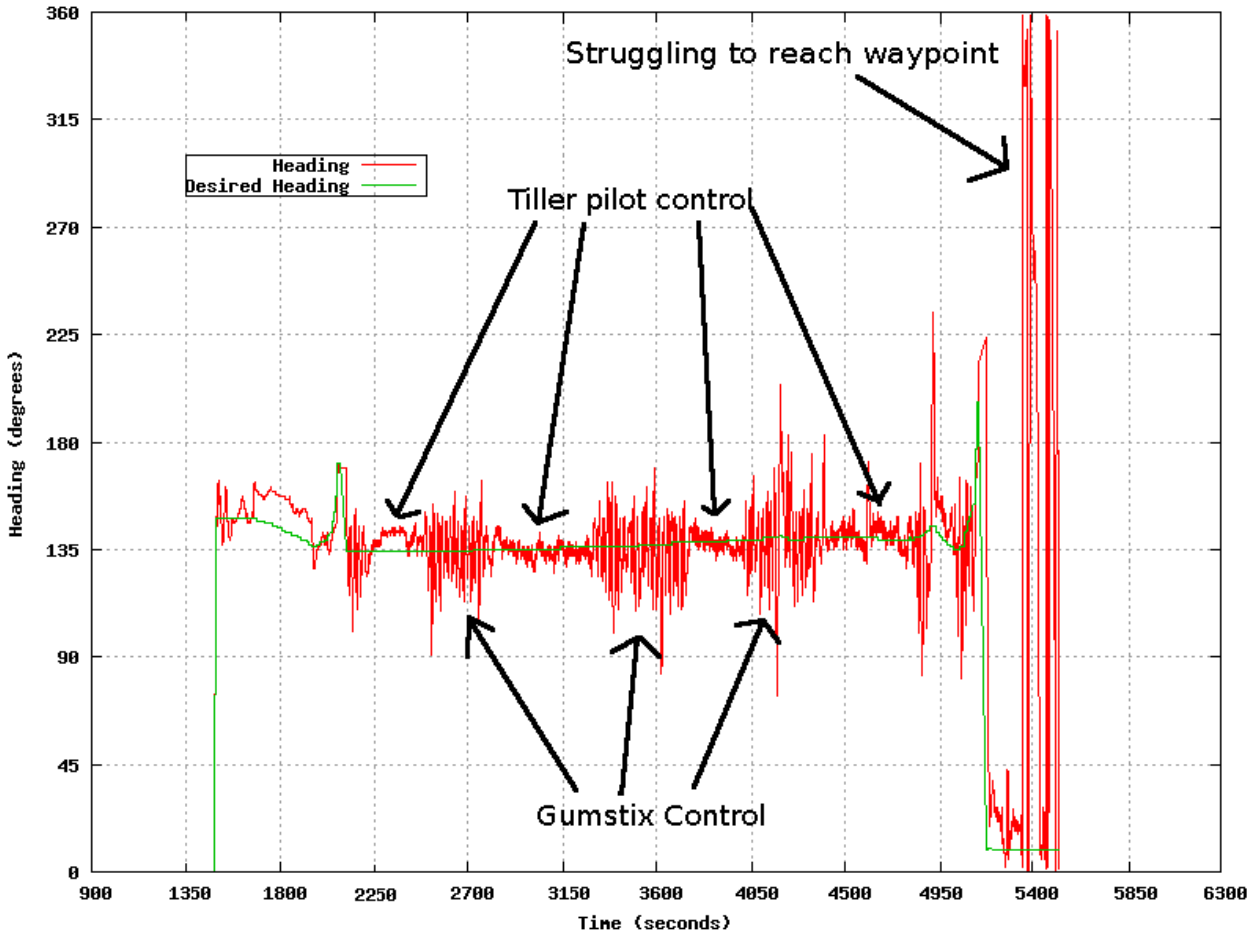

Figure 4.21: Pinta's compass heading during its first test run on May 21st 2008 (see figure 4.20 for corresponding GPS trace).

May 22nd was spent securing the wind sensor and making various refinements to the control system, a series of short tests were made. On May 23rd Pinta was entered into the first short race of the WRSC and initially sailed very well until a leak in the electronics compartment destroyed the GSM modem (which also doubled as the GPS). A graph of the compass heading can be seen in figure 4.24 and a map in figure 4.22. The circling around waypoint 2 in figure 4.22 was due to an attempt to keep Pinta the correct side of the buoy (in accordance with race rules) at waypoint 2 by actually specifying waypoint 2 as three separate waypoints, unfortunately this plan back fired when one of the waypoints was missed despite having passed the buoy correctly resulting in the behaviour shown in the map.

In figure 4.24 the flat green line representing the target heading after 8500 seconds is the point at which the GPS failed. Until the failure of the GPS Pinta had sailed well running down wind, beating up wind and reaching across the wind.

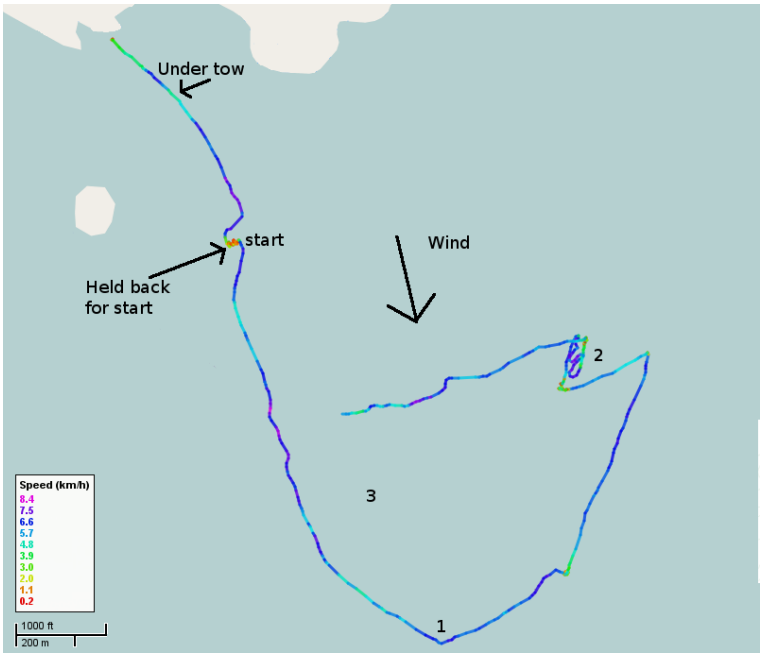

Figure 4.22: A map of Pinta's course during the first WRSC race on May 23rd 2008. Map courtesy of gpsvisualizer.com and OpenStreetMap.

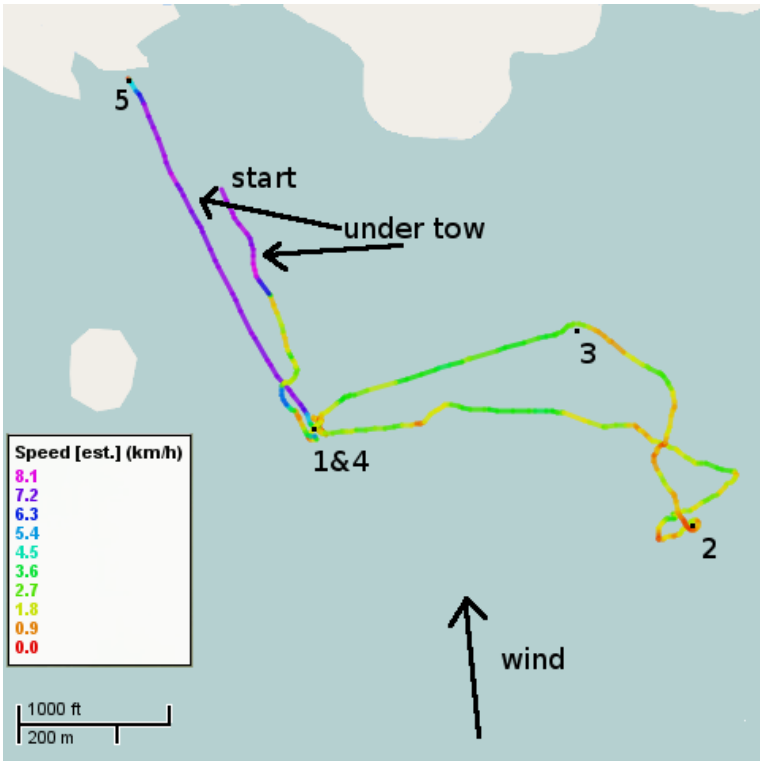

Figure 4.23: A map of Pinta's course during the final WRSC race on May 24th 2008. Map courtesy of gpsvisualizer.com and OpenStreetMap.

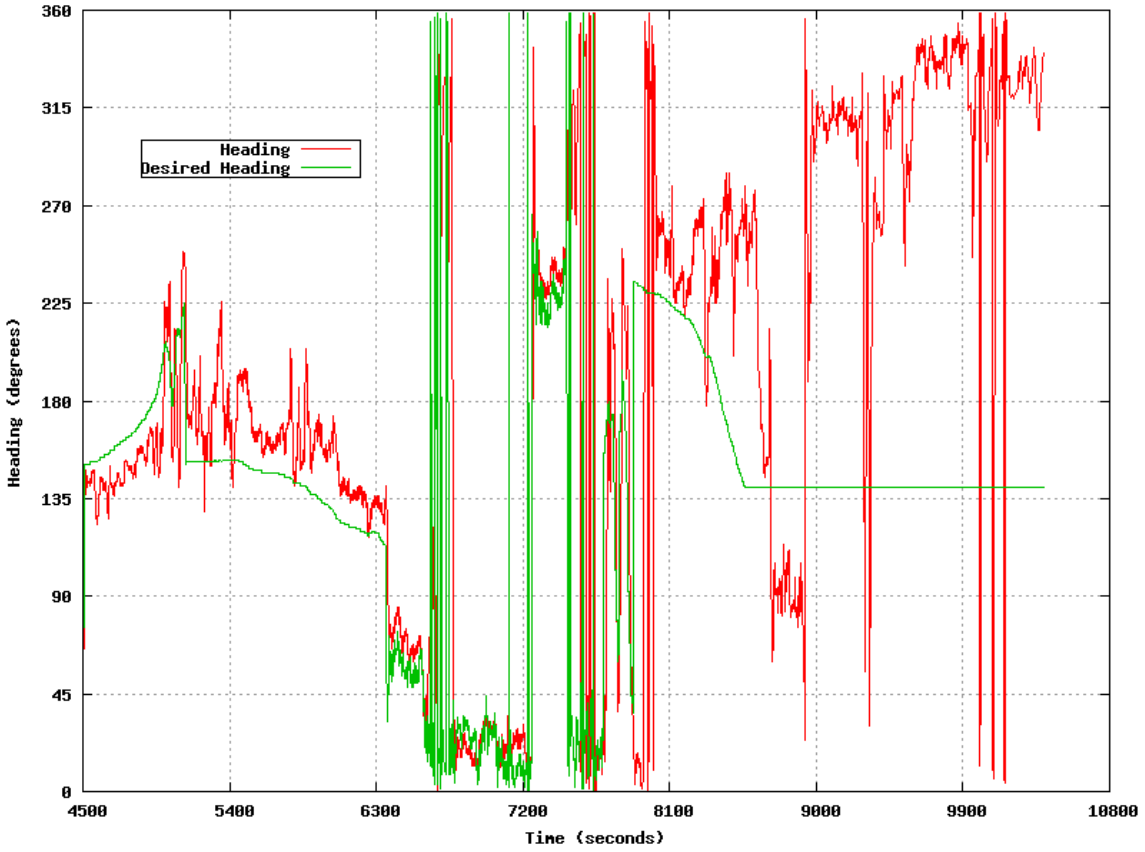

Figure 4.24: A graph of Pinta's compass headings and target heading during the first WRSC race on May 23rd 2008 (see figure 4.22 for corresponding GPS trace). The cause of the offset between the target and actual heading between 4500 and 6300 seconds is likely to be caused by the deadband set by the control system. .

At this point the GSM modem/GPS GM862 unit was replaced with a stand alone Garmin eTrex handheld GPS (the only spare GPS unit available at the time). By the evening of the 23rd Pinta was ready to sail again. The final race of the WRSC took place on the 24th and despite a lack of wind Pinta completed the entire course. A map and graph of compass headings are shown in figures 4.23 and 4.25.

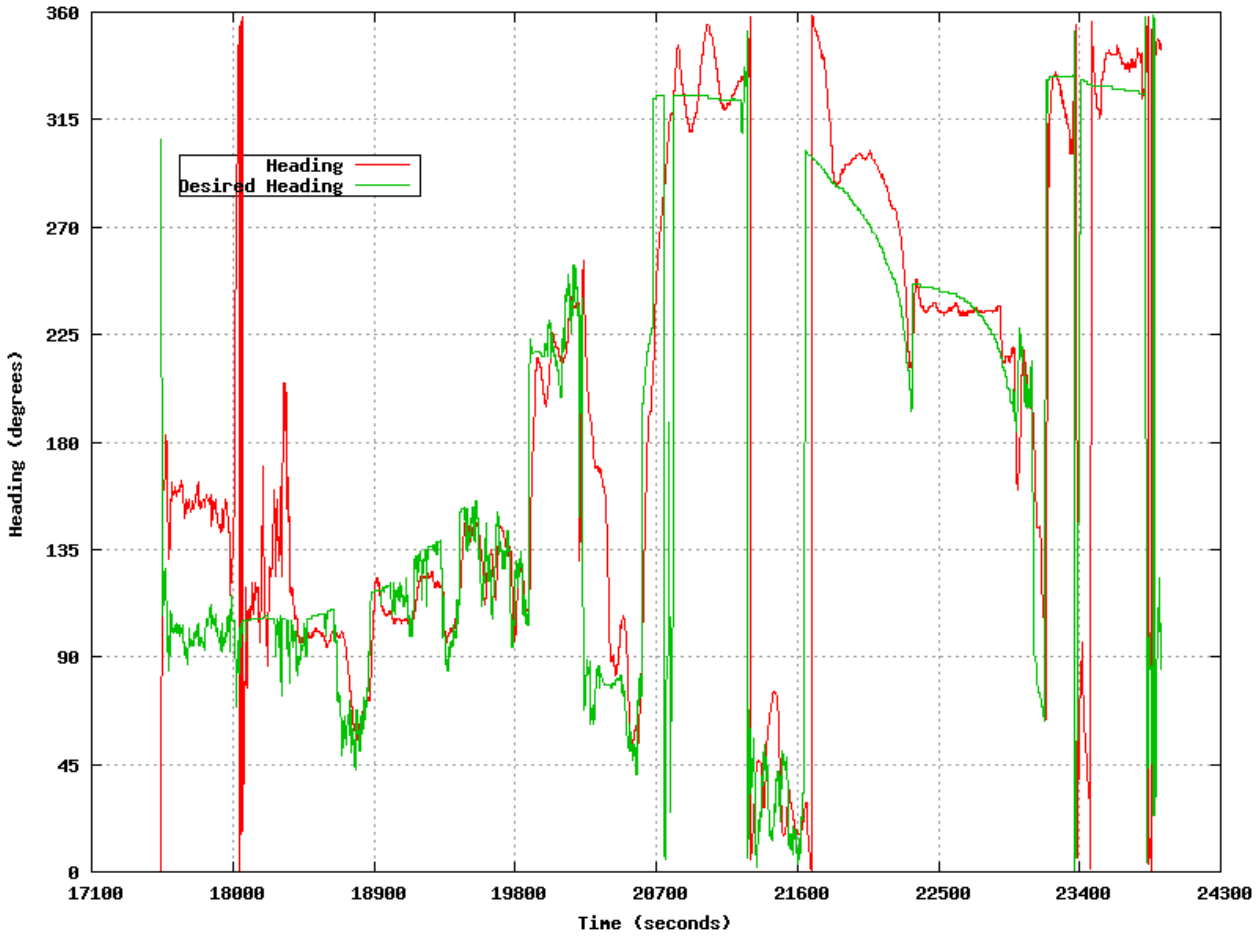

Figure 4.25: A graph of Pinta's heading during the final WRSC race on May 24th 2008 (the corresponding GPS trace can be seen in figure 4.23).

4.2.4.2 2010 Microtransat

Pinta's Microtransat attempt began at approximately 2pm GMT on September 11th 2010 from a point about 7km Northwest of Knightstown, Valentia, County Kerry, Ireland. The final target point was Martinique in the Caribbean, she was programmed with a series of waypoints to take her on a South Westerly course until 20 degrees North, at which point she would hopefully enter the trade winds and be able to turn east towards the Microtransat finish line (60 degrees West) and Martinique. The journey was expected to take around 2-3 months. Her first waypoint was approximately 28km South West of the start point. A west north west wind prevented her sailing directly to the waypoint and she instead sailed to the south west and came very close to land only two hours after launch. To add to the problems the waves were between two and three metres high with only a six or seven second period and the wind was between 17 and 20 knots. Figure 4.27 shows the course she took (from

Figure 4.26: Pinta departing for the Microtransat 2010.

hourly Iridium messages) until 3pm on September 13th.

By 3am on the 12th she had turned onto a north westerly course with the waypoint now sailable. Then for some unknown reason at around 5am she turned east , no longer indicating the waypoint as sailable. By 10am she completed a loop crossing over the previous days course. This pattern was then repeated again over the following 18 hours. At 3:43am an Iridium message indicated that the Gumstix had rebooted. There were no processes running that were supposed to reboot the Gumstix and it is believed this was a symptom of water entering the electronics compartment. Despite this Pinta continued to transmit telemetry messages and continued to sail south south east (and report that the waypoint was sailable despite sailing away from it). This trend continued until 3pm on the 13th, after which the messages from the Iridium modem ceased. After this time only position reports from the Spot were received.

The failure of the Iridium modem could have been due to a total failure of the main electronics system including the Iridium modem and Gumstix. If this were the case then all control of the sail would cease. It was possible that the tiller pilot was still working and would continue to try and follow the last target heading it was given (which Iridium data indicated to be 325 degrees). The next Spot message is shown at the bottom of figure 4.28 (which shows the entire route taken by Pinta from September 11th-29th) and indicates that Pinta had tacked and was now sailing towards the waypoint. This raised the hopes that maybe only the Iridium modem had failed and the control system was still running. These hopes were squashed when at 7am on the 14th it placed the boat 14km directly west of the waypoint. As the wind was blowing from the west the waypoint should have been directly sailable, even

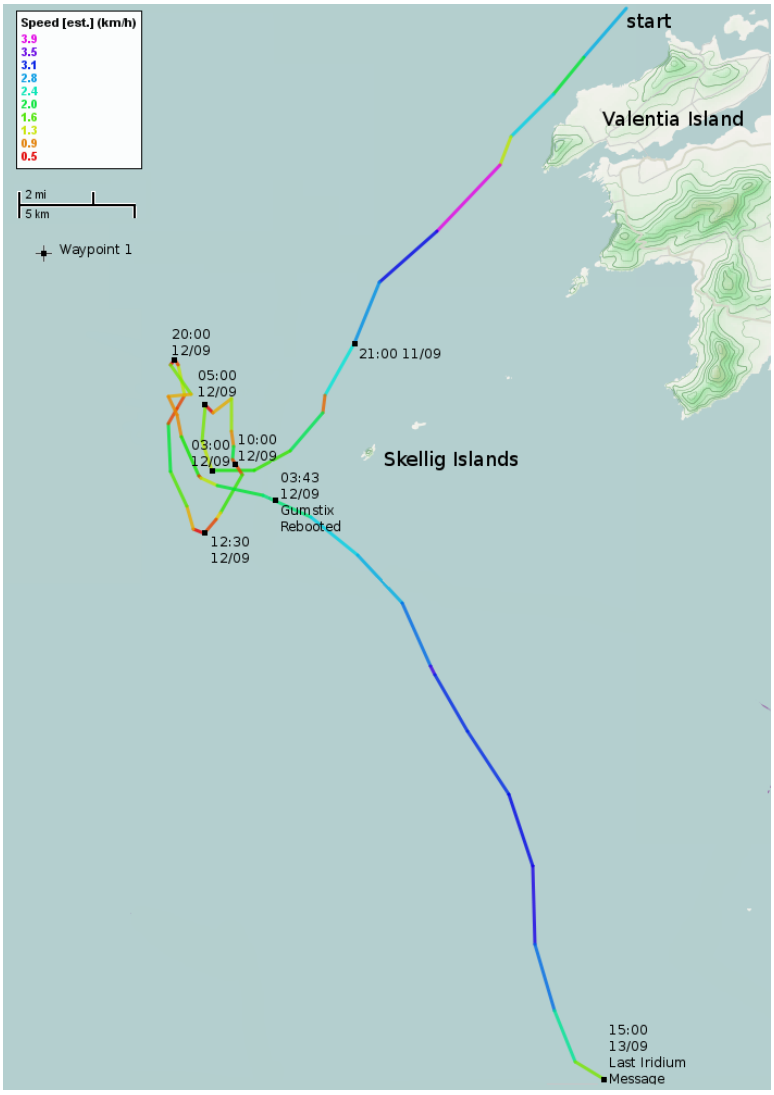

Figure 4.27: Pinta's path as generated from hourly Iridium telemetry messages, between September 11th and 13th 2010. Map courtesy of gpsvisualizer.com and OpenStreetMap.

if the rudder had failed just letting the sail out would have been enough to trigger a change of course in the right general direction. By 6am on the 16th Pinta had sailed nearly 150km to the north and was now about 50km west of the Aran islands which mark the entrance to Galway bay. This removed any hopes that the control system was still operational and it appeared that the sail was probably setup for a beam reach or broad reach with the rudder roughly centred. Pinta continued to follow an approximate beam reach course for a further 13 days and eventually ended up around 80km offshore to the west of Westport. In the final few days the Spot messages became increasingly infrequent, suggesting the possibility that the Spot may have become dislodged and no longer facing upwards or that the boat was spending significant amounts of time upside down.

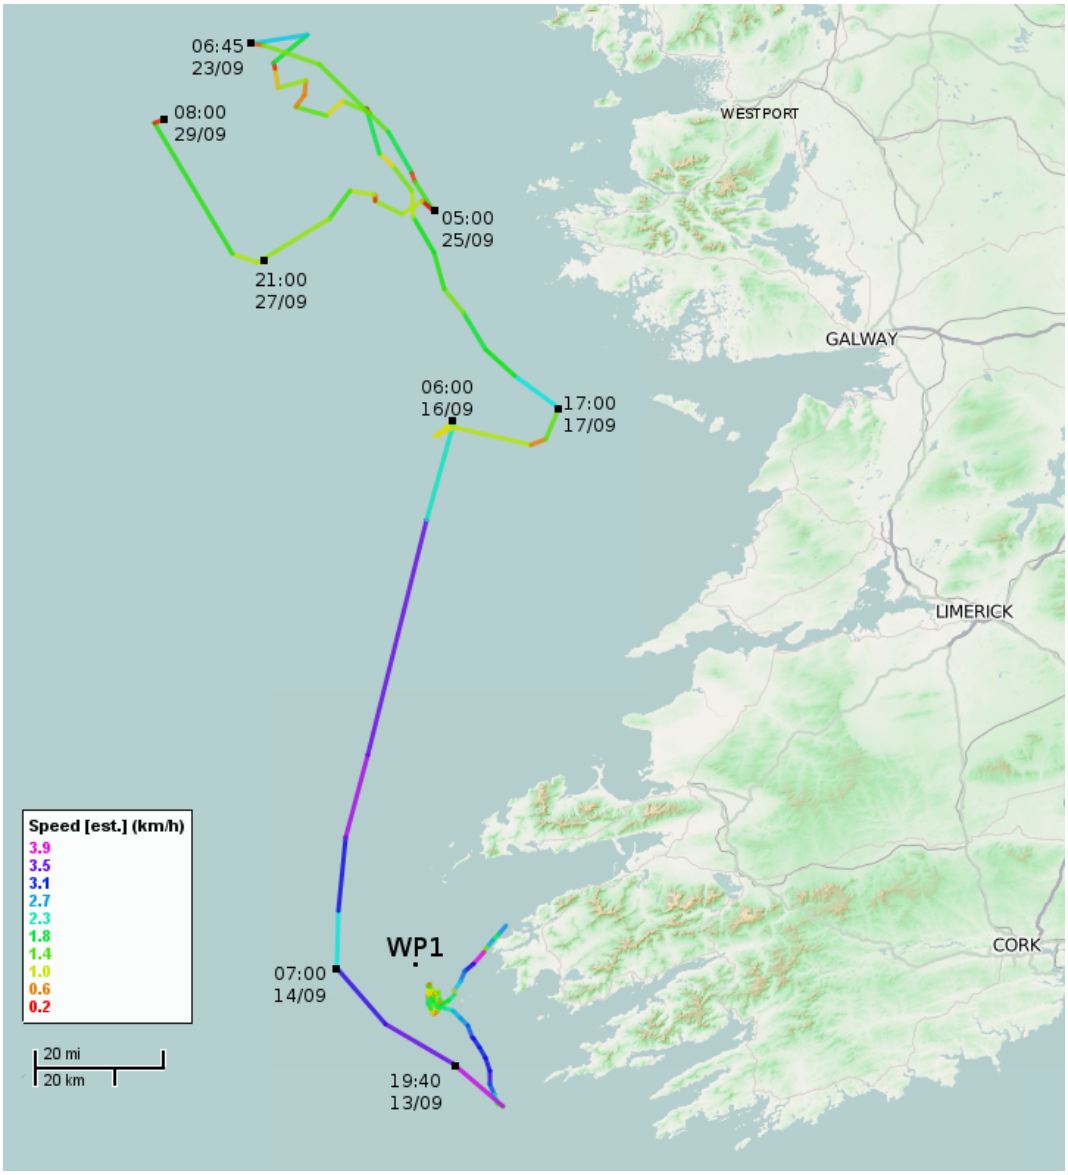

Figure 4.28: The whole of Pinta's journey between September 11th and 29th 2010. Map courtesy of gpsvisualizer.com and OpenStreetMap.

4.2.4.3 Analysis of Pinta's Power Consumption

The hourly telemetry sent via the Iridium modem included a battery voltage reading. This was obtained by running the battery output through a resistor bridge and then sampling it through an analogue to digital converter on the PIC microcontroller. Each reading was based on 10 sequential samples taken at the time the message was sent. It was possible for the rudder actuator to be moving while the sampling was undertaken, the extra load of actuator movement would most likely result in a temporary voltage reduction. During testing the readings were also shown to suffer from noise levels of approximately 1 volt.

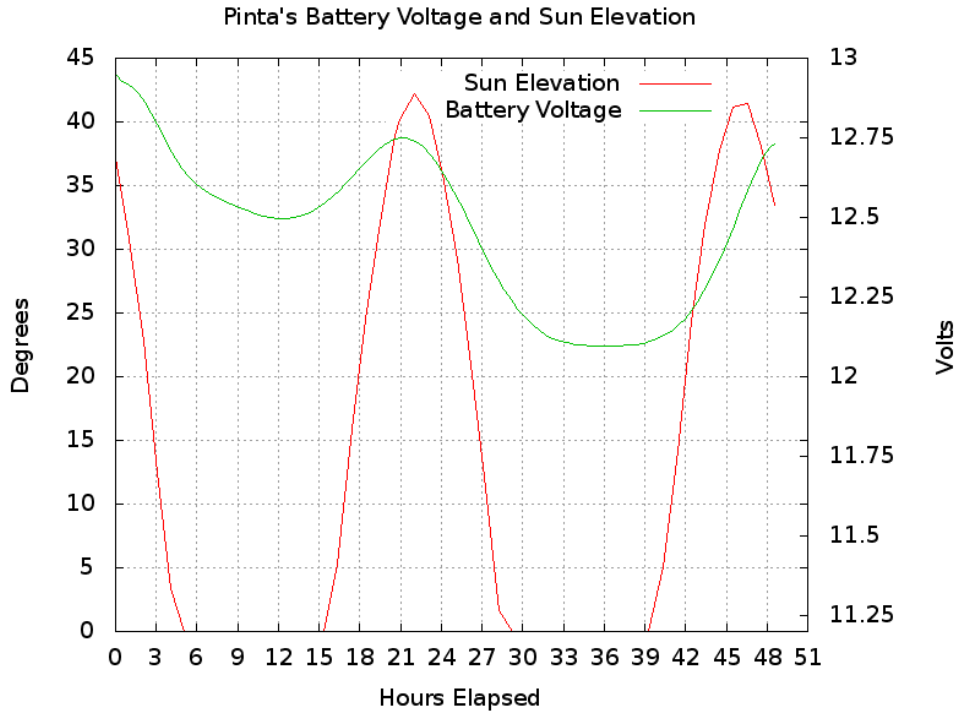

Figure 4.29: Pinta's hourly reports of battery voltage plotted against the sun elevation for the first 4 hours of its 2010 transatlantic attempt. The battery data has been smoothed to a Bezier curve.

Figure 4.29 shows a plot of battery voltage during the 49 hours that Iridium data was sent. A clear diurnal trend can be observed and the average battery level on the second day has dropped be nearly 0.5 volts. This could either be due to noise, lower sun light levels on the second day, depletion of the batteries or some combination of all three. Although, according to data from the Valentia weather observatory[20] there was only 0.2 hours of sunlight on September 13th, 0.9 hours on September 12th and 9.1 hours on the 11th. The heavy sea state may have also been taking its toll on the rudder actuator causing it to operate at very high duty cycles in comparison to those that would have been experienced in calmer seas. The separation of the rudder control into a separate system prevented any data from being gathered to confirm if this was the case or not. However despite what was believed to be a massive over engineering of Pinta's power systems this data suggests that even this over engineering may not have been sufficient for such conditions. Therefore, either more power efficient actuators are required, more attention must be paid to using power efficiently or more power needs to be generated in order to sustain a sailing robot at sea. Given the deck size of Pinta, it would be difficult to increase total power generation capabilities from the current 120 watts peak to much beyond 200 watts peak. Leaving the only viable option to reduce power

[20] `http://www.met.ie/climate/daily-data.asp` accessed 20/3/2011

consumption through better software management and more efficient actuators. Perhaps if some intelligent power management techniques had been applied, then Pinta would have been able to sustain its battery at a higher level. A neuro-endocrine controller (as discussed in chapter 2) could have been applied to suppress actuator movement and maintain the battery level within certain homeostatic bounds of viability.

4.2.4.4 Conclusions on Pinta's performance

Despite being built as a low cost "hack" Pinta proved to be a stable and capable sailing robot platform. It survived over 2 weeks off the Irish coast in waves ranging from 2 to 5 metres and winds typically in excess of 20 knots with gusts often reaching over 30 knots. By the time Spot messages stopped, she had sailed over 650km, although only 80km of that was with a working control system (which had obviously struggled to reach a waypoint less than 30km from the start point). The 49 hours of autonomous sailing are also believed to have exceeded any previous record (the previous best having been 24 hours, set by the INNOC team from Austria in the 2007 Microtransat in Aberystwyth).

It showed that only using 5 sail positions was sufficient to sail adequately enough to complete a triangular course in calm waters, although it is debatable if this success was transfered to the rough seas of the Atlantic. This could go a long way to reducing the requirements for the accuracy and repeatability (and potentially the complexity and cost) of sail actuators. However, the use of a fabric sail and a main sheet still felt very fragile in comparison to a wing sail. Offloading rudder control to a dedicated unit was initially shown to be a promising strategy and reduced the processing loads on the main computer but took away fine grain control over the rudder and reduced the ability of the control system to manage the power consumption of the rudder actuator. Ultimately Pinta was only built to sail in one race (the Microtransat) and was not built for continued use or mass production. Its power systems are heavily over-engineered, ignore many of the principles outlined in section 3.12 and do not need to support any scientific payload. Despite this over-engineering the limited evidence available suggests that this may have still been insufficient. This further strengthens the case for some kind of autonomous power management as discussed in section 3.12.

4.3 Discussion

4.3.1 Engineering of Sailing Robots

A number of conclusions regarding the engineering of both the hardware and software of sailing robots can be drawn from analysis of the work of others and from our own development of sailing robots. The first and perhaps most obvious conclusion is that the robot must be designed in a robust manner. It should be designed to be submerged temporarily, to be knocked around and to endure this for prolonged periods of time. The electronics system should be as simple as possible to minimise the number of potential points of failure. Where possible low cost off the shelf components should be utilised as they can enable faster repair times. Waterproof enclosures for electronics help in the event of a breach in the hull but can also trap condensation. The software systems must also be designed for robustness and should be subject to code reviews. Analysis tools which detect potential code faults (for example static analysers such as Lint [62]) should be used. Code must undergo extensive testing to eliminate issues such as memory leaks and parser bugs. The control software should be able to automatically restart itself and restore the previous state should a crash occur. Any string parsers should be robust, capable of handling invalid input caused by electrical noise or faulty devices and should include the checking of error correction information where appropriate (for example verifying the checksum of NMEA sentences as used by most GPS devices).

Second the systems should not be developed in isolation to each other but should be seen as a combined system. Decisions regarding hull and sail design will affect the requirements of the software in terms of the response rates required by the control system and the maximum durations that the boat can travel without any actuator movements. Faster and less stable hulls, will generally require faster response rates from their control systems which in turn will require faster moving actuators, higher frequency sensor data and faster microprocessors. Even in more stable designs the rudder typically requires a faster response rate than the sails (even if the sails are being used in co-operation with the rudder to improve steering). As the rudder control is usually in proportion to heading error, the compass must also be sampled at high frequencies. However, the GPS and wind sensor can be sampled at much lower rates.

Third, the hull should be designed to be passively stable and capable of continuing on its initial course once the sail and rudder are placed in the correct position. This reduces power consumption by reducing actuator movement and can allow for the complete shutdown of the control system when some distance from the next waypoint. The use of wind vane's to position sails as those used by Elkaim and Harbor Wing [39, 37, 38] or to maintain a rudder position with respect to wind [139] could achieve sail and rudder positioning without

requiring any electrical power (provided the wind direction does not change), although the mechanical complexity of such systems maybe an issue.

Finally, to enable rapid development the boat should be kept small and simple, it is best to scale up the design only after it has been perfected on a smaller boat. The deployment of larger boats, such as BeagleB and Pinta required a team of at least two (and often as many as four) people. Therefore, any testing required finding a day when everyone involved was free. Launching and recovering larger boats typically required towing them some distance off shore before beginning any tests and this represented a significant portion of the time involved. Smaller boats such as the MOOPs or ARC can be operated by one person, they can easily be chased with a canoe or small inflatable dinghy and they can be deployed and recovered in minutes not hours. They also tend to be easier and quicker (and cheaper) to repair should any damage occur. Observations at the various sailing robot competitions suggest that boats up to two metres in length can realistically be handled by one person, but larger boats require multiple people and proper chase boats.

4.3.2 The Case for Autonomous Power Management

The review of existing work in chapter 3 shows that relatively little attention has been paid to autonomously controlling power consumption in sailing robots to date. Many designers have simply opted for over-engineering of their power systems which face considerable financial, complexity and scalability problems. However, in a short term one off production this may be the simplest strategy. If sailing robots are to be used for long term missions, such as ocean monitoring, then it is likely to be desirable that as little power as possible is used to actually sail the robot and as much as possible is used to run scientific equipment. In order for this goal to be achieved, far greater attention needs to be given to autonomous control of power consumption. Data gathered by Pinta also suggests that previous estimates of the required power consumption are too low.

At this stage having developed a neural network controller for the MOOPs the neuro-endocrine controller strategy defined by Neal, Timmis and Mendao [76, 92] in the previous chapter, would appear to offer the potential to create a system which varies the magnitude and/or frequency of actuator movements in proportion to the available energy level to control power consumption. This system would allow for an almost continuously variable approach to controlling power consumption rather than having a few very distinct power saving modes such as suggested by Blair [18]. Such a system could be further integrated with artificial hormone signals, that represent both past and predicted future solar energy inputs, to provide a system which can modify its behaviour not just around its current state, but around past and future energy levels. This creates a method for exploiting the opportunity to

perform additional non-essential tasks such as improving sailing performance, transmitting extra telemetry data or gathering extra scientific data when there is excess power available. Through these mechanisms a neuro-endocrine inspired architecture also has the potential to adapt to problems across varying timescales. On the second or sub-second scale a neural network could be moving the rudder and sail actuators. On a minute by minute basis the duty cycle or magnitude or actuator movements can be adjusted in response to battery levels. Across the course of a day adjustments can be made in response to solar panel levels and seasonal adjustments to solar power levels can also be made. Finally at a "life long" level adjustments could be made to changes in actuator behaviour due to wear and tear or biofouling of control surfaces and solar panels.

Chapter 5

Neuro-Endocrine Control of Sailing Robots

5.1 Introduction

This chapter outlines the design of an artificial neuro-endocrine controller architecture suitable for controlling a sailing robot and manging its power consumption. It makes some modifications to the methods proposed by Neal and Timmis (2004) [92] and Mendao (2004) [76] which were discussed in section 2.6.3.7. These modifications are required to reduce the size of the parameter space and to create an architecture which is suitable for application to a sailing robot. This chapter also discusses the implications of which neurons in a neural network should have endocrine modulation applied to them. Finally this chapter discusses some other behaviours which it might be possible to create in a sailing robot using neuro-endocrine control and discusses a potential approach to implement collision avoidance.

5.1.1 Advantages of Neuro-Endocrine Controllers

As discussed in the previous chapters an artificial neuro-endocrine controller offers the potential to take a neural network control system and allow it to be modulated by a number of hormones to adjust its behaviour. A neural network can be trained upon data generated from a variety of different controllers. Once encoded, endocrine modulation can be easily applied to them regardless of the original type of controller. In many situations, for example attempting to add modulation to a rule based controller this can simplify the process of modulating the controller output. Neural networks also offer the ability to generalise a solution, allowing them to be trained on incomplete or low resolution data sets (for example data of a human remote controlling a robot).

The addition of an artificial endocrine controller allows a robot to adjust its behaviour in response to both internal and external stimuli on varying timescales. These behaviour transitions can take place gradually with the possibility of two (or more) behaviours being partially exhibited at any time. This is very much in contrast to many other action selection mechanisms traditionally employed in robotics, which follow "winner takes all" approaches, in which only a single behaviour is ever selected at any given time. This approach can help to ensure that any action has truly completed before another is selected, as only a single action is ever selected there is no worry about two or more selected actions competing. However, a winner takes all approach, by definition limits us to only performing a limited number of actions with little chance of a smooth transition between them. For example, the three power save modes described by Blair [18] in the chapter 3, the robot must explicitly be in one of these modes. Should none of the modes entirely suit the current situation then there is no in-between setting. Although it could be argued that if switching between different modes occurred fast enough it might give the illusion of an intermediate mode. In practice this may not be practical if any kind of penalty is suffered from switching modes, for example, if a device must be powered on or off (taking several seconds) when changing mode. A neuro-endocrine controller, on the other hand, has no distinct modes of operations and can change gradually between different actions. In a robotic system the potential candidates to vary are the frequency and/or magnitude of the actuator movements or the sensor sampling rates. These changes in frequency/magnitude could be in response to one or more environmental stimuli, such as battery state, sun light levels or actuator temperatures.

5.1.2 Disadvantages of Neuro-Endocrine Controllers

There are also a number of potential problems with using a neuro-endocrine controller. A number of these are problems which emerge from stand alone neural networks and others result from combination of the neural network and endocrine controller. As discussed in section 2.6.1.5 neural networks do not perform online learning and can suffer from over fitting training data and loosing their ability to generalise. Artificial endocrine controllers do not yet have a standard method for determining how sensitive a receptor should be to a hormone, although some work has been undertaken to perform online learning of hormone sensitivities [136]. There are also several approaches to integrating neural networks and artificial endocrine controllers. Neal and Timmis (2003) [92] and Mendao (2007) [78] had hormones interact with every neuron in a neural network, while Henley and Barnes (2004) [56] applied the hormone only to selected neurons. Given that artificial endocrine and neuro-endocrine controllers are a relatively new development, there are still likely to be many as yet undiscovered issues regarding their optimal use.

5.2 A neuro-endocrine Architecture for a Sailing Robot

This section outlines modifications to existing neuro-endocrine strategies which make them suitable for controlling a sailing robot. Previous approaches to produce an artificial endocrine system have been attempted by Neal and Timmis (2004) [92] and Mendao (2004) [76]. An overview of their work was described in section 2.6.3.7. Their work modified a traditional multi-layered perceptron and added an artificial hormone gland which released a hormone that modified the behaviour of the neural network. When applying it to a sailing robot some problems emerge due to the large number of parameters used. If we wish to test a neuro-endocrine controller across a range of parameter values on a real robot then reducing the number of parameters to a minimum is of key importance to allow for these experiments to be carried out in a reasonable amount of time. Additionally, with respect to power management in a sailing robot, power consumption can be modulated by changing either the number of actuator movements or the magnitude of those movements.

5.2.1 Combining Receptor Matching Distance with Hormone Sensitivity

In Neal and Timmis' architecture, the receptor matching distance (see equation 2.5 in section 2.6.3.7), hormone concentration and hormone sensitivity together control the level of behavioural change which will be seen in the neural network in the presence of a given hormone. As the receptor matching distance and hormone sensitivity both provide unchanging values they could be combined into a single number. By doing this we reduce the size of our parameter space. Therefore a simplified formula for calculating the output of a single perceptron becomes:

$$\sum_{i=0,j=0}^{n} w_i \cdot x_i \cdot C_j \cdot S_j \{w_i | w_i \in \mathbb{R}\}\{x_i | x_i \in \mathbb{R}\}\{C_j | C_j \in \mathbb{R}\}\{S_j | S_j \in \mathbb{R}\} \tag{5.1}$$

Where w represents the weights, x the input values, C the hormone concentration and S the hormone sensitivity. i is the current input from a set of inputs and j is the current hormone from a set of hormones. w,x,C and S are all real numbers. However this encounters some shortcomings when the hormone concentration (C) is zero as the end result will be zero. This would imply that if no hormone is present then we wish the neural network to have no output. A more useful scenario would be that if the hormone concentration were zero then we would want the network to behave normally, as if no hormonal modulation was being applied. Therefore, the formula was modified so that one was added to the hormone sensitivity and concentration as shown in equation 5.2. This ensured that in the event of there

being zero hormone the output of the perceptron remained unchanged. Hormone sensitivities are restricted to between zero and one, when set to zero a hormone will be ignored regardless of it its concentration level. When set to one it will have its maximum level of change over the neural network.

$$\sum_{i=0,j=0}^{n} w_i \cdot x_i(1 + C_j \cdot S_j)\{w_i|w_i \in \mathbb{R}\}\{x_i|x_i \in \mathbb{R}\}\{C_j|C_j \in \mathbb{R}\}\{S_j|S_j \in \mathbb{R}\} \tag{5.2}$$

5.2.2 Hormone Decay Function

In the Neal and Timmis' system the hormone is released and decayed in a geometric fashion (see equations 2.6 and 2.7 in section 2.6.3.7). To simplify this process and further reduce the number of variables in the system the hormone decay function was redesigned to incorporate a single variable release and decay rate. This lead to the formula:

$$C_{t+1} = C_t - r(C_t - q) \tag{5.3}$$

Where C_t represents the current hormone concentration, r represents the response rate and q represents the amount of hormone being created (in response to external stimuli). The response rate r determines how quickly the hormone concentration will change (both in release and decay) and is similar in function to the α and β value in equations 2.6 and 2.7 shown in section 2.6.3.7. There are two potential classes of hormone in this system: excitory and inhibitory hormones. Excitory hormones will increase the values of the weights in the neural network and thus lead to greater levels of activity in the network. Inhibitory hormones will reduce the weights in the neural network leading to lower levels of activity and eventually causing the network to output zeros. The class of hormone is determined by the value of q, if positive then it will lead to excitory hormone being produced, if negative it will lead to inhibitory hormones. This allows any hormone to transform from being inhibitory to excitory (or vice-versa). In earlier versions of the architecture this was instead controlled by the hormone sensitivity (S) used in formula 5.2, but it was decided that allowing a hormone to transition from being excitory or inhibitory depending on external conditions would result in a more flexible system.

A simple test of the hormone decay function (equation 5.3) is shown in figure 5.1. It shows the level of hormone concentration (C) varying over time with values of r of 0.1 and 0.05. At iteration 500 the gland output (q) rises from zero to 0.5 and remains at 0.5 until iteration 1500. During this time the hormone concentration rises more gradually. When r is 0.1 the hormone concentration has matched the gland output by iteration 1250, whereas when it is 0.05 the hormone concentration doesn't quite reach 0.5 before the gland stops producing

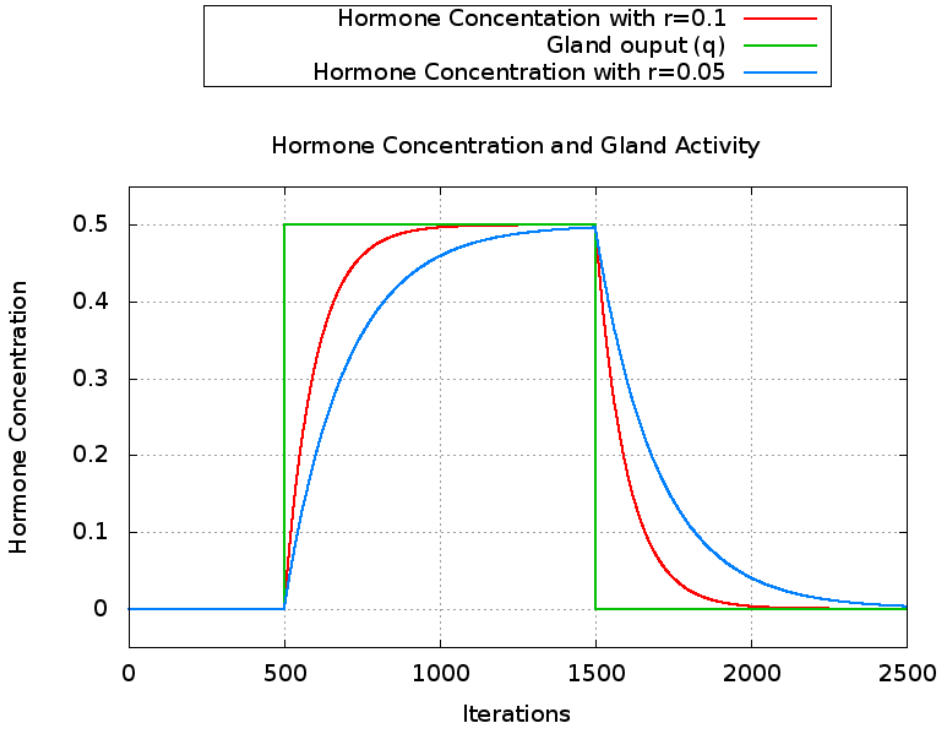

Figure 5.1: A graph showing the effects on hormone concentration over time of different values for r from equation 5.3. The red line represents an r value of 0.1 and responds to changes faster than the blue line ($r = 0.05$).

hormone at iteration 1500. After iteration 1500 the gland stops producing hormone and the hormone concentration begins to fall.

By using smaller r values the hormone will have a greater inertia and will continue to act for sometime after the gland has stopped producing it. This could be of use to create noise tolerance, when the inputs to the gland are noisy. Setting r too low could cause potential problems as the hormone may continue to act for too long after it is required. Setting r to a higher value will mean that the hormone level will mirror that of the input to the gland. This will remove noise tolerance but decrease response times to changes in the gland inputs.

5.2.3 Implementing and Testing the Neural Network

An example neural network controller was created for a simulated sailing robot. This was based upon the MOOP sailing robot described in section 4.1.5. The neural network was trained using a custom written Java implementation of the back-propagation algorithm described in section 2.6.1.2. The actual execution of the network is performed using the NNF open source neural network library[1] written in C, modifications were applied to allow hor-

[1] `http://nnf.sourceforge.net/nnf1/` accessed 19/03/2011

monal modulation using the formula described in section 5.2.1 and 5.2.2. NNF was selected as a lightweight, well tested and mature C implementation of a neural network, which would be suitable for use on embedded systems such as the PIC Microcontroller or Gumstix Single Board Computer used in the MOOPs and BeagleB.

Neural networks were created for setting rudder and sail positions, initially they featured a single input of the heading error (difference between target and current heading) for the rudder and wind direction for the sail. The output of the network was the new actuator position. It was later realised that the effect of inhibitory hormonal modulation would be to eventually move the actuator to one end of its range of motion, while an excitory hormone would move it the opposite way. This could result for example, in full inhibition moving the rudder completely to the right and full excitation moving it completely to the left. This would not result in a viable control strategy, therefore the neural network was redesigned so that the output value became the change in actuator position from the current position. This would result in inhibitory hormones reducing the range of motion, while excitory ones increased it. This also necessitated changing the inputs so that both the current actuator position and the heading error are supplied as inputs to the neural network.

The MOOP sailing robot encodes all actuators on integer positions between -5 and +5 (as discussed in section 4.1.5). This required the neural network outputs to be of the range -10 to +10. As the NNF library only allowed operation on value ranges between 0 and 1.0 a linear remapping of input and output values took place before and after executing a neural network operation. Rudder training data was generated through a proportional controller, the training algorithm is shown in algorithm 5.1. The sail training algorithm was based on a series of rules and is shown in algorithm 5.2. Both of these control algorithms were based upon controllers previously used in ARC, BeagleB and the MOOPs.

Algorithm 5.1 A rudder controller for a simulated sailing robot, which can be used to train a neural network.

```
pgain=0.1 //the proportional gain setting
/*loop through possible heading errors, use a jump of 36 to reduce
the size of training data*/
for heading_err = -180 ; heading_err <= 180 ; heading_err = heading_err + 36
    //loop through the range of rudder positions


    for rudder_pos = -5 ; rudder_pos<=5 ; rudder_pos++
         new_rudder_pos = (int)(heading_err * pgain)
         //stop positions exceeding the real maximum
         if new_rudder_pos < -5 then
             new_rudder_pos = -5
         else if new_rudder_pos > 5 then
             new_rudder_pos = 5
         rudder_change = new_rudder_pos - rudder_pos
```

Algorithm 5.2 Controller for sail setting on the simulated robot.

```
for relwind = 0 ; relwind <= 360 ; relwind = relwind + 5


    for curr_sail_pos = -5 ; curr_sail_pos <= 5 ; cur_sail_pos++
        if relwind < 180 then
            if relwind < 70 then new_sail_pos = 1
            else if relwind < 80 then new_sail_pos = 2
            else if relwind < 90 then new_sail_pos = 3
            else if relwind < 110 then new_sail_pos = 4
            else new_sail_pos = 5
        else
            if relwind >= 290 then new_sail_pos = -1
            else if relwind >= 280 then new_sail_pos = -2
            else if relwind >= 270 then new_sail_pos = -3
            else if relwind >= 250 then new_sail_pos = -4
            else new_sail_pos = -5
        /*calculate the difference between the desired
        position and current position*/
        sail_diff=new_sail_pos-curr_sail_pos
```

The relative wind directions for each sail positions used by the sail training algorithm are illustrated in figure 5.2. For both the sail and rudder 18 degree increments are used between each position, these are shown in table 5.1 and figure 5.3.

Position	Angle
-5	270
-4	288
-3	306
-2	324
-1	342
0	0
1	18
2	36
3	54
4	72
5	90

Table 5.1: A comparison of actual sail and rudder angles and their position numbers for the simulated boat.

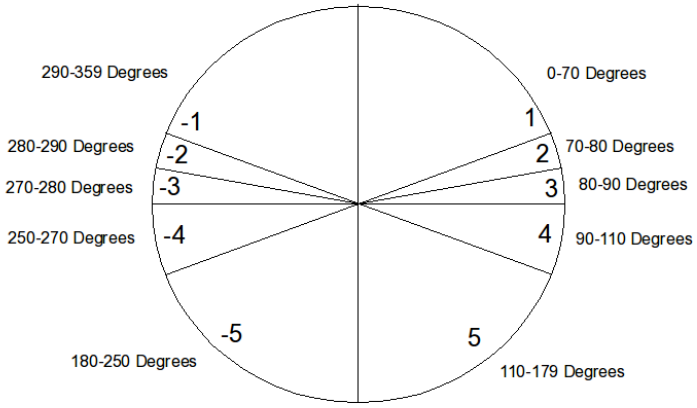

Figure 5.2: A digram comparing sail positions and relative wind directions for the simulated boat.

A variety of neural network configurations were tested with between two and eight nodes in the hidden layer. It was found that with fewer than eight nodes (for both the rudder and sail) the back-propagation algorithm struggled to train the network and became stuck in local minima. In all cases the network was "fully connected", meaning that each node in each layer is connected to all nodes in the neighbouring layers. Each neuron in the hidden and output layer also features a bias input, who's value is fixed at 1 but which has a weight which can be modified during back-propagation. These are present because they help to ensure back-propagation training algorithm will converge on a solution should all inputs be zero.

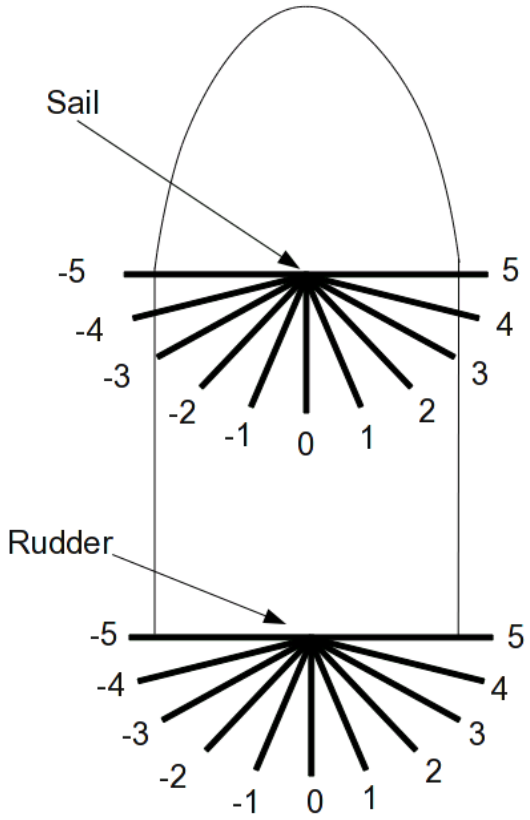

Figure 5.3: A diagram illustrating the sail and rudder positions for the simulated boat.

5.2.3.1 Where to apply modulation?

When applying hormonal modulation to a three (or more) layer multi-layer perceptron there are several different approaches that could be taken. Modulation could be applied to every connection within a neural network or just to a limited set of connections. Neal and Timmis (2003) [92] and Mendao (2007) [78] applied modulation to the entire network, while Henley and Barnes (2004) [56] describe difficulties in this approach and instead only applied the hormone to selected connections in their neural network.

To test the effect of various strategies a test network was built, this network is based upon the rudder neural network described in the previous section. For the purpose of these tests the heading error input was error set to 50 degrees and current actuator position was set to zero. The hormone concentration was then varied between zero and -1.0, as hormone concentration approaches -1.0 the network will become totally suppressed and should always

output zero change in actuator position. With a concentration of zero the network should behave normally and as we are 50 degrees off course the actuator will be placed in position 5 to generate the maximum amount of course correction.

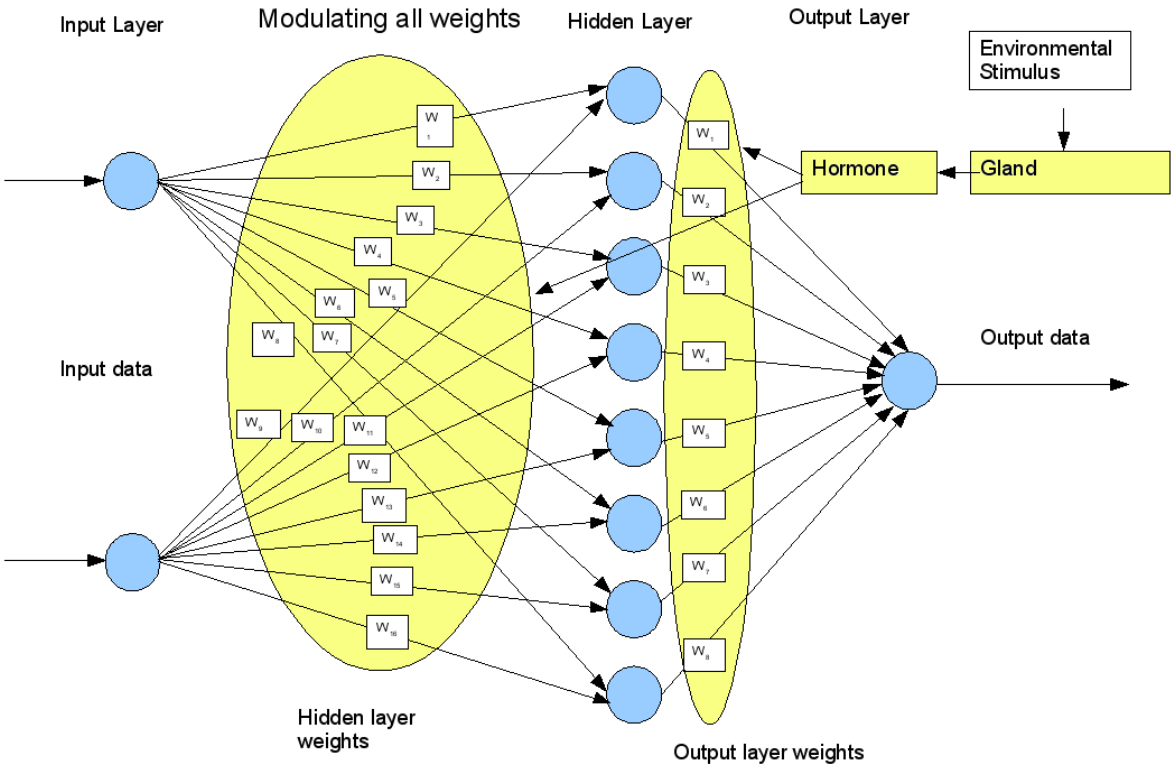

Figure 5.4: A diagram of a neural network with hormonal modulation applied to the hidden and output layers.

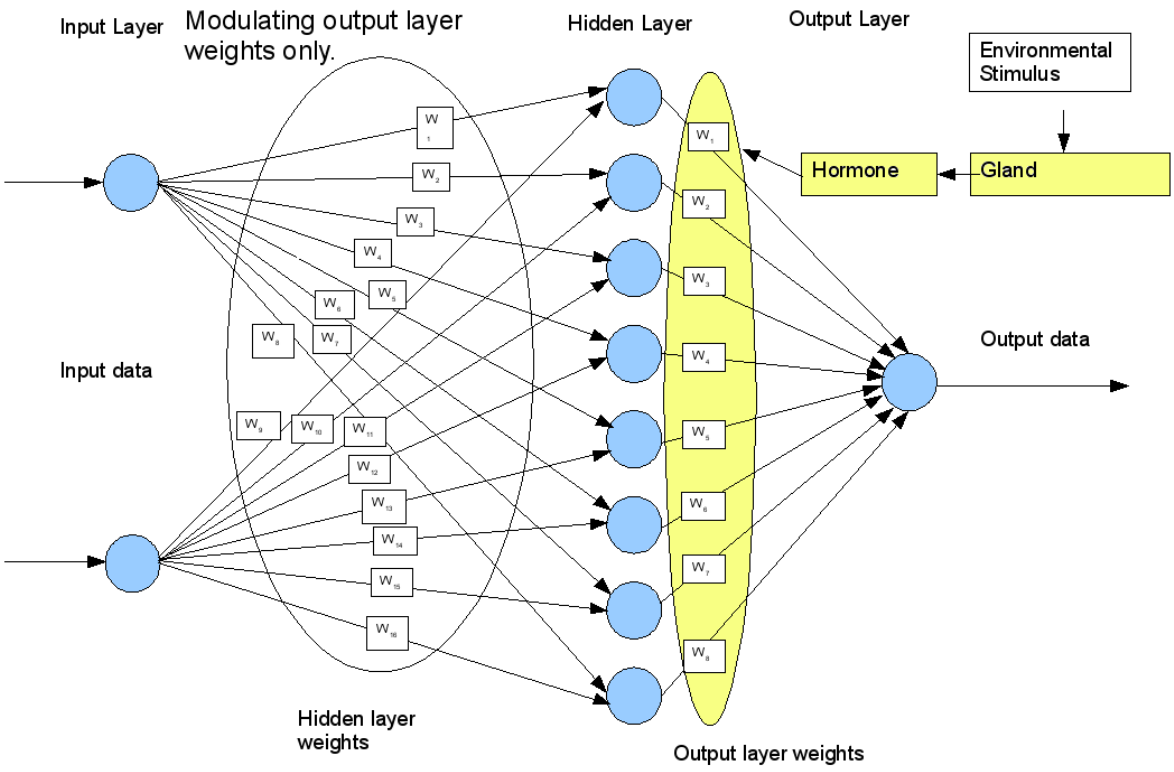

Figure 5.5: A diagram of a neural network with hormonal modulation only applied to the output layer.

Two different strategies were tested on the example network. The first was to apply hormonal modulation to the weights of all connections (input layer to hidden layer and hidden layer to output layer), a diagram of this network is shown in figure 5.4. The second strategy was to only apply it to the connections between the hidden and output layers as shown in figure 5.5. Figure 5.6 shows the output of the example network (actually the output rescaled and added to the current actuator position) for both modulation strategies. We can see that only applying modulation to the output layer produces a linear response while applying to all layers produces a non linear (although not dramatically dis-similar) response. Figure 5.7 shows the output of the network when the values are rounded as only whole numbers that actually make sense as actuator positions. As modulating only the output layer generates a linear response it will be easier to predict the effects of any modulation and it will require less computation time. Therefore only output layer modulation will be implemented in the final architecture used for any future experiments.

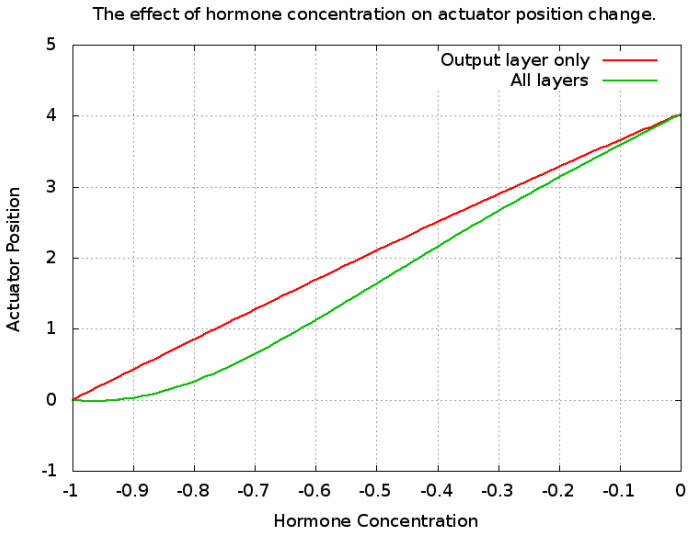

Figure 5.6: A graph showing the differing response of a hormonally modulated neural network depending on where modulation is applied. The X axis represents a changing level of hormone concentration and the Y axis the change in actuator position calculated by the neural network. The network has been given an input which should result in the actuator moving 4 positions, as is the case when no modulation is applied (hormone concentration=0). When hormone concentration is -1 the network is completely suppressed and no movement takes place. By applying modulation to only the output layer the effect of varying hormone concentration produces a linear change to the network output (the red line), while applying to all layers produces a non-linear change (the green line).

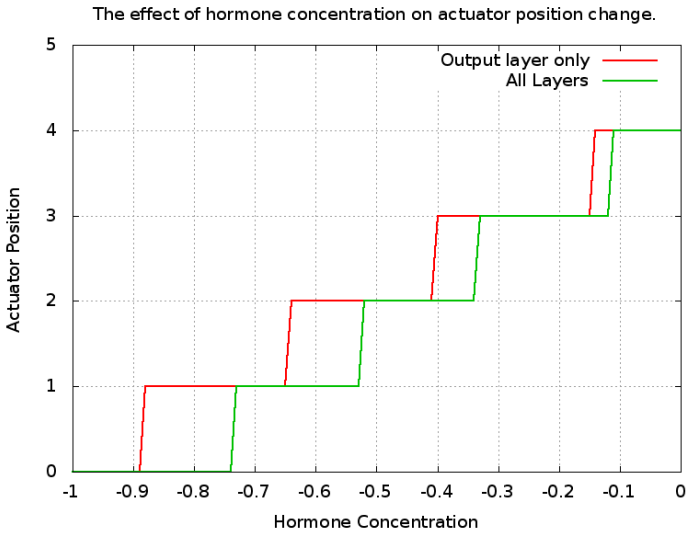

Figure 5.7: A graph showing the outputs of figure 5.6 rounded to the nearest whole number. This is because actuator positions must be a whole number.

5.2.3.2 Duty Cycle Modulation

We now have a framework in which we can modulate the behaviour of a neural network in response to one or more hormones. When the outputs of these neural networks represent the movements of a robot's actuator, then our hormonal modulation will effectively control the magnitude of actuator movements. What this method does not allow us to do is to (at least easily) vary the duty cycle of an actuator. Traditionally this is probably the more common method for controlling power consumption of actuators or sensors. Duty cycle modulation could possibly be achieved in the current architecture through a rapidly oscillating hormone which alternates between total inhibiting and zero inhibition of actuator movement. However, given that we can now control an actuator's power consumption by varying its magnitude, it is questionable as to whether there is any need to vary its duty cycle as well, since both methods can control an actuators power consumption. It is also not immediately obvious whether reducing actuator duty cycle is more or less detrimental to the control of a robot, than reducing the magnitude of the actuator's movement.

5.3 Managing Additional Behaviours with a Neuro-Endocrine Controller

For a robot such as a sailing robot, which must operate for a long period of time without any human contact, we can establish three essential objectives that the robot should attempt to achieve.

The first objective is that the robot must perform the mission that it has been sent to perform. After all, this is the whole reason that the user of the robot will have paid for the robot and deployed it into the target environment. This is likely to involve either remaining within a given area for a period of time or travelling between a set of waypoints. Typically it is envisaged that a sailing robot will be used to perform oceanographic monitoring, it will be replacing the role of either a buoy moored at a fixed location or of a survey ship sailing a transect between two (or more) locations. If the robot cannot continue to perform these tasks for a reasonable amount of time then there is little point in it being there. Therefore, any control system must ensure that the robot can continue to perform its mission for as long as possible. It is likely that some kind of sensing equipment will need to be run periodically and that in the event of insufficient power being available this need could be temporarily sacrificed in order to keep the robot's essential systems running.

The second objective is that the robot should use energy in a sustainable manner in order to remain operational. This may require the robot to adjust its behaviour to reduce power

consumption when energy is scarce. This might, at times require the robot's mission to be temporarily aborted in order to reduce power consumption. As this creates a conflict with the first goal the control system would be required to produce a sensible compromise. The compromise should still attempt the mission when possible but not at the expense of the robot completely running out of energy reserves. The end user of the robot may also wish to have the ability to configure the threshold point at which the mission is temporarily aborted. It is also desirable that the control system is able to exploit the opportunity of abundant energy when it is available. The amount of energy which the robot can spare is likely to vary from day to day. This depends upon how much energy is used for movement (which may be affected by environmental factors such as the weather), the amount used to run any mission equipment, such as scientific sensors and the amount of energy which can gather from the environment (e.g. from photo-voltaic solar panels). Although electrical power derived from photo-voltaic solar panels can be stored in batteries, it is more efficient to use this power directly from the solar panel as the total capacity of the batteries is limited and there are losses incurred when storing electricity in batteries. Therefore, it might be sensible for our robot to perform optional tasks or increase actuator movements to sail more accurately on a sunny day.

The third objective is that the robot should avoid danger. This might be achieved by steering clear of dangerous areas (e.g. the coastline), man made obstacles (e.g. other vehicles) or possibly even other robots. Depending on the exact wishes of the robot operator, the robot may need to temporarily ignore its mission and the need to reduce power consumption in order to achieve this. Operators of large fleets of low cost robots may be less concerned about this as the penalty for loosing a robot is much lower than those operating fewer expensive robots. The robot must also not be too "afraid" of danger. It is likely that at least part of the reason for deploying the robot is to operate in areas considered too dangerous for humans. If the robot is constantly avoiding danger then it is most likely failing to achieve its mission.

5.3.1 Hormone Roles

Given the requirements specified in the previous section we can start to consider the potential roles of the various artificial hormones that might allow those behaviours to be achieved through the use of an artificial neuro-endocrine system.

5.3.1.1 Energy Level Hormone

In order to manage power consumption the obvious choice is an energy hormone who's function is analogous to that of insulin in biological systems. This hormone could operate

either to promote power consuming behaviours when energy is abundant or to suppress them when energy is scarce. In biological systems insulin is produced when blood glucose (sugar) levels are high in order to trigger cells to increase their uptake of glucose. Glucagon is produced when blood sugar levels are low to signal that stored glycogen should be converted into glucose. However, in an artificial system this should work equally well as a suppressive mechanism to reduce power consumption when energy is scarce or as an excitory one to promote it when energy is plentiful. Its is only an implementation detail as to which way round we decide our system should operate. Sailing robots bring with them an additional complication to energy management. As they are dependent upon the wind for locomotion a lack of wind will prevent the robot from moving and a sailing robot has no ability to store wind power for future use (unless equipped with a wind turbine, large batteries and a propeller). Therefore the energy hormone could also be sensitive to wind speed (or in the absence of a wind speed sensor it could be sensitive to average boat speed) and could act to suppress the control system when there is no wind.

5.3.1.2 Day/Night Oscillating Hormone

A hormone which oscillates in phase with sunlight levels could act as an artificial analogue to the hormone melatonin. This could create a form of circadian rhythm for the robot, releasing more hormone when sunlight levels are at their highest. This could be useful when trying to optimise the efficiency of an electrical system so that power is used directly from a photo-voltaic solar panel instead of being stored inside a battery. The amount of sunlight could either be sensed through a light sensor, by reading the power output of a solar panel or calculating the elevation of the sun given the current time and the robot's position. If the hormone is triggered through calculation of the sun's elevation rather than direct sensing it would also be possible to apply a phase shift to it. This could be useful as instead of predicting the elevation of the sun it would predict its future elevation. If decisions must be taken about whether or not to begin a task which will take several hours to complete then it might be useful to consider what the sun will be doing by the time the task completes, not just what is happening when it starts. The time scales could also be adjusted to cover seasons instead of days. This shifts decisions from an hour by hour view (where the average power availability from one day to next is unlikely to change dramatically) to a longer term view. The end result could be that a robot operating over the course of a year will be able to attempt more ambitious power consuming tasks during the summer months. Hormones representing tides (and any other cyclical factor affecting the robot) could also be introduced using similar principles.

5.3.1.3 Mission Hormone

This hormone creates a desire for the robot to perform its scientific mission and may override the requirements to save power. Ideally the robot operator would have some ability to configure which would have highest priority, performing the mission or preserving the robot. This decision would ultimately depend upon the operator's need to obtain data quickly and at an increased risk to the robot or whether they wished to extend the robot's lifetime in exchange for occasionally loosing data. This might be, at least in part determined by the cost of each robot and how prepared the operator is to risk the robot's "life" in order to obtain data.

5.3.1.4 Activity Hormone

This hormone indicates the need for the robot to perform some sort of activity such as changing course due to imminent arrival at a waypoint. This hormone could be used to wake the robot from a state of inactivity when conditions have changed and a greater level of activity is now required.

5.3.1.5 Danger Hormone

Analogous to adrenaline, released when the robot is considered to be in danger, raising the weights within the steering and sailing neural networks causing them to react more dramatically. In biology adrenaline is often associated with the "fight or flight" response. In robotics this may translate to suppressing scientific data gathering and other behaviours which are not related to the immediate survival and promoting behaviours relating to avoiding danger and reaching safety. This hormone would typically be released in response to sensing dangerous conditions, for example proximity to hazards such as the physical obstacles (e.g. the coastline), another vehicle (manned or robotic) or storms. This approach also has a resemblance to work in the field of Artificial Immune Systems [53], based upon the "Danger Theory" [71] which states that the immune system responds to infection by acting upon danger signals given off by damaged cells.

5.3.1.6 "Pain" Signals for Component failure

In biology a number of hormones such as cortisol (often referred to as the "stress hormone") or epinephrine (adrenaline) are associated with responses to pain and other stressful situations. In an artificial system a single hormone will probably be sufficient. This will be released in response to component failures or difficulties (such as an actuator overheating). It provides a mechanism in which a failing component can be switched off or have its duty cycle reduced.

In order for the robot to continue to be fully operational some kind of redundant hardware architecture is required. Section 4.2.1.2 describes an attempt at building such an architecture for the ARC sailing robot. This architecture used two redundant motor controllers for each actuator. A temperature sensor detected overheating of the motor controller and a hormone controlled the duty cycle of each controller. However, this system was later abandoned for more reliable motor controllers which had far less tendency to overheat!

5.3.2 Artificial Hormone Cascades

In biological systems hormones often control their own production through a feedback loop involving several hormones as well as environmental feedback. In many cases the process is initiated by the neural system and production of the hormone is only stopped when a second (or third or nth) hormone signals that production should cease. Common examples are the feedback loops between the Hypothalamus, CRH, pituitary, ACTH and cortisol or between the Hypothalamus, TRH, TSH and Thyroxine. This process is known as a hormone cascade and was discussed in detail in section 2.6.2. The advantage of an artificial hormone cascade to a robot could be as a method of introducing a certain amount of lag inside the artificial endocrine system or ensuring a given behaviour is active for a minimum amount of time. As each hormone takes sometime to build up and trigger the next one in the cascade. This could be used as a method to ensure that a new behaviour is selected for enough time for it to have an effect. This could also act as a noise filter as small amount of the hormones which make up the early stages of the cascade may not be sufficient to continue the entire cascade and complete the feedback loop.

5.4 A Collision Warning "Danger Signal"

As discussed in section 5.3.1.5 a danger signal could be used to activate behaviours which avoid that danger. This section presents a potential mechanism to achieve this in relation to avoiding collisions with the coastline or other fixed obstacles.

A technique known as raycasting has been selected to determine the heading and distance to an obstacle and to discover obstacle free routes to sail. Raycasting [111] is a technique often used in early 3D computer graphics and games, it works by tracing the path of a series of rays originating from the players current position on a map of objects. When one of these rays hits an object (such as a wall) the distance is registered and the object is rendered at a size proportional to its distance. This can be applied to a robot by using the robot's GPS position to place the robot on a map and then casting the rays from the robot's position

Algorithm 5.3 The raycast algorithm for detecting the nearest coastline to the robot.

```
for angle = 0 ; angle < 360 ; angle++
    for dist = 0 ; dist < max_dist ; dist++
        x = sin(angle) * dist_y = cos(angle) * dist
        if getpixel(x,y) != 0 then
            if dist < nearest_dist then
                nearest_x = x
                nearest_y = y
                nearest_dist = dist
```

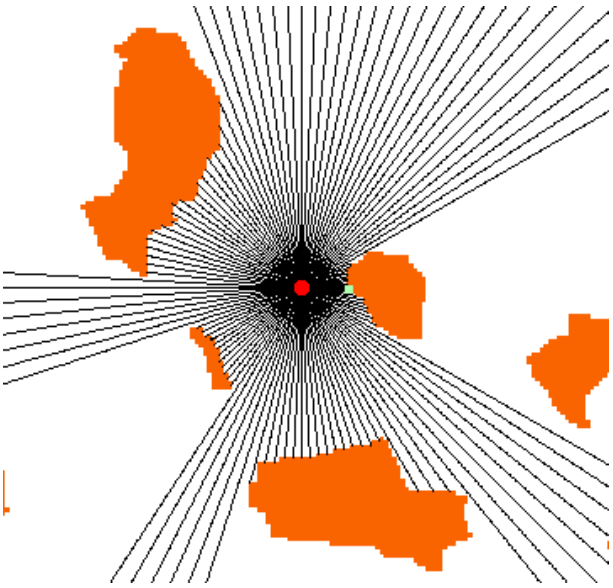

Figure 5.8: An example raycast in all directions. The robot is the red dot in the middle. The black lines are the rays. The orange blobs are islands the robot needs to avoid.

until they reach an obstacle (such as a land mass) on the map. This will result in the robot knowing the distance to the coastline in every direction. Based upon this information the robot can take appropriate action to avoid a collision with the coastline. A pseudo code algorithm is shown in algorithm 5.3. Figure 5.8 shows the path of rays emanating from the robot and hitting terrain. To prevent the robot from trying to avoid land that is behind or to the side of it the beam of rays can be restricted to only point ahead of the robot as shown in figure 5.9.

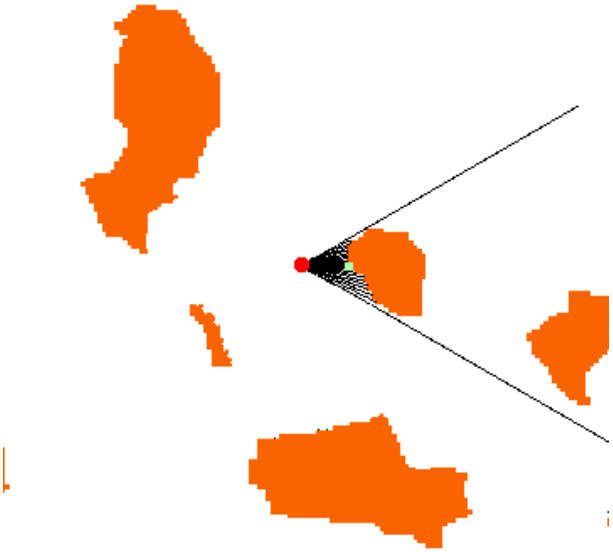

Figure 5.9: A raycast where the beam has been narrowed to only include obstacles ahead of the robot.

5.4.1 Switching Between the Collision Avoidance and Sail to Waypoint Behaviours.

After avoiding a collision, the robot needs to eventually return to sailing towards its waypoint. Once a new course has been decided upon, that course must be followed for enough time to allow the robot to actually complete the change of course and sail far enough to avoid the collision. There is a need to commit to the avoiding action and stay with it for sufficient time that it becomes useful. Conversely this time cannot be too long or the robot may become at risk of colliding with another obstacle that was not previously accounted for.

A artificial hormone can be used to represent the level of danger posed by nearby coastlines. When the hormone reaches a certain threshold it could trigger a change to a collision avoidance behaviour. It might be possible to gradually alter the course depending upon the hormone level with the default sail to next waypoint and collision avoidance behaviours essentially competing with each other. However a gradual switching between sailing towards the waypoint and avoiding an obstacle (when the obstacle is first detected) may not be as desirable as an early and decisive action is more likely to avoid a collision than a gradual one which leaves things too late.

To achieve these properties a hormone could be pooled using the method defined by Mendao (2006) [77]. When an obstacle is detected hormone is not released directly but instead it is stored in a pool until a threshold value is reached, then all of the stored hormone is released at once triggering a sudden change in behaviour. This hormone then gradually decays

to return to a normal behaviour. This is achieved by setting the heading that the robot follows to be determined by multiplying the hormone quantity (which is always between 0 and 1) by the difference between the robot's original heading and the new heading that was determined by the obstacle avoider. As the hormone decays the target heading gradually returns to its original value. The end result is for the first few seconds when the robot approaches an obstacle nothing happens while the hormone pools, then the hormone is released and there is a dramatic change in heading and then the robot gradually returns to its original heading. If the robot is still in danger of collision then the process will repeat itself, it is vital that the hormone decay is quite slow to allow enough time for the obstacle to be passed and it is also vital that the initial hormone release is fast enough to allow action to be taken before a collision occurs. To find a new course the robot must search for a clear course to sail that is as close as possible to the original target heading to the waypoint. The reason for this is to minimise the change in course in order to keep the boat going in the correct general direction and to minimise rudder/sail movements which are relatively expensive in terms of power consumption.

5.4.2 Application of Collision Avoidance

A simulated implementation of this approach to collision avoidance was undertaken in Sauzé and Neal (2010) [117], however, it was never integrated with any additional hormones. Behaviour switching was therefore limited to between sailing the default course and avoiding obstacles. The simulated scenario involved sailing from the northern to southern end of Strangford Lough in Northern Ireland (54.5 degrees North, 5.6 degrees West), this location was chose because of the presence of over 70 small islands in the space of a $150km^2$of water. This presents a complex navigation challenge even for human sailors and offers many potential locations where a robot could become stuck between several islands. The simulated robot was able to successfully navigate its target course in most cases, but would sometimes become stuck if it entered a narrow bay or found itself surrounded by islands with no large escape gap. This work demonstrated the basic feasibility of this approach in simulation, but should be integrated with other endocrine modulated behaviours (such as power management) in future work.

5.5 Chapter summary

In this chapter we have identified the key mathematical formulae and algorithms behind an artificial neuro-endocrine controller for use in a sailing robot. Some refinements and

simplifications have been made upon existing approaches. These reduce the size of the parameter space. Given the limitations of running experiments on a real robot, the parameter space needs to be small in order to achieve sufficient coverage in a limited time.

This chapter has shown how endocrine modulation can be applied to modulate the magnitude of actuator movements in a manner which is easily controllable and that should be able to reduce power consumption. Key objectives of the control system and potential hormones have been identified. These are that the robot must perform its mission, that it must use energy sustainably and that it must avoid danger. A method for detecting and avoiding fixed obstacles by switching of behaviours has also been presented.

The neural networks defined in this chapter will form the basis of a robot control system. This will be applied to a simulated sailing robot in chapter 6 and to an actual robot in chapter 7. Endocrine modulation will be applied to these in response to a number of factors. This will test if using an artificial endocrine controller to vary the magnitude of actuator movements will actually reduce power consumption and to see what effect it has upon sailing performance.

Chapter 6

Power Management Experiments and Results using Simulation

6.1 Introduction

Building upon the highlighting of power management as an issue in chapters 3 & 4 and the artificial neuro-endocrine controller in the previous chapter, this chapter presents the results of a series of experiments using this controller to manage the power consumption of a simulated sailing robot. These experiments are intended to test the hypothesis set out in section 1.3, that an artificial neuro-endocrine controller can manage power consumption in a sailing robot. A simulated environment was selected to perform an initial test to demonstrate the feasibility of this approach before proceeding to implementing it on a real sailing robot. It was thought that it would be easier and faster to develop the system and check its feasibility in simulation rather than on an actual robot. The simulator removes many of the complexities of real hardware such as component failures and physical damage to the robot and is far easier to deploy and run multiple times. However, it also over simplifies many of the real world control system issues.

The simulator not only covers the physical movement of the boat, but the electrical properties of its actuators, batteries and solar panels. Through this the effect of experiments upon power consumption can be evaluated. A set of four experiments using an artificial neuro-endocrine controller are undertaken. The first fixes the hormone at set levels to test the responses of the boat are sensible and to discover the limits of hormonal modulation. The second experiment varies the concentration of a single hormone in response to the battery level. The third adds a simulated solar panel to recharge the battery and the final experiment brings in a second hormone linked to the sun light level.

6.2 The Simulator

The use of simulator was employed primarily to test the feasibility of algorithms before they were deployed on a real robot. The selected simulator was developed based on the open source game Tracksail [1]. This provides a basic sailing physics model of a small sailing boat with a single mainsail. Simple wind direction shifts and gusts can be simulated, but they have not been used in this work to allow for conditions to be exactly recreated. Although it suffers many shortcomings, the simulator is sufficient to demonstrate the potential feasibility of the control algorithms for controlling a real robot, or at the very least to reject any unworkable approaches. An algorithm which does not function correctly in simulation is unlikely to work on a real robot, one which does work in simulation stands at least some chance of operating correctly on the robot.

The shortcomings of the simulator include the of waves, tides or currents and some behaviours which are not representative of a real sailing boat. It was previously found that [115] the response speed of the simulated robot was much faster than a real one with a 360 degree turn being possible in under one second in light winds. Although it would have been possible to improve the quality of the simulator, vast amounts of time could have easily been spent on this rather than focusing on the development of real robots. Little previous work exists to produce a accurate simulators which can be integrated with autonomous control systems. Some physics simulations such as those by Roncin (2004) [109] have developed more sophisticated and accurate simulations which consider many more variables than the Tracksail system. However, none of these were available as off the shelf programs which can be easily adapted for the purpose of this work.

The modified version of Tracksail used for this work has been forked from the original project and entitled Tracksail-AI and is available as an open source project. It is available online[2] for any readers of this thesis who may wish to undertake their own sailing robot simulations.

6.2.1 Simulation of Electrical Properties

Some additional modifications were made to the simulator in order to track the electrical power generation, storage and consumption of the simulated robot. A plausible set of figures for battery capacity, solar panel capabilities and power consumption were devised based upon a MOOP robot (as discussed in section 4.1.5). The battery capacity was calculated to

[1] `http://tracksail.sourceforge.net` accessed 15/04/2011

[2] `http://microtransat.svn.sourceforge.net/viewvc/microtransat/tracksail-AI/` accessed 7/04/2011, version 2 was used for all work in this thesis

Name	Power
Batteries	55 Watt Hours (198000 Joules) - equivalent to 20 AA NiMH rechargeable batteries (1.2 volts at 2300mAh each)
Solar Panels	4.75 Watts Peak, based on fitting 12x 4.4volt, 90mA solar panels
Rudder Actuator	0.3814 Joules per position (11 positions total) of actuator movement
Sail Actuators	1.1192 Joules per position (11 positions total) of actuator movement
Wind Sensor	75 mW Average
Compass	3.3 Volts, 4.5mA = 14.85 mW Average
GPS	3.3V, 44mA = 145.2 mW Average
Computer	750 mW Average

Table 6.1: The power consumption, battery capacity and solar panel figures used by the simulator.

be 55 watt hours, based upon the assumption that the robot would have 20 AA batteries setup in 4 parallel packs each containing 5 batteries. This assumption was made early in the development of the MOOPs when it was believed that 5 1.2V cells was sufficient to power a MOOP, this was later found to provide insufficient voltage for the sail servo when under load. The battery configuration was later changed to 18 batteries in 3 parallel packs each containing 6 batteries.

The power consumption of the rudder and sail servo were estimated by connecting them to a variable power supply with a voltage and current meter, moving them from end to end and observing the time taken, average current and voltage. Each actuator was moved across its full range of motion approximately 50 times to measure the average time taken and the power consumption. The sail actuator required an average of 10.25 watts (1.35 amps at 7.6 volts) and needed 1.2 seconds to move once through its full range. This gives a total power consumption of 12.312 joules to move through its full range of motion or 1.1192 joules to move from one of the 11 positions to the next. The rudder actuator required an average of 4.56 watts (0.6 amps at 7.6 volts) and needed 0.9 seconds to move through its full range. This gives a total power consumption of 4.9152 joules to move through its full range of motion or 0.3814 joules to move from one of the 11 positions to the next. These figures are of limited accuracy as the actuators were not under realistic load (the sail and rudder were connected, but the test was done indoors and out of the water so wind and water resistance is not factored in), although it is hoped they will give a rough approximation of real world power consumption.

Component	Power Consumption Range (per day)
Actuators	12.8kJ (1% duty cycle) to 128kJ (10% duty cycle)
Computers	6.48 kJ (10% duty cycle) to 64.8 kJ (100% duty cycle)
GPS	0.125 kJ (1% duty cycle) to 12.5 kJ (100% duty cycle)
Compass	0.125 kJ (10% duty cycle) to 1.25 kJ (100% duty cycle)
Wind Sensor	0.648 kJ (10% duty cycle) to 6.48kJ (100% duty cycle)
Total	20kJ to 212kJ

Table 6.2: The estimated power consumption ranges for each component of a MOOP sailing robot.

It was assumed that the boat had a small on-board computer, GPS, wind sensor, compass, solar panels and a battery. The figures used are shown in table 6.1. Solar panel capacity was estimated to be 4.75W peak by placing 12 4.4V, 90mA (peak) 6cm x 6cm mono-crystaline solar cells[3] on the deck. The output current of the solar panel was considered to be equivalent to the sine of the angle of the sun for the simulated time of day, it was assumed that the solar panels would be facing directly upwards and there would be no cloud. Sun elevation was calculated by the freely available library Solpos [4] from the United States National Renewable Energy Laboratory. The latitude and longitude for all experiments were assumed to be 52.4 degrees North and 4 degrees East (the approximate location of Aberystwyth). On a cloudless day at 52.4 degrees north in June the solar panels will provide around 150 kilo Joules or 41 watt hours of energy per day. At the equinox (March/September) it will provide around 82.5 kilo Joules or 23 watt hours per day and in late December it will only provide 20 kilo Joules or 5 watt hours per day.

Table 6.2 provides estimated daily power consumption ranges for each component in the robot. This assumes that it is possible to run the computer, compass and wind sensor at between 10% and 100% duty cycle and the GPS and actuators at between 1% and 10% duty cycle. These figures are likely to be an underestimate as they do not include the inefficiencies of batteries or voltage regulators. They show that it is plausible to run such a robot from solar power, but that behaviour will have to be modified to meet available energy demands.

Each iteration of the simulator is considered to be one second. The battery level is recalculated during each iteration to consider the amount of power generated by the solar panel and used by the actuators. This is specified by the formula:

$$b_{t+1} = s/3600 + b_t - c/3600$$

Where b is the remaining battery capacity in watt hours. s is the amount of power

[3] `http://www.mutr.co.uk/product_info.php?products_id=312` accessed 6/4/2011
[4] `http://rredc.nrel.gov/solar/codesandalgorithms/solpos/` accessed 13/4/2011

generated by the solar panel in watts and c is the power consumption of the actuators in watts during the current iteration of the simulator. The value for c is calculated by multiplying the distance moved by power consumption per position (as specified in table 6.1).

This model simplifies several aspects of the power consumption process as it assumes that 100% of the power which is collected by the solar panels is transferred into the battery and that the battery was able to return 100% of it at a later time. Actuators are always assumed to draw the same amount of power for each degree of movement, actuator load is not modelled and power consumption is assumed to be constant throughout the entire movement.

6.2.2 Simulator Integration

The control system including the neural network and endocrine controllers were implemented as part of a C program. As Tracksail was implemented in Java the two were integrated through a TCP/IP socket operating over a loopback network. This also allowed the possibility of running the simulator and the control system on different computers if required. The TCP/IP socket allowed the control system to set the rudder and sail position or to obtain the boat's heading, position, wind direction, current waypoint number or current rudder and sail positions. A series of waypoints are specified to Tracksail through a configuration file. The default algorithm in Tracksail requires each waypoint to be a pair of points and the boat must sail between these and in the correct direction. The control systems used by all the Aberystwyth University sailing robots operate by having the robot reach a threshold distance to a single point. This threshold distance is typically at least 15 metres to allow a GPS receiver suffering accuracy problems to still place the robot within the threshold. To simulate this behaviour both parts of the Tracksail waypoint are set to the same coordinate and getting closer than 15 metres to them will trigger it to move to the next waypoint.

6.3 Metrics for Evaluating Experiments

Metric Name	Calculated By	Potential Shortcomings
Total Power Consumption	Tracking power consumption to complete a given course.	Influenced by wind speed and direction. Either requires specialist hardware or software to track actions which discharge the battery.
Average Speed	Distance covered divided by time taken.	Influenced by GPS accuracy. Wind speed/direction and point of sail may not create fair comparisons.
Course Efficiency	Dividing the course length by the distance covered.	Influenced by GPS accuracy and threshold distance for reaching a waypoint.
Time to discharge Battery	Measuring the amount of time elapsed to completely discharge a fully charged battery.	Influenced by wind speed or direction and ability to measure battery state of charge.
Distance to discharge Battery	Measuring the amount of distance covered to completely discharge a fully charged battery.	Influenced by wind speed and direction, GPS accuracy and ability to measure battery state of charge.
Power Consumption per kilometre travelled	Dividing total power consumption by distance travelled.	Influenced by wind speed and direction, GPS accuracy and ability to measure battery state of charge.
Power Consumption per hour of travel	Dividing total power consumption by the total time.	Influenced by wind speed and direction and ability to measure battery state of charge.
Distance to Waypoint	Recording the distance to the next waypoint. If constantly repeating the same course this should never exceed the distance between two waypoints.	Influenced by GPS accuracy.

Table 6.3: A comparison of metrics for evaluating the success of sailing robot power management experiments.

In order to determine the effectiveness of a particular power management strategy, a consistent evaluation method is required. An ideal metric will be able to quantify the success of an experiment and should be applicable to both simulation and a real robot. The two key parameters which need to be compared are the power consumption and the sailing performance. A summary of potential metrics are shown in table 6.3.

Electrical power consumption can be measured either by keeping track of all actions which use battery power, by measuring the battery voltage or by measuring the flow of electrical

current. Sailing performance can be evaluated by the speed the boat travels, the accuracy with which it sails its course or the amount of progress it makes towards a goal or set of goals. In real world situations sailing robot performance will not be entirely consistent from day to day or even hour to hour, due to changes in wind speed and direction or sea state. Although it may be possible to normalise against these if they can be measured, it is not so easy to measure them frequently and accurately enough to perform a meaningful normalisation. An ideal metric will overcome or at least be somewhat resilient to these variations.

Power efficiency can be assessed by constantly repeating the same course and measuring the endurance time until the battery is completely discharged. Alternatively the amount of energy required to perform a given task (e.g. a sailing around a course a set number of times) can also be measured. By repeating the same course and varying control system parameters it may be possible to establish the effects of those parameters upon power consumption. However, power consumption rates alone do not reveal how efficiently the course was sailed. An insight into the accuracy with which the course was covered can be revealed by examining the total distance covered and the average speed. If the total distance travelled is divided by the total distance of the course then a course efficiency measure can be created, this is shown in equation 6.1, where c is the course efficiency, t is the distance travelled and l is the course length.

$$c = l/t \tag{6.1}$$

An additional metric to monitor course efficiency is to consider the distance between the robot and its next waypoint. Under ideal conditions this should never exceed the distance from one waypoint to the next waypoint. When this metric becomes most useful is when a robot significantly deviates from its intended course. It will work best when waypoints are close together and the deviation exceeds the distance between any two waypoints. It gives us an insight into how far the robot has deviated from the course and when viewed over time, how long it took to deviate and return to the course. If power consumption is expressed in terms of energy usage (Joules) for each hour of travel then this offers a metric which is not affected by the speed of travel but also loses any notion of work which has been achieved. It is possible that a robot which has used no energy, but travelled nowhere will be viewed favourably by this metric. Considering power consumption per kilometre gives an insight into the robot actually managing to travel somewhere, but will be susceptible to changes in wind speed which have an influence (but are by no means the sole factor) upon boat speed. Considering both of these metrics together should give a clearer picture of whether or not genuine changes in power consumption are occurring. The combination of these two metrics should be sufficient to answer the question of whether or not any hormonal modulation of

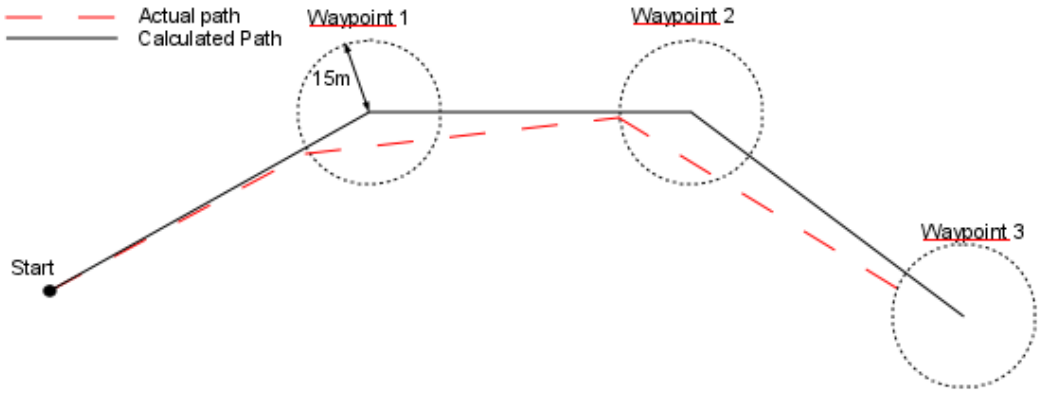

Figure 6.1: A comparison of the direct path between three waypoints and the actual path which will be taken due to the waypoint threshold algorithm.

the control system has reduced the power consumption significantly. Although they do not sufficiently answer the question as to whether or not sailing efficiency has been impacted upon to achieve these power savings.

Calculating these metrics creates some additional complications. Relying on GPS to provide the distance sailed is subject to the limitations of GPS fix accuracy and repeatability and the threshold distance around waypoints. As the boat is only required to be within 15 metres of a waypoint to be considered to have reached that waypoint. This can cause the total distance actually sailed to fall short of the actual distance between the waypoints, which can result in course efficiency numbers greater than 100%. This problem is illustrated in figure 6.1 where the dashed red line illustrates the actual course that can be sailed and the solid black line represents the course which will be calculated as the distance between the waypoints. However, as the problem consistently applies across all runs of the algorithm it does not reduce the comparative conclusions between multiple runs of the same algorithm. Therefore, there is no need to actually modify the method for evaluating course efficiency to compensate for this, but we must be aware that it is theoretically possible to see efficiency values over 100%. Metrics relying on the speed and distance a boat travels are heavily dependent on the wind direction relative to the boat. Typically the fastest point of sail is a beam reach, where the boat sails perpendicular to the wind, the slowest is when the boat attempts to sail closed hauled with approximately 45 degrees between the boat direction and the wind direction. Therefore, performing the same course (even with two identical algorithms and parameter sets) in different wind directions can result in differences in average speed. Differences in the distance travelled can also occur if the boat is required to sail up-wind.

As it is not possible to sail directly into the wind a boat must zig-zag (known as tacking up-wind) maintaining approximately 45 degrees between its heading and the wind direction. If one run of a particular course does not require upwind sailing and another does then even with no other differences the total distance travelled will vary. In a simulated environment it is easy to create many identical runs of the same course with the same wind direction, doing so in the real world is not so easy. Simply comparing average speed or distance covered does also not give any insight into whether or not the intended course was sailed or not. An ideal metric might also need to consider if a particular control algorithm has achieved the goals it was asked to as well as considering the speed, accuracy and power consumption. By limiting the course only to a beam reach (perpendicular to the wind direction) problems caused by varying distance when tacking up-wind are eliminated. This will require that the course waypoints are adjusted or that the possible wind directions are limited each time an experiment is carried out with a real robot. Under simulation this does not cause any major problems as consistent conditions can be applied each time. Although this does limit the coverage of the experiment since different behaviour may be experienced on other points of sail.

6.4 Simulator Experiments

6.4.1 Fixed Hormone Level

To test the effects upon power consumption and sailing performance of a neuro-endocrine controller, an initial experiment was devised to use fixed levels of hormone while a pre-determined course was sailed. Using the metrics established in section 6.3, it was decided to evaluate the power consumption both per kilometre travelled and per hour of operation. A fixed length course was setup and the aim was for each simulation run to complete the course. The course consisted of three waypoints in a 500m long straight line running East/West. There is one waypoint at each end of the 500 metres and one halfway between them. The middle waypoint was used to ensure that the robot stuck close to the course. It was primarily intended for the real robots than the simulator but was kept in the simulation for the purposes of keeping the experiments in simulation and the robot as similar as possible. The entire course (all 1km from the start point and back to it) was repeated 5 times in total. The wind was set to blow from the North at a consistent speed of 7 metres per second.

Power consumption was simulated using the method described in section 6.2.1. The battery level was fixed at 5, 25 and 55 watt hours remaining (from a total capacity of 55 watt hours), a zero battery level was not used as this was expected to produce a total inhibition

of the control system at higher hormone sensitivity levels. To calculate the hormone level a gland function producing the hormone used the formula:

$$0.023b - 1.030 \tag{6.2}$$

Where b is the amount of battery remaining in Watt hours (between 0 and 55 Watt hours). The output level of the gland function should always be between -1.03 and 0.235. This slight positive output when the battery level is above approximately 44.8 watt hours (about 82% of capacity) will have an excitory effect upon the neural network causing greater than normal actuator movements. This was intended to provide larger scale responses when the battery was full, which would exploit the plentiful supply of energy to sail the course more accurately. When the output of the gland is below 0 there will be an inhibitory effect upon the neural networks. So as the battery level drops the level of inhibition will increase, hopefully resulting in lower power consumption. Hormone sensitivities of 0, 0.25, 0.75 and 1.0 were used, a sensitivity level of zero will act as a control and should see no effect from the hormone system. As the input to the gland (the battery level) is constant for this experiment the resulting hormone level and the hormone sensitivity can be multiplied together to create a single number known as the effective hormone concentration. As hormone levels are constant it is this number which will be key to the amount of modulation applied to the neural networks. To ensure rapid changes in hormone level the response rate (r in equation 5.3) is set to 0.1.

In biological systems, the "normal" level of a hormone is not a zero value. This leaves two options for the implementation of an artificial system. Either the gland can constantly secrete hormone under normal circumstances or the normal situation can be setup so that no hormone is required. For the purpose of the experiments in this thesis the latter has been chosen and normal behaviour is considered to be when no hormone is secreted. The idea of a hormone being both suppressive or excitory is also present in biology. However, in biological systems a hormone can be suppressive to one group of target cells, while being excitory to another. At present the complexity of the control systems does not warrant a need for this. If at some point in the future the complexity increases to the point where this is required then a more complex approach maybe required. A linear gland function has also been selected here for the sake of simplicity. In biology hormonal reactions may often be non-linear and there is no reason why gland functions in our artificial system cannot be non-linear. This may be a potential area for future work.

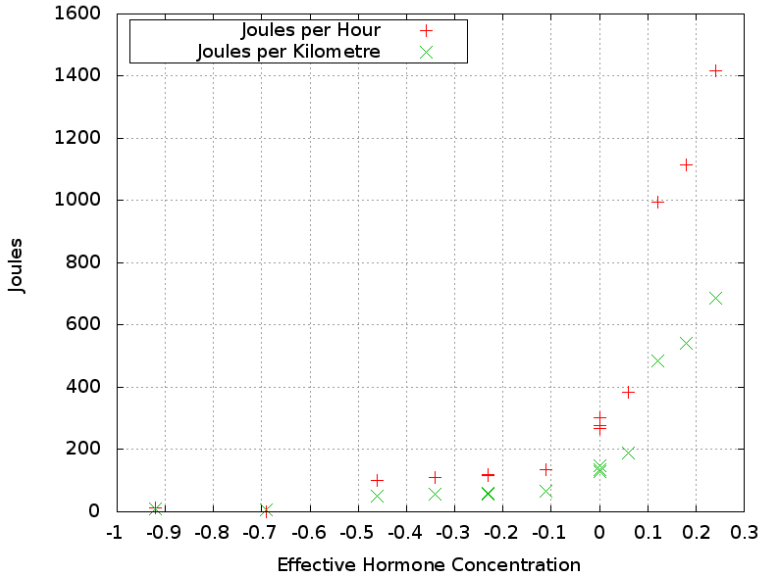

Figure 6.2: A graph showing the effects of varying hormone concentration in simulation against the number of Joules used for each kilometre travelled and each hour of travel.

6.4.1.1 Results

A summary of the experiments is shown in table 6.4, the two with the lowest effective hormone concentrations (the highest level of inhibition) both failed to complete the course and were terminated by manual intervention when it was realised that there was no possibility of them completing the course. A complete set of results including actual distances travelled, total times and average speeds is shown in Appendix D.

Figure 6.2 compares hormone concentration and the amount of energy in Joules used for each kilometre and each hour of travel. It suggests there maybe a correlation between the concentration of the hormone and amount of energy used. The correlation can be tested using Spearman's rank correlation coefficient. This is a non-parametric measurement and makes no assumptions about the underlying data unlike the Person's correlation coefficient which relies upon the data being linear. This leads to the following statistical hypothesis for the Spearman's rank correlation (rho):

H_0: Reducing hormone concentration does not alter energy use.

H_1: Reducing hormone concentration reduces energy use.

This forms a one-tail test as we are only testing against a decrease in energy use as the hormone concentration drops. The correlation coefficient was calculated using the R cor.test command[5], as R cannot calculate significance values for a Spearman's correlation so these

[5] `http://www.gardenersown.co.uk/Education/Lectures/R/correl.htm#correlation` accessed 15/4/2011

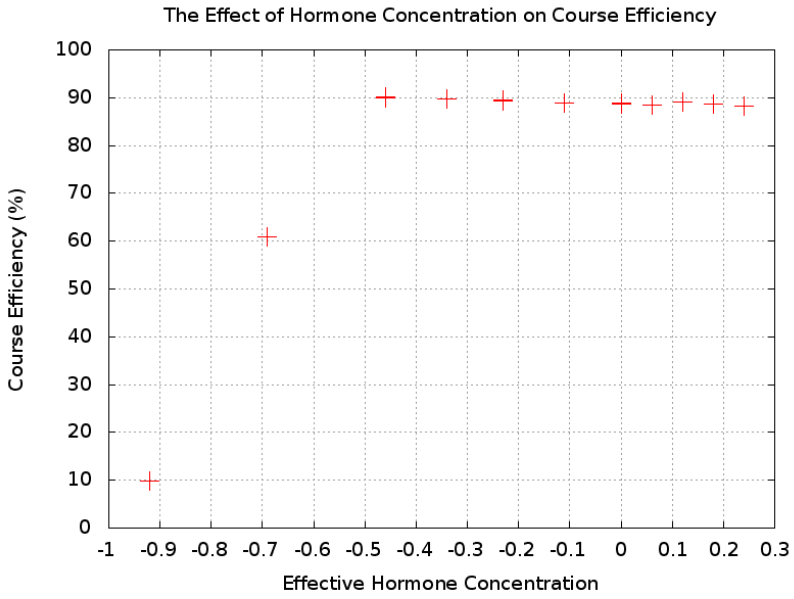

Figure 6.3: A graph showing the effects of varying hormone concentration against course efficiency.

were looked up in a significance table in Sprent (1990) [125]. Note that both power usage per hour and per kilometre are the same in this case because the ranks of the respective results are the same

$P = 0.9910377$
Using 99% significance level
$n = 15$
$P > 0.604$
Accept H_1

Figure 6.3 shows the course efficiency and effective hormone concentration and shows that, with the exception of hormone concentrations of -0.69 and -0.92, there is almost no change in course efficiency with all values being in the range of 88.5% to 90%. The two values which differ dramatically are the two runs which failed to complete the course. The high numbers suggest that the simulated robot has not heavily deviated from its intended course and that the level of deviation does not vary until hormone concentrations become exceedingly low. Examining the raw data of the course taken by the boat during the failing cases, they were not able to follow the course at all, but instead sailed off in totally the wrong direction. It appears that in both these cases the level of inhibition was high enough to prevent the control system from operating effectively.

Effective Hormone Concentration	Hormone Sensitivity	Battery Level (Watt hours, max 55)	Joules per km	Joules per hour	Course Efficiency (course length/distance travelled)	Completed course yes/no
0.00	0	5	135.941	278.738	88.58%	yes
-0.23	0.25	5	59.5	119.371	89.42%	yes
-0.46	0.5	5	50.259	99.816	90.11%	yes
-0.69	0.75	5	6.185	0.927	60.98%	no
-0.92	1.0	5	8.919	11.973	9.77%	no
0.00	0.0	25	147.592	303.778	88.98%	yes
-0.11	0.25	25	65.7912	134.598	88.99%	yes
-0.23	0.5	25	58.127	115.372	89.48%	yes
-0.34	0.75	25	55.388	109.979	89.79%	yes
-0.46	1.0	25	51.147	101.816	89.88%	yes
0.00	0.0	55	130.317	267.658	88.86%	yes
0.06	0.25	55	189.383	383.942	88.52%	yes
0.12	0.5	55	486.149	996.499	89.03%	yes
0.18	0.75	55	542.858	1114.477	88.72%	yes
0.24	1.0	55	686.686	1418.5	88.33%	yes

Table 6.4: Summary results from the simulator experiment.

6.4.1.2 Conclusions

These initial results demonstrate that hormonal modulation can have a dramatic effect upon the power consumption of a neural network controlled simulated sailing robot. Despite this dramatic change in power consumption there is little change upon the accuracy of sailing a fixed course until the effective hormone concentration level goes below something between -0.46 and -0.69. Once hormone concentration drops below between -0.46 and -0.69 then too much modulation is applied for effective sailing to be possible. This conclusion may simply suggest that for the task at hand the control system is far too over active and that up to a point this activity can be reduced when operating in light winds and flat seas offered by the simulator. However, in real life the control system needs to cope with a wide variety of conditions and some level of tolerance to harsher conditions is required. This experiment suggested that when conditions are calm this margin of error can be taken away without impeding performance. This therefore goes some way to confirming that this can also be used as a method for controlling power consumption. We may however, wish to create some kind of counteracting hormone which depends upon sea state.

6.4.2 Variable Level Battery Hormone

The previous section showed that fixed levels of hormonal modulation modulation can achieve control over a neural network controller and affect power consumption. The next stage is to apply a constantly varying hormone which represents feedback from the environment. For this example the hormone will be produced in response to battery levels using formula 6.2 that was used in the previous experiment, but this time each movement of the actuators will reduce the overall battery level. There will be no solar panels or any other power source being simulated to recharge the batteries. As this experiment is expected to take considerably longer than the previous one and due to limitations on computational resources the decision was taken to multiply the power consumption figures by 100. As computer and sensor power consumption is not being considered this figure is probably closer to real power consumption numbers than the figures used in the previous experiment.

The boat will, as in the previous experiment, sail back and forth on a beam reach, but in this experiment it will sail until the battery is flat instead of trying to complete a short course. The wind will, as before be set to a 7 metre per second Northerly. This will generate a key result of the amount of time the robot lasts and this can be compared with the hormone sensitivity. This leads us to the hypothesis that the greater the hormone sensitivity the longer the boat should be able to sail for as this will increase the modulation level and thus reduce power consumption. It was found in the previous experiment that during a 500 metre course the boat spent nearly the entire length of the course attempting to settle down from the turn it had initially made. This is believed to be due to a slow turning rate built into the simulation. To give the boat some period of calm time when it is not attempting to adjust its heading the 500 metre course has been lengthened to 5 kilometres for this experiment. As a hormone sensitivity of 1.0 resulted in the boat becoming totally unsailable in the previous experiment only the sensitivities 0.0, 0.25, 0.5 and 0.75 have been used in this experiment. Each of these is repeated five times to ensure that the results are not a chance occurrence. As with the previous experiment the hormone response rate (r in equation 5.3) is set to 0.1.

6.4.2.1 Results

Figures 6.4, 6.5, 6.6 and 6.7 show the battery level over time during these experiments. In figure 6.4 (which has zero hormone sensitivity so the behaviour is consistent across the entire experiment) a "staircase" effect can be seen . The sudden drops are caused by reaching waypoints and using considerable amounts of energy to turn the boat around, while in between waypoints the line is nearly flat as we are on course and there is nothing in the simulator to cause the boat to go off course. Figures 6.5, 6.6 and 6.7 all show a noticeable "step" effect

after around 1000 minutes of run time. This is at the point where the remaining battery capacity crosses the 44.8 watt hour threshold and the hormone switches from being excitory to inhibitory. It was also noticed that sometimes the sail would constantly oscillate between two positions, when on a boundary condition. This was most common while the hormone was still excitory. Added to this is that when reaching a waypoint and turning the boat 180 degrees requires both the sail and rudder to move dramatically, which (especially when the hormone is excitory or only slightly suppressive) uses considerable amounts of energy. This behaviour calms down as the battery level drops and the level of hormonal suppression on the actuators increases. This does bring into question if anything useful is being gained by having this hormone operate in an excitory manner and simply having it operate as a suppressive hormone would be sufficient. Even if it only became suppressive when the battery level dropped below a threshold amount, rather than beginning to suppress the actuators as soon as the battery dropped below full.

There appears to be a clear correlation between sailing time and hormone sensitivity. Summary statistics in table 6.5 show that the median time the boat is able to sail increases from 568.95 minutes for a sensitivity of zero (no modulation), to 2552.28 minutes for a sensitivity of 0.25, 4653.73 minutes for a sensitivity of 0.5 and 7642.9 minutes for a sensitivity of 0.75. A box and whisker plot is shown in figure 6.8, this indicates that there is quite a small spread and few outliers to this data. To ensure these changes were not simply a result of the boat slowing down and covering less distance the summary statistics for distance covered are shown in table 6.6 and in a box and whisker plot in figure 6.9.

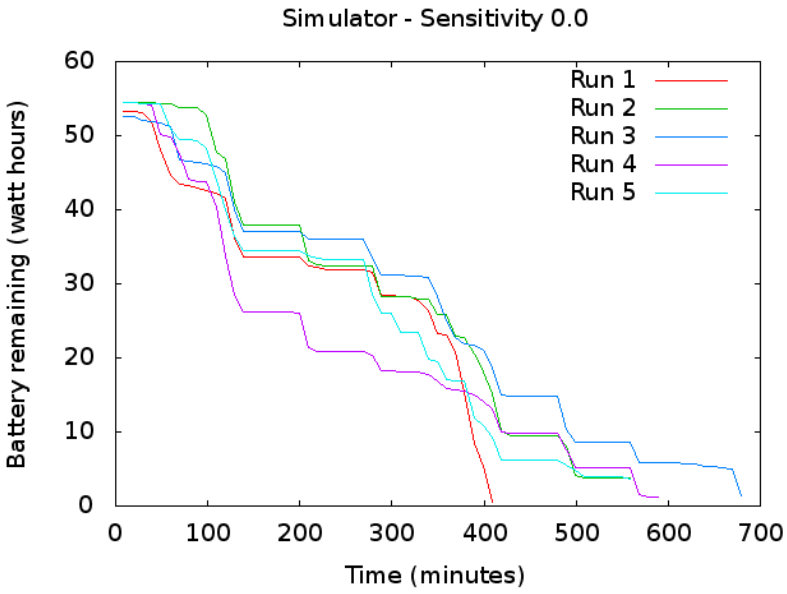

Figure 6.4: Graph of time and battery level for the simulator's running the variable hormone experiment with a sensitivity of 0.0.

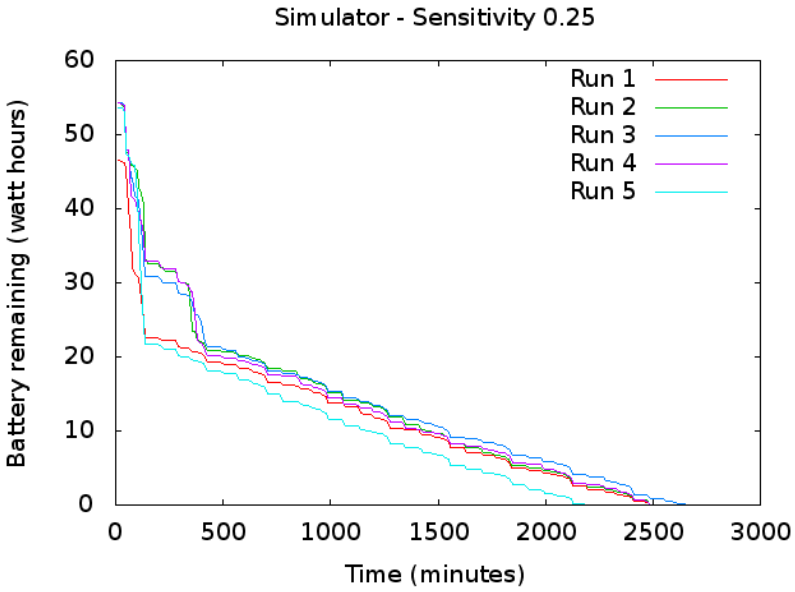

Figure 6.5: Graph of time and battery level for the simulator's running the variable hormone experiment with a sensitivity of 0.25.

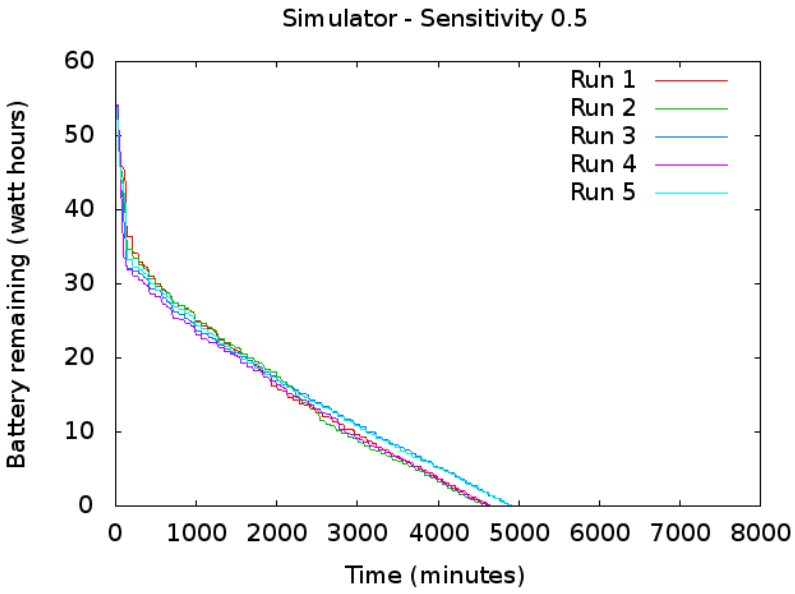

Figure 6.6: Graph of time and battery level for the simulator's running the variable hormone experiment with a sensitivity of 0.5.

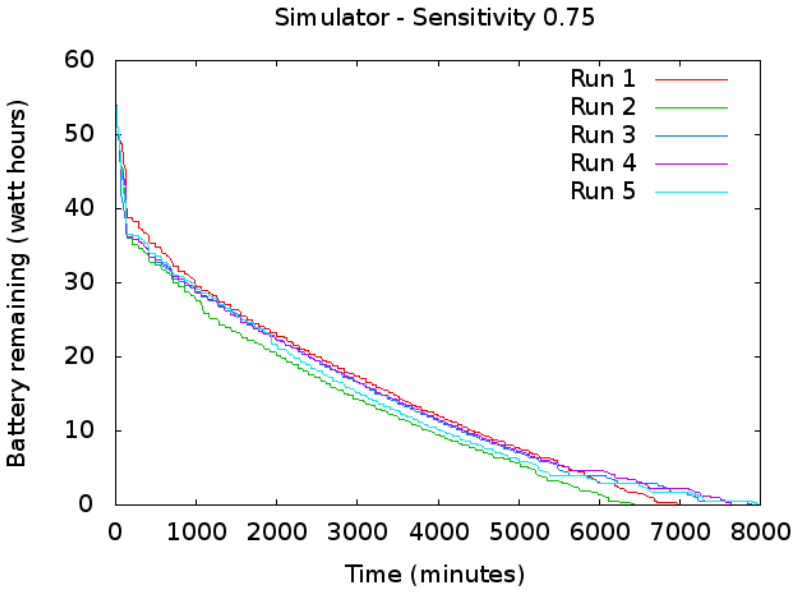

Figure 6.7: Graph of time and battery level for the simulator's running the variable hormone experiment with a sensitivity of 0.75.

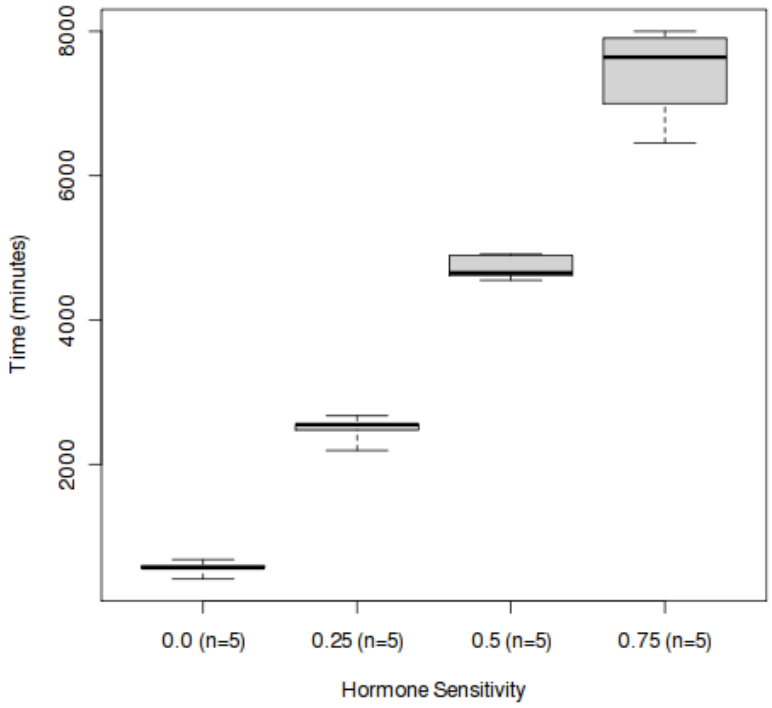

Figure 6.8: A Box and Whisker plot of the simulator's battery discharge time and hormone sensitivity during the variable battery hormone experiment.

Hormone Sensitivity	Mean	Median	Min	Max	Lower Quartile	Upper Quartile
0.0	565.01	568.95	411.7	682.82	567.9	594.22
0.25	2490.67	2552.28	2193.53	2674.23	2477.85	2555.45
0.5	4726.41	4653.73	4552.73	4912.2	4619.92	4893.48
0.75	7401.14	7642.9	6454.8	8000.27	7000.63	7907.12

Table 6.5: Summary statistics for the effect of the battery hormone on the time taken to discharge the battery. All values are in minutes.

Hormone Sensitivity	Mean	Median	Min	Max	Lower Quartile	Upper Quartile
0.0	20.15	20.23	14.88	24.48	20.03	21.14
0.25	88.21	90.25	77.68	94.84	87.89	90.39
0.5	169.24	166.29	161.87	175.94	166.24	175.86
0.75	231.09	235.2	215.96	238.63	228.78	236.88

Table 6.6: Summary statistics for the effect of the battery hormone on the distance travelled to discharge the battery. All values are in kilometres.

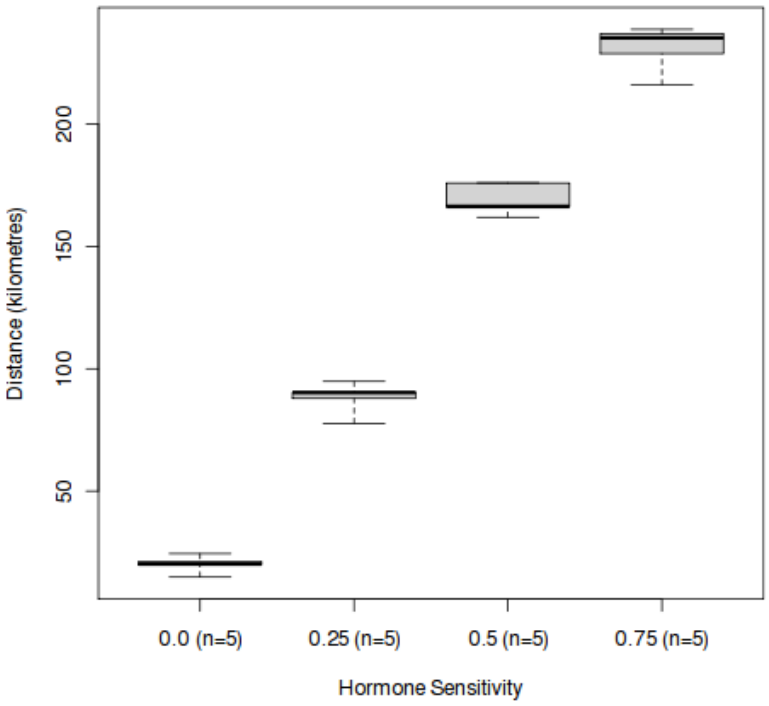

Figure 6.9: A Box and Whisker plot of the simulator's distance covered and hormone sensitivity during the variable battery hormone experiment.

A Kruskal-Wallis test (non-parametric analysis of variance) was performed to test if the differences caused by changing the hormone sensitivities was statistically significant or not. As this test operates on ranked data the results for both distance covered and time taken were identical as the ranks of the data were the same. The results were calculated in R [6] using the kruskal.test command. This leads us to generate the following statistical hypotheses:

> H_0 : That varying hormone concentration does not affect the time or distance that a robot can sail on a single battery charge.
>
> H_1 : That varying hormone concentration does affect the time or distance that a robot can sail on a single battery charge.

$H = 17.8571$, $df = 3$, $p = 0.0004707$

Given that $p = 0.0004707$ (0.4707%) we can accept H_1 at the 99% significance level, there is a statistically significant change in both time and distance that can be achieved on a

[6] `http://www.r-project.org` accessed 15/04/2011

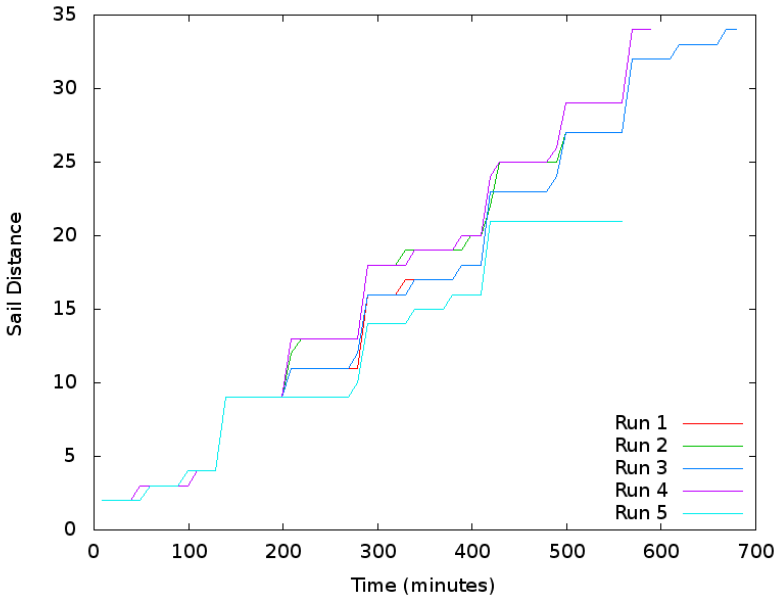

Figure 6.10: The cumulative distance moved by the sail actuator during the simulated variable hormone experiment with a hormone sensitivity of 0.0.

single battery charge. This conclusion can be reinforced by examining the box and whisker diagrams in figures 6.8 and 6.9 where the median time and distance can both be seen to significantly increase as hormone sensitivity is increased. There is also only a small spread of values suggesting that these trends are consistently appearing and not a random occurrence.

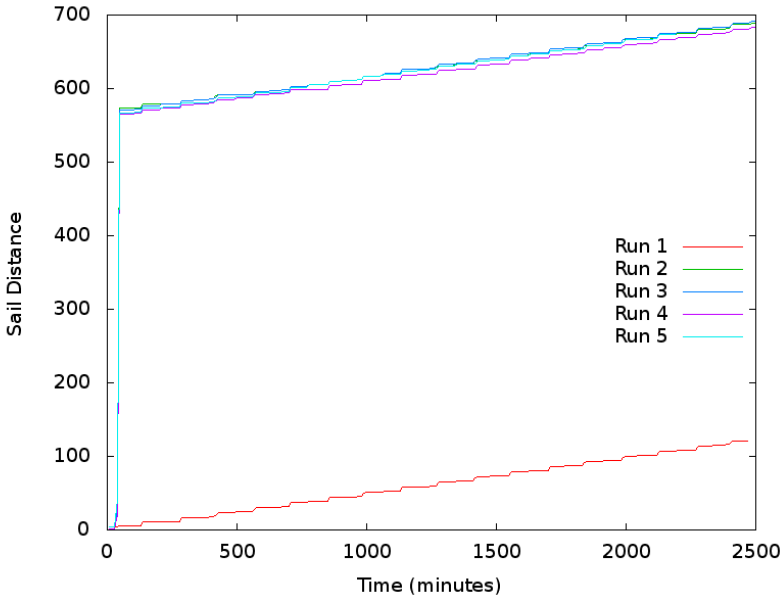

Figure 6.11: The cumulative distance moved by the sail actuator during the simulated variable hormone experiment with a hormone sensitivity of 0.25.

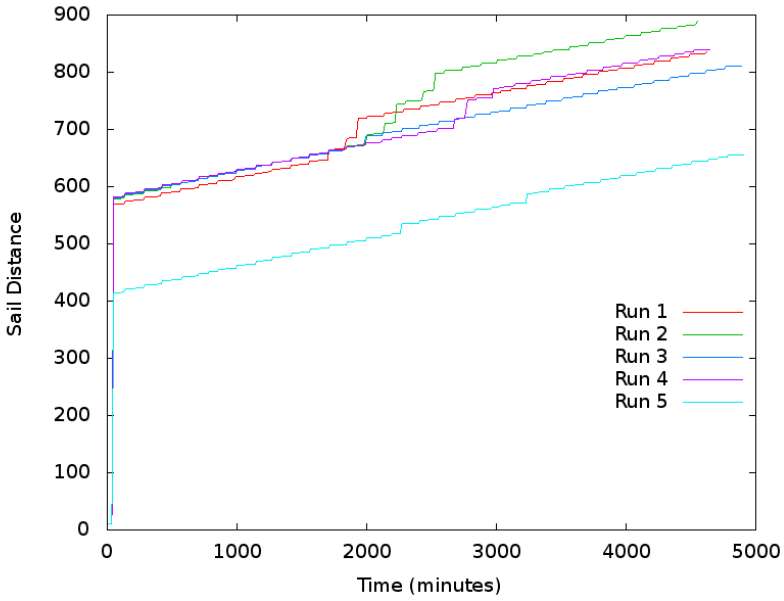

Figure 6.12: The cumulative distance moved by the sail actuator during the simulated variable hormone experiment with a hormone sensitivity of 0.5.

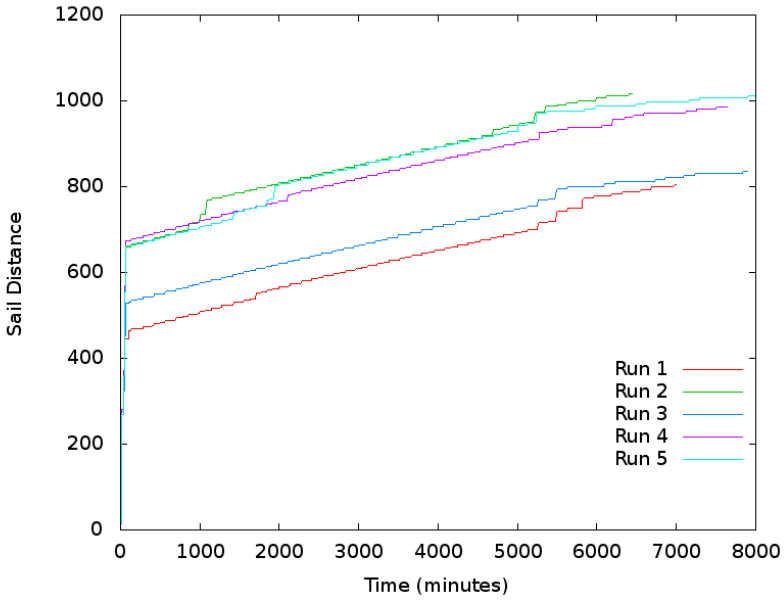

Figure 6.13: The cumulative distance moved by the sail actuator during the simulated variable hormone experiment with a hormone sensitivity of 0.75.

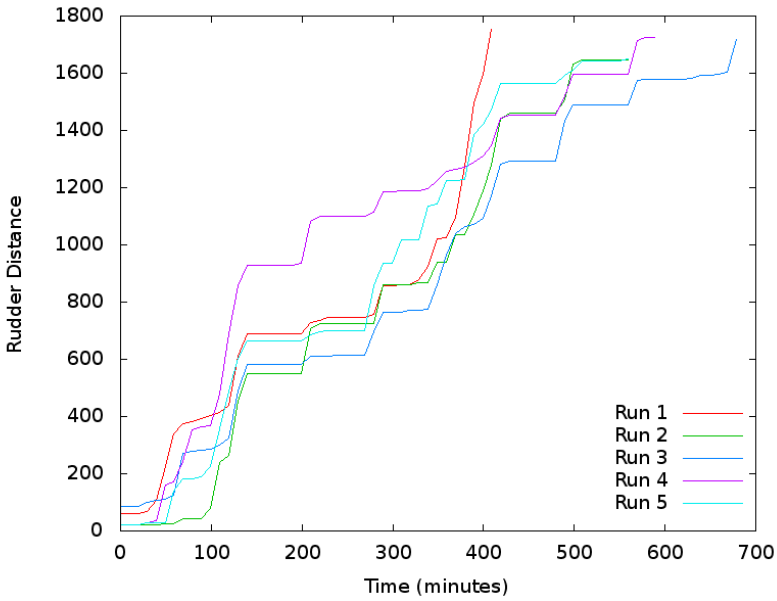

Figure 6.14: The cumulative distance moved by the rudder actuator during the simulated variable hormone experiment with a hormone sensitivity of 0.0.

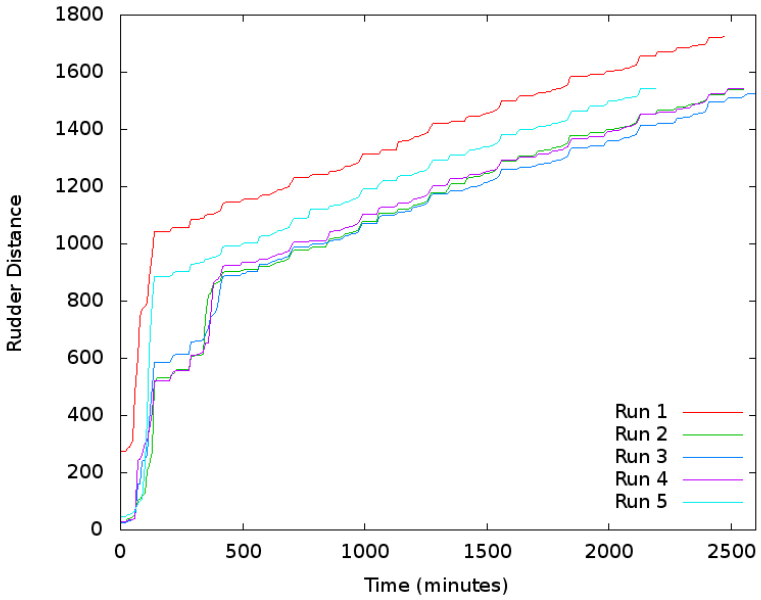

Figure 6.15: The cumulative distance moved by the rudder actuator during the simulated variable hormone experiment with a hormone sensitivity of 0.25.

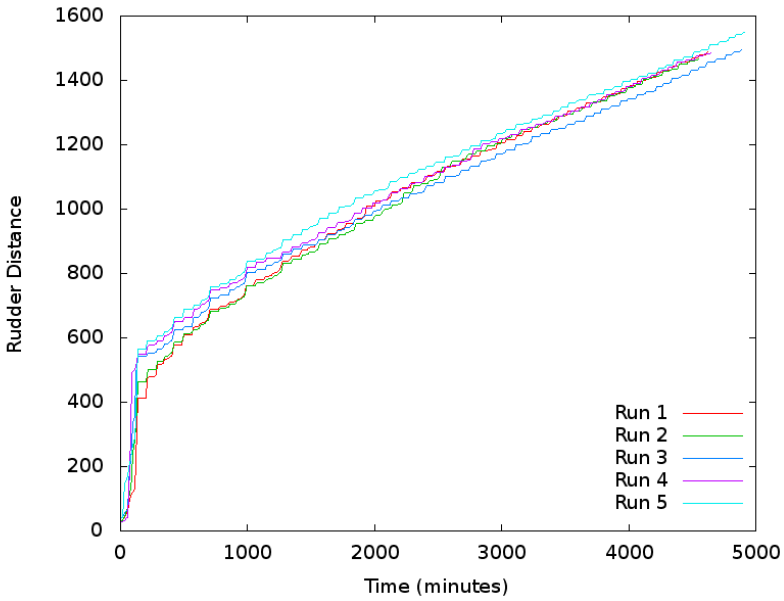

Figure 6.16: The cumulative distance moved by the rudder actuator during the simulated variable hormone experiment with a hormone sensitivity of 0.5.

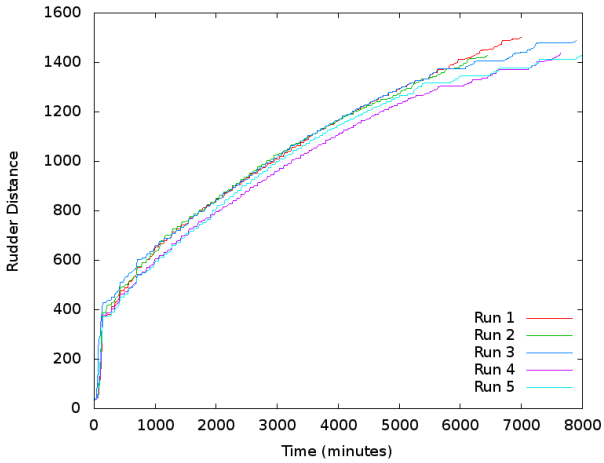

Figure 6.17: The cumulative distance moved by the rudder actuator during the simulated variable hormone experiment with a hormone sensitivity of 0.75.

Hormone Concentration	Mean	Median	Upper Quartile	Lower Quartile
0.0	28	28	34	26
0.25	573	683	689	673
0.5	787.4	812	820	804
0.75	909.4	973	980	828

Table 6.7: The summary statistics showing the number of sail movements during the variable battery hormone experiment.

Hormone Concentration	Mean	Median	Upper Quartile	Lower Quartile
0.0	1746.3	1746	1746	1743
0.25	1518.3	1481	1494	1475
0.5	1397.4	1387	1400	1384
0.75	1362.6	1346	1399	1332

Table 6.8: The summary statistics showing the number of rudder movements during the variable battery hormone experiment.

Tables 6.7 and 6.8 show the mean, median, upper quartile and lower quartile of sail and rudder movements against the hormone concentration. The full data is shown in appendix D. From this data it can be seen that the median number of sail movements increases with hormone concentration, while rudder movements decrease with hormone concentration. This data does not present a full picture though. It does not distinguish between small and large actuator movements, so although higher hormone sensitivities may see a different number of movements the size of those movements may also have changed. To overcome this limitation figures 6.10, 6.11, 6.12 and 6.13 show a graph of the cumulative distance moved by the sail

over time and figures 6.14, 6.15, 6.16 and 6.17 show the cumulative distance moved by the rudder. These values are in terms of the actuator positions (-5 to +5) discussed in section 5.2.3, so moving an actuator through its full range of motion will add 11 to the cumulative total. These graphs reveal when actuator movements occur and the size of the movements as well as the total level of movement. Figures 6.10 and 6.14 show the control runs where hormone sensitivity was set to zero. In these a "staircase" effect can be seen where most of the actuator activity occurs when the robot reaches a waypoint. By comparison figures 6.11, 6.12, 6.13, 6.14, 6.15, 6.16 and 6.17 which show the experiments where there is some sensitivity to the hormone all show a sudden jump near the start. This is while the battery level is still high and the hormone is slightly excitory causing larger actuator movements and higher power consumption. Once the initial jump in actuator movements has taken place the sail movement (figures 6.11, 6.12, 6.13) continues to increase at an almost linear rate for the rest of the experiment. Figure 6.11 appears to behave differently from all other cases and does not have the initial step, but immediately begins the linear phase giving a much lower overall sail use than any other run. However, this appears to be compensated for by additional rudder movement shown in figure 6.15. The rudder graphs for hormone sensitivities of 0.25, 0.5 and 0.75 all show a more curved response after their initial step. Ultimately in this simulated scenario battery state of charge is only affected by actuator movements, so the curves shown in figures 6.4, 6.5, 6.6 and 6.7 are entirely a product of actuator movements.

6.4.2.2 Conclusions

This experiment has helped to reinforce the conclusion of the previous experiment, that hormonal modulation can have a dramatic effect upon power consumption. By increasing the hormone sensitivity from 0.0 to 0.75 we have been able to achieve over a 13 fold improvement in the amount of time the simulated robot can sail for. Also, in the previous experiment this may simply be taking advantage of an excessive margin of error which has effectively been built into the control system. It does seem reasonable that a potential power saving strategy is to reduce this margin when running low on battery power. This experiment has also highlighted how excitory behaviours cause the battery to drain much faster. It is questionable as to whether or not this is a useful behaviour. If the neural network, operating at its "normal" level (i.e. no hormonal modulation of any kind) is sufficient for all situations then is there really any need to excite it when energy is plentiful?

6.4.2.3 Swapping Rudder and Sail Power Consumption Figures

At some point during the modelling actuator power consumption the sail and rudder figures were swapped, this mistake was only realised after a number of experiments had taken place. This gave the rudder a higher power consumption figure of 1.1192 joules per movement and the sail a figure of 0.3814 joules per movement. This section presents an attempt to rectify this situation and to test if it is likely to have affected the conclusions reached in the previous section. It could be argued that this is not a totally unrealistic scenario as in boats which steer the wing sail (as in Atlantis and Harbor wing boats discussed in section 3.2) the rudder actuator would be the larger actuator. The discovery of the error led to concern that the "step" behaviour shown in figures 6.5, 6.6 and 6.7 might be the result of this error. To investigate if this was the case the data from the experiment was reprocessed and battery levels recalculated as if the correct power consumption figures were being applied. This does not entirely compensate for the problem, as the rate at which the hormone changed during the experiment was still determined by the incorrect power consumption figures. However, it should give enough of an insight to determine whether or not it is necessary to rerun the experiment. Figures 6.18, 6.19, 6.20 and 6.21 show the results of this. Compared with figures 6.4, 6.5, 6.6 and 6.7 we can see very similar trends in the battery levels. Figure 6.18, like figure 6.4 still shows a staircase effect. While figures 6.19, 6.20 and 6.21 like figures 6.5, 6.6 and 6.7 show a steep drop in the battery level at the beginning and a levelling out later on. The only noticeable change is that a greater separation exists between each run in figures 6.19, 6.20 and 6.21 than in figures 6.5, 6.6 and 6.7. This is could be due to oscillations in the rudder occurring during some simulations and not others. As the results in the previous section exaggerated sail power consumption and understated rudder power consumption this difference was not exposed as clearly. With this reversed, changes in rudder movement are exaggerated and the difference appears. Apart from this there does not appear to any major change in behaviour and no suggestion that a rerun of the experiment is required.

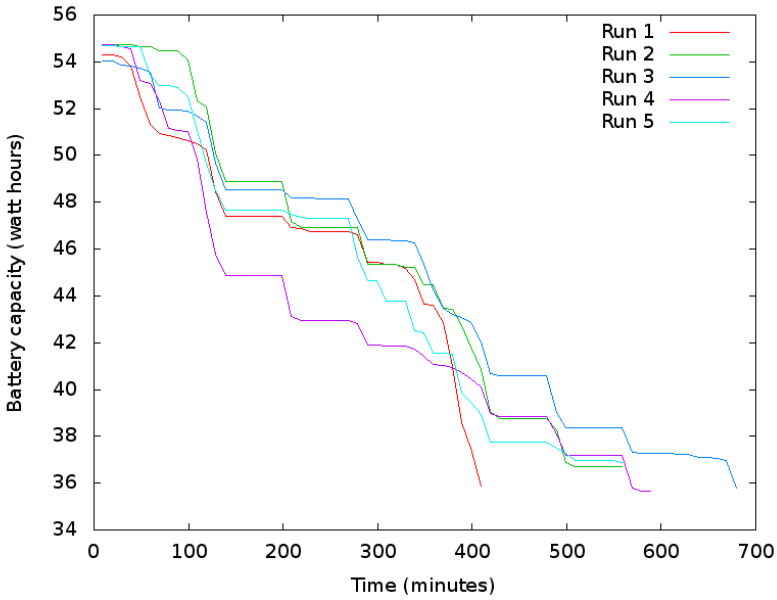

Figure 6.18: A graph showing the simulator's battery discharge over time, after data was reprocessed to swap rudder and sail power consumption figures, when the hormone sensitivity was set to 0.0.

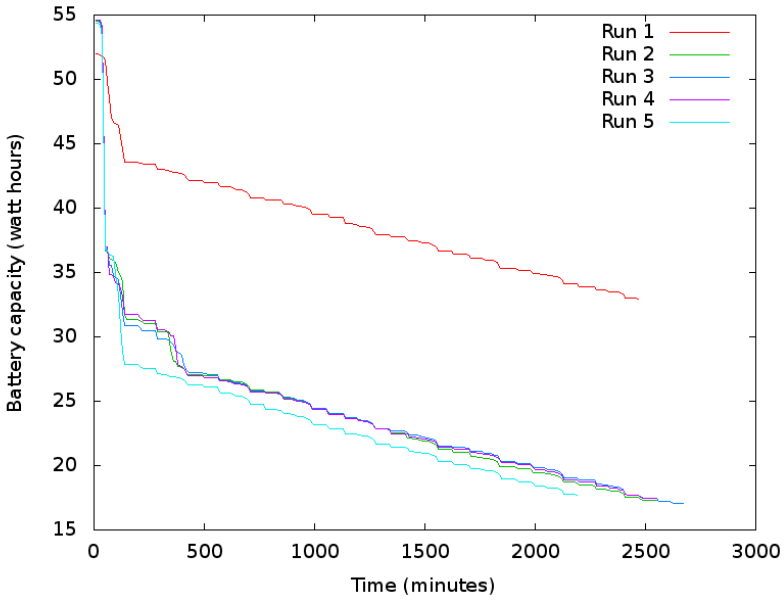

Figure 6.19: A graph showing the simulator's battery discharge over time, after data was reprocessed to swap rudder and sail power consumption figures, when the hormone sensitivity was set to 0.25.

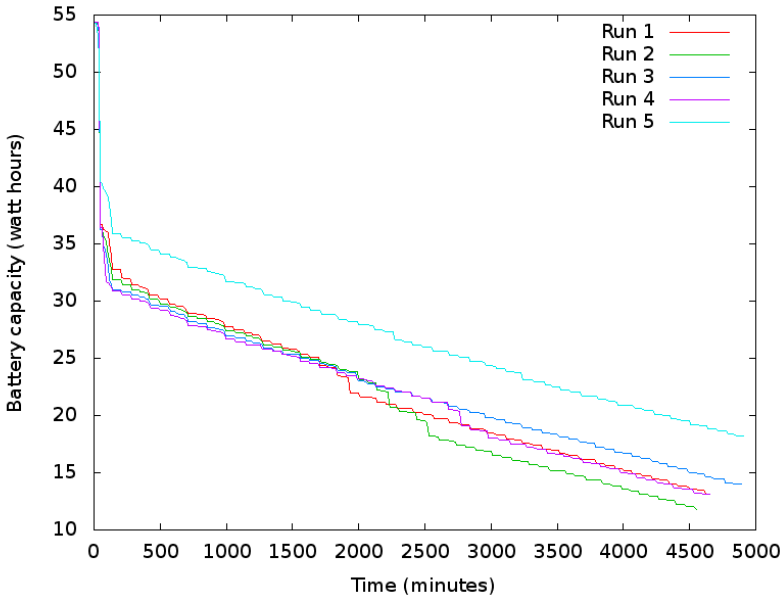

Figure 6.20: A graph showing the simulator's battery discharge over time, after data was reprocessed to swap rudder and sail power consumption figures, when the hormone sensitivity was set to 0.5.

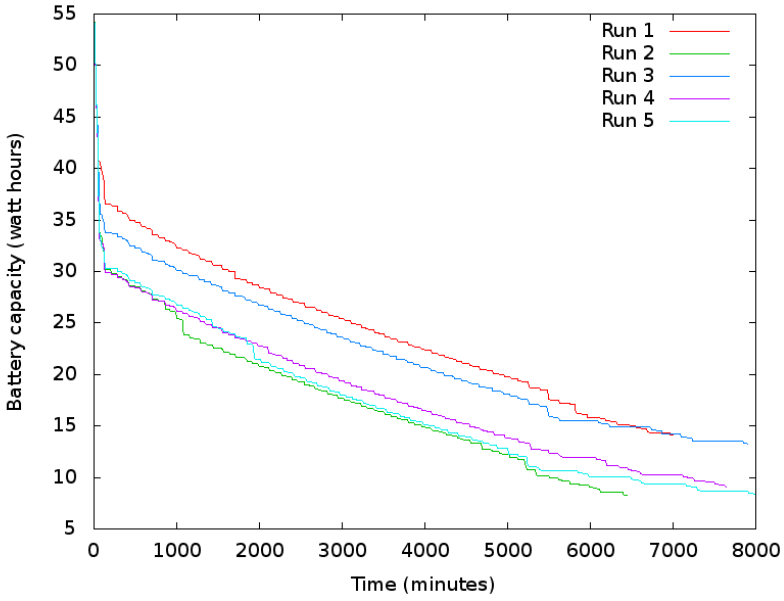

Figure 6.21: A graph showing the simulator's battery discharge over time, after data was reprocessed to swap rudder and sail power consumption figures, when the hormone sensitivity was set to 0.75.

6.4.3 Solar Power Simulations

We have now established a simulated electrical and battery system for a sailing robot and have investigated how changes in hormonal modulation will affect power consumption. What has not yet been considered is that the robot's battery can be recharged from a photo-voltaic solar panel. This has the potential to give perpetual operation to the robot, however as discussed in section 6.2.1 the simulated robot is only able to generate a maximum of 4.75 Watts from its solar panels. On a cloudless day at 52 degrees north in June this will provide around 150 kilo Joules or 41 watt hours of energy. At the equinox (March/September) it will provide around 82.5 kilo Joules or 23 watt hours and in late December it will only provide 20 kilo Joules or 5 watt hours. The previous experiment in section 6.4.1, showed that simulated actuator power consumption will be between 100 and 1400 Joules per hour or 2.4 to 33.6 kilo Joules per day. However, evidence in section 6.2.1 suggests that, using the figure multiplied by 100 as in the previous experiment maybe closer to the reality of a real small scale sailing robot. As this will be in excess of the solar power budget some portion of the day will have to be spent not sailing properly.

There are two potential approaches which can be taken to a hormonally modulated and solar powered system. The first is not to assign an additional hormone to the sun light level and simply have the control system feedback through the battery levels. These are likely to lag behind the solar output by a few hours and thus power wasting behaviour may continue after the sun has gone down. The alternative is to have a second hormone track the level of the sun or the future level of the sun (perhaps one hour into the future). So if we enter a situation where battery is low but we have several hours of sunlight to come then it might be acceptable to go ahead and use lots of power. Conversely if we have significant amounts of battery left but little sunlight coming up then we might wish to be more conservative. This approach could also be extended to cover seasonal variations rather than daily ones which might be of more use for robots with larger batteries that can easily survive several days without recharging. Such robots would see greater benefit from varying their behaviour seasonally rather than daily. However, smaller robots such as the MOOPs or the simulated robot will require daily adjustment since their batteries cannot last more than a day or two.

As there is the potential for some of these experiments to result in perpetual operation of the simulated robot, some termination mechanism is required. So, a course of 400 waypoints each 250 metres apart (100km total course length) was setup. Simulations will terminate either when the battery level reaches zero or when all 400 waypoints have been visited. As with previous experiments the waypoints are setup on a beam reach with the wind perpendicular to the course. The course just causes the boat to sail back and forth across the same area.

Figure 6.22: A graph showing sun elevation and battery levels during the course of the solar power simulation in June.

6.4.3.1 Solar Power Simulations with only Battery Hormone

For the first of the solar experiments no extra hormones were created. This leaves only feedback through the battery level to inform the control system of the effects of sunlight falling upon the solar panels. The battery hormone sensitivity was set to 0.75. This value had been shown in the previous experiment to apply the maximum level of modulation without causing adverse problems for the boat's ability to sail. Five runs of each experiment were carried out to help ensure that any conclusions were not due to a chance event. The experiment was repeated on the simulated dates of June 22nd, September 22nd and December 22nd to approximately cover the summer solstice, autumnal equinox and winter solstice.

It is also of interest that figure 6.22 bears a striking resemblance to figure 4.29 from the sailing robot Pinta which attempted a transatlantic crossing in September 2010, which was discussed in section 4.26. Both the Pinta data and this simulation cover a similar time span. Although this similarity is not a scientific comparison and the Pinta data was subject to high levels of noise and variations due to cloud cover, it does suggest that the gap between the simulation and real world data is quite small.

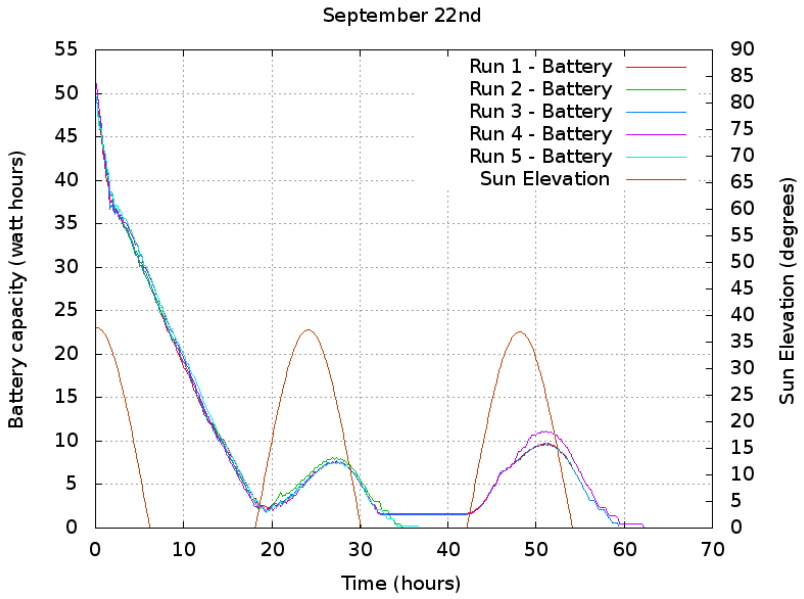

Figure 6.23: A graph showing sun elevation and battery levels during the course of the solar power simulation in September.

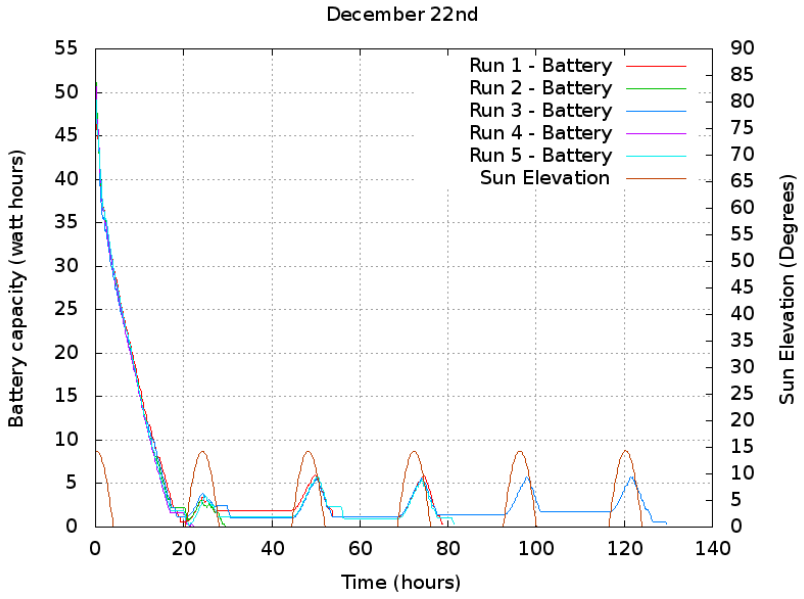

Figure 6.24: A graph showing sun elevation and battery levels during the course of the solar power simulation in December.

Figures 6.22, 6.23 and 6.24 show graphs of the sun elevation and the battery level for these simulations. In the June simulation the course was completed in all cases in around 50 hours. It appears that, given the level of sunlight in June this would allow this course to be repeated indefinitely (or at least until the days get short enough to not provide enough

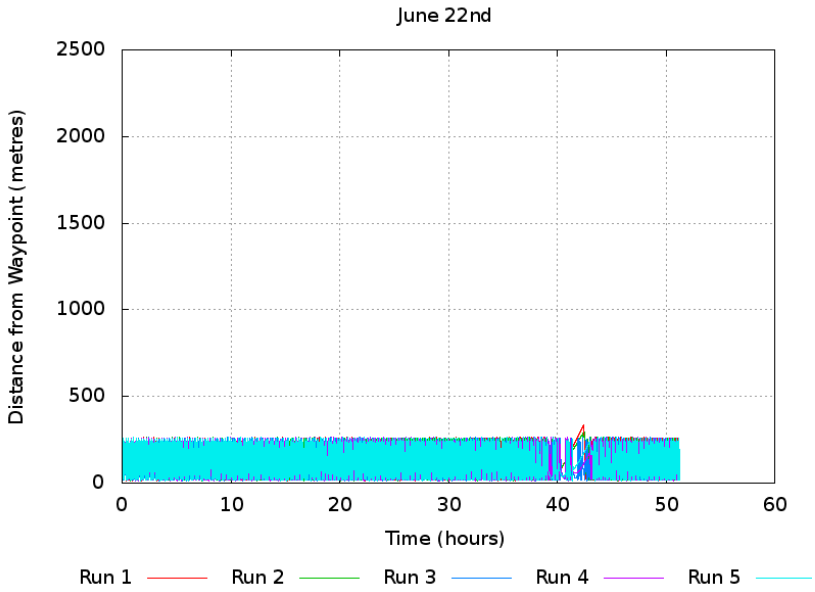

Figure 6.25: A graph showing the simulated robot's distance from its waypoint over time, running in June.

solar power). In both the September and December the course was not completed and a flat battery was the cause of the simulation terminating. The September graph shows that two of the experiments ran out of battery after only approximately 35 hours. While the other 3 have suppressed their control systems sufficiently to consume no power all night and only begin to consume any power again the next morning. This same behaviour is seen, to an even greater extreme in December, where two simulations end within 30 hours, another two manage to last 80 hours by performing no activity at night and another manages to last nearly 130 hours. While this may appear at first to be a desirable situation, it should be noted that the reason they are consuming no power during the night is because the battery is so low that the hormone has totally suppressed any actuator movement and therefore, the boat is just drifting.

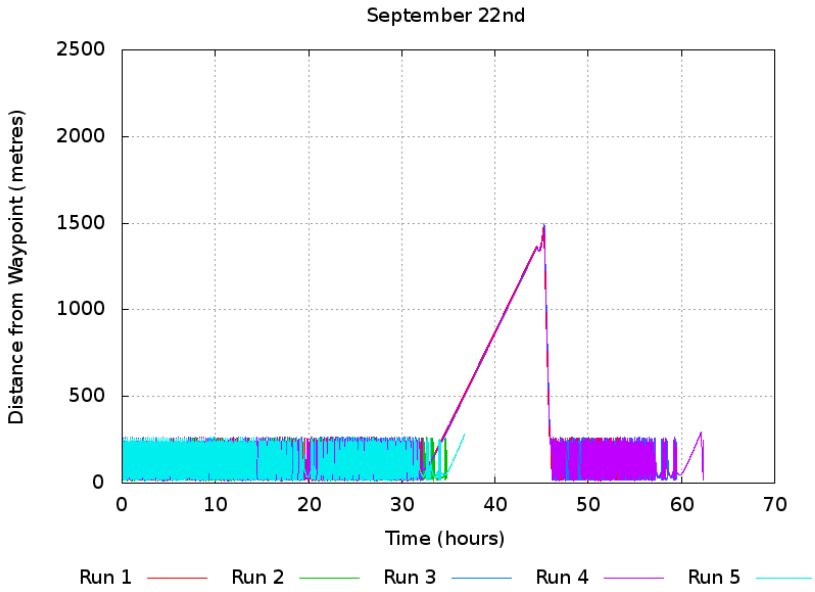

Figure 6.26: A graph showing the simulated robot's distance from its waypoint over time, running in September.

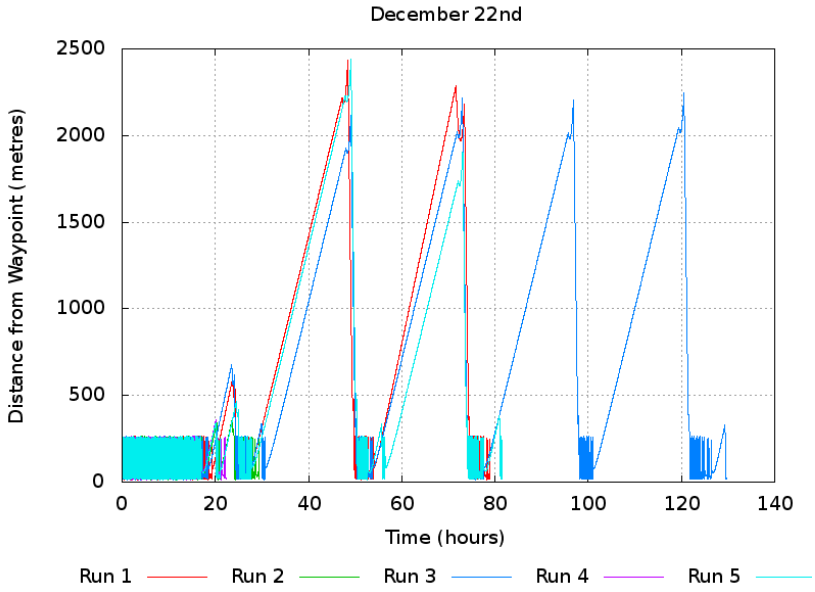

Figure 6.27: A graph showing the simulated robot's distance from its waypoint over time, running in December.

This can be confirmed in figures 6.25, 6.26 and 6.27, which show the distance of the robot from its waypoint over time. As the distance between waypoints is only 250 metres, under normal circumstances this distance should never be significantly more than 250 metres. In cases when it does dramatically exceed this it is an indication that the control system has

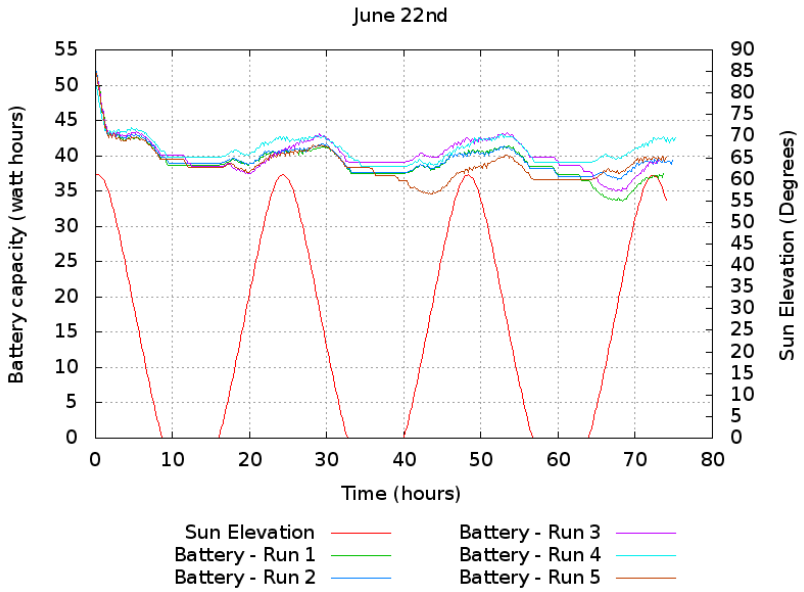

Figure 6.28: A graph showing sun elevation and battery levels during the course of the solar power simulation, running in June with an additional solar hormone.

not succeeded in keeping the robot sailing towards the waypoint and has instead let it drift away or overshoot its waypoint.

In the June simulation the robot is never more than approximately 300 metres from the waypoint and this value is constantly fluctuating. In the September simulation some of the experiment runs manage to travel as far as 1500 metres from the waypoint. In the December simulation the robot spends large amounts of time overshooting the waypoint and gets over 2000 metres from the waypoint on several occasions. These overshoots are because it was left in a sailing state when the control system stops performing any action due to a low battery. This behaviour is not ideal and what needs to be done is to lower power consumption during the day so that battery levels are higher at night. Although for many applications a robot which can remain within 2km of a target point is still usable.

6.4.3.2 Solar Power Simulations with Sunlight Level and Battery Hormone

A second hormone was introduced which was produced in inverse proportion to the elevation of the sun. This hormone is only inhibitory, and will produce its maximum level of inhibition when the sun is below the horizon and no inhibition when the sun is directly overhead. The previous experiment was rerun with this extra hormone. The sensitivity to the new hormone was also set to 0.75.

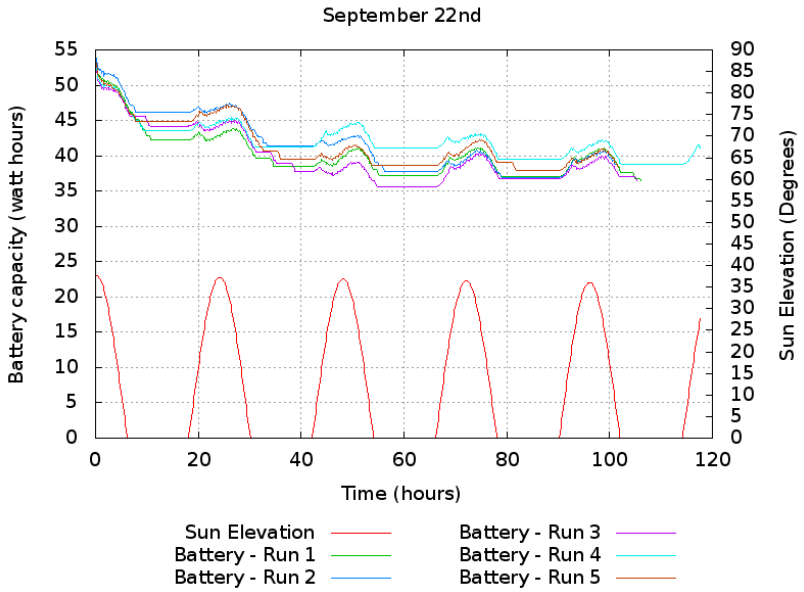

Figure 6.29: A graph showing sun elevation and battery levels during the course of the solar power simulation, running in September with an additional solar hormone.

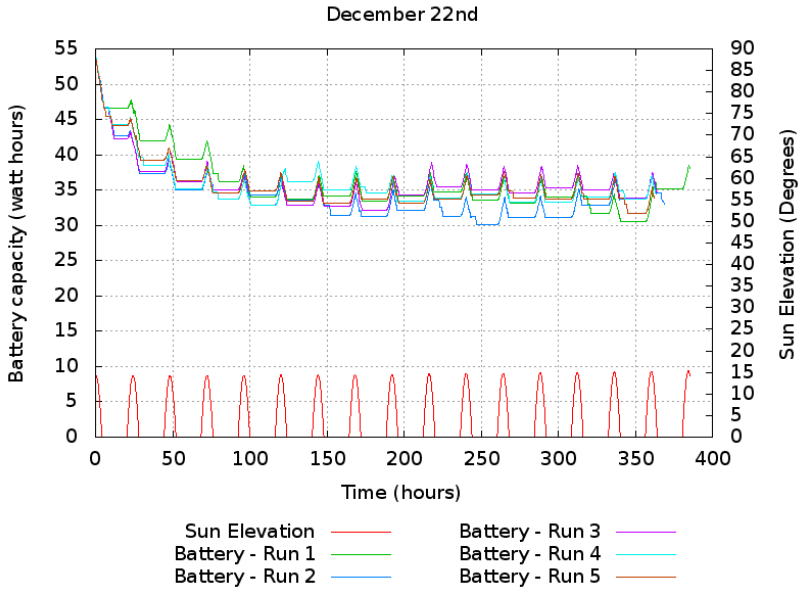

Figure 6.30: A graph showing sun elevation and battery levels during the course of the solar power simulation, running in December with an additional solar hormone.

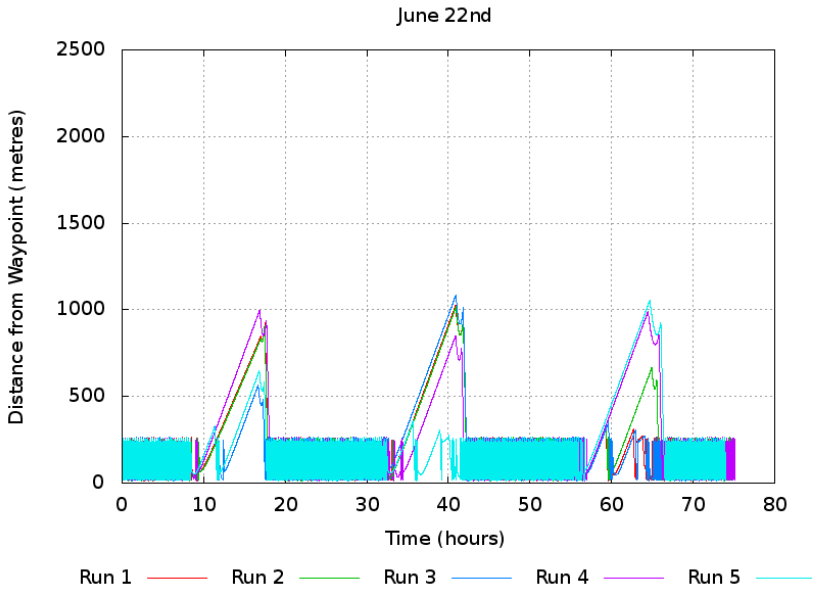

Figure 6.31: A graph showing the simulated robot's distance from its waypoint over time, running in June with an additional solar hormone.

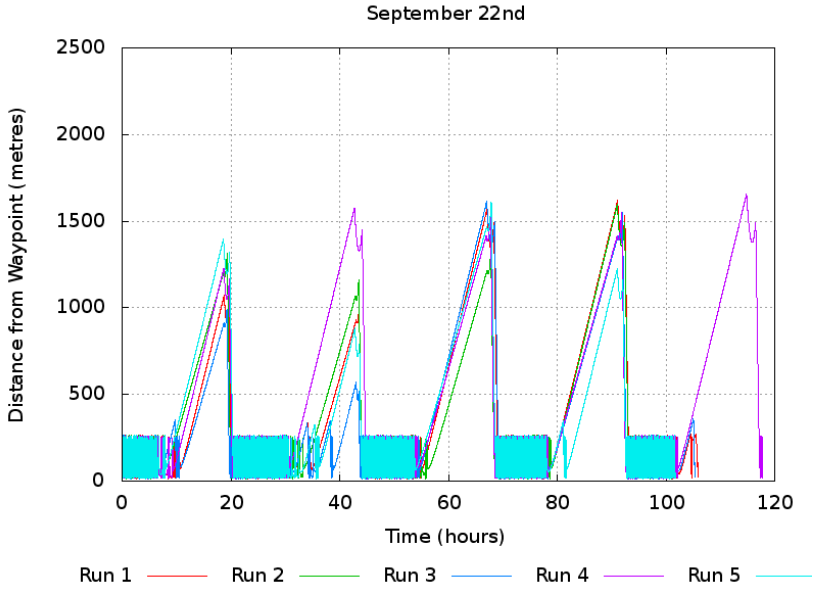

Figure 6.32: A graph showing the simulated robot's distance from its waypoint over time, running in September with an additional solar hormone.

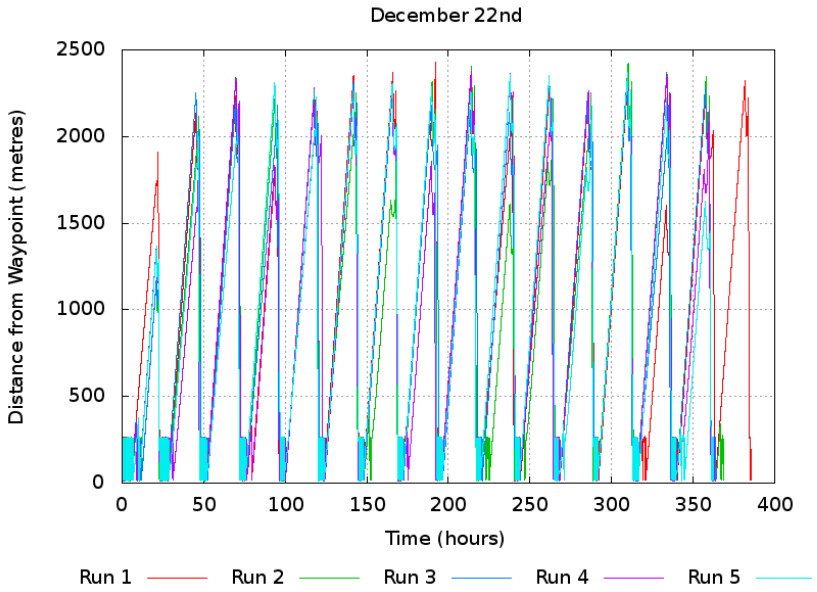

Figure 6.33: A graph showing the simulated robot's distance from its waypoint over time, running in December with an additional solar hormone.

Figures 6.28, 6.29 and 6.30 show the battery and sunlight levels for these experiments. In all cases the robot was now able to complete the course, with the June simulations taking around 75 hours, the September simulations between 80 and 120 hours and the December simulations over 350 hours. However, in all cases this was achieved at least in part by having the robot make no movement at all during the night. Figures 6.31, 6.32 and 6.33 show the distance from the waypoint. Compared with figures 6.25, 6.26 and 6.27 we can see that distances from the waypoint have increased. In all cases the amount of time spent more than 250 metres away from the waypoint has increased substantially, due to the amount of time the control system has spent being suppressed, to the point where the boat cannot sail. However, it is only by doing this that the robot has been able to complete the course.

Given the power budget constraints it would be impossible to complete the course without having to totally suppress the control system during September or December. However, it may have been possible during June suggesting that the hormone sensitivity level for the sun elevation hormone is too high.

These results show an example of how two hormones can help to reinforce a given behaviour. They also show an example of a sustainable homeostatic system which is able to maintain battery levels perpetually. Although this is essentially achieved through a "deep sleep" which totally suppresses the control system for much of the day. Despite the obvious dangers of doing this the boat never strays further than 2.5km from its waypoint. In a number of off-shore scenarios, if the boat could remain within 2.5km of its intended course

then this might be sufficient for safe operation. However, this figure of 2.5km is derived from simulations which maybe operating at unrealistic speeds for a real boat and in an overly simplistic world in which there are no tides or currents and no obstacles to avoid. If sailing longer distances between waypoints (providing the boat is reasonably stable and able to stay on course without constant correction) then this might be sufficient to suppress the control in this way when some distance from the waypoint and only "wake up" and begin moving actuators once a significant distance has been travelled.

6.5 Chapter Summary

This chapter presented the results of four experiments run on a simulated sailing robot. The first of these experiments (section 6.4.1) fixed the hormone levels at preset values, it was designed to identify the limits of hormonal modulation and to test that it was plausible to apply it at all. This experiment demonstrated that it was possible to apply hormonal modulation to the simulated robot. It showed that decreasing the effective hormone concentration (which makes the hormone suppressive) was correlated with both the amount of energy used for each hour and each kilometre of travel. The effect on sailing performance was less conclusive with little difference being shown until a certain threshold was reached, after which performance suddenly dropped to the point where the robot could no longer sail in a useful manner.

The second experiment (section 6.4.2) varied hormone levels in response to battery levels and tested what effect this had upon the time and distance which could be sailed on a single battery charge. It showed that by using a hormone sensitivity of 0.75 the robot could sail over 13 times further/longer than with a sensitivity of zero (no hormonal modulation).

The third experiment (section 6.4.3.1) added a simulated solar panel to the robot. This allowed the simulated battery to be recharged. A simulation was undertaken in June, September and in December. Despite the presence of an inhibitory hormone being produced as battery levels fell in most of the September and December simulations the robot was unable to complete a 100km course without totally discharging its battery. In the final experiment (section 6.4.3.2) a second hormone was introduced. This was produced in proportion to the sunlight levels and by suppressing the control system at night was able to allow the robot to complete its mission in each case.

6.6 Conclusions

Revisiting the research question from chapter 1 we can conclude that in simulation an artificial neuro-endocrine system can manage power consumption of a sailing robot through

actuator movement. The effects upon sailing performance appear to be a form of binary step function. Few effects are seen up until a critical point, but after this point performance rapidly deteriorates to the point at which a course cannot be followed. There is a possibility that during the period where inhibition levels are increased we are effectively eating into a safety margin created by the use of an excessively high gain in the network training. Therefore, sailing in stronger winds or heavier seas may result in sailing performance deteriorating earlier. A possible extension for future work might be to include sensing of wind speed and sea state and representing these with their own hormones which excite the control system in more difficult conditions and overcome the suppression of power saving systems. An alternative method of detection might be to analyse heading error over time rather than directly using sea state or wind speed, this would also not require additional sensors on a real robot.

The work in this chapter has also demonstrated that multiple hormones can be used to reinforce a behaviour, in this case a hormone linked with the expected level of energy being received by a photo-voltaic solar cell. Through this method we are able to demonstrate the potential to achieve indefinite operation (at a latitude of 52 Degrees with only sunny days). This shows that an artificial neuro-endocrine controller is able to maintain a constant internal state with respect to battery charge remaining, despite external conditions changing from 8 hours to 16 hours of sunlight and from 20 to 150 kilo joules of energy input per day. However, when faced with a power budget that can clearly not sustain round the clock operation of the robot, the result is to effectively have the robot sleep in a diurnal cycle. This could be seen as an emergent property rather than an intended design feature although the capabilities of the resulting system bear a resemblance to sleeping robot systems discussed in section 2.6.3.4. The user of any real robot undertaking such an approach would need to be aware of the potential danger of having the entire robot effectively shut down for several hours per day. Potentially other excitory hormones could be introduced which could wake the robot from this state should a dangerous situation be encountered, although this will still require enough power to run systems which can detect the danger situation.

To test the scalability of this work these experiments now need to be repeated using a real robot. This will establish if they can scale from a relatively tame simulated environment without waves, wind shifts or gusts to a much harsher real environment. This will be discussed in the next chapter.

Chapter 7

Power Management Experiments and Results using Sailing Robots

7.1 Introduction

This chapter presents the results of repeating both the fixed hormone level and variable hormone level experiments from the previous chapter using a pair of real sailing robots. It also describes the process of adapting the neuro-endocrine controller described in chapter 5 to operate on a real robot, some of the difficulties encountered in setting up appropriate experimental conditions and operating robots outside during winter for prolonged periods of time.

A pair of MOOP robots (known as MOOP0 and MOOP1) described in section 4.1.5 were used for these experiments. These robots are controlled by a relatively power hungry Gumstix single board computer and constantly stream telemetry data over an 802.11 wireless network. The Gumstix simplifies the process of developing software as it runs a reasonably complete Linux based operating system which provides useful facilities such as networking and file system support. However, the Gumstix combined with the wireless network uses between 3 and 4 watts continuously compared with a microcontroller such as a PIC which would use less than 0.1 watts. The ease of use advantages were considered to be significant enough to accept this tradeoff. This choice reduced the operational lifetime of the MOOP to around 4 or 5 hours and meant that computers and communications equipment actually represented the bulk of the power budget. This is not necessarily the choice that somebody wishing to operate a MOOP in a real world situation would wish to opt for, but for the purpose of these experiments it was seen to be sufficient.

As in the simulator experiments in the previous chapter only actuator power consumption was actually considered in these experiments. Due to these limitations it was only possible

Algorithm 7.1 The algorithm for producing the MOOP sail training data.

```
if relwind >= 0 and relwind < 15 then new_sail_pos = 0
else if relwind >= 15 and relwind < 50 then new_sail_pos = -1
else if relwind >= 50 and relwind < 80 then new_sail_pos = -2
else if relwind >= 80 and relwind < 120 then new_sail_pos = -3
else if relwind >= 120 and relwind < 145 then new_sail_pos = -4
else if relwind < 180 then new_sail_pos = -5
else if relwind >= 180 and relwind < 215 then new_sail_pos = 5
else if relwind >= 215 and relwind < 240 then new_sail_pos = 4
else if relwind >= 240 and relwind < 280 then new_sail_pos = 3
else if relwind >= 280 and relwind < 310 then new_sail_pos = 2
else if relwind <= 345 then new_sail_pos = 1
else new_sail_pos = 0
```

to recreate the fixed hormone and the variable battery hormone experiments that were performed in the simulator. It was not possible to perform the solar power experiments. Section 7.4 contains the results of the fixed hormone experiment and section 7.5 contains the variable hormone experiments.

7.2 Adapting Neural Network Controllers to the robot

For these experiments, the same neural network software was deployed on the robot as had been used in the simulator although, some low level differences existed as the API to the MOOP is not quite identical to the simulator's. The neural networks were retrained based on algorithms generated during previous testing of the MOOPs, both MOOPs used the same neural network despite small differences in their rudder systems (see section 4.1.5 for more details). These are shown in algorithms 7.1 and 7.2.

As with the simulator power consumption was tracked only through the movement of actuators. Although it is theoretically possible to measure the power draw on the batteries using an ammeter circuit there are difficulties in obtaining precise measurements. Adding this extra complexity to the robot's electronics only creates further points of failure and increases the weight of the robot. Therefore, the same method used for power tracking in simulation (described in section 6.2.1) was used. This allows for close comparison with the simulator output, but there will be some difference between these numbers and the real level of power consumption.

Algorithm 7.2 Proportional rudder setting algorithm for producing MOOP rudder training data.

```
pgain=1.0
if abs(heading_err)<15 then
   new_rudder_pos = 0
else
   new_rudder_pos = heading_err * pgain
if new_rudder_pos<-5  then
   new_rudder_pos=-5
else if new_rudder_pos>5  then
   new_rudder_pos=5
rudder_change = new_rudder_pos - curr_rudder_pos
```

7.3 Location and course design

Experiments were under taken at Llyn-Yr-Oerfa (52.4 degrees North, 3.87 degrees West), a lake approximately 12 miles east of Aberystwyth. This provided a safe environment where there was no traffic on the water and if something went wrong with a robot it would usually be blown onto an easily accessible shoreline. Concerns were raised that if experiments took place on the sea or a larger lake that finding the robot in the event of a failure would be very difficult (especially if it was dark). Operating on bodies of water with other traffic was also considered dangerous as the MOOPs small size meant they were difficult to see even at close range. A collision between a MOOP and any other craft would almost certainly leave the other craft without significant damage however, the MOOP might not fare so well.

The useable area of the lake (portions of the lake are covered in reeds and not suitable for sailing on) is approximately 120 x 200 metres. The local topography favoured testing during winds from either the south west, west, northwest, north east or east. During the development of the robots attempts were made to sail triangular courses on Llyn-Yr-Oerfa but this proved to be difficult for a number of reasons. Firstly the size of the lake tended to create small but very high frequency waves which would slow the boat down especially when sailing to windward. When combined with small wind shifts caused by local topography the robot tended to accidentally tack every few minutes. The small size of the lake often resulted in the boat washing up on the leeward (downwind) bank of the lake within a few minutes if a tack failed and the boat end up pointing into the wind. This caused a number of experiments to end prematurely. Generally it was considered that if a robot was ashore for more than one minute then the experiment needed to be terminated. The reason for this is that when the robot runs aground, it is no longer able to correct any course error as rudder movements have not effect. It will also typically tip the robot onto its side which will eventually skew wind

sensor readings (wind sensor readings are heavily dampened but a few minutes of changes will lead to a change in the dampened direction) and lead to the sail also being incorrectly moved. This will most likely skew power consumption figures to the point where it could effect the conclusions of the experiment. If the robot was reached within one minute (only possible when standing near to it at the time it ran into difficulty) then it would be gently pushed back into deeper water. If the robot then sailed back into the land immediately after, then the experiment was terminated.

The wide profile keel of the MOOPs also caused drag reducing speed and overall sailing efficiency. As sailing to windward is the slowest point of sail this will be affected worst by these inefficiencies. The MOOPs also struggled with down wind sailing and would jibe (moving the back of the boat through the wind) violently more than once per minute. To give a MOOP reasonable chances of continuing to sail a stable course for a prolonged period, it was decided just to reach back and forth with the wind perpendicular to the boat. This required the course to be adjusted depending on conditions, however under most conditions a course which started at the southern end of the lake and sailed approximately 150 metres north was sailable. A second waypoint was placed about 20 or 30 metres north of the start point, the boat would then sail back and forth between these two points. A photograph of the lake and this course are shown in figure 7.2. The location this was taken from is marked as "Photo Location" in figure 7.1. A typical course is show in figure 7.1, the boat would be launched at the point labelled "start", then sail to waypoint 1, then to waypoint 2 and then repeat the course between waypoint 1 and 2 a number of times before attempting to sail to the point labelled "end". This end point is deliberately set some way inland so that the boat will sail itself into the shore when the course was completed.

The robot was considered to have reached a waypoint when according to the on-board GPS it was within 20 metres of that waypoint. This was to ensure that despite any degradation in the quality of the GPS signal (which has a typical accuracy of around 5 metres) the boat would still be able to consider that it had reached a waypoint. An unintended consequence of this is to shorten the actual distance the boat sails, if for example there is 130 metres between waypoint 1 and waypoint 2, the boat will only sail approximately 90 metres on average when sailing back and forth continuously. Travelling from the start point to waypoint 1 and back to waypoint 2 typically took between 10 and 20 minutes depending on the wind speed.

7.3.1 Suitability of the Robot and Choice of Test Location

Due to the limitations of the MOOP robot it was not possible for an experiment to last more than 4 or 5 hours. The limited battery life which was primarily due to the power hungry Gumstix and 802.11 Wifi Access Point prevented any chance of true perpetual operation.

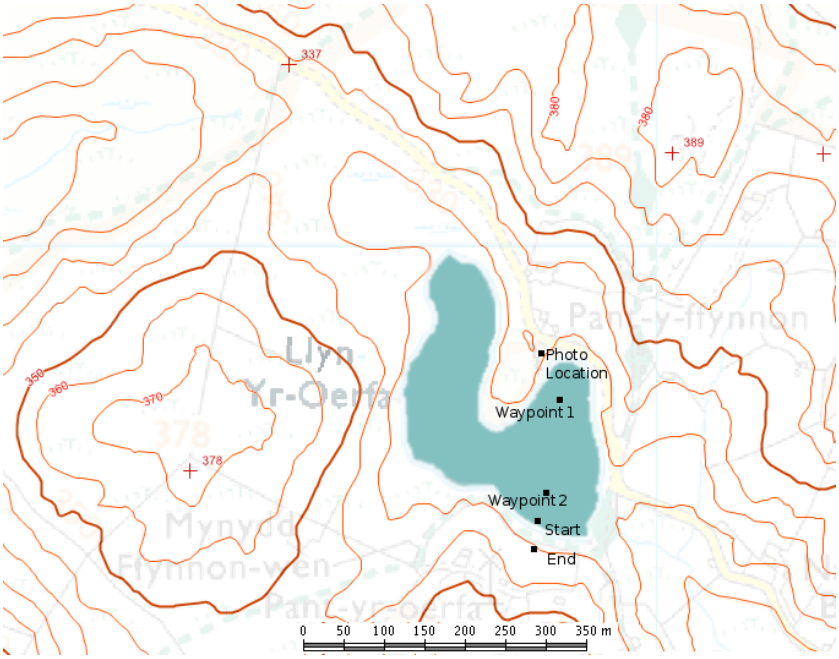

Figure 7.1: A map of Llyn-Yr-Oerfa showing the surrounding terrain. Contours are at 10 metre intervals. © Crown Copyright/database Right 2009. An Ordnance Survey/(Datacentre) supplied service.

If the MOOP were re-engineered with only a small lower power microcontroller then this might just be a possibility. Although, developing software that will fit within these memory constraints adds to the challenge. Despite these limitations it was only due to the small size and flexibility of the MOOPs that allowed these experiments to be undertaken through such varied conditions during a particularly cold winter. To have used a significantly larger sea going boat would have required chase boats, weather conditions suitable for humans to sail in and at least two people to launch and recover the boat. Due to the sensitivity to minor adjustments in weight distribution or control system parameters, developing the MOOPs to the point where their control system was sufficiently stable and fine tuned took over 4 months after the first "successful" autonomous sailing, a larger boat would have most likely been less sensitive to these adjustments and faster to develop.

This worked also highlighted a number of strengths of the MOOP robots. They demonstrated their resilience in sailing often for several hours per day across three or four subsequent days. No major faults occurred as a result of wear and tear although some minor problems occurred with unsecured wires breaking and batteries dislodging. The boats were subjected to gusts of wind thought to exceed 40 km/h (no wind speed monitoring was available) and to have sailed for a total in excess of 100 hours. They could also typically be deployed within 10-15 minutes of arrival at the lake, when compared with the hours involved in launching

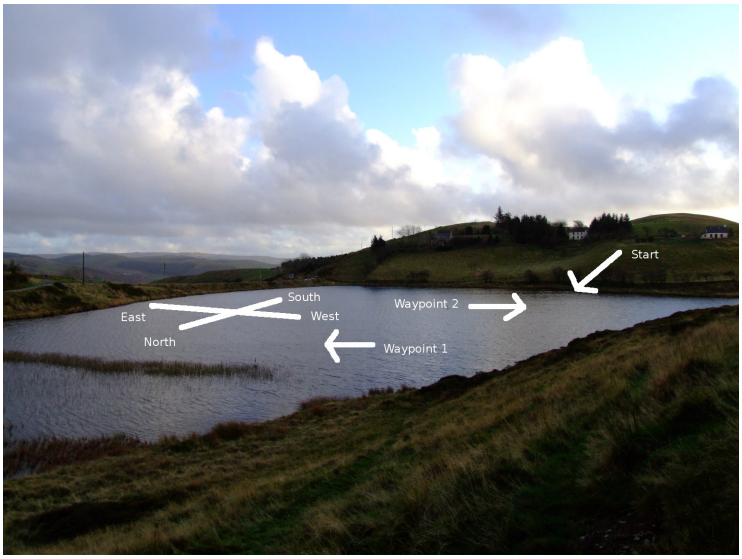

Figure 7.2: A photograph of Llyn-Yr-Oerfa showing the approximate course that was sailed. The location this was taken from is marked in figure 7.1 as "Photo Location".

larger boats such as BeagleB or Pinta this is a major advantage.

Due to the topography around the lake and its relatively coastal location relatively (12 miles inland) often meant that wind speeds would drop significantly soon after sunset. This caused the failure of a number of experiments which had been started late in the day (especially between November and February). This also presented the challenge of recovering a small robot from the middle of a dark lake at night. A series of small LEDs were added to the boats to help locate them in the dark and an inflatable dinghy used to recover them in the event of failures or lack of wind during darkness. Working with small custom made robots at a relatively remote location also brought with it a number of challenges. No mains electricity was available, so if recharging of robots or laptops was required it had to be done from car batteries. Although not required to keep an experiment running, the usual procedure was to frequently monitor telemetry data received over a wireless network using a laptop computer. On many occasions it was the laptop battery life rather than the robot battery which was the limiting factor. The simplest solution to this was to bring along a selection of laptops. Additional problems occurred on rainy days as the main laptops in use were not waterproof. When hardware faults did occur it was not usually possible to fix them at the lake side. Anyone wishing to undertake similar work should bear in mind the immense amount of time that will be lost to hardware failures, inappropriate weather and travelling. They should also consider the difficulties of managing a robot which is essentially a network of floating computers in a harsh, sometimes dark, cold and wet environment.

7.4 Fixed Hormone Level Experiments

The fixed hormone experiment that was carried out in simulation (as described in section 6.4.1) was recreated on the robots. A single hormone is produced in response to battery level. Its concentration was calculated using the formula:

$$0.023b - 1.03 \tag{7.1}$$

Where b is the amount of battery remaining in Watt Hours (between 0 and 55 Watt Hours)

The hormone will have an excitory effect upon the neural networks controlling the sail and rudder when the battery level is above 44.8 watt hours and an inhibitory effect below this. Sensitivities to the hormone of 0, 0.25, 0.5 and 0.75 were used. By multiplying the sensitivity by the hormone concentration we can create an "effective hormone concentration", which accounts for the level of hormone and the neural network's sensitivity to it. As with the simulation work three metrics were used to evaluate the success of each experiment. These metrics compared the power consumption for each hour of sailing, for each kilometre travelled and evaluated the course efficiency by dividing the course length by the total distance travelled.

The experiments were carried out over 7 days between the 3rd and 23rd of September 2009. This was during a period of variable weather with winds varying from Beaufort force 1 to 6 and blowing from the West, South West and North East. This represents almost the full range of conditions in which it was practical to sail at Llyn-Yr-Oerfa. A summary of the weather conditions are shown in table 7.1.

The course consisted of two waypoints approximately on a beam reach course and around 150 metres apart. The target course length was around 3-400 metres so this required one or two loops of the course. Due to wind variations during this time the course had to be altered several times to ensure a beam reach was still sailed. As a result course lengths were not always identical. A mistake was made entering waypoints on the experiments on September 11th 2009 and the course was accidentally several times longer than it was supposed to be giving a total length of approximately 1.5km instead of 3-400 metres. Originally it was intended to run the experiments with hormone sensitivities of 0, 0.25, 0.5, 0.75 and 1.0 as had been done with the simulator in section 6.4.1. However, it was found that some of the experiments with a sensitivity of 0.5 and all of those with a sensitivity of 0.75 ended up with the robot sailing into the shore before the experiment was completed. Therefore it was decided not to run experiments with a sensitivity of 1.0 as they were all expected to fail. To confirm that these were not random chance events several short experiments with

Date	Hormone Sensitivity	Battery Level (Watt hours, max 55)	Wind Observation
03/09/2009	0	5	Force 3 W
04/09/2009	0.25	5	Force 3-4 W
04/09/2009	0.5	5	Force 3-4 W
04/09/2009	0.75	5	Force 3-4 W
11/09/2009	0	25	Force 1-2 NE
03/09/2009	0.25	25	Force 3 W
03/09/2009	0.5	25	Force 3 W
03/09/2009	0.75	25	Force 3 W
08/09/2009	0	25	Force 5-6 W
22/09/2009	0.25	55	Force 4-5 SW
21/09/2009	0.5	55	Force 3-4 SW
22/09/2009	0.75	55	Force 4-5 SW

Table 7.1: Wind observations during the fixed hormone experiments.

a sensitivity of 0.8 were also run and these too all resulted in the robot sailing into the shore before it could complete the course. It was intended to run half of the experiments on MOOP0 and the other half on MOOP1, with two robots simultaneously on the water and running with different parameters at any one time. Unfortunately MOOP0 suffered a series of faults during this time and only took part in one experiment on September 21st.

7.4.1 Results

A summary of the results from these experiments is presented in table 7.2. Graphs comparing effective hormone concentration with course efficiency, energy use per hour and per kilometre are shown in figures 7.3 and 7.4.

There appears to be a very wide spread in the results for experiments with a zero sensitivity to the hormone. As shown in table 7.2 and figure 7.3 there was a spread between 1767.25, 5716.6 and 11140.46 Joules per hour and 2260.25, 4610.76 and 8040.17 Joules per kilometre despite no change in parameters. This level of variation is greater than that seen between any of the other experiments. This suggests that there is a possibility that the other results are all subject to such levels of variation. A larger sample size might help to overcome this, although obtaining such large sample sizes is difficult as these results alone represent nearly three weeks of work during relatively favourable weather.

Figure 7.3 suggest there maybe a correlation between the concentration of the hormone and amount of energy used. As with the simulator data a test of the correlation will be

Effective Hormone Concentration	Hormone Sensitivity	Battery Level (Watt hours, max 55)	Robot	Joules per km	Joules per hour	Course Efficiency (course length / distance travelled)	Completed course yes/no
0	0	5	MOOP1	2260.05	1767.25	96.25%	Yes
-0.23	0.25	5	MOOP1	2928.96	3859.55	123.90%	Yes
-0.46	0.5	5	MOOP1	1926.26	2398.57	129.15%	No
-0.69	0.75	5	MOOP1	661.54	613.34	132.01%	No
0	0	25	MOOP1	4610.76	5716.6	111.03%	Yes
-0.11	0.25	25	MOOP1	3831.18	4538.13	117.82%	Yes
-0.23	0.5	25	MOOP1	3271.92	4122.2	97.49%	Yes
-0.34	0.75	25	MOOP1	2853.87	3412.55	141.09%	No
0	0	55	MOOP1	8040.17	11140.46	91.01%	Yes
0.06	0.25	55	MOOP1	3442.04	4333.32	85.71%	Yes
0.12	0.5	55	MOOP0	5294.21	3829.29	112.26%	Yes
0.18	0.75	55	MOOP1	3663.45	4824.05	215.73%	No

Table 7.2: Summary results from the robot fixed hormone experiment. Raw data including total distance covered, time taken, total power consumption, date and weather conditions can be found in Appendix D.

calculated using a one tail Spearman's rank correlation coefficient (rho) run in the R statistics program. This leads us to generate the following statistical hypothesis:

H_0: Reducing hormone concentration will not reduce energy use.

H_1: Reducing hormone concentration will reduce energy use.

7.4.1.1 The correlation between hormone concentration and energy use per hour:

$P = 0.5608684$

Using 99% significance level

$n = 12$

$P < 0.678$

Reject H_1

Using 95% significance level

$P > 0.503$

Accept H_1

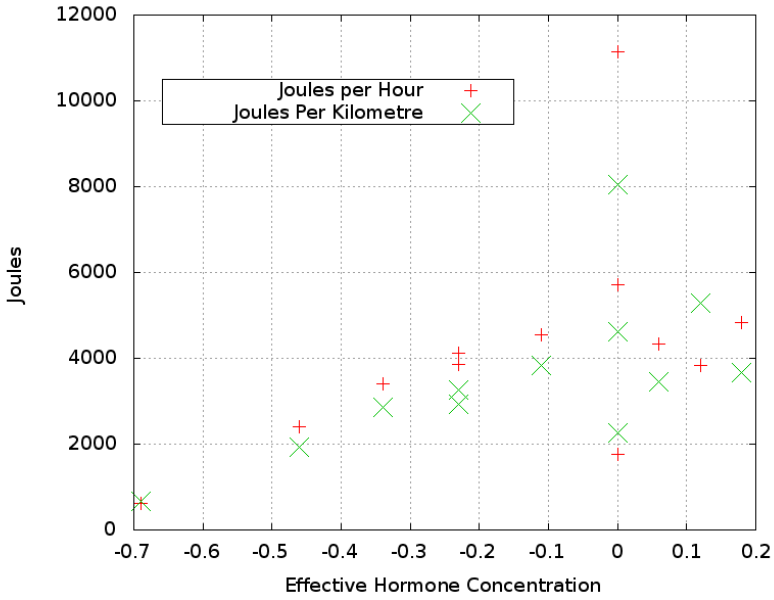

Figure 7.3: Robot Hormone concentration vs Power consumption per kilometre and per hour.

7.4.1.2 The correlation between hormone concentration and energy use per km:

$P = 0.7090223$

Using 99% significance level

$n = 12$

$P > 0.678$

Accept H_1

This leads us to accept the hypothesis that increasing the hormone concentration will increase the power consumption. From these correlations it would appear that reducing hormone concentration in the robot still has an effect upon energy usage. Although, the lack of results with a hormone sensitivity of 1.0 reduces the sample size and thus reduces the probability of accepting the alternative hypothesis.

7.4.2 Conclusions

These experiments have indicated that it is possible to modulate a neural network that is acting as a robot control system, this reduces the magnitude of actuator movement in order to reduce power consumption while still managing to complete the required course. There is a point at which the level of modulation prevents the robot from achieving its objectives, in these results this begins to occur when the effective hormone concentration reaches -0.34 and consistently occurs beyond a level of -0.46. It is likely that this threshold may change

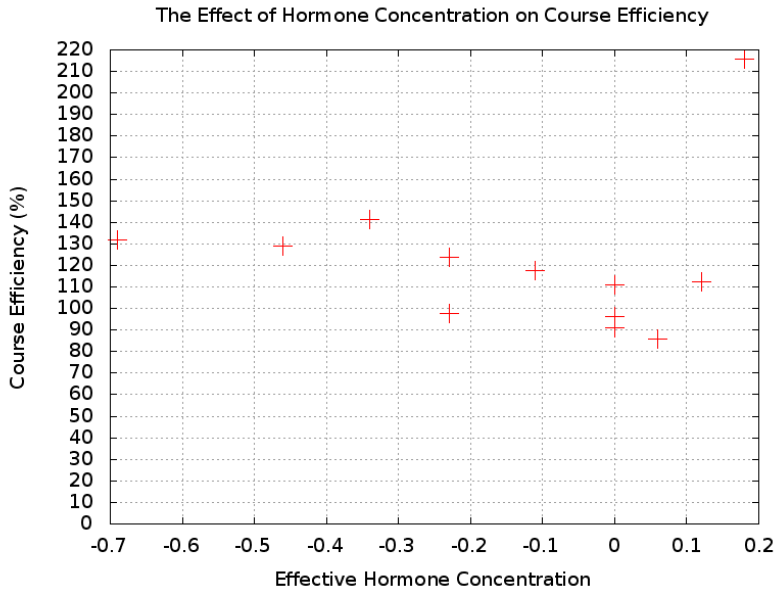

Figure 7.4: A graph comparing course efficiency and hormone concentration in the robot version of the fixed hormone experiment.

depending upon the conditions the robot is operating in. This still suggests that a hormone based controller can modulate the power consumption of a real sailing robot to an extent and that while this does ultimately impact upon sailing performance it still allows some level of useful modulation. This also suggests that it should be possible to repeat the variable hormone experiment as performed on the simulator in section 6.4.2 upon the robot and obtain similar results.

7.4.3 The effects of Over Stimulation

The neural networks in the control system output a value which represents how far to change the current actuator position. Suppressing the output of the neural network (by using a negative hormone concentration) will reduce the probability that the neural network will decide a change is required, exciting the network (using a positive hormone concentration) will increase the probability that the actuator will move. With regards to rudder movements increasing the magnitude of movements will cause more sudden and dramatic changes in direction while, decreasing the magnitude will result in smaller more subtle changes. With the sail increasing the magnitude can actually result in the sail moving beyond the required position and actually causing the boat to sail the wrong course. This was observed in some of the experiments with positive hormone concentrations. This can be considered the equivalent of letting out the sheet in a traditional sailing boat even when the sail was correctly set. This

has the effect of causing the boat to turn away from the wind. As the sail setting neural network acts only to react to the current wind direction and not the target heading situations can occur where this actually counter productive and stops the boat sailing its desired course. It is thought to have caused a number of experiments to fail to complete the entire course that they had been set. This can be seen in table 7.2 where several of the experiments with positive effective hormone concentrations are seen to not complete. This exposes a fundamental shortcoming in our control system design, that the sail has an influence upon steering and that the steering neural network needs someway to influence the sail position network to help with steering. It also raises the question as to whether or not it makes any sense to increase stimulation of the sail setting network beyond its "normal" level.

7.5 Variable Hormone Experiments

The battery hormone experiment which was performed in simulation in section 6.4.2 was repeated using the robots. This experiment produces a hormone in proportion to the battery level instead of leaving it at a fixed value as in the previous experiment. The battery level is artificially represented by a variable in the control system instead of being read through sampling the actual battery. This is because, as explained in section 7.1, there are a number of difficulties in accurately determining battery level and we are not considering the power consumption of the on-board computers, sensors and telemetry equipment. As with the simulator each rudder or sail movement will reduce the remaining battery level. In section 6.2.1 it was calculated that each rudder movement would take 0.3814 Joules and each sail movement would take 1.1192 Joules per position and that both had 11 positions in total. A mistake in the simulator experiments accidentally reversed these two numbers for consistency and to make comparisons that reversal has been preserved. In the simulator experiment in section 6.4.2 these values were multiplied by 100 to speed up the experiment's run time. It was estimated that if this level were applied to the robot that a typical experiment would complete in between 10 and 30 minutes and that this would typically only cover crossing the lake and coming back to the start point. As a longer experiment was sought to give a more realistic impression of the MOOP performance the actuator power consumption values were only multiplied by 10 instead of 100. This gave a typical course length of between 100 and 300 minutes. This was an ideal length as it would require nearly the full real battery capacity of the MOOP to travel for this long.

Some minor hardware modifications were made at this point to improve the MOOPs sailing ability. The number of batteries was increased from 12 to 18 AA 2700 mAh Nickel-Metal-Hydride cells, these were connected as three sets of six batteries each providing a

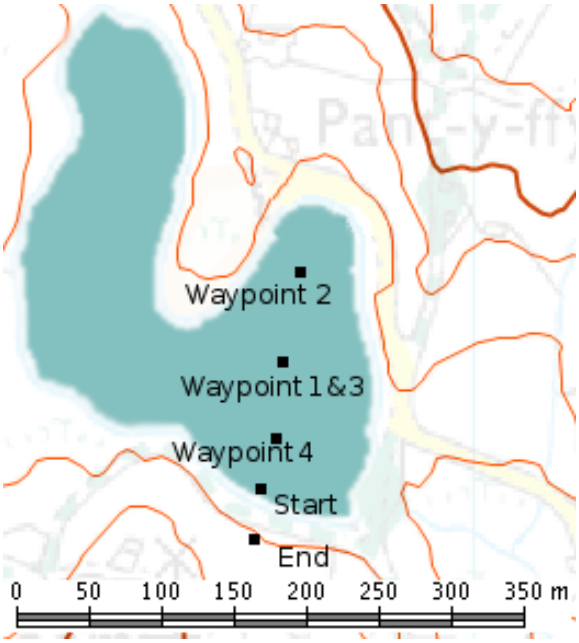

Figure 7.5: The three waypoint course used for the robot variable hormone experiment.

nominal voltage of 7.2 volts and 2700 mAh. The wind sensors were also upgraded from a continuous rotation potentiometer to an Austria Microsystems AS5040 magnetic encoder. The potentiometers had suffered from a small deadband and non-linear response while the magnetic encoder suffered from neither of these. The rudders of both boats were also extended to approximately double their surface area, this was expected to help allow the MOOPs to tack better and to handle sailing in stronger winds.

A beam reach course similar to that used in the previous experiment was used, although the course was to be repeated until the simulated battery level fell below zero. A third intermediate waypoint was added halfway along the course, this was to reduce the chance of deviation from the course while travelling between two waypoints. An illustration of this is shown in figure 7.5. Once the battery level reached zero the experiment ended, the robot stopped logging data, all hormone levels were reset to zero and the boat would sail back towards a recovery point which was intentionally placed about 30 metres inland to ensure the robot would beach itself. Had the waypoint been placed on the shoreline then the control system would by default hold station at its final waypoint (or 20 metres from the waypoint), previously this had resulted in it holding station just off shore beyond arms reach.

As with the simulation version of this experiment hormone sensitivities of 0.0, 0.25, 0.5 and 0.75 were used. Each was repeated a total of six times, three times on each boat. The experiment was run between November 2009 and April 2010, poor weather during much of December and the whole of January (the lake was frozen) prevented any experiments being

run. Wind speeds between Beaufort force 1 and 5 were observed. Of these experiments there were 18 attempts which resulted in failure (the robot's battery level never reached zero). A number of these were due to hardware faults or the robot becoming stuck on submerged rock near the shore or in reeds at the Northern end of the lake.

7.5.1 Results

This section reviews the results from these experiments. An analysis is provided of both the amount of time taken and the distance covered before the battery discharged. This is because increasing the time the robot can sail for might come at the expense of speed and thus reduce the distance covered. If this is the case it becomes questionable as to whether anything useful has been achieved if all we have done is make the boat sail for longer at a slower speed. This could result in the same distance being covered for each joule of energy being used. If applied to a solar powered robot, it is more useful to reduce power consumption and slow down at night and continue operating all night than it is to discharge the battery and drift until sunrise.

Some natural level of variation should be expected with these results due to changes in wind speed. Had it been possible to accurately record wind speed throughout the course of each experiment it might be possible to normalise against wind speed. However as the boat's wind sensors had no ability to measure wind speed, only highly generalised wind speed recordings based on visual observations are available.

7.5.1.1 Weather Observations

The weather observations during the course of these experiments are shown in table 7.3 (at the end of this chapter). The wind speeds are recorded according to their speed as a Beaufort force. Wind directions are based on their 16 point compass headings. These observations were taken through manual observation rather than with any wind recording equipment (as none was available). Wind estimates were always taken at the start of an experiment, in cases where a dramatic change in wind speed or direction was noticed this was also noted down. The topography of the terrain around the lake will have meant that wind speeds were not consistent across the entire lake.

7.5.1.2 Analysis of Battery Discharge Time

The battery level over time for each of the experiments is shown in figures 7.6, 7.7, 7.8 and 7.9. It had been expected that the shape of the lines would not be linear straight lines, but would appear more like an inverse exponential function (as seen in the simulator in section 6.4.2.1),

especially at higher hormone sensitivities. There are two exceptions to this, run 3 with the sensitivity of 0.0 (figure 7.6) and run 4 with the sensitivity of 0.5 (figure 7.8) although these appear to be outliers rather than normal behaviour. It is possible these trends might be due to a change in wind speed rather than any action of the control system. As the boats did not have any means to record wind speed, sufficiently detailed records of wind speed are not available to test this hypothesis. These graphs also do not reveal the same 'kink' that was seen in the simulator results in section 6.4.2 (figures 6.4, 6.5, 6.6 and 6.7). One possible explanation for this is that the compass used on the MOOPs is prone to giving quite noisy readings, especially while sailing. This is believed to be because the compass uses an accelerometer to sense its orientation. This is perfectly adequate for applications where the compass is relatively still (such as handheld devices), but when the boat is constantly being vibrated by waves it leads to quite a large amount of noise. As the control system directly uses the compass reading, without averaging at a rate of over 10hz this can lead oscillation of the rudder. Even when the level of hormone suppression was quite high this constant oscillation of the rudder was still observed. By comparison, in the simulator actuator movement rarely occurs once the robot is on course. This could explain the reason for the near linear power consumption despite increasing levels of hormonal suppression.

The simulator experiments in section 6.4.2 had seen a much steeper drop in charge level during the initial phase of the experiment and this had then levelled out. This trend does not appear to be repeated in the robot experiments which drop at an almost constant rate. The simulator results also showed a step effect as most power was consumed when turning around at waypoints with few course corrections being needed between waypoints. However, in the robot the rudder actuator in particular seems to be almost constantly moving (especially when hormone concentration is positive or near zero) as the MOOPs have a tendency to always try and sail upwind the control system must constantly fight this. The end result is that power consumption rates are far more constant than they were under simulation.

Despite the discharge graphs not matching expected trends and showing some dis-similarity to the simulator experiment, the amount of time taken still increases with hormone concentration. Table 7.4 shows that the median time of operation increases from 99.63 minutes with zero sensitivity to 111.68 minutes at 0.25 sensitivity, 143.04 minutes at 0.5 sensitivity and 195.43 minutes at 0.75 sensitivity. Figure 7.10 shows a box and whisker diagram illustrating the increase in median times and the inter-quartile range although several outliers still defy the overall trend. As the time required will be at least partially dependent upon wind speed this is not unexpected.

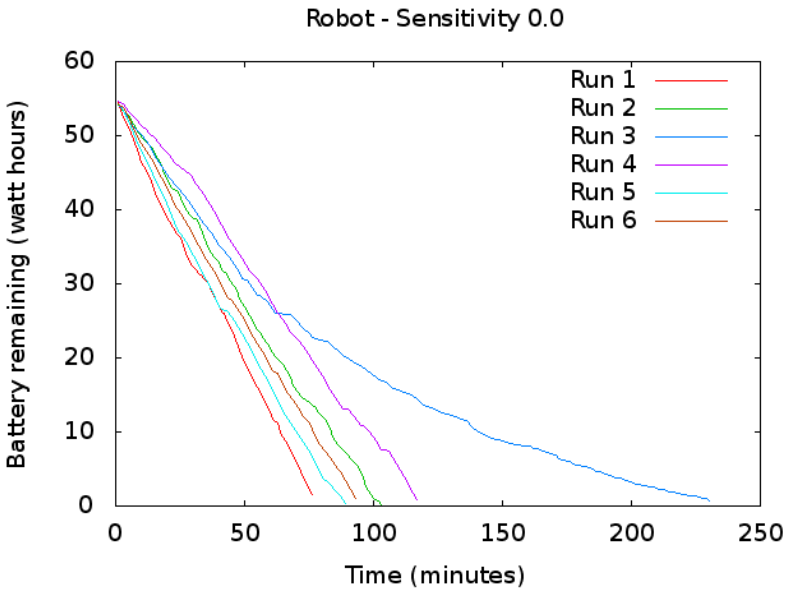

Figure 7.6: A graph showing battery level over time during the robot variable hormone experiment with a sensitivity of 0.

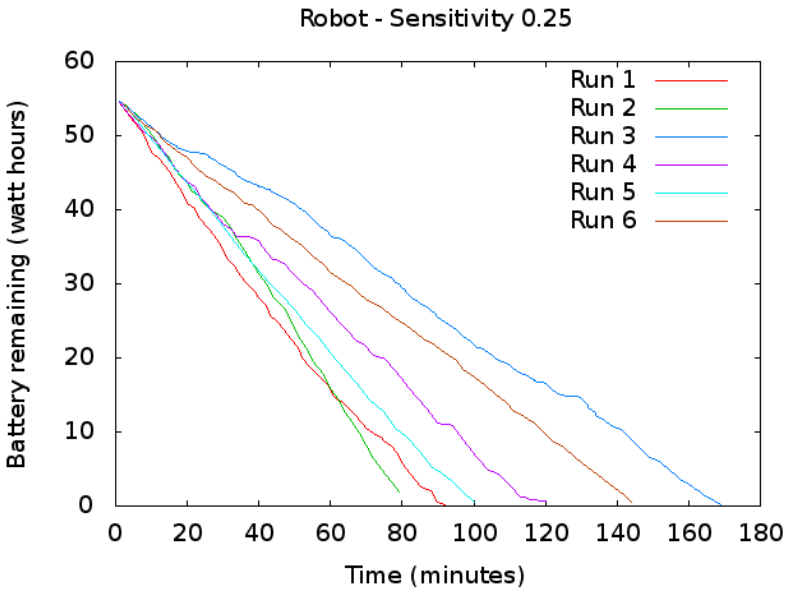

Figure 7.7: A graph showing battery level over time during the robot variable hormone experiment with a sensitivity of 0.25.

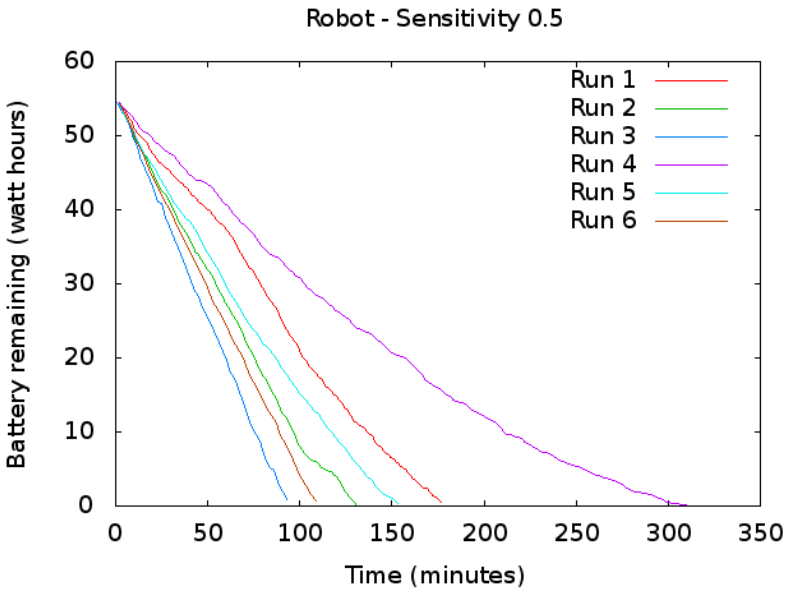

Figure 7.8: A graph showing battery level over time during the robot variable hormone experiment with a sensitivity of 0.5.

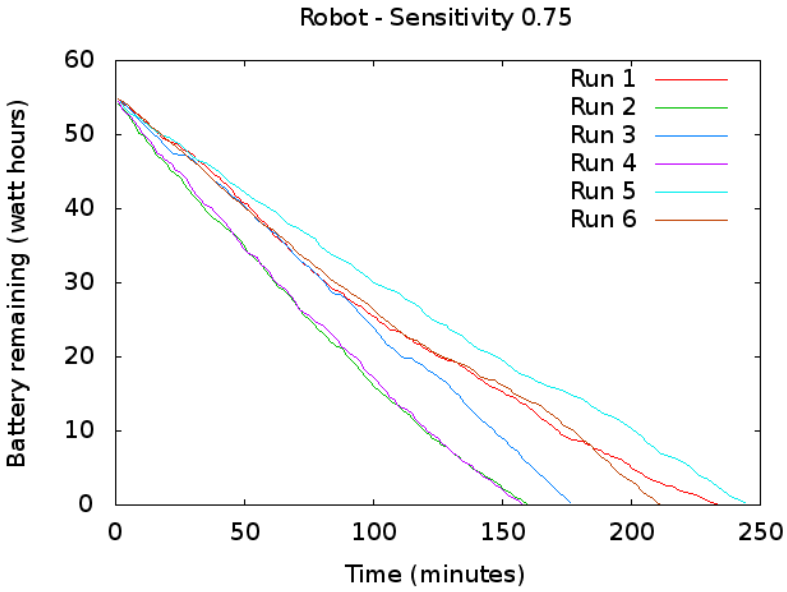

Figure 7.9: A graph showing battery level over time during the robot variable hormone experiment with a sensitivity of 0.75.

Hormone Sensitivity	Mean	Median	Min	Max	Lower Quartile	Upper Quartile
0.0	119.9	99.63	78.65	231.48	91.5	115.83
0.25	118.73	111.68	81.23	169.5	94.63	139.91
0.5	164.12	143.04	95.42	313.18	116.31	172.73
0.75	198.86	195.43	159	245.35	165.37	229.88

Table 7.4: Summary statistics for the battery discharge time for the robot's running the variable hormone experiment. All values are in minutes.

To test for a statistically significant difference between the experiments a Kruskal-Wallis test (non-parametric analysis of variance) was performed. This led to the statistical hypothesis that:

> H_0: Varying hormone concentration does not affect the amount of time that it is possible to sail on a single battery charge.
>
> H_1: Varying hormone concentration does affect the amount of time that it is possible to sail on a single battery charge.

Kruskal-Wallis $H = 9.0067$, $df = 3$, $p = 0.0292$

Given the p value of 0.0292 (2.92%) we cannot accept the alternative hypothesis at the 99% significance level although we can accept it at the 95% level. This result shows that a statistically significant difference is occurring between the experiments, but it does not show where those differences lie. The Box and Whisker diagram shown in figure 7.10 reinforces this conclusion by showing the increase in median value, but it also illustrates the high level of noise with a few outliers far exceeding the interquartile range. In particular an outlier in the control group (sensitivity of 0.0) and the 0.5 sensitivity group both exceed the median times in any other group. This suggests the possibility that had the experiment been repeated more then there might be a far wider spread to the data. However it is not sufficient evidence to reverse the conclusions reached from the Kruskal-Wallis test or the analysis of the median and box and whisker diagram.

7.5.1.3 Analysis of the Distance Sailed

An analysis of the distance sailed was also performed, the results are shown in table 7.5 (at the end of the chapter). This was to confirm that increasing the amount of time the robot sailed for also increased the distance it was actually able to cover. As with the analysis of time taken a Kruskal-Wallis test was performed to test for a statistically significant difference between the samples. This leads to the following statistical hypotheses:

> H_0: Varying hormone concentration does not affect the distance that it is possible to cover on a single battery charge.

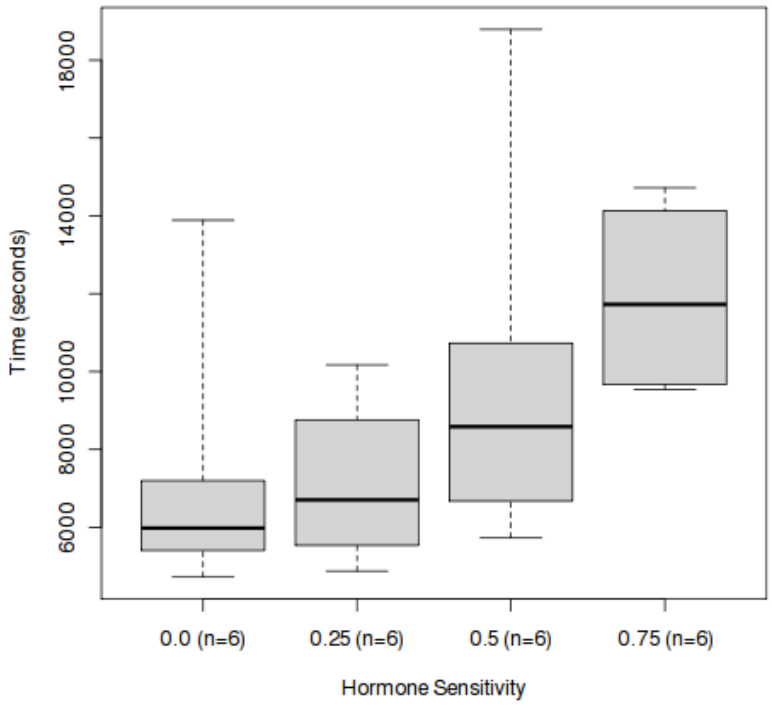

Figure 7.10: Box and Whisker plot for robot battery discharge times during the robot variable hormone experiment.

H_1: Varying hormone concentration does affect the distance that it is possible to cover on a single

Kruskal-Wallis $H = 8.183$, $df = 3$, $p = 0.0424$

Given the p value of 0.0424 (4.24%) we cannot accept the alternative hypothesis at the 99% significance level, although we can accept it at the 95% level. As in the analysis of distance covered a box and whisker plot is provided in figure 6.9 to illustrate where these differences lie. When comparing this to figure 6.8 the medians follow a similar trend, increasing with hormone concentration. However the outlier previously seen on the zero sensitivity run does not reoccur. This suggests that despite taking longer during this outlier, no extra distance was covered, this might have been due to a reduction in wind speed and/or a shift in wind direction. The outlier in the 0.5 sensitivity group does still exist. The medians are also not as close to the centre of the interquartile range suggesting that this data is not as close to a normal distribution as the time data.

In general, this figure supports H_1, so we can conclude that as well as increasing the amount of time spent sailing we are also increasing the distance covered. This suggests

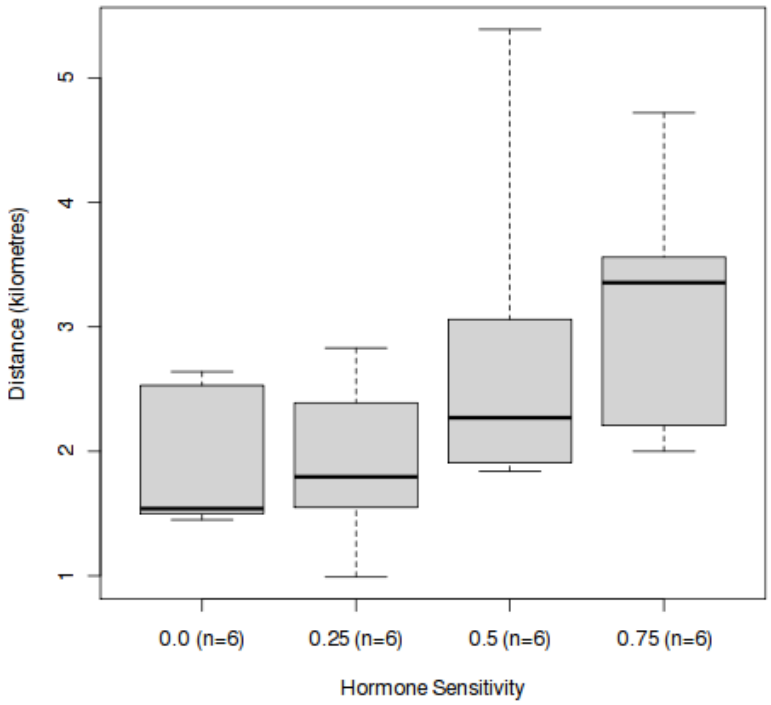

Figure 7.11: Box and Whisker plot of the distances achieved by the robot in the variable hormone experiment.

that applying hormonal modulation in this way will be particularly useful in reducing power consumption and that the effect on performance is not particularly dramatic. This reinforces the conclusion of the fixed hormone experiments and simulated experiments that power consumption can be reduced with little effect upon sailing performance.

7.5.1.4 Analysis of Actuator Usage

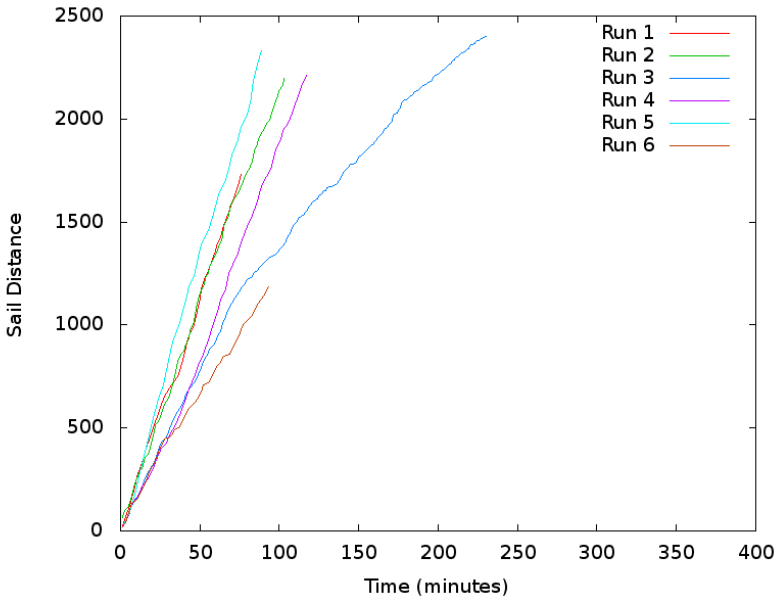

Figure 7.12: The cumulative distance moved by the sail actuator during the robot variable hormone experiment with a hormone sensitivity of 0.0.

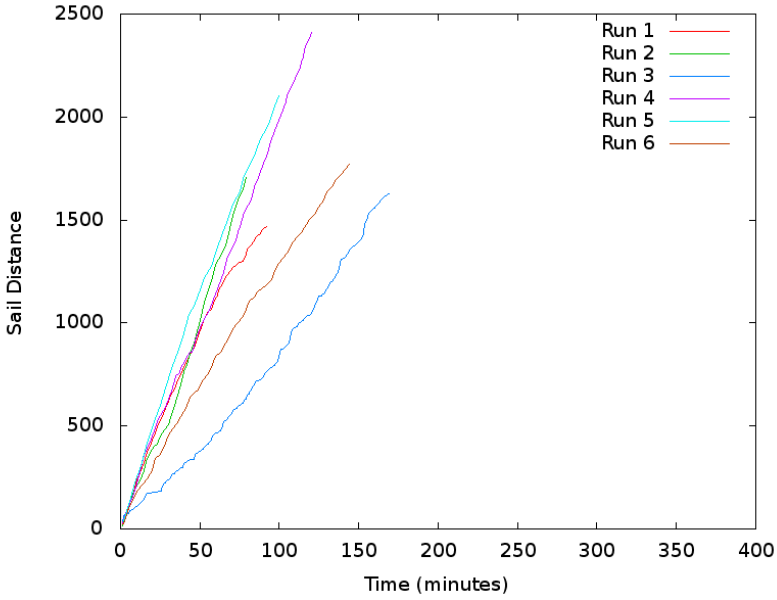

Figure 7.13: The cumulative distance moved by the sail actuator during the robot variable hormone experiment with a hormone sensitivity of 0.25.

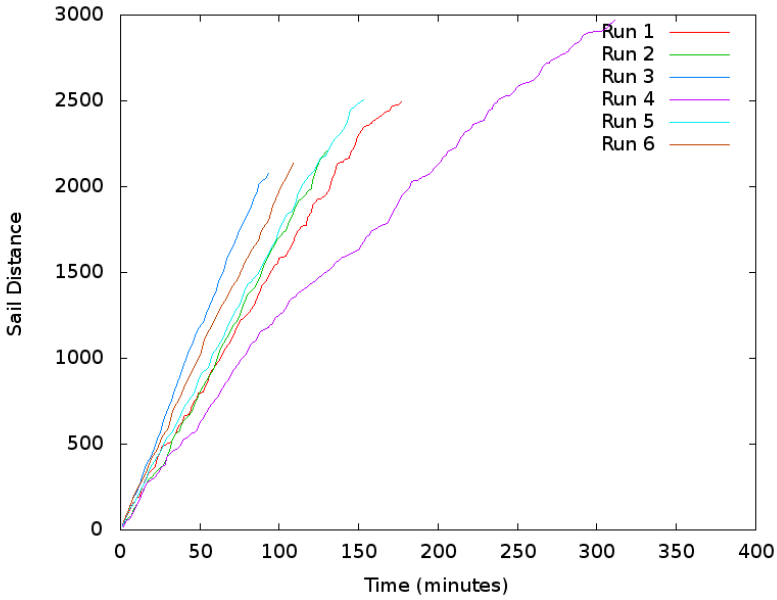

Figure 7.14: The cumulative distance moved by the sail actuator during the robot variable hormone experiment with a hormone sensitivity of 0.5.

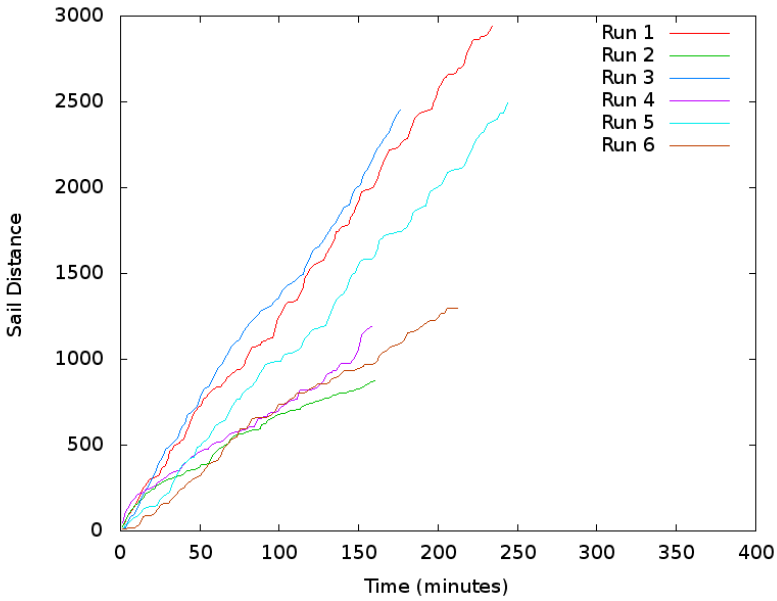

Figure 7.15: The cumulative distance moved by the sail actuator during the robot variable hormone experiment with a hormone sensitivity of 0.75.

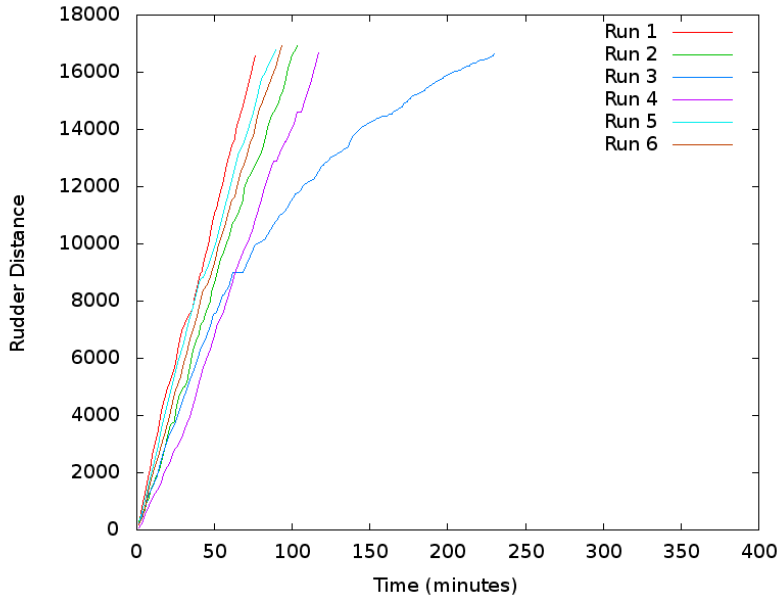

Figure 7.16: The cumulative distance moved by the rudder actuator during the robot variable hormone experiment with a hormone sensitivity of 0.0.

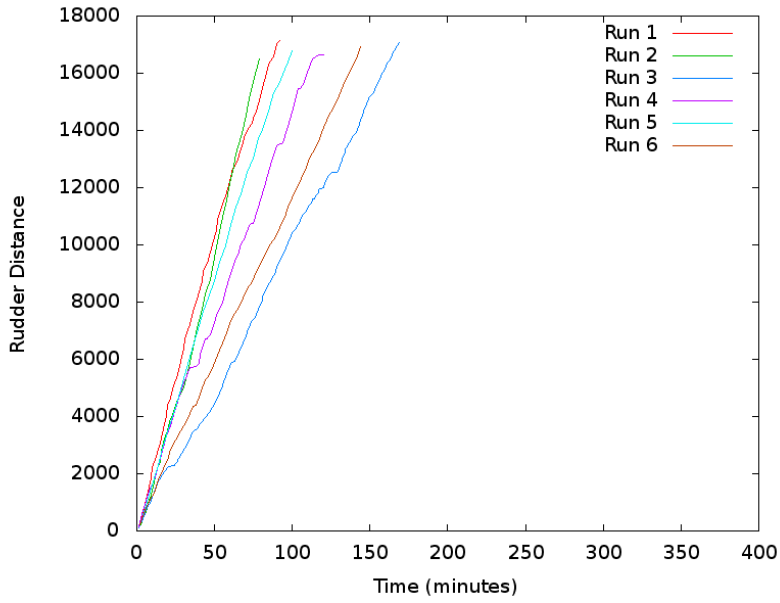

Figure 7.17: The cumulative distance moved by the rudder actuator during the robot variable hormone experiment with a hormone sensitivity of 0.25.

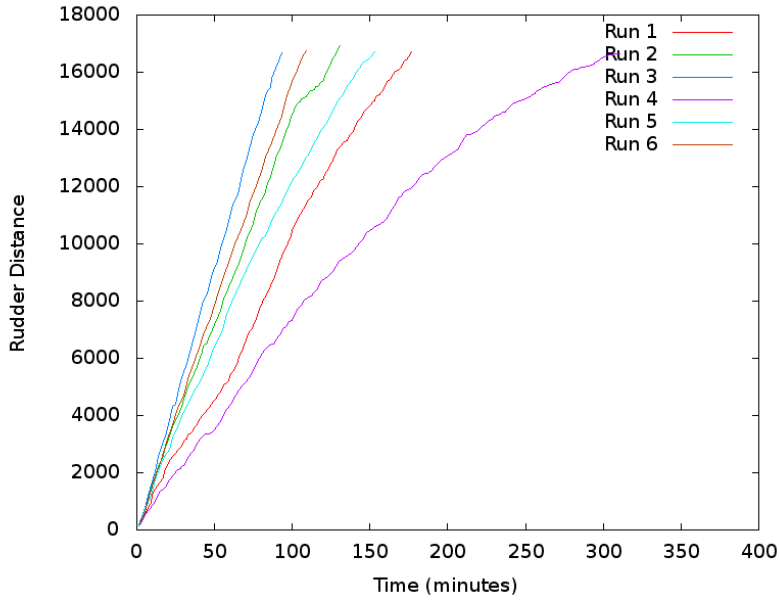

Figure 7.18: The cumulative distance moved by the rudder actuator during the robot variable hormone experiment with a hormone sensitivity of 0.5.

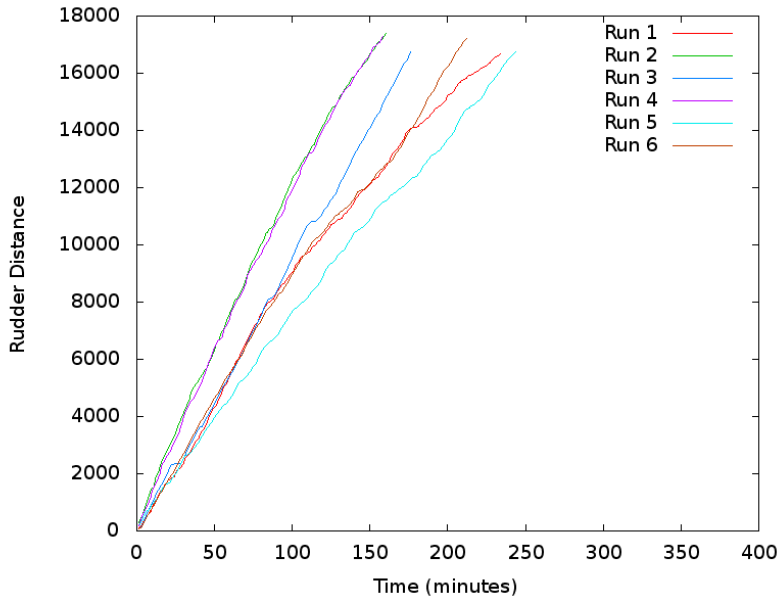

Figure 7.19: The cumulative distance moved by the rudder actuator during the robot variable hormone experiment with a hormone sensitivity of 0.75.

Hormone Concentration	Mean	Median	Upper Quartile	Lower Quartile
0.0	1821.83	1889	1991.5	1701.75
0.25	1670.33	1600	1640	1501.5
0.5	2267.67	2222.5	2388	2045.75
0.75	1772.83	1699.5	2378.5	1118.75

Table 7.6: Sail actuator use during the robot variable battery hormone experiment.

Hormone Concentration	Mean	Median	Upper Quartile	Lower Quartile
0.0	10585.5	10826	11406.75	10183
0.25	10236.17	9734.5	11644.5	8984.75
0.5	11780.67	11474	12448.5	10676.5
0.75	14580.67	14562.5	15183	14098.75

Table 7.7: Rudder actuator use during the robot variable battery hormone experiment.

Table 7.6 shows the summary statistics for the number of movements made by the sail actuator and table 7.7 shows the same for the rudder actuator. The full data can be found in appendix D. Figures 7.12, 7.13, 7.14, 7.15, 7.16, 7.17, 7.18 and 7.19 show graphs of the cumulative actuator usage over time. These show the cumulative size of all movements made by the sail and rudder actuators. As discussed in section 5.2.3 there are 11 possible actuator positions, so moving through the full range of motion will increase the cumulative distance by 11. From these graphs the distribution of actuator movements over time can be seen as can the relative contributions of the rudder and sail to battery discharge (as these are the only things which influence battery levels). These graphs show relatively little difference in the overall trends between both rudder and sail power consumption, unlike the simulated version of this experiment in section 6.4.2. In cases of the two outliers, (run 3 with sensitivity of 0 and run 4 with a sensitivity of 0.5) both the sail and rudder still demonstrate similar trends with a long curve rather than the near linear trends seen in all other runs.

7.5.1.5 Conclusions on Variable Hormone Experiments

This experiment has demonstrated that through hormonal modulation it is possible to nearly double the distance and sailing time achieved by a robot operating under real world conditions. This doubling appears to be possible without severely reducing the ability of the robot to continue to manoeuvre. This result is not as strong as in simulation, where a 13 fold improvement was seen when performing the same experiment and the discharge rates of the robot also appear to differ from those of the simulation. This is likely to be due to the need of the robot to constantly move its rudder to maintain a stable course, the noise of using

real sensors and from variations caused by winds and waves. This appears to be a reasonable strategy which can be applied to reducing power consumption in a sailing robot as it provides a mechanism which is almost continuously variable. This experiment has not managed to exceed the limits at which the boat stops sailing effectively as was seen in the previous fixed hormone experiment. This is believed to be because the amount of time spent with effective hormone concentrations comparable to those which caused failures in the previous experiment was relatively short. The modifications made to the rudder may have also helped prevent a repeat of these failures.

7.5.1.6 Swapping Sail and Rudder Power Consumption Values

As discussed in section 6.4.2.3, at some point during the calculation of the sail and rudder power consumption an error was made and these figures were accidentally reversed. This section attempts to rectify this error by recalculating the battery discharge using the correct power consumption figures. There was a suspicion that the error might have accounted for some of the unexpected trends seen in figures 7.7, 7.8 and 7.9. In particular the way in which these graphs are extremely linear compared with those from the simulator, with the exception of run 3 in figure 7.6 and run 4 in figure 7.8. The data from these experiments was reprocessed to recalculate the battery level over time based on the counts of actuator movements and the correct actuator power consumption figures. The results are shown in figures 7.20, 7.21, 7.22 and 7.23. Since the rudder moves more frequently than the sail, it previously accounted for most of the total power consumption. By swapping the rudder and sail power consumption values the total power consumption has been reduced. As a result the corrected experiments terminate with approximately 30 watt hours of battery capacity remaining. However, the shape of the graphs appears very similar to those in section 7.5.1.2 implying that swapping the rudder and sail power consumption does not have any effect on the shape of the graph.

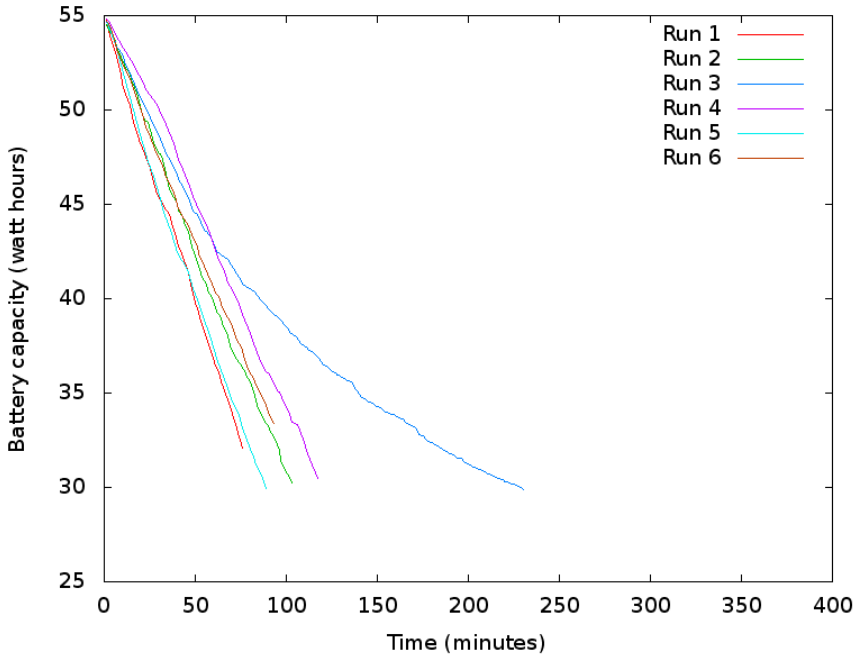

Figure 7.20: A graph comparing time and battery discharge time after reprocessing results to swap rudder and sail power consumption figures, using a hormone sensitivity of zero.

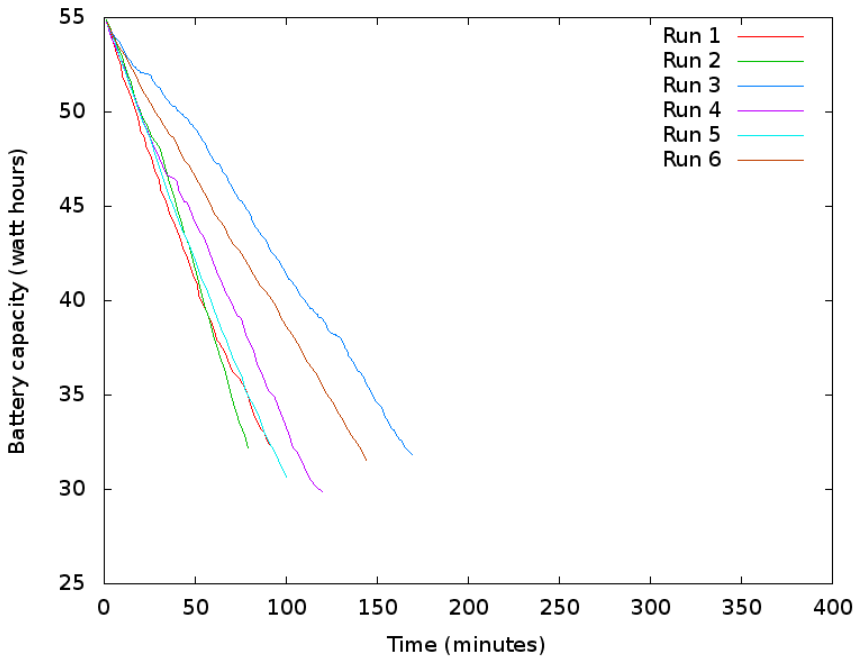

Figure 7.21: A graph comparing time and battery discharge time after reprocessing results to swap rudder and sail power consumption figures, using a hormone sensitivity of 0.25.

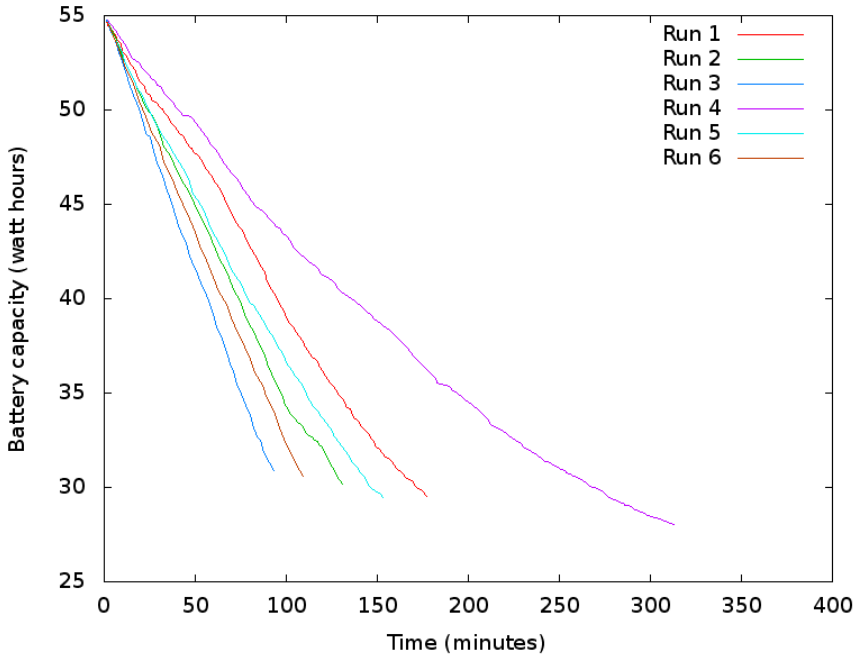

Figure 7.22: A graph comparing time and battery discharge time after reprocessing results to swap rudder and sail power consumption figures, using a hormone sensitivity of 0.5

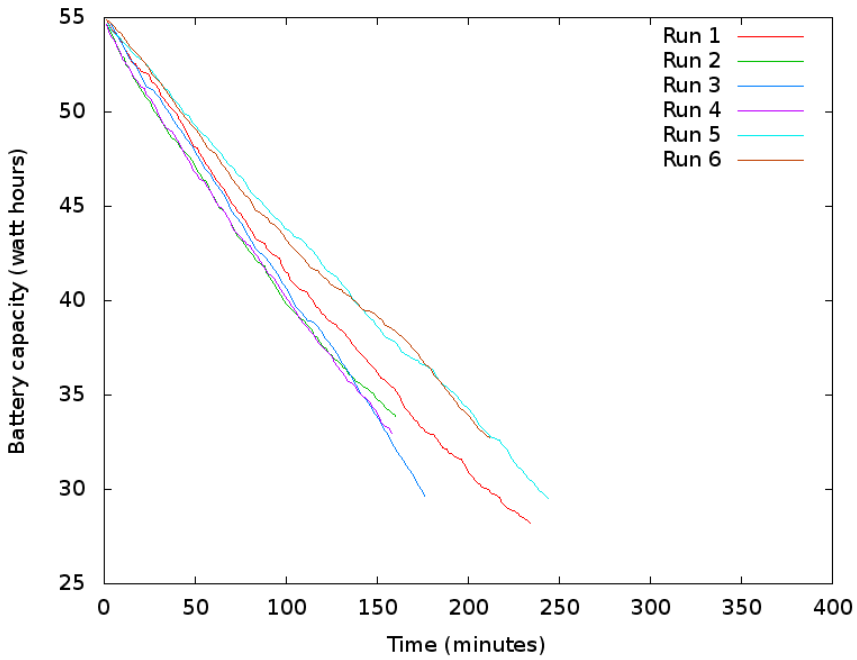

Figure 7.23: A graph comparing time and battery discharge time after reprocessing results to swap rudder and sail power consumption figures, using a hormone sensitivity of 0.75

7.6 Chapter Summary and Conclusions

In this chapter two sets of experiments have been conducted using a real sailing robot. These repeat the first two experiments which were simulated in the previous chapter, in sections 6.4.1 and 6.4.2. The first experiment used a set of fixed hormone levels to establish if the robot could still sail effectively when a hormone excited or suppressed the control system and to see what effect this had upon power consumption. A similar trend to that of the simulator experiment was identified although the effect was not quite as consistent and sailing performance became unacceptable at a lower level of (inhibitory) hormone concentration. The second experiment varied the hormone in response to the battery level. It showed a similar trend to the simulations although, when using a hormone sensitivity of 0.75 had yielded a 13 fold improvement in the distance covered and time spent sailing over a control with a sensitivity of zero in the simulated robot it only yielded a doubling in the real robot.

These experiments have demonstrated the feasibility of an artificial neuro-endocrine controller for modulating the power consumption of a sailing robot operating in a real world environment. They have demonstrated a robustness in the face of varying wind directions, wind speeds between a Beaufort force of 1 and 6 (approximately 1 to 45 km/h) and wave heights between zero and approximately 20cm. Although compared with open ocean sailing these conditions could still be considered relatively benign, a larger more suitable robot will be required to sail in more hostile conditions.

The general conclusions of these experiments match those of the simulator experiments in the previous chapter. However, the magnitude of the differences seen by varying hormone sensitivity and concentration have been significantly reduced. Although the savings achievable through the use of a neuro-endocrine controller are still large enough to make a considerable difference to potential lifespan of a real robot.

From the data available it is not entirely clear to what extent modulating the magnitude of actuator movements is reducing sailing performance. As course lengths were restricted due to the size of the lake, the metrics which might have revealed the impact on sailing performance have failed to illustrate any meaningful differences. As with the simulation experiments it is possible that by increasing the level of modulation we are taking away some kind of safety margin in the control system. However, the range of conditions these experiments ran in did not reveal this. Additionally during the development of the controllers which were used to train the neural networks a considerable amount of parameter tweaking was undertaken to find an optimal setting and relatively small variations were found to prevent well controlled sailing.

These results show that a neuro-endocrine controller is a promising method for controlling

power consumption in a real sailing robot. Future experiments need to take place on a larger body of water and with a robot that can sail for longer. Ideally such a robot will have its own solar panels and will allow for a recreation of the solar power experiments in section 6.4.3. It should also have the ability to track actual power consumption and battery state, rather than estimating them from actuator movements. The experiment should take place at sea or on a large lake to provide a more challenging operational environment. Re-running the experiments at different times of year, different latitudes or in differing levels of cloud cover could provide further insight into adapting behaviours to varying levels of energy availability.

Date	Robot	Hormone Sensitivity	Run Number	Weather Observations
06/11/2009	MOOP0	0	1	Force 2 SW
17/11/2009	MOOP0	0	2	Force 2-3 SSW
04/04/2010	MOOP0	0	6	Force 2 WSW
25/11/2009	MOOP1	0	4	Force 4 WSW
02/02/2010	MOOP1	0	5	Force 4-5 W
20/11/2010	MOOP1	0	3	Force 2 SW, later force 1 variable
20/11/2009	MOOP0	0.25	1	Force 2 SW, later force 1 variable
25/11/2009	MOOP0	0.25	2	Force 4 WSW
10/12/2009	MOOP0	0.25	3	Force 2 WSW, force 1 variable later
02/02/2010	MOOP1	0.25	4	Force 4-5 W
26/02/2010	MOOP1	0.25	5	Force 4 W
26/04/2010	MOOP1	0.25	6	Force 3 W
01/02/2010	MOOP0	0.5	1	Force 3 W
15/02/2010	MOOP0	0.5	2	Force 4 WSW
26/02/2010	MOOP0	0.5	3	Force 4 W
10/03/2010	MOOP1	0.5	4	Force 2-3 ENE
15/03/2010	MOOP1	0.5	5	Force 3 W
22/03/2010	MOOP1	0.5	6	Force 4-5 WSW
10/03/2010	MOOP0	0.75	1	Force 2-3 ENE
15/03/2010	MOOP0	0.75	2	Force 3 W
07/04/2010	MOOP0	0.75	4	Force 3-2 NW
04/04/2010	MOOP1	0.75	3	Force 2 WSW
14/04/2010	MOOP1	0.75	5	Force 2-3 NE
17/04/2010	MOOP1	0.75	6	Force 2 W

Table 7.3: Table showing the weather observations, which of the two robots were used, the experiment parameters and run number during the variable hormone experiments. The run numbers correspond to the numbers used in section 7.5.1.2.

Hormone Sensitivity	Mean	Median	Min	Max	Lower Quartile	Upper Quartile
0.0	1.86	1.54	1.45	2.64	1.5	2.28
0.25	1.89	1.8	0.99	2.83	1.56	2.3
0.5	2.79	2.27	1.84	5.39	1.92	2.94
0.75	3.2	3.35	2	4.72	2.49	3.51

Table 7.5: Summary statistics for the distance covered during the variable hormone experiment running on the robot. Distances shown are in kilometres.

Chapter 8

A Comparison of Simulator and Robot Results

8.1 Introduction

This chapter discusses the results presented in chapters 6 and 7, examining the gap between the simulator and robot results. It also discusses the limitations of both platforms and what effect this may have had upon the results.

8.2 Comparison of Results

This section compares the robot and simulator results from both the fixed hormone experiments in sections 6.4.1 & 7.4 and the variable hormone experiments in sections 6.4.2 & 7.5.

8.2.1 Fixed Hormone Experiments

In these experiments (described in detail in sections 6.4.1 and 7.4), hormone levels were fixed for the duration of each experiment run. Hormone concentration was varied between each run and the effects upon power consumption per hour and per kilometre travelled were compared. This is shown in figures 8.1 and 8.2, where the robot and simulator results are plotted together. In these figures it can be seen that there is a strong correlation between power consumption and hormone concentration in both the robot and simulation experiment. This is demonstrated by the Spearman's rank correlation coefficients shown in tables 8.1 and 8.2. Overall, the same correlations are shown in both datasets, but the robot results show a greater spread and an order of magnitude higher average power consumption.

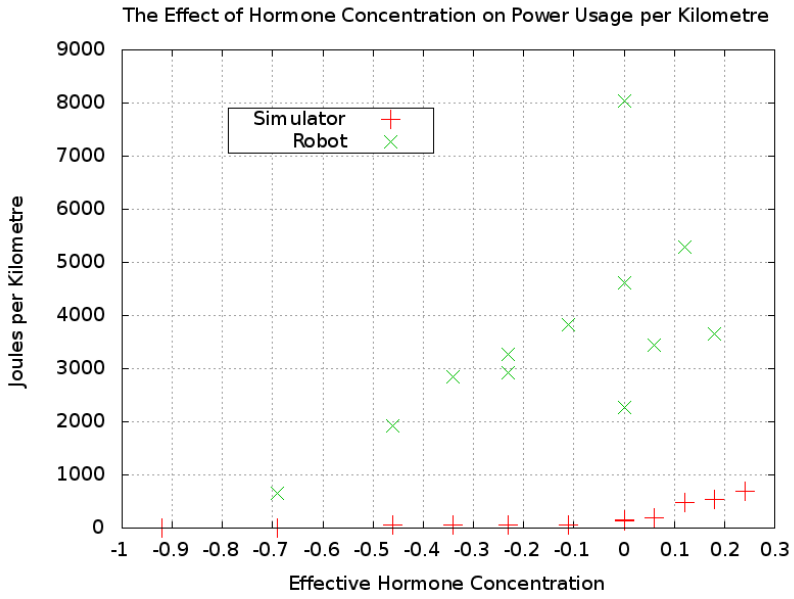

Figure 8.1: A comparison of the hormone concentration and power consumption per kilometre for both the robot and simulator during the fixed hormone experiment.

	Robot	Simulator
Hormone Concentration and Energy use per KM	0.7090223	0.9910377
Accept at 95% significance?	Yes	Yes
Accept at 99% significance?	Yes	Yes

Table 8.1: Comparison of the robot and simulator correlation between hormone concentration and energy use per kilometre.

The course efficiency metric (defined in section 6.3) was designed to test how efficiently the robot followed its target course. It is calculated by dividing the course length by the distance travelled. The robot's course efficiency results shown in figure 8.3 are less consistent than the simulator results. This is in part caused by differences in what happens when the experiment fails due to excessive suppression of the control system. Under simulation the robot has an essentially infinite area of water to drift in until the simulation terminates. As simulations are generally run unattended, there is unlikely to be any manual intervention to terminate them. Instead, they will continue to run until a timeout value terminates them. In the real robot there is at best a few hundred metres of water until the robot washes up on the shore and the experiment is immediately terminated, this can occur within just a few minutes of the control system becoming suppressed (or a failed tacking manoeuvre). The result is that when hormone concentrations are at their lowest (approaching -1.0) the simulated robot will just drift slowly for several hours (until the simulation reaches its maximum allowed time or is terminated by the user) and will move a distance that is significantly longer than

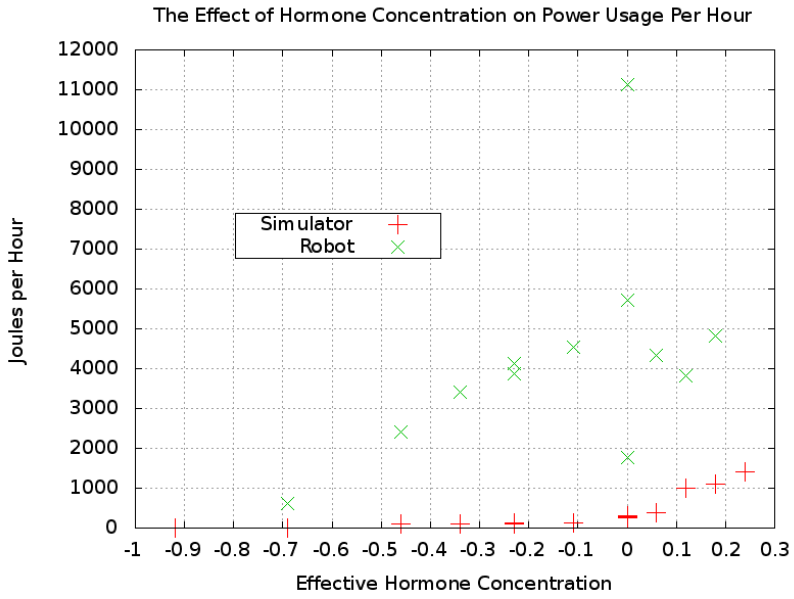

Figure 8.2: A comparison of the hormone concentration and power consumption per hour for both the robot and simulator during the fixed hormone experiment.

	Robot	Simulator
Hormone Concentration and Energy use per Hour	0.5608684	0.9910377
Accept at 95% significance?	Yes	Yes
Accept at 99% significance?	No	Yes

Table 8.2: Comparison of the robot and simulator correlation between hormone concentration and energy use per hour.

the entire course length. This causes a low course efficiency number to be reported. As the robot experiment ends very quickly this results in only a very short distance being covered which will be less than the total course length resulting in a very high (over 100%) efficiency number. Even robot experiments which did complete reported very high course efficiency values, this is believed to be caused by a combination of short courses and waypoint threshold problem (described in detail in section 6.3). The threshold problem is caused because the robot is considered to have reached a waypoint when it is within 20 metres of that waypoint. With distances between waypoints of only 150 or 200 metres loosing 40 metres represents a significant portion of the distance and this has adverse affects on this metric. The difference between the robot and simulator is exacerbated as simulator waypoints are 1 kilometre apart, while the robot waypoints are only 150-200 metres apart. These differences in course layout were a result of the behaviour of the limitations of both systems. The robot course was limited by the size of the lake, although it was capable of turning 180 degrees within a few seconds and would settle down on a new course very quickly. The simulated boat had a much slower

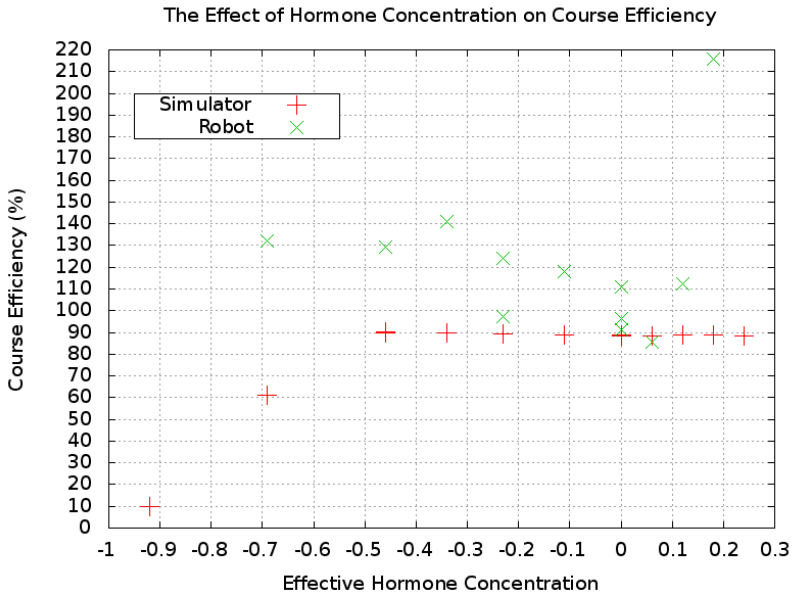

Figure 8.3: A comparison of the course efficiency between the simulator and robot during the fixed hormone experiment.

rate of turn and would take several hundred metres to settle down onto a course, requiring relatively long courses to achieve some stable sailing between waypoints.

8.2.2 Variable Hormone Experiments

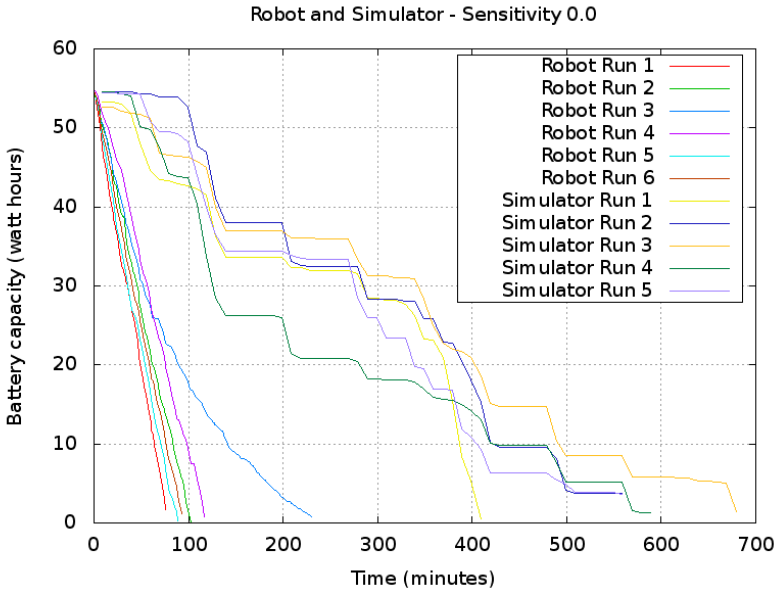

Figure 8.4: The simulator and robot battery discharge over time for the control (no sensitivity to hormone) in the variable hormone experiment.

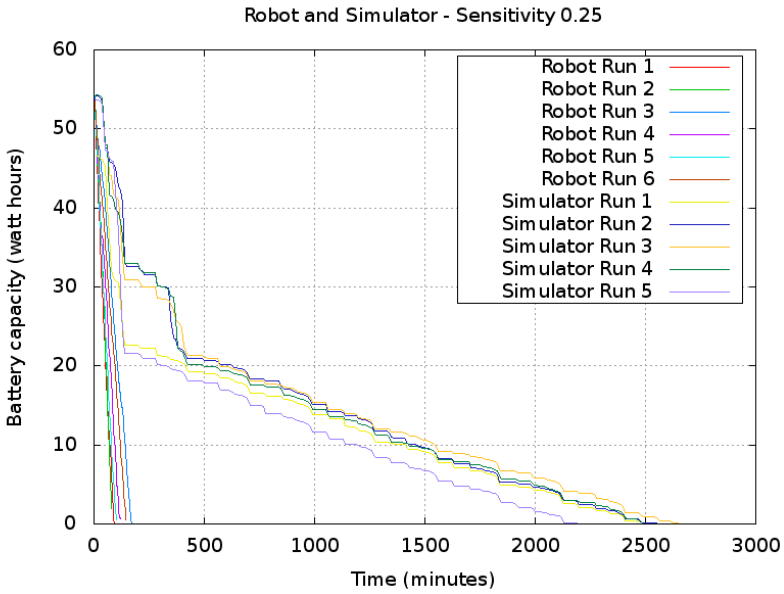

Figure 8.5: The simulator and robot battery discharge over time during the variable hormone experiment, with a hormone sensitivity of 0.25.

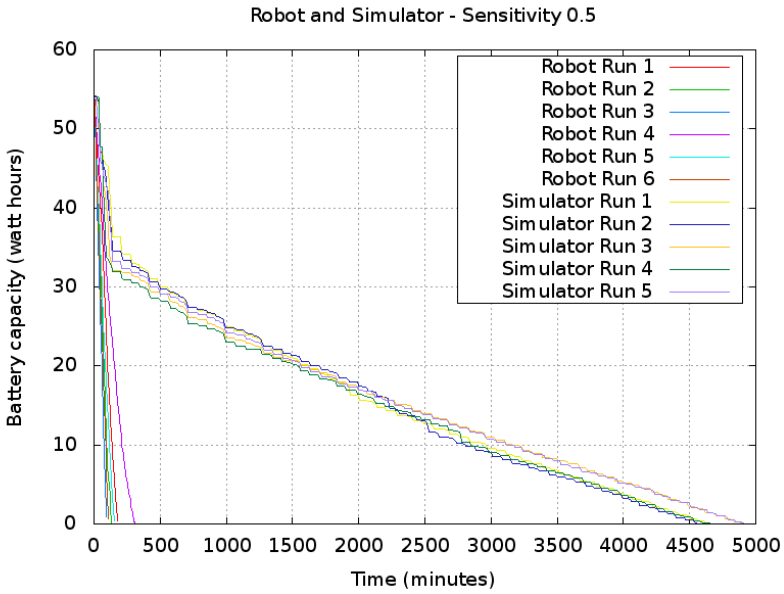

Figure 8.6: The simulator and robot battery discharge over time during the variable hormone experiment, with a hormone sensitivity of 0.5.

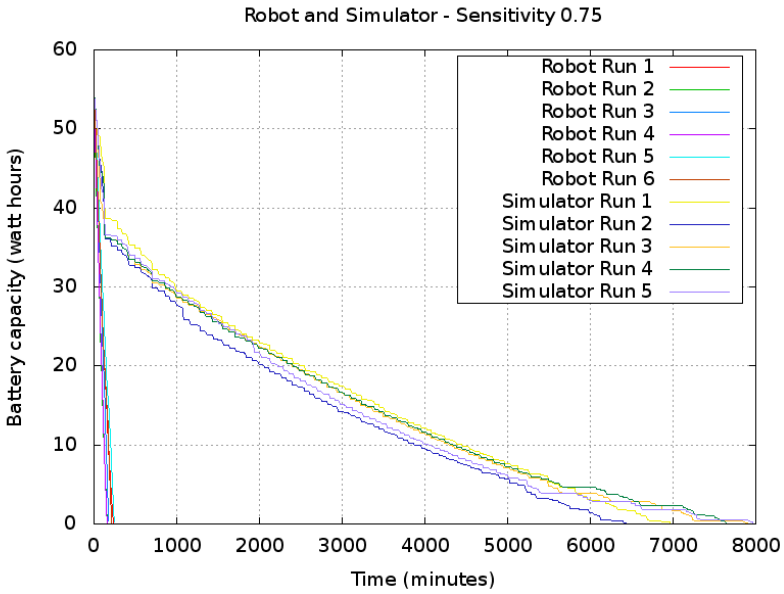

Figure 8.7: The simulator and robot battery discharge over time during the variable hormone experiment, with a hormone sensitivity of 0.75.

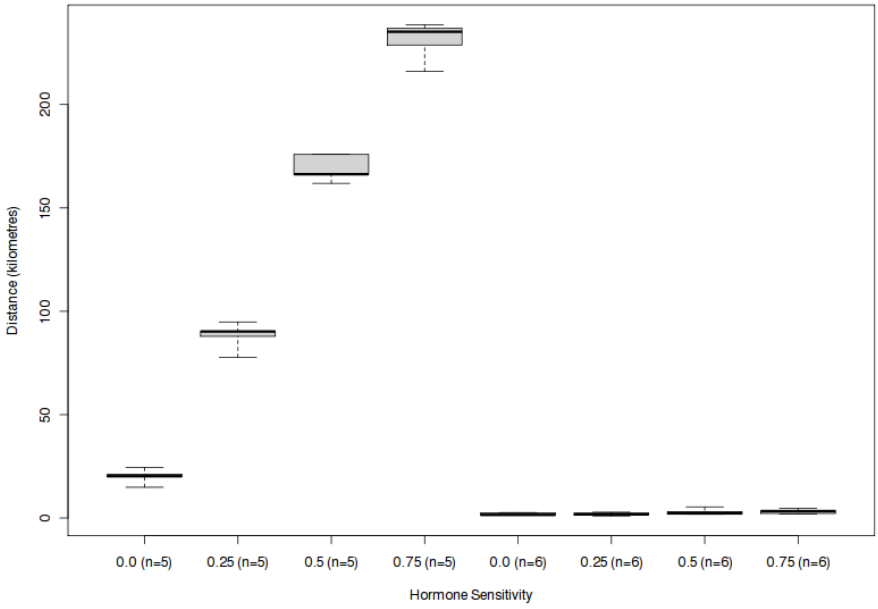

Figure 8.8: A box and whisker plot showing the distance covered by the simulator (left) and robot (right) during the variable hormone experiment.

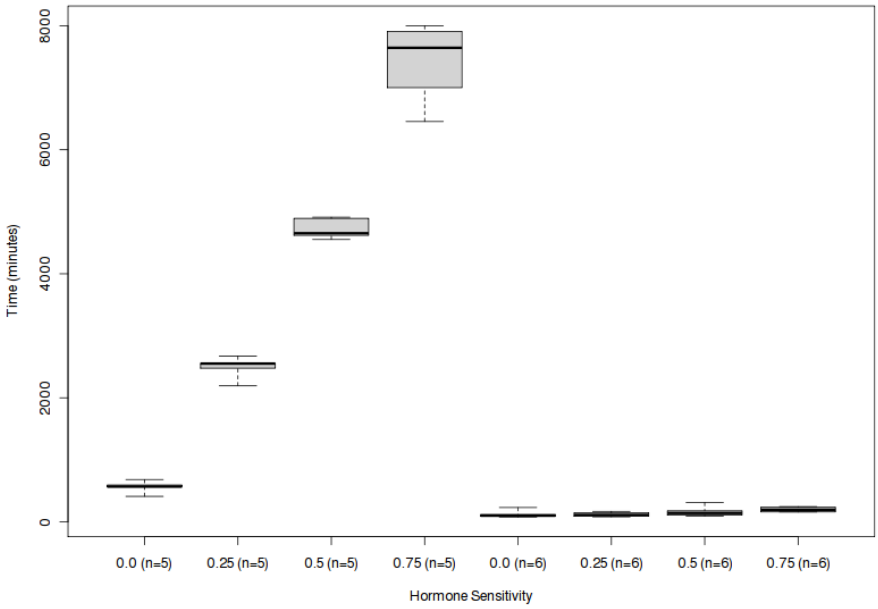

Figure 8.9: A box and whisker plot showing the time taken by the simulator (left) and robot (right) during the variable hormone experiment.

This section compares the results of the variable hormone experiments from sections

Hormone Sensitivity	Robot Median Distance Sailed (Kilometres)	Simulator Median Distance Sailed (kilometres)	Robot Median Time Taken (Minutes)	Simulator Median Time Taken (Minutes)
0.0	1.54	20.23	99.63	565.01
0.25	1.8	90.25	111.68	2490.67
0.5	2.27	166.29	143.04	4726.41
0.75	3.35	235.2	195.43	7401.14

Table 8.3: The median time and distance covered by the simulator and real robot on a single battery charge in comparison to hormone concentration.

6.4.2 and 7.5. In these experiments hormone concentration was allowed to vary in response to battery levels and the experiment ran until the battery was discharged. The battery charge level was reduced every time an actuator moved. As battery charge level dropped the hormone would transition from exciting the control system to perform larger movements to suppressing it to making smaller ones, with the maximum level of suppression occurring as the battery level neared zero.

The median distance sailed and time taken on a single battery charge are shown in table 8.3 and the discharge over time is shown in figures 8.4, 8.5, 8.6 and 8.7. These figures combine the data shown in figures 6.4, 6.5, 6.6, 6.7 and 7.6, 7.7, 7.8, 7.9. From them it can be seen that the robot discharge rates are much faster and follow a near linear relationship, while the simulator has a change in the rate at around 35 watt hours and the rate of discharge continues to slow as time progresses. The difference in the total time elapsed and distance covered are explained by the simulator course being longer (5km instead of 150m), but this does not explain the difference in the shape of the curves. A potential cause for the difference in curve shape is that the simulated robot is stable once on course and does not experience any noise from sensors, which can trigger rudder and sail movements when the robot is actually on a stable course. The MOOP sailing robots are also far less stable when left on a given course and have a tendency to always turn towards the wind and therefore requires constant action from the control system to remain on course, while the simulated robot will typically remain on any given course without constant action. Despite these differences in the robots and the courses, both exhibit the same underlying trend with regards to power consumption and hormonal modulation: the more the hormonal modulation suppresses the control system, the lower the average power consumption causing the boat to sail for longer and travel further on a single battery charge. The box and whisker diagrams for both simulator and robot results are shown together in figure 8.8 (comparing distance covered) and figure 8.9 (comparing

time taken). These demonstrate the different orders of magnitude between the simulator and the robot results although make it difficult to see how the spread of data compares as the simulated results extend the scale so dramatically that the robot results are dwarfed.

Overall from these comparisons we can see that the same underlying trends in power consumption are shown in both the simulator and the robot. Although, the simulator results are less prone to noise and exaggerate the trends.

8.2.3 Future Improvements to the Simulator

A number of modifications could be made to improve the simulator in order to close the gap between simulation and real experiments. Noise could be added both to the sensor data (so that position, heading and wind direction information varies despite no actual change taking place) or to the wind direction. This could be extended so that real data recorded by actual sailing robots (or even a weather station) was replayed to form the basis of the simulated data. The simulator physics model could also be improved to better match the performance of a real boat. One possibility would be to implement the work of Roncin (2004) [109], who created a detailed simulation physics model which incorporated far more parameters than the physics model of Tracksail does. However, even with modifications such as these a large gap is still likely to exist between the performance of a real sailing robot and a simulated one. Given that these limitations, future experiments need to be run on real robots as well as simulations before their results can be fully accepted.

8.3 Limitations of Results

The results presented in this chapter and chapters 6 & 7 are subject to a number of limitations due to the environment in which the simulator and robot operate. The world of the simulator in which the water is always calm and the wind consistent is not realistic when compared to that of the robot. This appears to have exaggerated many of the simulated results. The robot was restricted to operating in a single location, travelling only on a beam reach (sailing across the wind) and going back and forth along the same 150m long course. This restricted the set of wind strengths and directions in which it was practical to sail. These environmental conditions are very different to those which would be experienced by an ocean going sailing robot, which is likely to experience a wider range of wind speeds and on average much higher speeds, waves which are often higher than the robot (but who's wavelength is at least several seconds), tides and currents. It is not entirely clear from the data gathered in this thesis how an ocean going robot who's control system is being suppressed by an artificial endocrine

system, would behave in such conditions. The experience of Pinta shows that creating an experiment to test this would be difficult as there is a high risk of not getting the robot (and more importantly the data) back. As a result the work in this thesis was very much a compromise between operating the system in the harshest conditions possible and the practicalities of getting the boat back at the end and having safe conditions for humans to launch and recover it. However, this work still managed to obtain results across a wide range of weather conditions and across a long time period.

The results shown in this thesis were not just influenced by the environmental conditions that the robot operated in, but also by the design of the robot itself. Altering the size and shaped hull may have caused them to interact very differently with the waves and the control system. Using a fabric sail or multiple wing sails would have also resulted in different behaviours. In future it would be useful to repeat these experiments on robots of different designs, to increase the sample size, to use a more complex course (e.g. a triangular shape) and to operate at sea in rougher conditions in order to test if the same trends are shown. It would also be useful to repeat the solar power simulations from section 6.4.3 upon a real robot operating under realistic ocean conditions.

8.4 Chapter Summary

This chapter has compared and contrasted the differences between the simulator and robot results from the fixed hormone and variable battery hormone experiments. Both sets of results presented similar trends, although there is an order of magnitude difference in average power consumption during the fixed hormone experiment. In the variable battery hormone experiment the battery discharge is much faster on the robot and is almost perfectly linear, while the simulator battery lasts longer and the relationship between battery charge and time is curved as increased levels of hormonal suppression decrease power consumption. It has discussed potential future improvements to the simulator and how the operating environments of both the robot and simulator may have affected and limited the results. Both the simulator and robot environment are thought to have an impact on the final result and further work is required to scale up to a more realistic simulation or an ocean going sailing robot.

Chapter 9

Conclusions and Future Work

9.1 Introduction

This thesis has shown that through the use of an algorithm inspired by the mammalian neural and endocrine systems it is possible to control the power consumption of both simulated and real sailing robots. This is achieved by modulating the magnitude of actuator movements in response to battery and sunlight levels. This chapter reviews the contents of this thesis, presents a list of contributions, summarises the answer to the research question and outlines some of the limitations of this work and possible directions for future work.

9.2 Overview

The neuro-endocrine controller developed in chapter 5 was applied to a simulated sailing robot in chapter 6 and a pair of real robots in chapter 7. In chapter 6 several experiments were performed using a simulated sailing robot controlled by two neural networks. The behaviour of these networks could be either excited or suppressed through the action of hormones in an artificial endocrine system. The first experiment fixed a hormone for the duration of the experiment and compared the relative performance of the control system between several preset hormone levels. It was shown that the hormone could be used to suppress actuator movement and reduce power consumption. The concentration of a suppressive hormone could be increased to a considerable extent without causing any noticeable reduction in performance. Although eventually the level of suppression began to affect performance and prevented the simulated robot from sailing effectively. A second experiment produced a hormone in proportion to a simulated battery charge level, reducing this every time an actuator moved. This resulted in a gradual suppression of actuator movement as the battery level dropped. As with the fixed hormone experiment the effect upon sailing performance

was not particularly noticeable until the battery reached extremely low levels and the highest hormone sensitivities were applied. A simulated solar panel was then added to the simulated robot which would recharge the battery when the sun was above the horizon. When combined with the battery hormone system this allowed the robot to sail for longer. However, in simulations with lower sun light levels it still did not allow for perpetual operation without causing the battery to become entirely drained. A second hormone was added to represent the level of sun light. This prevented the battery from becoming completely drained but achieved this by suppressing all actuator movement for much of the day leaving the robot to drift uncontrolled.

The fixed hormone and battery hormone experiments were repeated on the robots in chapter 7. The experimental results showed similar characteristics to the simulator, as it was still possible to reduce power consumption while not impacting sailing performance until the highest levels of hormone concentration were used. However, the overall power savings were significantly smaller than in the simulator but still large enough to be useful to an autonomous power management system. The difference between these results was examined in chapter 8, which showed a large gap between the results but the same overall trends in response to hormonal modulation.

To the best of the author's knowledge, this work represents the only attempt to date to experimentally verify an autonomous power management algorithm on a real sailing robot. In the process of achieving this, it also probably represents one of the only attempts to rigorously test the performance of a sailing robot over a prolonged period of time. It also represents one of the longest and most in depth studies into autonomous power management on any robotic platform. Of the relatively small amounts of prior work in this area, most have been limited to simulation or simple and short term experiments using laboratory robots. Even the timescales of the simulation work of this thesis (covering between approximately 6 and 400 hours) cover longer timescales than many of those in previous autonomous power management work.

9.3 Conclusions

This section revisits the original hypothesis and research question specified in section 1.3.

The hypothesis was:

> "That an artificial neural and endocrine system can manage power consumption in a sailing robot operating in a real environment through the modulation of actuator movement. This will result in an improvement in power management when compared to an unmodulated neural network."

And the research question was:

> "To what extent will a control system based upon an abstracted artificial neuro-endocrine control system, be able to manage power consumption in a sailing robot through the modulation of actuator movement and what effects will this have upon sailing performance, in comparison to an unmodulated neural network controller?"

From this work it has been shown that an artificial neuro-endocrine controller is able to manage power consumption in a sailing robot through the modulation of the magnitude of actuator movements. This has been demonstrated in both simulation and in a real robot.

There is a large amount of data to support the original hypothesis through both simulation and real world experimentation. The first piece of evidence to support the hypothesis is through the application of fixed levels of hormonal modulation to a neural network controlling a robot to sail a pre-defined course. The impact of this upon power consumption was evaluated in sections 6.4.1 and 7.4. It was shown that when a suppressive hormone was applied, that the power usage was reduced with regards to both power used for each kilometre travelled and each hour that elapsed. A strong correlation was shown between the hormone concentration and the power consumption suggesting that the hormone concentration was the cause of this behaviour. To further investigate this a hormone which varied in proportion to battery level was created. Sections 6.4.2 and 7.5 show that adding a variable battery hormone to a neural network control system, increased the median operation time from a single charge of a battery by up to 13 times in simulation and nearly two fold in a real robot. Attempts to measure the effects of hormonal modulation upon sailing performance have not been as conclusive with no identifiable changes in the sailing efficiency metrics being observed. The experiments in sections 6.4.3 and 6.4.3.2 demonstrate the ability of an artificial neuro-endocrine controller to work over longer time scales, adjust to external environmental factors and for multiple hormones to be combined together. This experiment demonstrates the feasibility of an artificial neuro-endocrine controller to manage power consumption, enable greater power autonomy and can enable a robot to adapt to variable conditions.

9.3.1 Key Contributions

A number of key novel contributions have been made by this thesis:

- The development of an architecture inspired by the mammalian neural and endocrine systems, which can steer and manage the power consumption of a sailing robot.
- The development of a series of metrics and a methodology for testing the power consumption of a sailing robot.
- Refinements to existing artificial neuro-endocrine control algorithms.
- The demonstration through both simulation and robot experiments, that it is possible to control power consumption in a sailing robot through the modulation of actuator movement using an artificial neuro-endocrine controller. This offers a (near) continuous method for modulating power consumption in contrast to existing fixed mode systems.

This work has the potential to enable sailing robots to adapt their behaviour to cope with unpredictable supplies and demands for energy. It allows for small and gradual changes to be made to the robot's behaviour, although large and sudden changes are also possible. This could assist sailing robots in operating autonomously for long periods of time, to cope with the degradation of batteries and solar panels and to enable them to operate with narrow margins of error within their power budgets. Ultimately this could increase the length of missions achievable by a sailing robot and reduce the need to over-engineer their power systems, reducing financial costs.

9.4 Limitations

There are several limitations of the work presented in this thesis. These primarily centre around the lack of any online learning mechanism. Should the control system take an incorrect decision there is no mechanism to identify this, rectify it or prevent it from happening again. There is also no mechanism to learn optimal hormone sensitivities. Section 6.4.3 demonstrated some of the difficulties involved in choosing a suitable sensitivity. Scaling up the system to include more hormones is likely to worsen this. In particular with hormones which are only released occasionally choosing appropriate sensitivities and testing their responses becomes difficult. This could potentially be solved through a mechanism which can learn the optimal sensitivities. It is possible that as new situations are encountered the optimal sensitivity will change and therefore any learning algorithm should be run online rather than in advance. Some potential strategies for performing this are discussed in the next section.

9.5 Future Work

There are a number of potential directions for future work based upon this thesis. These include increasing the number of hormones, longer term experiments, improvements to the neuro-endocrine architecture and using neuro-endocrine architectures in new applications.

9.5.1 Expansion of Current Architecture to Include Additional Hormones

The current neuro-endocrine architecture could be expanded to cover a number of additional variables represented as additional hormones. At present only battery level and sun elevation have been considered. Wind speed information could be used to produce an inhibitory hormone, which will suppress the entire control system in the event of there being no wind. This information could be derived from the speed the robot has achieved, or from a wind speed sensor. If a passively stable robot, such as BeagleB (described in section 4.1.3), were used then it would also be possible to suppress almost all control system activity when sailing on a fixed course between two waypoints.

Another possibility is to produce an excitory hormone in response to heading errors (the difference between current heading and target heading). This hormone would excite the steering neural network to increase the magnitude of steering movements. Such a system would also be able to adjust its behaviour to changes in sea state, as these would cause increases in heading errors. This could be combined with a system which suppresses steering between waypoints so that any heading error will excite the control system to produce remedial action if the heading error increases.

A shortcoming of the current system is that there is no interaction between the sail and rudder neural networks. The sail network reacts directly to the current wind direction without consideration for the target heading of the boat even though the sail will affect the steering of the boat. An additional hormone could be used to allow the rudder setting neural network to secrete a hormone or hormones, which are able to influence the sail setting neural network in order to help with steering.

This work has already hinted at the possibility of adjusting behaviour in response to degradation of components. For example photo-voltaic solar panels, fuel cells and batteries all degrade in performance over their lifetime which will reduce the available power budget to the system. A neuro-endocrine architecture could be used to gradually adjust the behaviour of a system as these components degrade. This could allow for a trade off between reducing the length of a mission or reducing performance, while keeping the mission length the same.

To test the scalability of this work, it needs to be reapplied in longer term studies in

robots which are capable of operating for at least several days. The 2011 Microtransat race from Europe to the Caribbean presents a potential opportunity to apply these algorithms in an experiment which could be expected to last several months. It would also be interesting to compare the performance of a neuro-endocrine power management system against a multi-mode system such as that proposed by Blair (2010) [18] or Frey (2009) [44].

9.5.2 Modifications to the Current Architecture

A number of further modifications to the artificial neuro-endocrine architecture developed in this thesis could be applied to increase the resilience and capabilities of a neuro-endocrine controller. These include online learning of hormone sensitivities and through artificial hormone cascades. Potentially the Hebbian learning approach proposed by Timmis, Neal and Thorniley (2009) [136] or an artificial immune systems approach such as that used by Timmis, Neal and Murray (2010) [135] or Singh and Nair (2005)[122].

In the current architecture hormone sensitivities are fixed at preset values, the two hormone system discussed in section 6.4.3 demonstrated the difficulties in deciding upon sensible values for these. This problem could potentially be solved with some form of online learning system which can adjust the hormone levels. An artificial immune system or genetic algorithm could potentially perform this role. If a genetic algorithm were used the early generations are likely to produce highly undesirable results and some kind of simulated bootstrap process would probably be required.

The artificial endocrine system could be extended to create an artificial hormone cascade (as discussed in section 2.6.2 and 5.3.2) where one hormone triggers the production of another. In biological examples this can often involve three or more hormones all acting as part of a larger feedback loop. This process will allow a more gradual response with the secondary (and tertiary etc) hormones being able to either reinforce certain behaviours or to act as a form of timeout to limit the production of the initial hormone. If glands were also implemented as neural networks, which themselves can be modulated by hormones, then this offers a method to implement this process. By using hormones with slow response times (small r values) this delayed response can be created. By combining several hormones with linear response rates it is possible that a system which is non-linear in its behaviour will emerge under certain situations.

9.5.3 Alternative Applications

Artificial neuro-endocrine controllers could be applied to power management problems in many systems other than sailing robots. Potential candidates are any system which must

contend with limited, variable and unpredictable energy sources and that offer some continuous mechanisms through which power consumption can be modulated (even if these reduce performance). It is also useful to possess the ability to exploit the opportunity of excess power availability to improve performance and make up for poor performance during periods of low energy availability.

Potential applications include robotic systems requiring long term autonomy, portable battery powered electronics, sensor networks and even electricity grids.

Robotic systems which must operate unattended for long periods of time could benefit from the use of neuro-endocrine controllers, examples might include autonomous underwater vehicles (AUVs) or unmanned aerial vehicles (UAVs). In particular those systems which must operate from unpredictable power sources such as solar power or from power sources which degrade over their lifespan. A prime example is hydrogen polymer electrolyte membrane fuel cells [1] which degrade by approximately 5% per hour for every 100 hours of operation after an initial 500 hour period.

Applications could also be made to portable electronic devices such as mobile phones. While creating hormones in response to battery levels, user demand and usage profiles (e.g. minimum power consumption or maximum performance) it is possible to modulate CPU clock speeds, LCD back light brightness and timeout periods or by switching off non-essential auxiliary functions such as GPS and bluetooth.

It might also be possible to apply an artificial endocrine controller to a "smart" electricity grid [104] in which the grid is able to communicate with devices connected to it and have them switch off or reduce their power consumption. Target devices would be high power demand, non-critical devices which can be safely switched off for short periods of time. Examples of such devices might include fridges and freezers, washing machines and dish washers, electric car chargers or heaters. By modulating the power consumption of these devices demand peaks can be smoothed out and unpredictable spikes in renewable energy can be exploited. Users of these appliances could specify certain criteria such as "I must have my electric car charged by 8am tomorrow" allowing the grid flexibility in when to supply the user with power while still fulfilling the users requirements. Power demand can be modelled as a single hormone, when demand exceeds supply it becomes inhibitory and begins to shutdown devices, when supply exceeds demand it becomes excitory and begins to switch on devices. Additional hormones can also be used to act as supply and demand predictors based upon previous usages patterns, weather forecasts and known future availability of power stations. As coal and nuclear power stations in particular take a long time to reach full power from idle this prediction would be required in order to achieve a fully autonomous system.

[1] `http://www.hes.sg/files/AEROPAK_Technical_Data_Sheet.pdf` accessed 14/04/2011

Bibliography

[1] J. Abril, J. Salmon, and O. Calvo. Fuzzy control of a sailboat. *International Journal of Approximate Reasoning*, 16(3-4):359–375, April - May 1997.

[2] J. C. Alves, T. M. Ramos, and N. A. Cruz. Ocean sampling and surveillance using autonomous sailboats. In *Proceedings of the 1st International Robotic Sailing Conference, May 23rd-24th 2008, Breitenbrunn, Austria*, pages 13–20, 2008.

[3] J. C. Alves, T. M. Ramos, and N. A. Cruz. A reconfigurable computing system for an autonomous sailboat. In *Proceedings of the 1st International Robotic Sailing Conference, May 23rd-24th 2008, Breitenbrunn, Austria*, pages 13–20, 2008.

[4] N. Ammann, R. Biemann, F. Hartmann, C. Hauft, I. Heinecke, P. Jauer, J. Kruger, T. Meyer, R. Bruder, and A. Schlaefer. Towards autonomous one-design sailboat racing: navigation, communication and collision avoidance. In *Proceedings of the 3rd International Robotic Sailing Conference*, pages 44–48, 2010.

[5] N. Ammann, F. Hartmann, P. Jauer, R. Bruder, and A. Schlaefer. Design for a robotic sailboat for the wrsc/sailbot. In *Proceedings of the 3rd International Robotic Sailing Conference*, pages 41–43, 2010.

[6] M. Arbib. *The Metaphorical Brain 2 - Neural Networks and Beyond*. Wiley, 2nd edition, 1989. ISBN 0-471-09853-1.

[7] R. C . Arkin. *Survivable Robotic Systems: Reactive and Homeostatic Control in Robotics and remote systems for hazardous environments*, volume 1, chapter 7, pages 135–154. Prentice Hall, 1993.

[8] R. C. Arkin. Homeostatic control for a mobile robot: Dynamic replanning in hazardous environments. *Journal of Robotic Systems*, 9:197–214, 1992.

[9] R.C. Arkin. Motor schema based mobile robot navigation. *The International Journal of Robotics Research*, 8(4):92–112, 1989.

[10] R.C. Arkin. Towards the unification of navigational planning and reactive control. In *Proceedings of AAAI Spring Symposium on Robot Navigation*, 1989.

[11] W. R. Ashby. *Design for a Brain: The Origin of Adaptive Behavior*. Springer, second edition, 1960. SBN 412-04930-9.

[12] Avila-García and D. Cañamero. Using hormonal feedback to modulate action selection in a competitive scenario. In S. Schaal, A.J. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J. A. Meyer, editors, *Proceedings of the 8th International Conference on Simulation of Adaptive Behaviour (SAB'04)*, pages 243–252. MIT Press, 2004.

[13] T. Balch and H. Yanco. Ten years of the aaai mobile robot competition and exhibition looking back and to the future. *AI Magazine*, 23(1):13–22, 2002.

[14] W. H. Bayliss and E. H. Starling. The mechanism of pancreatic secretion. *The Journal of Physiology*, 28(5):325–353, 1905.

[15] L. Benini, A. Bogliolo, and G. De Micheli. A survey of design techniques for system-level dynamic power management. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 8(3):299–317, 2000.

[16] P. J. Bentley. *Comparative Vertebrate Endocrinology*. Cambridge University Press, 2nd edition, 1982. ISBN 0-521-24653-9.

[17] H. Besedovsky and E. Sorkin. Network of immune-neuroendocrine interactions. *Clinical and Experimental Immunology*, 1:1–12, January 1977.

[18] C. Blair. Autonomous sailboat power system design. In *Proceedings of the 3rd International Robotic Sailing Conference*, pages 49–56, 2010.

[19] V. Braitenberg. *Vehicles - Experiments in Synthetic Physcology*. MIT Press, 1996. ISBN 0-262-52112-1.

[20] Y. Briere. Iboat: an unmanned sailing robot for the microtransat challenge and ocean monitoring. In *Proceedings of TAROS 2007, September 3rd-5th 2007, Aberystwyth, UK*, 2007.

[21] Y. Briere, F. Bastianelli, and M. Gagneul. Challenge microtransat. In *Proceedings of CETSIS 2005, Nancy, France, October 25th-27th 2005*, 2005.

[22] Y. Briere, F. L. Cardoso Ribeiro, and M.A. Vieira Rosa. Design methodologies for the control of an unmanned sailing robot. In *Proceedings of the 8th IFAC International Conference on Manoeuvring and Control of Marine Craft, September 16-18, 2009, Guarujá (SP), Brazil*, pages 58–65, 2009.

[23] R. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2:14–23, 1986.

[24] R. Brooks. *Elephants Don't Play Chess.* MIT Press, 1990.

[25] R. Brooks. Integrated systems based on behaviors. *SIGART Bulletin*, 2(4):46–50, 1991.

[26] R. Brooks. Intelligence without reason. Technical report, MIT, 1991.

[27] R. A. Brooks and A. M. Flynn. Fast, cheap and out of control: A robot invasion of the solar system. *Journal of the British Interplanetary Society*, 42:478–485, 1989.

[28] R. Bruder, A. Schlaefer, and B. Stender. Model sailboats as a testbed for artificial intelligence methods. In *Proceedings of the 2nd International Robotic Sailing Conference*, pages 37–42, July 2009.

[29] M. Campbell, E. Garcia, and D. Huttenlocher et al. Team cornell: Technical review of the DARPA urban challenge vehicle. Technical report, Cornell University, 2007.

[30] W. B. Cannon. *Bodily Changes In Pain, Hunger, Fear And Rage - An Account Of Recent Researches Into The Function Of Emotional Excitement.* Cannon Press, 1927. ISBN 1-406-75539-7.

[31] W. B. Cannon. Organization for physiological homeostasis. *Physiological Review*, 9:339–431, 1929.

[32] K. Capek. R.U.R(Rossum's Universal Robots). Play, January 1921.

[33] D. Cañamero. Modeling motivations and emotions as a basis for intelligent behavior. In *Proceedings of the First International Symposium on Autonomous Agents (Agents'97)*, pages 148–155. ACM Press, 1997.

[34] L. Crapo. *Hormones: The Messengers of Life.* W. H. Freemany and Company, 1985. ISBN 0-7167-1753-0.

[35] L.N. de Castro and F. J. Von Zuben. Learning and optimization using the clonal selection principle. *IEEE Transactions on Evolutionary Computation, Special Issue on Artificial Immune Systems*, 6(3):239–251, 2002.

[36] M. Ebb. Zero handed transatlantic race. *Latitude 38*, October 1988.

[37] G. Elkaim. *System Identification for Precision Control of a Wingsailed GPS Guided Catamaran*. PhD thesis, Stanford University, December 2001.

[38] G. Elkaim and C.O. Boyce. An energy scavenging autonomous surface vehicle for littoral surveillance. In *Proceedings of ION Global Navigation Satellite Systems Conference*, 2008.

[39] L. Elkins, D. Sellers, and W. Reynolds-Monach. The autonomous maritime navigation (amn) project: Field tests, autonomous and cooperative behaviors, data fusion, sensors, and vehicles. *Journal of Field Robotics*, 27(6):790–818, 2010.

[40] C. Reinholtz et al. DARPA urban challenge technical paper - team victor tango. Technical report, Virginia Tech, 2007.

[41] J. Firby. *Adaptive Execution in Complex Dynamic Worlds*. PhD thesis, Yale University, 1989.

[42] S. Forrest, A.S. Perelson, L. Allen, and R. Cherukuri. Self-nonself discrimination in a computer. In *Proceedings of the 1994 IEEE symposium on Research in Security and Privacy*, 1994.

[43] C. Francheschi and E. Ottaviani. The invertebrate phagocytic immunocyte: clues to a common evolution of immnue and neuroendocrine systems. *Immunol Today*, 18:169–174, 1997.

[44] D. Frey. Power aware communication platform for an autonomous sailboat. Master's thesis, Swiss Federal Institute of Technology Zurich, 2009.

[45] S. Gadanho and J. Hallam. The role of emotions: Exploring autonomy mechanisms in mobile robotics. *Adaptive Behaviour*, 1997.

[46] S. Gadanho and J. Hallam. Exploring the role of emotions in autonomous robot learning. In D. Ca namero, editor, *Emotional and Intelligent: The Tangled Knot of Cognition*, pages 84–89, 1998.

[47] S. Gadanho and J. Hallam. Emotion triggered learning in autonomous robot control. *A Special Issue of Cybernetics and Systems*, 32:531–559, 2001.

[48] J. A. Gally, R. Montague, G. N. Reeke, and G. M. Edelman. The no hypothesis: Possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system. *Neurobiology*, 87:3547–3551, May 1990.

[49] E. Gat. *Three Layer Architectures*, chapter 8, pages 195–210. AAAI Press / MIT Press, 1998.

[50] L. Giger, S. Wismer, S. Boehl, H. Erckens G. Busser and, J. Weber, P. Moser, P. Schwizer, and R. Siegwart C. Pradalier. Design and construction of the autonomous sailing vessel avalon. In *Proceedings of the 2nd International Robotic Sailing Conference, July 6th-12th 2009, Matosinhos, Portugal*, pages 17–22, 2009.

[51] D. G. Hardie. *Biochemical Messengers: Homrones, Neurotransmitters and Growth Factors*. Chapman and Hall, 1991. ISBN 0-412-30350-7.

[52] S. Harnad. The symbol grounding problem. *Physica*, 42:335–346, 1990.

[53] E. Hart and P. Ross. Improving sosdm: Inspirations from the danger theory. In *Artificial Immune Systems*, Lecture Notes in Computer Science, pages 194–203. Springer Berlin / Heidelberg, 2003.

[54] D. O. Hebb. *The organization of behavior : a neuropsychological theory*. John Wiley & Sons Inc., 1949.

[55] J. Henley. *An Artificial Neuroendocrine Kinematics Model for Legged Robot Obstacle Negotiation*. PhD thesis, Jan 2006.

[56] J. Henley and D. Barnes. An artificial neuro-endocrine kinematics model for legged robot obstacle negotiation. In *Proceedings of the 8th ESA Workshop on Adavanced Space Technologies for Robotics and Automation*, 2004.

[57] O. Holland. Exploration and high adventure: The legacy of grey walter. *Philosophical Transactions*, 361:2085–2121, 2003.

[58] B.P. Howell, S. Wood, and S. Koksal. Passive sonar recogition and analysis using hybrid neural networks. *Proceedings of IEEE Oceans 2003*, 2(22-26):1917–1924, 2003.

[59] R. Humza, O. Schloz, M. Mokhtar, J. Timmis, and A. Tyrrell. Towards energy homeostasis in an autonomous self-reconfigurable modular robotic organism. *Comptation World*, pages 21–26, 2009.

[60] P. Husbands, T. Smith, N. Jakobi, and M. O'Shea. Better living through chemistry: Evolving gasnets for robot control. *Connection Science*, 10(3 and 4):185–210, 1998.

[61] X. Jiang, J. Polastre, and D. Culler. Perpetual environmentally powered sensor networks. In *Proceedings of Fourth International Symposium on Information Processing in Sensor Networks*, pages 463–468, April 2005.

[62] S. C. Johnson. A c program checker. Technical Report 65, Bell Labs, September 1977.

[63] A. Kansal, D. Potter, and M. B. Srivastava. Performance aware tasking for environmentally powered sensor networks. *ACM SIGMETRICS Performance Evaluation Review*, 32(1):223–234, 2004.

[64] C. Kilian. *Modern Control Technology: Components and Systems*, chapter 11, pages 361–409. Delmar Thomson Learning, 2002. ISBN 0-766-82358-X.

[65] A. Klesh and P. Kabamba. Solar-powered aircraft: Energy-optimal path planning and perpetual endurance. *Journal of Guidance, Control and Dynamics*, 32:1320–1329, 2009.

[66] C. R. Kube and H. Zhang. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–218, September 1993.

[67] J.H Lever, L.R. Ray, A. Streeter, and A. Price. Solar power for an antarctic rover. *Hydrological Processes*, 20(4):629–644, 2006.

[68] M. Mila Macchi and J. N. Bruce. Human pineal physiology and functional significance of melatonin. *Frontiers of Neuroendocrinology*, 25:177–195, 2004.

[69] P. Maes. The dynamics of action selection. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 991–997, 1989.

[70] O. Matan, R. K. Kiang, C. E. Stenard, B. Boser, J.S .Denker, D. Henderson, R. E. Howard, W. Hubbard, L. D. Jackel, and Y. Le Cun. Handwritten character recogition using neural network architectures. In *proceedings of the 4th USPS advanced Technology Conference*, pages 1003–1011, November 1990.

[71] P. Matzinger. The danger model: A renewed sense of self. *Science*, 296(12):301–305, 2002.

[72] W. S. McCulloch and W. Pitts. A logical calculus of the ideas imminent in nervous activity. *Bulletin of Mathematical Biophysics*, 5:115–133, 1943.

[73] D. McFarland. *Autonomy and Self-Sufficiency in Robots. In: The 'Artificial Life' Route to Artificial Intelligence*, chapter 6, pages 187–213. Lawrence Erlbaum Associates, 1995. ISBN 0-8058-1519-8.

[74] D. McFarland and E. Spier. Basic cycles, utility and opportunism in self-sufficient robots. *Robotics and Autonomous Systems*, 20(2):179–190, June 1997.

[75] Y. Mei, Y.H. Lu, Y.C Hu, and C.S.G. Lee. A case study of mobile robot's energy consumption and conservation techniques. In *12th International Conference on Advanced Robotics, 2005. ICAR '05.*, pages 492–497, July 2005.

[76] M. Mendao. Hormonally moderated neural control. In *Proceedings of the 2004 AAAI Spring Symposium.*, AAAI Press, Menlo Park, CA, 2004.

[77] M. Mendao. A neuro-endocrine control architecture applied to mobile robotics. 2006.

[78] M. Mendao. David meets goliath: A bio-inspired architecture measured against a traditional architecture. 2007.

[79] P. Miller, B. Beal, C. Capron, R. Gawboy, P. Mallory, C. Ness, R. Petrosik, C. Pryne, T. Murphy, and H. Spears. Increasing performance and added capabilities of usna sail-powered autonomous surface vessels (asv). In *Proceedings of the 3rd International Robotic Sailing Conference*, 2010.

[80] P. Miller, O. Brooks, and M. Hamlet. Development of the usna sailbots (asv). In *Proceedings of the 2nd International Robotic Sailing Conference, July 6th-12th 2009, Matosinhos, Portugal*, pages 9–16, 2009.

[81] M. Minsky and S Papert. *Perceptrons*. MIT Press, Cambridge, MA, USA, 2nd edition, 1989. ISBN 0-262-63111-3.

[82] M. Mirolli and D. Parisi. *Artificial Organisms That Sleep in Advances in Artificial Life*, volume 2801/2003 of *Lecture Notes in Computer Science*, pages 377–386. Springer Berlin/Heidelberg, 2003.

[83] Team MIT. Technical report: DARPA urban challenge. Technical report, Massachusetts Institute of Technology, 2007.

[84] R. C. Moioli, P. A. Vargas, F.J. von Zuben, and P. Husbands. *Advances in Artificial Intelligence - SBIA 2008*, volume 5249/2008, chapter Evolving an Artificial Homeostatic System, pages 278–288. Springer Berlin / Heidelberg, 2008.

[85] R.C Moioli, P.A Vargas, F.J. Von Zuben, and P. Husbands. Towards the evolution of an artificial homeostatic system. In *Evolutionary Computation, IEEE World Congress on Computational Intelligence 2008, June 1st-6th 2008, Hong Kong*, pages 4023–4030, 2008.

[86] F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klaptocz, S. Magnenat, J. Zufferey, D. Floreano, and A. Martinoli. The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions*, volume 1, pages 59–65, 2009.

[87] F. Mondada and D. Floreano. Evolution and neural control structures: Some experiments on mobile robots. *Robotics*, 16(2-4):183–196, 1995.

[88] H. Moravec. Visual mapping by a robot rover. In *Proceedings of the 6th IJCAI*, 1979.

[89] K. Murphy, P. Travers, and M. Walport. *Janeway's Immunobiology*. Garland Sceince, 7th edition, 2008. ISBN 0-8153-4123-7.

[90] M. Neal. A hardware proof of concept of a sailing robot for ocean observation. *IEEE Transactions on Oceanic Engineering*, 31, 2:462– 469, 2006.

[91] M. Neal, J. Feyereisl, R. Rascunà, and X. Wang. Don?t touch me, i?m fine: Robot autonomy using an artificial innate immune system. In *Proceedings of the 5th International Conference on Artificial Immune Systems (ICARIS2006)*, pages 349–361. Springer Berlin / Heidelberg, 2006.

[92] M. Neal and J. Timmis. Timidity: A useful emotional mechanism for robot control? *Informatica*, 27:197–204, 2003.

[93] M. Neal and J. Timmis. Once more unto the breach? towards artificial homeostasis. In L.N. De Castro and F.J. Von Zuben, editors, *Recent Developments in Biologically Inspired Computing*, pages 340–365, 2005.

[94] U. Nehmzow. *Mobile Robotics - A Practical Introduction*. Springer-Verlag, 2003. ISBN 1-85233-726-5.

[95] T. D. Ngo and H. Schioler. Sociable robots through self-maintained energy. *International Journal of Advanced Robotic Systems*, 3(4):313–322, 2006.

[96] D. Nguyen and B. Widrow. The truck backer-upper: An example of self-learning in neural networks. In *proceedings of the International Joint Conference on Neural Networks*, volume 2, pages 357–363, 1989.

[97] N. J. Nilsson. A mobile automation: An application of artificial intelligence techniques. In *Proceedings of the IJCAI*, 1969.

[98] S. Nolfi and D. Floreano. *Evolutionary robotics : the biology, intelligence, and technology of self-organizing machines.* MIT Press, 2000. ISBN 0-262-14070-5.

[99] D. Parisi. Internal robotics. *Connection Science*, 16(4):325–338, 2004.

[100] A. S. Perelson, , and G. F. Oster. Theoretical studies of clonal selection: Minimal antibody repertoire size and reliability of self-non-self discrimination. *Journal of Theoretical Biology*, 81(4):645–670, December 1979.

[101] N. Petrovsky. Towards a unified model of neuroendocrine-immune interaction. *Immunology and Cell Biology*, 79:350–357, 2001.

[102] P. Piction. *Neural Networks.* Palgrave, Houndmills, Basingstoke, Hampshire, RG21 6XS, UK, 2000. ISBN 0-333-80287-8.

[103] D. Pomerleau. *Neural Network Vision for Robot Driving.* MIT press, 1995.

[104] C. W. Potter, A. Archambault, and K. Westrick. Building a smarter smart grid through better renewable energy information. In *Proceedings of Power Systems Conference and Exposition*, pages 1–5, 2009.

[105] L. Ray, J. Lever, A. Streeter, and D. Price. Design and power management of a solar-powered "cool robot" for polar instrument networks. *Journal of Field Robotics*, 24(7):581–599, 2007.

[106] E. Rich and K. Knight. *Artificial Intelligence*, chapter 18, pages 492–509. McGraw Hill, 2nd edition, 1991. ISBN 0-071-00894-2.

[107] P. H. Richter. A network theory of the immune system. *European Journal of Immunology*, 5(5):350–354, 1975.

[108] C. Rocks and D. Barnes. Entrainment as a paradigm for modelling a planetary robot's circadian rhythm. In *Proceedings of the 8th ESA Workshop on Advanced Space Technologies for Robotics and Automation*, ESTEC, Noordwijk, The Netherlands, November 2004.

[109] K. Roncin and J. M. Kobus. Dynamic simulation of two sailing boats in match racing. *Sports Engineering*, 7(3):139–152, September 2004.

[110] F. Rosenblatt. The perceptron: A probabilistic model of information sotrage and organisation in the brain. *Pyschological Review*, 65:386–408, 1958.

[111] S. Roth. Ray casting for modeling solids. *Computer Graphics and Image Processing*, 18(2):109–144, 1982.

[112] D. Rumelhart, G. Hinton, and R. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, October 1986.

[113] D. Rumelhart and J. McCelland. *Parallel Distributed Processsing*, volume 1. MIT Pr, 1988. ISBN 0-262-18120-7.

[114] P. F. Rynne and K. D. von Ellenrieder. Unmanned autonomous sailing: Current status and future role in sustained ocean observations. *MTS Journal*, 43(1):21–30, 2009.

[115] C. Sauzé. Control software for a sailing robot. Master's thesis, University of Wales, Aberystwyth, 2005.

[116] C. Sauzé and M. Neal. A biologically inspired approach to long term autonomy and survival in sailing robots. In *Proceedings of the International Robotic Sailing Conference, May 23rd-24th 2008, Breitenbrunn, Austria*, pages 6–11, 2008.

[117] C. Sauze and M. Neal. A raycast approach to colliision avoidance in sailing robots. In *Proceedings of the 3rd International Robotic Sailing Conference, June 7th-10th 2010, Kingston, Ontario, Canada*, pages 26–33, 2010.

[118] W. Shen, C. Chuong, and P. Will. Digital hormone models for self-organization. *Artificial Life*, 8:116–120, 2002.

[119] W. Shen., P. Will, A. Galstyan, and C. Choung. Hormone-inspired self-organization and distributed control of robotic swarms. *Autonomous Robots*, 17:93–105, 2004.

[120] R. Shiffrin and W. Schneider. Controlled and automatic human information processing: Ii. *Psychological Review*, 84:127–190, 1977.

[121] K. Shillcutt. *Solar Based Navigation for Robotic Explorers.* PhD thesis, Carnegie Mellon University, 2000.

[122] C. T. Singh and S. B. Nair. An artificial immune system for a multiagent robotics system. *World Academy of Science, Engineering and Technology*, 11, 2005.

[123] J. Sliwka, P. Reilhac, R. Leloup, P. Crepier, H. De Malet, P. Sittaramane, F. Le Bars, K. Roncin, B. Aizier, and L. Jaulin. Autonomous robotic boat of ensieta. In *Proceedings of the 2nd International Robotic Sailing Conference, July 6th-12th 2009, Matosinhos, Portugal*, pages 1–8, 2009.

[124] B. Smith. Skamp - robot boat with rigid sails patrols ocean beat. *Popular Science*, pages 70–71,142, May 1970.

[125] P. Sprent. *Applied Non-Parametric Statistical Methods*. Chapman and Hall, 1990. ISBN 0-412-30610-7.

[126] R. Stelzer and K. Jafarmadar. A layered system architecture to control an autonomous sailboat. In *Proceedings of TAROS 2007*, September 2007.

[127] R. Stelzer and K. Jafarmadar. Communication architecture for autonomous sailboats. In *Proceedings of the 2nd International Robotic Sailing Conference, July 6th-12th 2009, Matosinhos, Portugal*, pages 31–36, 2009.

[128] R. Stelzer, H. Klinck, K. Jafarmadar, and D. Mellinger. AAS endurance: An autonomous acoustic sailboat for marine mammal research. In *Proceedings of the 2nd International Robotic Sailing Conference, July 6th-12th 2009, Matosinhos, Portugal*, pages 43–48, 2009.

[129] R. Stelzer and T. Pröll. Autonomous sailboat navigation for short course racing. *Robotics and Autonomous Systems*, 56:604–614, July 2008.

[130] Z. Sun and J. Reif. On energy-minimizing paths of terrains for a mobile robot. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2003.

[131] Ben Franklin Racing Team. 2007 DARPA urban challenge the ben franklin racing team. Technical report, University of Pennsylvania, Lehigh University and Lockheed Martin Advanced Technology Laboratories, 2007.

[132] Stanford Racing Team. Stanford's robotic vehicle "junior": Interim report. Technical report, Stanford University, 2007.

[133] S. Thrun, A. Bucken, W. Burgard, D. Fox, T. Frohlinghaus, D. Hennig, T. Hofmann, M. Krell, and T. Schmidt. Map learning and high-speed navigation in rhino. In S. Thrun et al, editor, *Artificial Intelligence and Mobile Robots*, chapter 1, pages 21–52. AAAI Press / MIT Press, 1998.

[134] J. Timmis, A.Tyrrell, M. Mokhtar, A. R. Ismail, N. Owens, and R. Bi. *An Artificial Immune System for Robot Organisms*, chapter 4, pages 268–288. 2010.

[135] J. Timmis, L.Murray, and M. Neal. A neural-endocrine architecture for foraging in swarm robotic systems. *Studies in Computational Intelligence*, 284:319–330, 2010.

[136] J. Timmis, M. Neal, and J. Thorniley. An adaptive neuro-endocrine system for robotic systems. In *IEEE Workshop on Robotic Intelligence in Informationally Structured Space. Part of IEEE Workshops on Computational Intelligence*, pages 129–136, 2009.

[137] The Open University. Learning space unit S377 - cell signalling. http://openlearn.open.ac.uk/mod/oucontent/view.php?id=398848. accessed 19/04/2011.

[138] C. Urmson and et. al. Tartan racing: A multi-modal approach to the DARPA urban challenge. Technical report, Team Tartan, 2007.

[139] W. Verity. I crossed the atlantic in a 12-foot sloop. *Popular Science*, pages 104–108, January 1967.

[140] J. Walker. Wingsail systems. US Patent 4,770,113, May 1986.

[141] J. Walker and M. Wilson. Hormone-inspired control for group task-distribution. In *Proceedings of TAROS (Towards Autonomous Robotic Systems) 2007*, 2007.

[142] G. W. Walter. *The Living Brain*. Gerald Duckworth & Co. LTD., 1953.

[143] A. Whitbrook, U. Aickelin, and J. Garibaldi. Idiotypic immune networks in mobile-robot control. *IEEE Transactions on Systems, Man, and Cybernetics*, 37(6):1581–1598, 2007.

[144] A. Whitbrook, U. Aickelin, and J. Garibaldi. Two-timescale learning using idiotypic behaviour mediation for a navigating mobile robot. *Applied Soft Computing*, 10(3):876–887, June 2010.

[145] N. Wiener. *Cybernetics: or Control and Communication in the Animal and the Machine*. The MIT Press, 2nd edition, 1965. ISBN 978-0-262-73009-9.

[146] K. Xiao, J. Sliwka, and K. Roncin L. Jaulin. A wind-independent control strategy for autonomous sailboats based on voronoi diagram. 2009.

[147] M. Yamamoto. Sozzy: A hormone-driven autonomous vacuum cleaner. Technical Report FS-93-09, Matsushita Research Institute and MIT Artificial Intelligence Laboratory, 1993.

Appendix A

Glossary of Terms

A.1 Sailing Terminology

This section provides a glossary of sailing terms for those readers who are not familiar with these.

Beam Reach - Sailing with the wind blowing towards the side of the boat. The direction of travel for the boat and the direction the wind is blowing from will be perpendicular to each other.

Bermudan Rig - A popular style of sail rig, found on most modern sailing boats. It features two triangular shaped sails, a jib and a mainsail.

Boom - A pole which extends backwards from the mast and is connected to the bottom of the main sail.

Bow - The front of the boat.

Broad Reach - Sailing with the wind halfway between straight behind and to the side of the boat. Halfway between a run and beam reach.

Close Hauled - Sailing towards the wind with an angle of about 45-60 degrees between the boat's direction of travel and the wind.

Jib - The front sail on a bermudan rigged boat.

Jibe - Turning the stern (back) of the boat through the direction of the wind.

Leeward - The downwind side.

Mainsail - The rear (and usually larger) sail on a bermudan rigged boat.

Mast - A pole extending upwards from the deck which supports the sails.

Point of Sail - The course the boat is sailing with respect to the wind. Common points of sail include Close Hauled, Beam Reach, Broad Reach and Running.

Port - Left side.

Running - Sailing with the wind blowing from directly behind the boat.

Sheet - A rope controlling the position of a sail.

Starboard - Right side.

Stern - Back of the boat.

Tack - The side of the boat the wind is blowing from.

Tacking - Turning the front of the boat or alternating from one tack to another in order to sail upwind in a zig-zag pattern.

Windward - The upwind side.

A.2 Biology Terminology

This section provides a glossary of biology terms used throughout this thesis for those readers who are not familiar with them.

Autocrine Signalling - A form of intra cellular signalling where the cell releases a messenger particle which then binds to a receptor on the cell surface. Cytokines sometimes communicate in this way.

Cytokine - Protein based messengers used primarily by the immune system. They trigger a behavioural change in cells with appropriate receptors for a particular cytokine. They can reach target cells either in an autocrine, paracrine or endocrine manner.

Endocrine Signalling - A form of inter cellular signalling where messenger chemicals (such as hormones or cytokines) are released into the bloodstream and can reach cells almost anywhere in the body.

Gland (endocrine) - A group of cells or an organ which secrete hormones into the bloodstream. This will be done in response to certain stimuli which can include neural stimuli, immune system signals or other hormones.

Glucose - A simple sugar which can be directly used by cells as a source of energy. Cells will take glucose from the blood when the hormone insulin is released.

Glycogen - A starch based substance which is used to store energy by the body. The process to convert it into glucose is controlled by the hormone glucagon.

Hormone - A chemical messenger which is released by an endocrine gland and travels through the bloodstream to virtually all cells in the body. They bind with a target cell to trigger a change in that cell's behaviour. This can either be by binding to receptors on the cell surface or by diffusing into the target cell.

Hormone Cascade - A process in which a hormone triggers the production of several others, which creates a feedback loop controlling the production of the original hormone.

Paracrine Signalling - An inter-cellular signalling process where the messenger particles diffuse through the extra cellular fluid between cells. This method of signalling is sometimes used by cytokines.

Peptide Hormone - A class of hormone consisting of a short chain of amino acid molecules. They are relatively large and must bind with receptors on the surface of a target cell. Upon binding they trigger the release of intra-cellular messengers which trigger behavioural changes within the cell itself. Their actions typically occur within a few minutes unlike steroid hormones which take much longer to act.

Receptor - A structure located on the surface of or inside a cell that binds only with specific substances such as hormones or cytokines.

Steroid Hormone - A class of hormone derived from cholesterol molecules. Physically much smaller than a peptide hormone. They enter the body of the cell and interact directly with the internal apparatus of the cell, triggering a change of behaviour in that cell.

Appendix B

ARC's Motor Controllers

Figure B.1 shows a block diagram of the components in ARC and the connections between the computers and motor controllers. Figure B.2 shows the circuit diagram for one of ARC's (see section 4.1.2) motor controllers. Each controller is capable of driving two stepper motors. Each stepper motor is driving by sequentially turning each one of its four pins on and off. The direction of the sequence determines the direction of travel for the motor. The PIC16F83P microcontroller is a Devantech GPIO14 [1] which acts as an IO extender for a Gumstix single board computer communicating via an I2C bus (SDA and SCL lines). The Gumstix is able to control the state of each of the 8 motor control pins via the GPIO14. Two DS1631 temperature sensors are attached to the heatsinks which back power transistors T1-T4 and T5-T8. These also use the I2C bus and provide feedback to the Gumstix as to how hot the heatsinks and power transistors are. The DIP switches on each DS1631 control its address on the I2C bus (which must be unique). Connectors at PAD7-12 provide analogue inputs to the GPIO14 from a feedback potentiometer and from feedback micro switches on PAD3-6.

ARC has a total of three stepper motors (two for sails, one for the rudder) and it was intended that each motor would be controlled by two independent controller boards. Should one begin to overheat or fail then the other could be used instead.

[1] `http://www.robot-electronics.co.uk/htm/gpio14tech.htm` accessed 15/04/2011

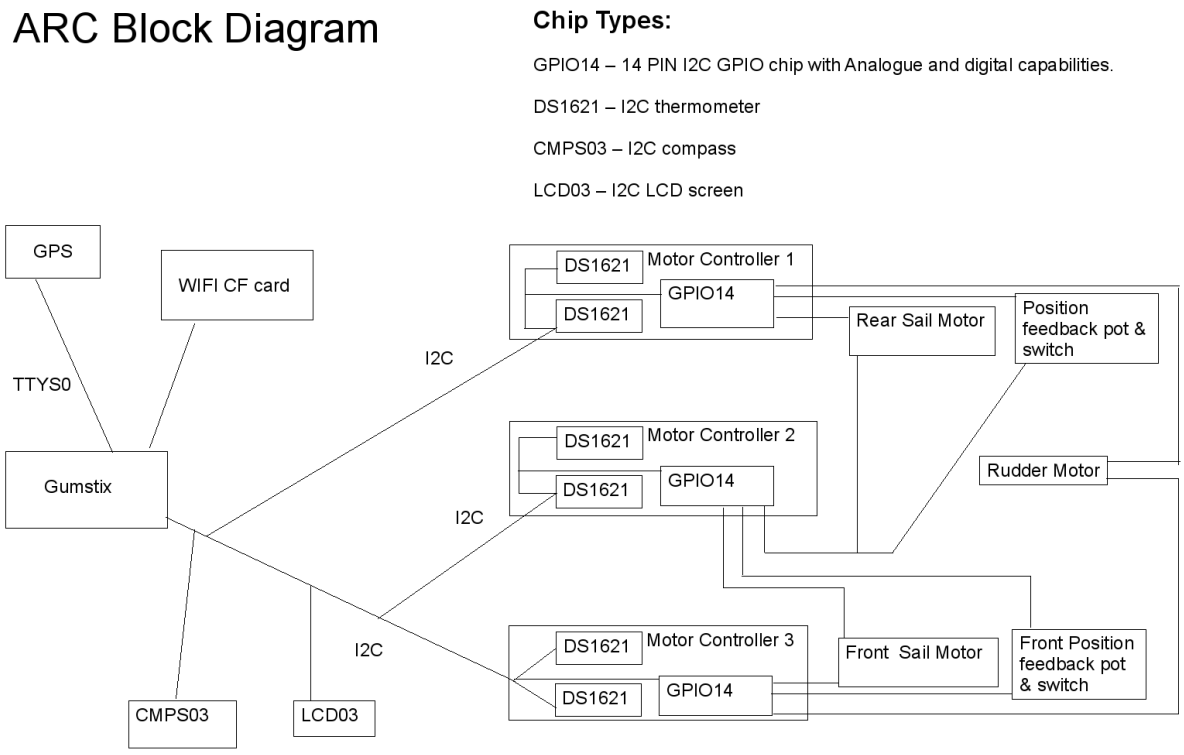

Figure B.1: A block diagram showing ARC's computers, sensors and actuators.

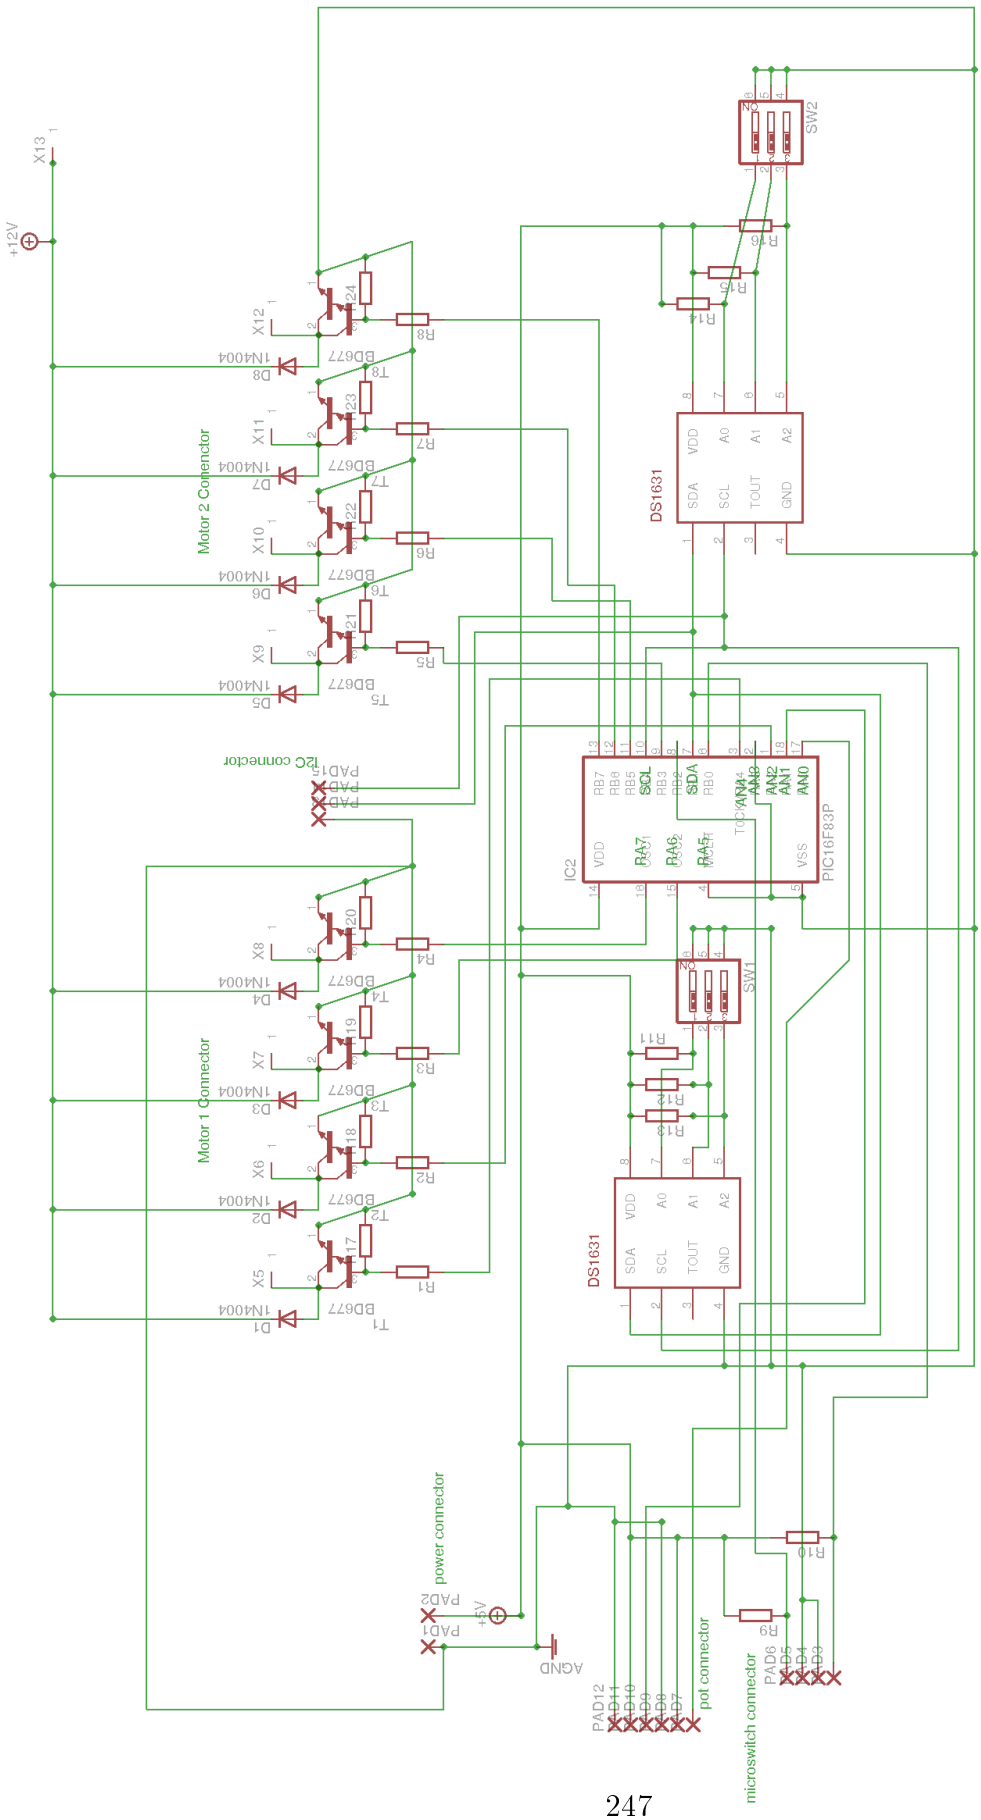

Figure B.2: A circuit diagram of ARC's redundant motor controller.

Appendix C

MOOP Diagrams

This appendix shows detailed annotated photographs of the MOOP1 sailing robot discussed in chapter 3 and used for experiments in chapter 7.

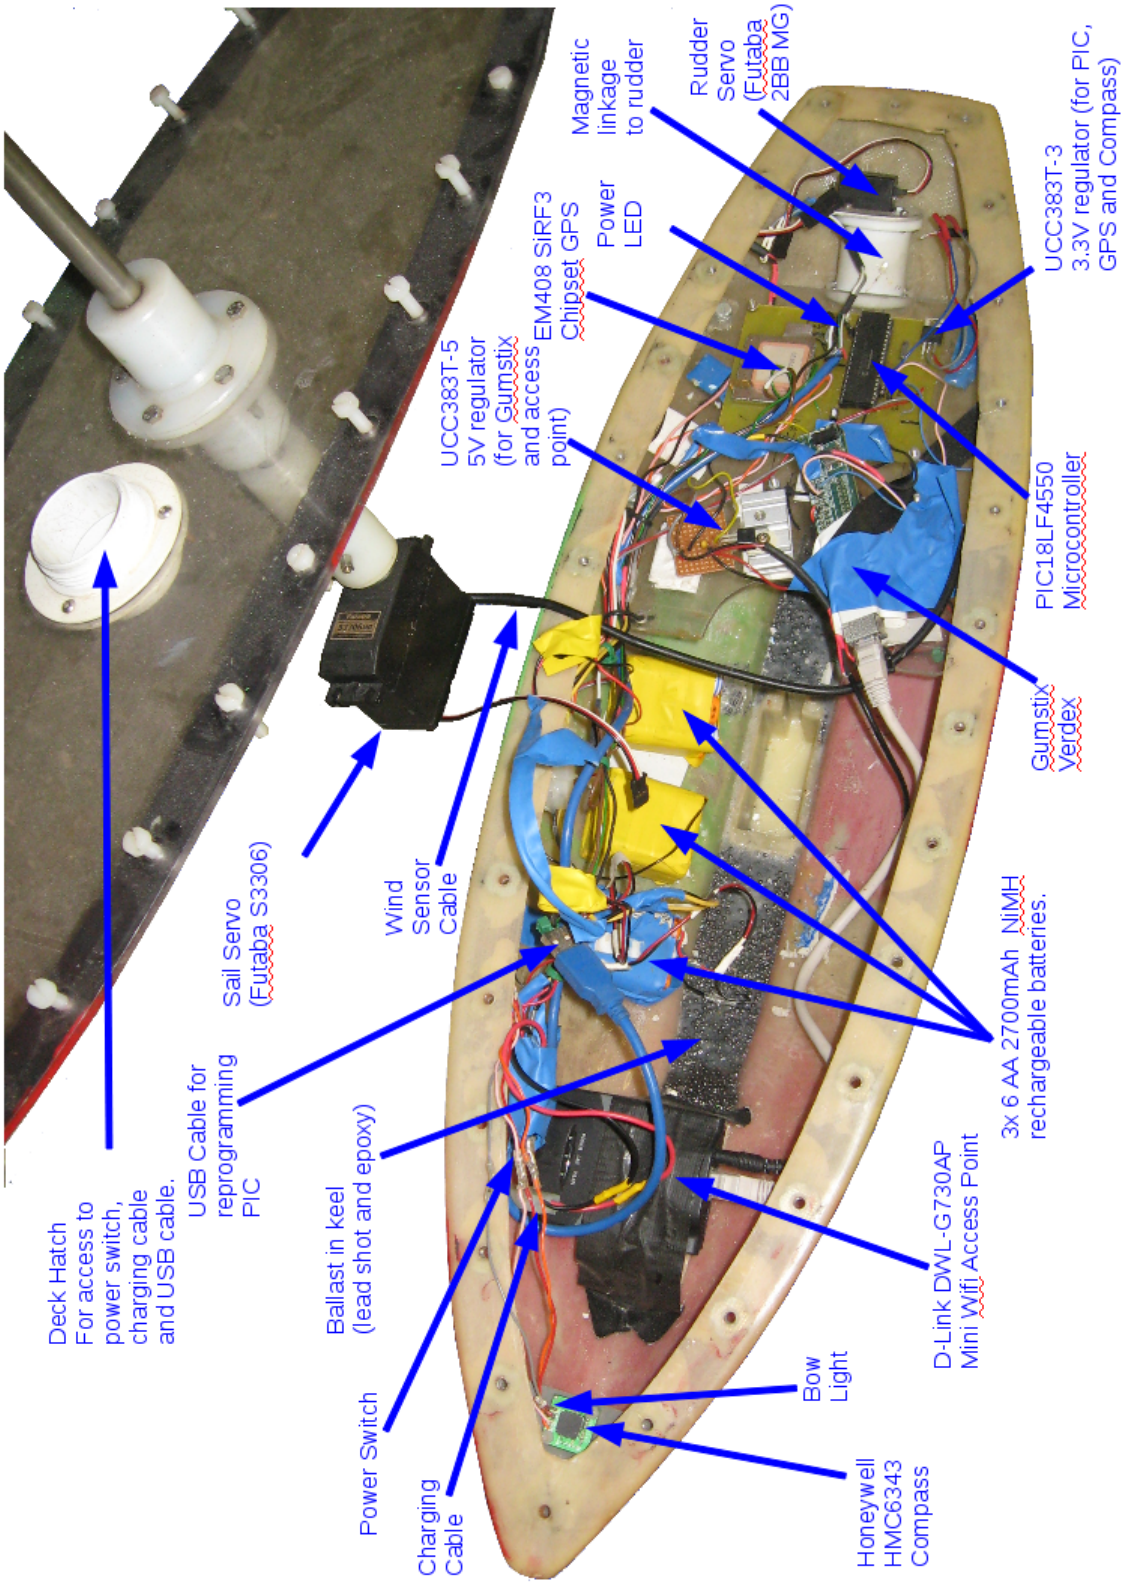

Figure C.1: A diagram showing the internal distribution of components in MOOP1.

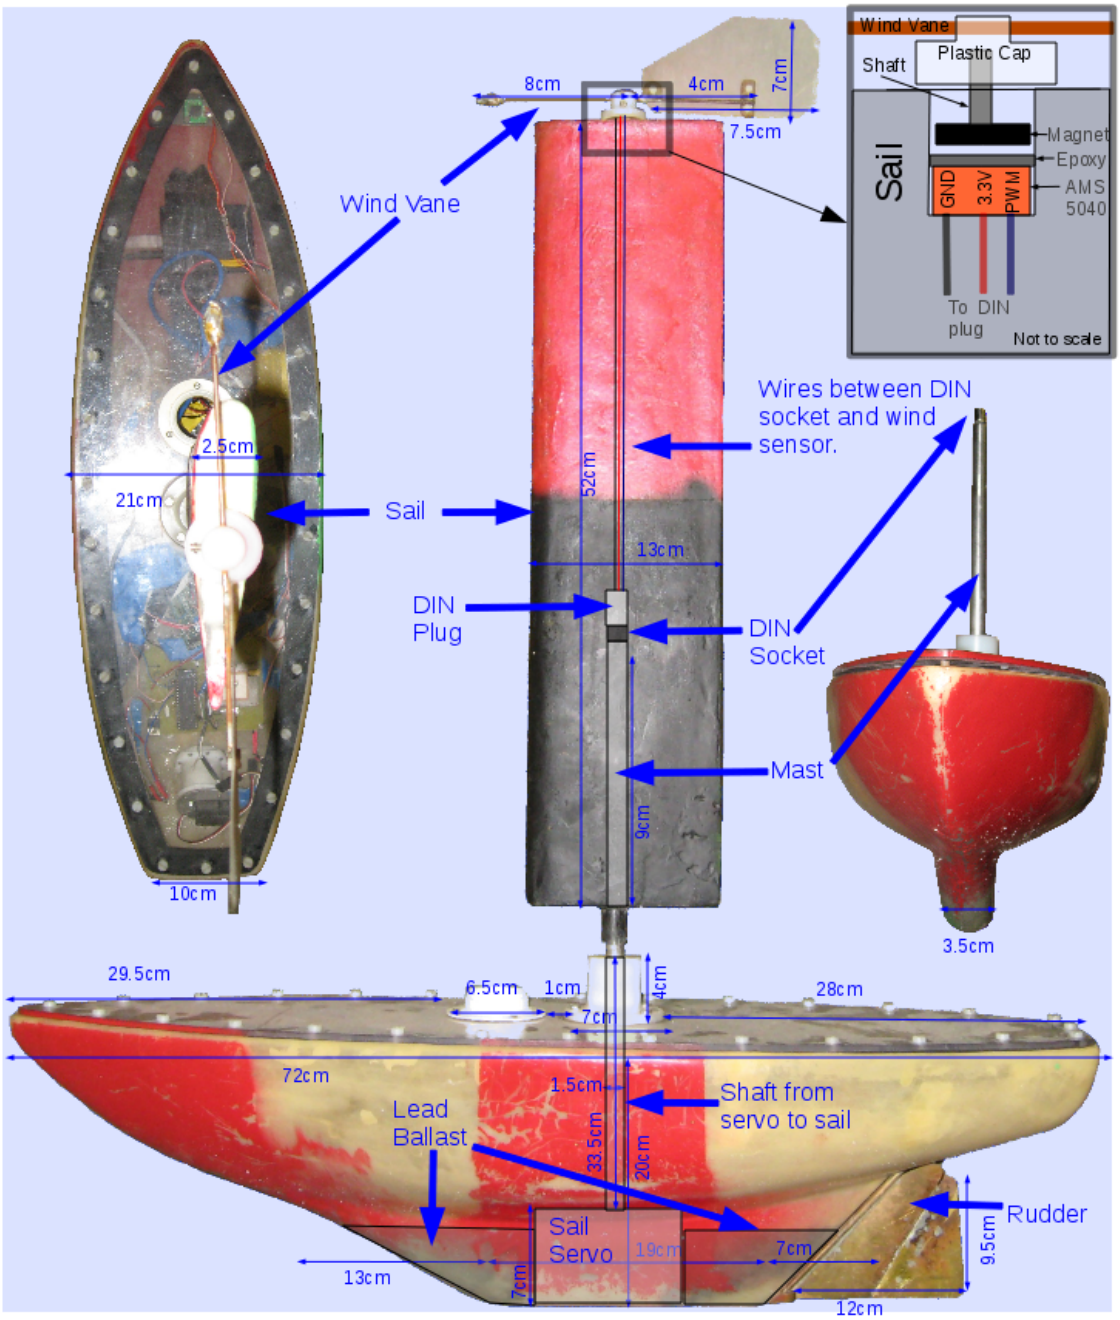

Figure C.2: A digram indicating the dimensions of MOOP1's hull and sail. Also shown is the wiring of the wind sensor.

Appendix D

Raw Data

D.1 Variable Hormone Experiment

D.1.1 Simulator Results

Hormone Concentration	0	0.25	0.5	0.75
Run 1	14.88	87.89	166.29	235.2
Run 2	20.03	90.39	161.87	215.96
Run 3	24.48	94.84	175.86	236.88
Run 4	21.14	90.25	166.24	238.63
Run 5	20.23	77.68	175.94	228.78

Table D.1: Distance covered for the simulated variable hormone experiment in section 6.4.2. All distances are in kilometres.

Hormone Concentration	0	0.25	0.5	0.75
Run 1	411.17	2477.85	4619.92	7000.63
Run 2	567.9	2555.45	4552.73	6454.8
Run 3	682.83	2674.23	4893.48	7907.12
Run 4	594.22	2552.28	4653.73	7642.9
Run 5	568.95	2193.53	4912.2	8000.27

Table D.2: Times taken during the simulated variable hormone experiment in section 6.4.2. All times are in minutes.

Hormone Concentration	0	0.25	0.5	0.75
Run 1	18	122	812	784
Run 2	28	689	859	980
Run 3	34	698	804	828
Run 4	34	683	820	973
Run 5	26	673	642	982

Table D.3: The number of sail movements during the simulated variable hormone experiment.

Hormone Concentration	0	0.25	0.5	0.75
Run 1	1770	1734	1491	1502
Run 2	1767	1541	1468	1430
Run 3	1764	1538	1499	1491
Run 4	1764	1543	1491	1439
Run 5	1767	1547	1553	1431

Table D.4: The number of rudder movements during the simulated variable hormone experiment.

D.1.2 Robot Results

Hormone Concentration	0	0.25	0.5	0.75
Run 1	2.53	2.39	2.6	4.72
Run 2	1.52	2.02	3.06	3.33
Run 3	1.5	0.99	1.91	2
Run 4	2.64	1.57	1.94	3.56
Run 5	1.56	1.55	1.84	3.38
Run 6	1.45	2.83	5.39	2.21

Table D.5: Distance covered for the robot variable hormone experiment in section 7.5. All distances are in kilometres.

Hormone Concentration	0	0.25	0.5	0.75
Run 1	119.95	146	154.48	245.35
Run 2	90.07	121.65	111.22	212.98
Run 3	95.82	169.5	178.82	161.2
Run 4	231.48	81.23	95.42	235.52
Run 5	103.45	92.27	131.6	177.88
Run 6	78.65	101.72	313.18	159

Table D.6: Times taken during the robot variable hormone experiment in section 7.5. All times are in minutes.

Hormone Concentration	0	0.25	0.5	0.75
Run 1	1656	1420	2412	2878
Run 2	1939	1631	2129	830
Run 3	1173	1479	2018	1142
Run 4	2009	2280	2936	2257
Run 5	1839	1643	2316	2419
Run 6	2315	1569	1795	1111

Table D.7: The number of sail movements during the robot variable hormone experiment in section 7.5.

Hormone Concentration	0	0.25	0.5	0.75
Run 1	10155	12110	10664	15286
Run 2	11414	10248	12234	14251
Run 3	11385	7705	10714	14048
Run 4	10267	8906	14293	13538
Run 5	7628	9221	12520	14874
Run 6	12664	13227	10259	15487

Table D.8: The number of rudder movements during the robot variable hormone experiment in section 7.5.

D.2 Fixed Hormone Experiment

D.2.1 Robot Data

Hormone Sensitivity	Battery level (Watt hours remaining)	Effective hormone concentration (sensitivity * battery level)	Completed yes/no	Course distance (kilometres)	Distance travelled (kilometres)	Time taken (seconds)	Total power used (Joules)	Course efficiency	Average speed (metres/hour)	Power usage Joules/hour	Power usage Joules/km
0	5	0.00	yes	4.615	5.644	9910	852.17	81.76%	2050	309.6	151.0
0.25	5	-0.23	yes	4.615	5.591	10034	317.83	82.53%	2006	114.0	56.8
0.5	5	-0.46	yes	4.615	5.548	10058	264.46	83.17%	1986	94.7	47.7
0.75	5	-0.69	no	4.615	8.199	196945	47.49	56.29%	149	0.9	5.8
1	5	-0.92	no	4.615	51.187	137269	448.65	9.02%	1342	11.8	8.8
0	25	0.00	yes	4.615	5.619	9829	931.35	82.12%	2058	341.1	165.8
0.25	25	-0.11	yes	4.615	5.618	9887	364.82	82.14%	2045	132.8	64.9
0.5	25	-0.23	yes	4.615	5.587	10135	309.92	82.59%	1984	110.0	55.5
0.75	25	-0.34	yes	4.615	5.568	10096	289.17	82.88%	1985	103.1	51.9
1	25	-0.46	yes	4.615	5.562	10060	270.02	82.96%	1990	96.6	48.5
0	55	0.00	yes	4.615	5.627	9863	806.63	82.01%	2053	294.4	143.3
0.25	55	0.06	yes	4.615	5.648	10030	1135.28	81.71%	2027	407.5	201
0.5	55	0.12	yes	4.615	5.615	9863	2424.33	82.18%	2049	884.9	431.7
0.75	55	0.18	yes	4.615	5.636	9883	2722.83	81.88%	2052	991.8	483.1
1	55	0.24	yes	4.165	5.66	9865	3441.78	81.53%	2065	1256.0	608.0

Table D.9: Complete results from the experiments in section 7.4

Appendix E

Actuator Power Consumption Estimates

Table E.1 contains power consumption estimates for the rudder and sail servos of a MOOP sailing robot. The current estimates were produced using a laboratory mains power supply to operate the servos and noting down the current while they were moving. This was performed under laboratory conditions and the servos were not under any load other than those caused by the friction of moving the rudder or sail. Therefore, this is likely to be a slight underestimate of the actual power consumption.

Actuator	Voltage	Current	Power	Time to move full range	Number of positions	Energy to move full range	Average Energy Used per position
Sail	7.6 Volts	1.35A	10.26W	1200 ms	11	12.312 Joules	1.1192 Joules
Rudder	7.6 Volts	0.6A	4.56W	920ms	11	4.1952 Joules	0.3814 Joules

Table E.1: Power consumption estimates for the sail and rudder of a MOOP sailing robot.

Appendix F

Paper's Published During this Thesis

This appendix lists papers published during the course of this thesis. My personal role and a brief summary of each paper is included in italics after the citation.

Papers Published to Date (April 2011)

- C. Sauzé and M. Neal. (2006) "An Autonomous Sailing Robot for Ocean Observation", in proceedings of TAROS 2006, Guildford, UK, Sept 4-6th 2006, pages 190-197. *I was the primary author for this paper which summarised work with AROO and ARC.*

- C. Sauzé and M. Neal. (2007) "Endocrine Inspired Modulation of Artificial Neural Networks for Mobile Robotics", Dynamics of Learning Behaviour and Neuromodulation Workshop, European Conference on Artifical Life 2007, Lisbon, Portugal, September 10th-14th 2007. *I was the primary author for this extended abstract which summarised the aims of a neuro-endocrine controller for a sailing robot.*

- C. Sauzé and M. Neal. (2008) "A biologically inspired approach to long term autonomy and survival in sailing robots.", Austrian Journal of Artificial Intelligence, Issue 2, Volume 27, Pages 11-17. *I was the primary author of this paper which detailed the work in section 4.2.1.2.*

- C. Sauzé and M. Neal. (2008) "Design Considerations for Sailing Robots Performing Long Term Autonomous Oceanography.", Austrian Journal of Artificial Intelligence, Issue 2, Volume 27, Pages 4-10. *I was the primary author of this paper which summarised the design issues and operating experinece of AROO, ARC and BeagleB.*

- M. Neal , J. C. Alves, C. Sauzé and B. Thomas (2009) "Technologies for Autonomous Sailing: Wings and Wind Sensors", in proceedings of the 2nd IRSC , Matosinhos,

Portugal, July 6th-12th 2009, pages 23-30. *I contributed a section on an averaging algorithm for wind sensors and put together the text from the other authors.*

- C. Sauzé and M. Neal (2010). "A raycast approach to colliision avoidance in sailing robots", In proceedings of the 3rd International Robotic Sailing Conference, June 7th-10th 2010, Kingston, Ontario, Canada, pages 26-33 *I was the primary author of this paper which covered work discussed in section 5.4.*

- C. Sauzé and M. Neal (2010), "A neuro-endocrine inspired approach to long term energy autonomy in sailing robots.", in proceedings of TAROS (Towards Autonomous Robotic Systems) 2010, August 31st-September 2nd 2010, pages 255-262. *I was the primary author of this paper which summarised much of the work in chapters 6 and 7.*

Papers Pending Publication

- C. Sauzé and M. Neal (2011), "Long Term Power Management in Sailing Robots.", in proceedings of IEEE Oceans 2011, June 6th-9th 2011, Santander, Spain, pending publication. *I am the primary author of this paper which repeats the work of chapters 6 and 7 on a different robotic platform.*

- C. Sauzé and M. Neal (2011), "MOOP: A miniature Sailing Robot Platform", in proceedings of the 4th International Robotic Sailing Conference, August 16th-20th 2011, Lübeck, Germany, pending review. *I am the primary author of this paper which describes the MOOP robot from section 4.1.5.*

- C. Sauzé and M. Neal (2011), "Simulating Sailing Robots", in proceedings of the 4th International Robotic Sailing Conference, August 16th-20th 2011, Lübeck, Germany, pending review. *I am the primary author of this paper which describes the gap between simulated and real sailing robot and is based upon chapter 8 and new work on in hardware simulation out carried out since the initial submission of this thesis.*

- C. Sauzé and M. Neal, "Pinta the sailing robot, the world's first autonomous shipwreck? (Working title)", Journal of Field Robotics, in preparation.

- C. Sauzé and M. Neal, "A Review of Artificial Neuro Endocrine Controllers for Robotic Systems (Working title)", Journal of Robotics Research, in preparation.