A computational study of the crayfish escape circuit

Steven C. Versteeg

Submitted in total fulfilment of the requirements
of the degree of Doctor of Philosophy

14th June 2005

Department of Computer Science and Software Engineering
University of Melbourne

(Produced on acid-free paper)

Abstract

The crayfish escape circuit is a well studied circuit in neurobiology. A massive and simultaneous stimulus, which is the signature of a predator's attack, triggers a rapid stereotyped tail-flip. The circuit is a case study of robust reflex behaviour in animals. Computational models of the electrical characteristics of the escape circuit, based on biological data, simulate the decision making of circuit. The models demonstrate that the rectifying junctions, which connect the sensory neurons to the command neuron, discriminate for highly coincidental stimuli. Furthermore, the sensory neurons are directly connected to form a lateral excitatory network, whereby stimulated neurons recruit non-stimulated neurons. The models show that the network sharpens the stimulus threshold required to trigger a tail-flip and enables the escape circuit to discriminate between different combinations of sensory inputs. Escape behaviour is also investigated in a broader context by simulating the interaction of escaping with other behaviours.

Declaration

This is to certify that

1. the thesis comprises only my original work towards the PhD except where indicated in the Preface,
2. due acknowledgement has been made in the text to all other material used,
3. the thesis is less than 100,000 words in length.

Steven Versteeg

Preface

The work presented in chapters 5, 6 and 7 was done in collaboration with the Donald H. Edwards and his laboratory at Georgia State University. The computational models in these chapters are based on biological data recorded by Dr Jens Herberholz and Dr Brian Antonsen.

All photos of crayfish neurons that are included in figures in this thesis (including those in figures 3.1, 5.1, 5.7 and 5.13) are the work and property of Dr Brian Antonsen.

The morphological model of LG presented in chapter 5 was created in collaboration with Dr Brian Antonsen and Jason Agran, at Georgia State University. The model was constructed from confocal stacks imaged by Dr Antonsen. Jason Agran was responsible for starting the project of constructing a morphological model of the lateral giant cell. The morphological model of LG presented in chapter 5 is a continuation of Jason Agran's work.

Acknowledgements

My sincere thanks to my supervisors Dr. Leon Sterling and Dr. Alan Blair for all their guidance, support and patience with me. Thank you to Leon for steering me towards this interesting topic, thank you for giving me the freedom to pursue it, thank you for all the guidance and encouragement in writing up my thesis. Thank you to Alan for working with me and guiding me for two years. Thank you for all your help in keeping me on track. Thank you for all the coffees and lunches.

I give a very warm thank you to Dr. Donald Edwards. Don was effectively my third supervisor. Much of this thesis is built on Don's ideas. Thank you to Don for inviting me to Georgia State University for a collaboration in simulating the crayfish escape circuit. Thank you for all the ideas, stimulating discussions and encouragement. Thank you for the feedback on my thesis. Thank you for the generous financial support that enabled me to live in the United States for 18 months.

Thank you to Dr. Brian Antonsen and Dr. Jens Herberholz, at Georgia State University, for all their hard work in uncovering the secrets of the escape circuit, on which the work in this thesis depends. Thank you Brian

and Jens for communicating with me your knowledge and understanding of how the cells in the LG function. Thank you Jens and Brian for all the discussions had and the beer consumed.

Thank you to Dr. Les Kitchen for his stimulating discussions.

Thank you to Gary Jolley-Rogers, at Zoology at Melbourne University, for introducing me to the crayfish escape circuit.

Thanks to Bernie, Inga, Kevin, Pete and Tim for keeping me sane while we procrastinated together. A special mention to Tim for donating a keg towards the completion of this thesis.

And a big thank you to my family for their patience and their support.

My PhD candidature was funded by the Australian Postgraduate Award (APA) scholarship.

The work on this thesis undertaken in the United States was supported by NINDS (NS 26457) grant to D. H. Edwards.

My 2000 visit to the United States was supported by a Postgraduate Overseas Research Experience Scholarship (PORES).

CONTENTS

CHAPTER 1

Introduction

The goal of Artificial Intelligence is to make artificial systems that are able to function with human level intelligence. Despite fifty years of research, Artificial Intelligence is a long way from achieving its goal. There is still a wide disparity between the intelligence we observe in nature and the artificial intelligence we can create. Biological systems display superior intelligence to the state of the art artificial systems, especially in regard to adaptiveness and robustness. Even 'simple' animals exhibit robust, adaptive behaviour that cannot be found in the most advanced artificial systems.

The observations above immediately suggest several questions: what do biological systems have that artificial ones do not? Which of these differences are relevant to our goals of building robust and adaptive artificial systems? In manufacturing, when a competitor has superior technology, a common approach is to reverse engineer the competitor's product. By understanding the advantages biological systems have over artificial ones we can try to reverse engineer animal systems. This thesis is in the spirit of 'reverse

engineering' biological systems.

Possibly the most advanced brain, and the most relevant to Artificial Intelligence is the human brain. If we had a complete understanding of the human brain we could surely manufacture an artificial one. Unfortunately the human brain is of a complexity that causes this to be a rather difficult task. The human central nervous system (CNS) is estimated to have 10^{11} neurons and at least 10^{14} synapses (Bear et al., 1996). These numbers give an impression that underestimates the complexity. Individual cells are themselves highly complex and cells have many different mechanisms for interacting with each other. With current techniques in Neuroscience we are only able to obtain a superficial understanding of how our brains think.

By comparison, invertebrate brains are simpler and better understood. The number of neurons in invertebrate brains is many orders of magnitude less. For example, the honeybee is estimated to have one million neurons and the flatworm, *Caenorhabditis elegans*, has only 302 or 381 (male or female) neurons. In invertebrates neuroscientists are able to identify individual cells. Complete circuits controlling identifiable behaviours have been mapped out.

Invertebrate brains are interesting in themselves. Invertebrate brains produce many interesting behaviours. Such behaviours are superior to the state of the art Artificial Intelligence systems in many aspects of cognition.

Fortunately, despite vastly differing in complexity, there are many similarities between invertebrate and vertebrate brains. The neurons that comprise these respective types of brains are very similar and function according to the same principles. The types of channels, receptors and neurotransmitters found in vertebrates are (for the most part) also found in invertebrates. Understanding invertebrate brains can yield many principles that are applicable also to vertebrates.

Current knowledge of biological brains is very far from a complete understanding. It is difficult to understate that fact. What we do not know is far greater than what we know. This is true even for the best understood systems. There are some specific neuronal circuits whose structure and cognition are relatively well understood. Even though it is impossible to completely reverse engineer a biological brain, it is possible to study parts of it in isolation. Understanding these specific circuits gives us concrete examples of how biological brains are organised and make decisions. An accurate understanding of specific neural circuits is a precursor to a general theory of how brains function.

Previous studies of modelling invertebrate systems have yielded success. (Beer, 1990) simulated the neuroethology of cockroach walking to demonstrate a robust and adaptive locomotive controller. Beer's model was used to implement walking in a robotic system (Beer et al., 1992; Espenschied et al., 1993; Espenschied et al., 1996).

This thesis is directed to investigate an example of a reflex behaviour. Understanding reflex behaviours will allow us to implement robust decision making in Artificial Intelligence. Many reflex behaviours in different animals are well studied in neurobiology. Reflex behaviours are relatively simple and interesting. When an animal acts reflexively it has made a decision. Reflex behaviour is a specialised form of decision making. The decision processes for reflexes are very robust. Reflex behaviour provides an example of robust decision making in nature. Arguably to understand reflexes is to understand how biological systems get their robustness.

Two approaches were taken to studying reflex behaviour. The first is at a high level that considers the behaviour but ignores the underlying neural mechanisms. The goal was to consider the role of escape behaviour in a wider context of competing behaviours. Chapter 2 describes an animat

in a simulated world that integrates an escape behaviour with the task of collecting food to survive. The work was presented at the Australian Joint Conference on Artificial Intelligence 2001 (Versteeg, 2001). We examine the combination of escape behaviour with other roles in an artificial system in a very simple case. We show that by having a robust escape response the agent is able to be bolder in the pursuit of primary goals.

The animat simulation treats the decision of whether to escape at a very high level. The simulation demonstrated the usefulness of escape behaviour but no really interesting complexities emerged. Furthermore, perhaps as a consequence of the logic based decision making being used, the results showed rigidity and brittleness, which are common attributes of artificial systems but absent from biological systems. This directed us to our second approach. In order to truly understand how animals make decisions it was decided to investigate at the neurophysiological level.

A good way of gaining an understanding of reflex behaviour is by simulating the neural circuit that controls it. In general, simulating neural circuits enables us to test theories of how things work. Neural circuits are extremely complex. Biologists are able to identify the specific components of the circuit but are only able to get a intuitive and qualitative understanding of how the system functions as a whole. Computer simulations enable us to test these ideas. If we can reproduce an interesting behaviour in simulation then we have understood it.

There is a positive feedback loop between experimental biology and computer simulations. Neural simulations depend on experimental results from biology. The simulation results provide feedback for biologists. Simulations verify theories of how a neural circuit functions. Simulations make predictions about the broader functions of the circuit.

The crayfish escape circuit is a particularly well studied invertebrate cir-

cuit that commands a reflex behaviour. The circuit controls a stereotyped short latency tail-flip that propels the animal from danger when unexpectedly attacked. The neurons in the circuit have been identified and well characterised.

The crayfish escape response is a rare example of a high level behaviour where the controlling neural circuitry is well described. The crayfish escape circuit is a system that is responsible for activating and controlling a specific behaviour. The neurons and connections involved in this decision are identified. The circuit is simple enough to be able simulate all of its components. The escape circuit is responsible for making a critical life threatening decision. It provides a concrete example of how a biological brain robustly makes a decision.

The escape circuit is also a case study of other interesting capabilities of biological brains. The escape circuit is involved in social behaviour. Animals flip their tail during agnostic (fighting) encounters to demonstrate their strength to assert their dominance. The crayfish escape circuit undergoes learning. The circuit is not static. The response is regulated by the conditions in the environment and the crayfish' internal state. The circuit is able to habituate to be sensitised or desensitised to frequently occurring stimuli.

The major component of the thesis is a model of the crayfish escape circuit. By building a quantitative model of the crayfish escape circuit, one that accurately reproduces its interesting characteristics, we can gain insight into qualities that make biological brains interesting.

A key underlying issue is the level of detail that should be included in the model. Ideally our model should be as simple as possible while still retaining the interesting characteristics that explains the decision making process in crayfish escape. We have chosen to model the electrical properties of the neurons in the escape circuit. The model includes neurophysiological

and neuroanatomical information about the cells. The model is far from complete. The model does not explain all the known phenomenon of the circuits behaviour nor does it include all the biological information. The work presented provides a foundation for a complete model of the escape circuit behaviour.

Chapters 3 and 4 give background. Chapter 5 describes the model. The model was used to perform a series of simulations to identify what aspects of the escape circuit's organisation assists its function.

Chapter 6 investigates how the lateral giant is able to detect a group of simultaneous excitation, which is the signature of a predator's attack.

Chapter 7 simulates a network of sensory neurons that are connected together. Stimulated neurons are able to recruit non-stimulated neurons and thereby bring the lateral giant to spiking threshold. The work in this chapter was presented at Neuroscience 2003 (Versteeg et al., 2003).

Escape behaviours are useful to study as an example of decision making. Escape and reflex behaviours also have a direct application to the design of robots and artificial systems. The fact that every animal has a startle response of some description provides compelling evidence that they have a critical role in the function of biological nervous systems. When an animal finds itself unexpectedly in a critical situation it must rely on fast, simple and robust reflexes to manœuvre out of trouble. Robots when they are placed in a real environment also have to deal with unexpected events. Studying escape reflexes in nature is instructive as to how to design for such situations in artificial systems.

Escape behaviours have obvious applications in RobocupRescue and Robocup. RobocupRescue provides direct parallels to the natural world. Robots will find themselves in physical danger. In Robocup, robots are also frequently

faced with unexpected events. Reflex behaviours may provide a suitable mechanism for desperately preventing goals.

The introduction is concluded by noting the great gulf that existsx between neuroscience and artificial intelligence. Neuroscience is the study of brains in animals. Artificial intelligence aims to engineer brains in robots and artificial systems. For the most part, the two communities are unaware of each others work. The two communities have a lot to learn from each other. This thesis is intended to contribute to crossing the gulf.

If you are interested to know how an animat can escape from a predator while collecting enough food and whether it will survive will you must read the next chapter.

CHAPTER 2

Testing escape behaviour in a simulated environment

This chapter is concerned with testing an animat with escape behaviour in a simulated environment. We consider how an escape response can be used by an agent in an environment and be integrated with its other behaviours. Sections 2.1-2.6 describes simulations that evaluate some specific questions in regard to how escape behaviour is used within an environment. These simulations are runnable online as Java applets at:

`http://www.cs.mu.oz.au/~scv/botsim/`

In section 2.7 we describe a framework for creating simulations of multiple agents interacting with each other in an environment.

2.1 Introduction

In nature, animals going about their daily routine need to avoid predators in order to survive. Many animals have evolved some kind of startle response (see section 3.5), which enables them to escape from dangerous situations. The term 'animat', for an artificial animal, either simulated on a computer or a physical robot, which must adapt to increasingly more challenging environments (Cliff et al., 1994). As robots are required to operate in more hostile environments, mechanisms analogous to startle responses may be critical to building robust systems. Animats may also need to escape from sources of danger, such as malevolent passers by and curious children (Miller and Cliff, 1994). The key is not only to have effective escape mechanisms but also to integrate escape with the animat's other activities. An animal that spends all its time running away without stopping to eat and find food will not survive very long.

This chapter presents some preliminary work exploring (1) how some reactive evasive behaviours can be added to an agent operating in a hostile environment, and (2) how evasive measures can be integrated with the agent's other activities.

(Miller and Cliff, 1994) argued for studying pursuit and evasion tactics as a useful problem domain for artificial intelligence. Pursuit-Evasion tactics are amongst the common, most challenging contests in animals. Their main reasons for studying pursuit and evasion included:

- Pursuit and evasion require highly robust forms of adaptive behaviour. Systems that are slow, brittle or easily confused do not do well. Traditional AI may not be suited. Newer reactive, behaviour based approaches may be more suitable.
- Studying pursuit and evasion may illuminate behavioural arms races

in general. Pursuit and evasion scenarios may be the simplest.

- Pursuit and evasion is well studied in 3 fields: behavioral biology, neuroethology and game theory.
- Pursuit and evasion is the simplest situations that can favour protan (adaptively unpredictable.)
- Understanding P-E may give better understanding of evolution and have other scientific implications.

Miller and Cliff used co-evolution to evolve predators and prey.

There have been many previous studies of evasion in isolation. A couple of examples of evolving optimal strategies in scenarios with fixed predator behaviour include (Koza, 1991) and (Grefenstette et al., 1990). Miller and Cliff (Miller and Cliff, 1994) co-evolved pursuer and evader tactics using noisy neural network controllers. The pursuer-evader problem has been reformulated as a one-dimensional, time-series prediction game. (Ficici and Pollack, 1998) There has been exploration of the evolution of evasion strategies when the game is made slightly asymmetric between the pursuer and evader (Wahde and Nordahl, 1998).

(Edwards, 1991) created a pursuit-evasion scenario for testing a hypothesis of how a crayfish may choose between different behavioural modes. Central to Edwards' model are mutually inhibiting command systems. Edwards created a scenario with a crayfish and a predator. The crayfish needs to eat food to survive and may retreat to a shelter for safety from the predator. The crayfish was given seven different behavioural modes: foraging, eating, hiding, defending, retreating, tail-flipping and swimming. Each behavioural mode has a separate command system. The command systems are excited by sensors in the environment. The command systems also mutually inhibit for each other. The command system with the greatest excitation, after

the inhibitory effects, has its behaviour executed. Edwards' simulation to showed that a mutually inhibiting command systems can produce robust and stable behaviour. The simulation is viewable as a Java applet at: `http://www.cs.mu.oz.au/~scv/sim/simcray.html`.

We create a simple scenario where an animat has a primary task, in this case collecting food. A predator is placed in the environment to make it hostile. Escape capabilities are incrementally added to the agent. As we build up the prey's evasive capabilities, we observe the impact on the agent's survival and the change to its approach to its primary task.

2.2 The Scenario

The set-up of the predator-prey scenario is similar to the crayfish-predator scenario described by (Edwards, 1991). There is one predator, one prey, a shelter and some food.

The prey has the task of collecting enough food to survive while being hunted by a predator. The predator has a greater maximum velocity (V_{pred}) than the prey (V_{prey}). Furthermore the predator's seeing distance (R_{pred}) is also greater than the prey's (R_{prey}). The prey has superior acceleration over the predator and may choose to hide in a shelter where it is safe from the predator.

The environment is a continuous two-dimensional plane of $n \times m$ units. It has wraparound edges (this is to avoid the artifact of the prey being trapped in a corner.) The world contains pieces of food located at random locations. New pieces of food are added and old ones are removed at random time intervals. Also situated in the environment is a shelter. When the prey is in the shelter the predator is unable to see it and unable to kill it.

The predator and prey are able to move within the world and are able to make some limited interactions with the other entities in the world. The prey is able to eat a piece of food if it is close enough. The predator is able to kill the prey if it is close enough. The simulation is updated in discrete time steps. At each time step the predator and prey are queried by the simulation engine about their intended movements and other actions they want to take. All updates to positions and interactions are executed simultaneously. If the predator or prey elect to change their velocity, they are not able to do it instantly but instead accelerate to new velocities. It may take several time steps for the predator or prey to reach their new velocity. The predator and prey each have maximum rates of acceleration.

The predator follows a very simple behaviour pattern which is fixed for all the experiments. The predator roams around at half-speed ($\frac{v_{\text{pred}}}{2}$) travelling in a straight line. At random time intervals it changes to a new random direction. The predator is continuously looking for the prey. If at any point the predator spots the prey it will immediately change its direction to head directly toward it and accelerate to its maximum velocity (v_{pred}). When the predator gets within a distance of k_{kill} units of the prey, the predator kills the prey and then eats its. The predator is present somewhere in the world for the entire duration of the simulation.

The prey has an internal energy level. To avoid starvation it must maintain its energy level above zero. At each time interval the prey consumes an amount of energy determined by equation 2.1.

$$\Delta E = -(B + Av^2) \tag{2.1}$$

The base energy consumption (determined by B) forces the prey to occasionally go and collect food. The other term is dependent on the square of the prey's velocity to penalise travelling at high speeds. The prey replen-

ishes its energy level by eating food. The prey needs to move close to a piece of food before it can eat it. When a piece of food is eaten the food is removed and the prey gains k units of energy. If E drops below zero the prey starves to death.

The simulation was implemented using the botsim framework (which is described in detail in section 2.7.)

2.3 Architecture of the prey

The prey uses a layered architecture which we build up incrementally (Pfeifer and Scheier, 1999; Brooks, 1986). The prey operates in distinct behavioural modes. It monitors the level of some simple stimuli, such as 'hunger', to determine which behavioural mode to operate in. At the most basic stage the prey ignores the predator completely and is solely focused on collecting enough food to avoid starving. At each stage, another behavioural mode or stimulus is added to the prey's repertoire to assist it in avoiding the predator. New behaviours are able to subsume or suppress behaviours introduced at previous levels.

2.3.1 The Hiding Bot

First we consider a very simple bot. The hiding bot looks for food when it is hungry and hides in the shelter when it is not. It does not detect an approaching predator. The hiding bot's survival strategy is basically to spend as much time in the shelter as possible, while avoiding starvation.

The hiding bot uses one stimulus with which to make its decisions: the internal energy level (E). It uses this information to choose between one of two behavioural modes in which it can operate:

Forage The prey searches for food and eats it. The prey follows the odour gradient emitted by the food until it reaches an item of food which it then eats. The prey travels at half speed to conserve energy.

Hide The prey moves to the shelter and hides there when it arrives. The prey travels at half speed to conserve energy.

Figure 2.1a shows the hiding bot's control architecture.

2.3.2 The Running Away Bot

The running away bot actively keeps an eye out for the predator while it is outside of the shelter. The bot operates in the same way as the hiding bot in that it ventures out of the shelter for food when it is hungry, but if while the bot is out of the shelter it detects the predator it will scurry back to the shelter for safety.

The running away bot may use the behavioural modes of the hiding bot: hide and forage, and in addition may use: *run away* mode.

Run Away The prey runs to the shelter at maximum velocity.

To determine when to run away, the bot uses the stimulus *predator fear*. It is dependent on the distance (d_{pred}) between the prey and the predator as shown in equation 2.2.

$$P = A_P e^{-\frac{d_{pred}}{L}} \tag{2.2}$$

The algorithm uses to control the running away bot is shown in figure 2.1b.

2.3.3 The Memory Bot

The hiding bot and the running away bot are stateless. The memory bot explores what advantage a simple piece of state information can give an agent. (Dorigo and Colombetti, 1994) demonstrated that adding even a few memory bits can give a significant improvement to the performance of a reactive object tracking system.

The memory bot remembers when it last saw the predator. This affects the stimulus *memory fear*. When the predator is seen memory fear instantly rises to the maximum. It decays exponentially with time (t_{pred}) from when the predator was last seen as shown in equation 2.3

$$M = A_M e^{-\frac{t_{pred}}{\tau}} \tag{2.3}$$

The memory bot operates in the same way as the running away bot but if its memory fear is still above a threshold T_M then it will continue to hide. Figure 2.1c shows the control algorithm.

2.3.4 The Dodging Bot

The dodging bot has a reflex action with which it attempts to evade the predator if it gets too close.

If the predator fear stimulus crosses a threshold T_D then the prey will go into *dodge* mode:

Dodge The prey immediately changes to a new direction which is orthogonal to the direction of the approaching predator. The prey very rapidly accelerates to maximum velocity. (It is in effect a jump to the side.)

This manœuvre is somewhat analogous to the escape reflex in the crayfish

(which is described in detail in chapter 3) in that (1) a simple test is used to determine when to activate the response, and (2) to be effective a special piece of hardware is needed. The crayfish makes use of special flexors in the abdomen which are only used in an escape tail-flip; this gives it extremely rapid acceleration. The dodging bot uses a higher amount of acceleration in the dodge manœuvre than it normally has available to it. The dodge manœuvre is also somewhat similar to the zig-zagging used by the evasive agents in the (Miller and Cliff, 1994) simulations.

The dodge behaviour takes precedence over all other behaviours. Figure 2.1d shows the control layer diagram of the dodging bot.

2.4 Evaluation methods

The four bots were placed in the environment to see how well they survived.

To evaluate how the prey modifies its approach to its primary task (collecting food) as its evasive capabilities increase, we measured the optimal hunger threshold (T_H) for each bot. This was done performing a linear search over the entire range of valid hunger thresholds. A total of 101 sample points were considered. For each hunger threshold one thousand simulations were performed. This was necessary to obtain reasonably precise statistical measures of the prey's expected life. One thousand trials was within the capability of the computer resources available and produced quite smooth curves of the recorded mean statistics versus the varied parameters.

For each configuration the following statistics were recorded:

1. The mean number of time steps that the bot survived.

2. The median number of time steps that the bot survived.

3. The standard deviation of the survival time.

4. The frequencies of the alternative survival outcomes. For each simulation the bot may have died because (a) it ran out of energy (*starved*), or (b) it was caught by the predator (*killed*), or (c) it reached the end of the simulation still alive (*survived*).

The conditions of the environment, the prey's energy costs and the constraints on the prey's sensory capabilities were kept constant for all bot configuration trials. Refer to appendix A for the values of these constants which were used in the simulations. Prior to the bot configuration trials being executed, some different environmental settings were experimented with by hand. Changing the environmental constants modifies the balance between the predator and the prey. The parameters in equation 2.1 reduce or increase the frequency at which the prey needs to forage for food. Modifying the predator's seeing range and maximum velocity significantly tilts the balance. If the seeing range is too great or the chasing velocity is too high, the predator becomes too efficient at hunting the prey to the extent that the prey always dies. The final values adopted for all these settings gave a fair balance; the predator has a reasonable chance of catching the prey and a well-adapted prey has a reasonable chance of survival.

2.5 Results

2.5.1 The Hiding Bot

Since the hiding bot is unable to detect the predator its behaviour is governed only by the hunger threshold, which determines when it will hide in the shelter and when it will venture out to look for food. Figure 2.2 shows how the survival time and cause of death vary as the hunger threshold (T_H)

changes. If the hunger threshold is very low, then the bot will wait until it is almost completely out of energy before venturing out to look for food. There is a very high chance that the bot will starve to death before finding the food. If the hunger threshold is very high, then the bot will spend most of its time out of the shelter looking for food and there is a greater chance of it being eaten by the predator. The bots that do the best are the ones that stay in the shelter as long as possible while still having a good chance of finding food in time before starving.

Figure 2.2 shows how the survival rate of the hiding bot varies by adjusting the hunger threshold. A curious feature of the graph in figure 2.2 is that as the hunger threshold goes from 0 to about 2, the median declines slightly but the mean rises. The explanation of this phenomenon is at very low hunger thresholds when the bot ventures out it is more likely that the bot will starve than that it will find food. While most of the bots die of this cause the median will stay low. However the lucky few who find food live significantly longer so push up the mean.

2.5.2 The Running Bot

The running bot survives significantly better than the hiding bot as can be seen in figure 2.3.

After seeing the predator the running bot is generally effective at running back to the shelter in time before the predator overhauls it. Provided that the hunger threshold is set at a reasonable level the death of the running bot generally occurs under one of three circumstances: (1) the bot is too far away from the shelter when the predator sees it and is therefore rundown by the predator in pursuit, (2) the predator approaches so that it is between the bot and the shelter; in this case the bot is cut off while running to safety, and (3) the bot is chased before it reaches any food so it is even shorter of energy

when it gets back to the shelter; this causes the bot to die of exhaustion rather than being killed by the predator directly. Furthermore, the running bot suffers from having no memory and a shorter seeing distance than the predator. If the bot is chased to the shelter and is still hungry it will venture out as soon as it can no longer see the predator. But it is very likely that the bot will be still within the predator's seeing distance and therefore is immediately chased again.

The optimal threshold at which escape is triggered, T_P (see equation 2.2), was found to be quite low. If the threshold was too high, the prey would have an insufficient head start on the predator to safely get back to the shelter. On the other hand, if the threshold was extremely low, then the prey would continuously hide in the shelter (and starve to death.)

In the context of the limited visibility range given to the prey, its best strategy was to run as soon as it caught sight of the predator. This is because by the time the prey can detect the predator, the predator can expect to reach the prey within about four time steps if no evasive action is taken. We conducted some experiments where the prey's visibility range was made to be unlimited to determine what the optimal value of T_P should be in that context. Under those conditions the optimal value of T_P was about 0.05 which equates to the predator being about ten time steps away. This is also the distance at which the predator can spot the prey.

2.5.3 The Memory Bot

The simple state information provided by the memory fear stimulus gives the memory bot a big boost in survival chances. The memory bot is able to run back to the shelter and stay there long enough until the predator has passed. Since the predator moves randomly, the longer it stays in the shelter the greater the chance that the predator will have passed. Balanced

against this is that staying longer in the shelter reduces the chance of having enough energy to reach new food before it starves.

The bot will stay in the shelter until the memory fear (see equation 2.3) drops below the threshold T_M. The effect of the memory threshold was evaluated using the method described in section 2.4 and was repeated for a number of different hunger threshold settings. If the bot has a very high threshold it will stay in the shelter only a short time. If the bot has a very low threshold it will stay in the shelter a very long time. Figure 2.4 shows how the survival rate is affected by the value chosen for the memory threshold. Note that the survival rate is fairly even for thresholds between 0.2 and 0.6, reflecting that the tradeoff between waiting out the predator and risk of starvation is fairly evenly balanced.

Figure 2.5 shows how the survival rate is affected by the hunger threshold for the memory bot. The optimal hunger threshold is greater for the memory bot than for the running bot. This can be explained by two factors: (1) the memory bot has got a better chance of surviving an encounter with the predator so it can risk spending more time out of the shelter and (2) if the memory bot collects more energy then it can spend more time in the shelter after it has been chased by a predator, giving it higher chance of the predator having left.

The memory bot is still killed if it is too far away from the shelter when pursuit begins or if the predator is in between it and the shelter. It is fairly successful in waiting in the shelter until the predator has passed. It sometimes starves while waiting for the predator to pass and sometimes is unlucky and finds the predator still waiting outside after it has waited.

2.5.4 The Dodging Bot

The dodge manœuvre of the dodging bot allows it to survive in one of the situations where the memory bot is frequently killed. If the predator attacks the prey coming from the direction of the shelter, the dodging bot is able to evasively side-step the predator. The predator coming at full speed is unable to adjust its direction in time to catch the prey. The dodge manœuvre is less successful in evading the predator chasing from behind. However as with the running bot and the memory bot, in most cases when the prey is being chased from behind it will be able to reach the shelter in time. It is usually only when the prey is very far from the shelter that it is caught in pursuit.

Figure 2.6 shows how the survival rate of the dodging bot varies with the hunger threshold. The dodging bot has a higher optimal hunger threshold again compared to the memory bot.

2.6 Discussion

A trend that emerges is the prey spends more time outside of the shelter foraging for food as it gets better at avoiding the predator. This phenomenon can be explained by a combination of two factors: (1) the evasion manœuvres require a lot of energy, if the prey accumulates greater energy reserves, it is better prepared for an encounter with a predator, as it is able to deploy more evasive actions, and (2) because the prey is more likely to survive an encounter with the predator it can afford to spend more time outside the shelter foraging for food, thereby reducing its risk of starvation.

To optimise its survival time the prey needs to balance its primary task, collecting food, with taking evasive action. If the predator were nonexistent the optimal survival strategy for the prey would be to spend all its time

collecting food to nullify any risk of starvation. In the presence of a predator, collecting food becomes a risky task. This causes the bots with poor predator avoidance to wait until they are really hungry before they will venture out to look for food. As the prey is equipped with more evasive capabilities the risk in collecting food diminishes. This increased confidence allows the prey to act more like it would if there was no predator. As more evasive capabilities are added, the prey's behaviour pattern (when it is not taking evasive action) approaches what it would be if the predator did not exist.

A more general implication of this result is that having separate subsystems to deal with dangerous situations allows an agent to be less obstructed in undertaking its primary activities. Robots presumably have a set of primary tasks which they have to perform. However some robots while undertaking their work robots may periodically have to face obstructions or even dangers. A hypothetical example is a rescue robot sent into a burning building that may have to dodge falling debris. Having separate subsystems to deal with obstructions may allow robots to be minimally affected in the way they achieve their primary tasks. Animals have dedicated neural circuits to deal with unexpected hazardous situations (Bennett, 1984; Hoy, 1993).

One of the biggest improvements given to the prey in this simulation is the addition of some simple memory information. In this memory model the time that the predator was last seen is used implicitly as a predictor of the predator's present proximity.

The fixed hierarchical model used for the agents seems suboptimal for this scenario. One kind of behaviour should not always have precedence over another. For example because hiding in a shelter after seeing a predator has a higher precedence than foraging, in some cases the bot would starve to death in the shelter. A better model would be more flexible: wait longer in the shelter if energy reserves are relatively high, and shorter if energy

reserves are low. Different actions need different precedence at different times.

There is biological evidence that in animals the precedence of actions is much more flexible. (Prescott et al., 1999) review biological findings about parts of the vertebrate brain and argue that the basal ganglia acts as a central decision making point for arbitrating between conflicting actions. They argue that a similar specialised switching mechanisms might be employed in layered robot architectures (such as (Brooks, 1986)) to provide more flexible action selection.

In the crayfish the giant command neurons responsible for triggering the escape response are modulated by other parts of the nervous system (Vu and Krasne, 1993a; Glanzman and Krasne, 1983; Edwards et al., 1991). The trigger threshold adjusts according to various circumstances, such as during feeding and restraint (Wine, 1984) and also adjusts according to longer term conditions such as the mating cycle and social dominance (Yeh et al., 1997).

Edwards (Edwards, 1991) proposes a model for behavioural choice in crayfish that uses mutual inhibition amongst the neural command centres. In Edwards' model there is one command neuron for seven different behavioural modes. Each neuron receives excitatory stimuli from sensors. Each neuron is able to inhibit other command neurons and also receives inhibitory signals from the excited command neurons. After summing the excitations and the inhibitions, the command neuron with the greatest excitation wins. Edwards' model is able to give actions different precedence at different times. An attempt was made to write a bot based on Edwards' architecture for the scenario described in this chapter. Preliminary results indicate that in this scenario it performs slightly better, but the results are inconclusive.

The scenario examined in this chapter is very specific. Previous pursuit-

evasion experiments (Cliff and Miller, 1996) have shown that effective evasion strategies are often very sensitive to the parameters of the environment. Future work may consider what kind of escape measures work best when faced with different kinds and variable sources of danger. The prey currently uses 'magic' perception to get the position of the predator. This is unrealistic. In future simulations the prey should have to infer the presence of a predator from noisy sensors.

The work presented in this chapter is a preliminary step in exploring the role of integrating evasive actions in the context of doing other activities. As robots move out of the laboratory into more hostile environments, handling evasive actions may prove an important component of the robot architecture.

2.7 The botsim framework

Botsim is a framework for creating high level simulations of robots or animal like agents. The design goals of botsim are:

- to provide a flexible and extendable framework for creating heterogeneous multi-agent simulations, botsim should be able to support many different kinds of agents;
- the simulation results should be recorded in detail, allowing the simulations to be 'played back' for detailed analysis.
- to have a command-line interface to facilitate large multiple-run experiments;
- to able to run through a graphical interface, that should be viewable in a web browser;
- to adhere to a design that excludes agents from compromising the integrity of other agents or the world; and

- the framework should be able to run on multiple operating system platforms.

Botsim is a library of classes for creating simulations. In a botsim simulation there is a world which contains entities. The entities in the world may be custom coded. The properties of the world and the types of entities which populate it are set in configuration files.

Botsim is written in the Java programming language (Joy et al., 2000; Chan et al., 1998), version 1.1.

The source code for botsim is available at:

`http://www.cs.mu.oz.au/~scv/thesis/botsim/botsim_src.jar`

A precise specification of the all botsim library classes is available at:

`http://www.cs.mu.oz.au/~scv/thesis/botsim/doc/`

2.7.1 The world and entities

A botsim simulation consists of a world and a population of entities that exist within it.

The world's properties and functions are encapsulated in the `World` class. Time is discrete, it elapses in time-steps. The world is a finite two-dimensional space in which the entities are placed. Space in the world is continuous. The world has a rectangular shape, its dimensions are determined by the configuration file settings (see section 2.7.3.) Botsim supports two alternative ways of handling the edges of the world: (1) they may be treated as boundaries which cannot be passed, or (2) they can be wrap-around edges (*i.e.* if an entity passes over the edge it moves to the other side of the world.) The `World` class also contains the code for the simulation engine.

Within the world exist entities. The `World` class keeps a list of all the

entities that exist within it and the positions at which they are located. The relationship between these the world and entities is described in UML notation in figure 2.7.

There are two top-level entity subclasses: static entities and active entities. These subclasses have many subtypes but each entity in the world is either static or active. The botsim library provides the Java interfaces `StaticEntity` and `ActiveEntity`. A botsim entity must implement one of these interfaces. Interfaces were used to allow the programmer maximum flexibility in how they implement an entity.

Static and active entities are fundamentally different in two respects. Firstly, active entities can move (although they do not necessarily have to) and static entities cannot. Secondly, active entities have an `act()` method which is called by the simulation engine at each time step of the simulation. Active entities may use the `act()` method to observe the conditions in the world, update their movements and initiate interactions with other entities. Static entities exist passively in the world.

As an illustration of the different entity types, imagine one wished to use botsim to create a simple simulation of an autonomous lawn mower which needs to avoid rocks. The lawn mower would be implemented as an active entity and the rocks as static entities.

Each entity is allocated a unique identifier by the simulation. Since entities are not given a direct reference to any other entity in the world, entity identifiers are used within the simulation as public addresses. When entities interact, entity identifiers are used to specify the entities that are being referred to.

Entities may belong to one or more *categories*. Categories are a loose mechanism for allowing different types of entities belong to the same group. Each

entity lists which categories it belongs to. For example in an artificial life simulation, a carrot entity may be in the categories 'carrot' and 'food'. Categories allow entities to search for certain kinds of entities within the world. For example, a rabbit entity may search for the closest 'food'.

In addition to entities there are also generators. Generators occupy no 'physical' space but are able to generate new entities in the world. Generators are polled by the simulation engine at each time step when they are given the opportunity to add new entities to the world.

2.7.2 Simulation properties

The world proxy

Entities are not given direct access to the `World` class. They are instead given access to a world proxy. All communication between entities and the world must go through the world proxy. The purpose of the world proxy is to allow entities to access the services of the world, without the entities being able compromise the integrity of either the world or other entities. The world proxy prevents entities gaining gain access to the properties of the world that should be hidden from the entities and prevents entities from getting direct references to other entities. It is a precaution to make it more difficult for entity code to compromise the integrity of the world or other entities.

An entity calls the `takeAction()` method in the world proxy to notify the simulation engine of the actions it intends to take. Actions are discussed in more detail in section 2.7.2.

Entities may use the methods provided by the world proxy to query about the conditions in the world. The kinds of information obtainable about the world include:

- whether there are any entities belonging to a specific category in the proximity. For example, an artificial life agent may query the world proxy whether there is any food within a fixed radius of the agent. The identifiers of the matching entities are returned to the querying entity.

- the relative distance from the entity to another entity. The querying entity must know the identifier of the target entity. The world or the entity may impose conditions on how close the target entity must be in order for its position to be 'visible'.

- general information about the world, for example the dimensions of the world and the duration of the simulation.

- odour information. Some entities emit an odour. An entity may request the world to calculate the odour gradient at a given point.

Entity interactions

The botsim simulation framework aims to give maximum flexibility to the ways in which entities may interact. Since the framework aims to allow many different kinds of custom coded entities, it must also allow for different kinds of custom coded entity interactions.

A botsim interaction occurs between two entities. One entity initiates the interaction, the other is the target. Both the initiator and the target may be affected by the interaction.

The framework provides an abstract class called `Interaction` as a generalisation of entity interactions. The interaction class is given a direct reference to both entities. An interaction implementation must define these three methods:

`validate()` This method should check the properties of the initiating and target entities and determine whether the interaction is valid.

`proximityNeeded()` Calculates the minimum distance between the initiating entity and the target entity for the interaction to be possible.

`doAction()` This method performs the interaction. It is executed by the simulation if the constraints imposed by the previous two methods are satisfied. The `Interaction` class has references to both entities in the interaction, giving it the power to directly call any of the public methods of either entity. This gives the interaction programmer a great amount of flexibility over how the interaction affects the entities.

The interaction class is sufficiently general to capture many different kinds of actions. Examples of the kinds of interactions that can be implemented include:

- eating food;
- communication between entities;
- one entity killing the another entity; and
- mating.

An interaction may be invoked by an entity through the `takeAction()` method in the world proxy.

Sensory processing

The botsim framework leaves sensory processing as the responsibility of the entity. This is to keep the simulation engine simple and to give maximum flexibility to the simulation framework. The botsim framework allows entities to know the precise relative positions of other entities in the simulation.

In what form this information is passed on to the entity decision making logic is up to the entity programmer:

- The programmer may opt to give the entity 'magic perception' and pass the locations of the other entities directly to the decision making code. In this case low level perception is ignored entirely.
- Alternatively, the programmer may wish to have more realistic sensory processing. In this case the programmer needs to simulate the entity's sensors. The positions of the other entities may be used to calculate what what the entity is able to sense.

2.7.3 Running a botsim simulation

The initial settings of a botsim simulation are described using configuration files. Botsim parses these files to create an initial world. The simulation is then executed and a recording is produced. The recording may be used to replay what happened in the simulation using the graphical interface or may be analysed by another program. Figure 2.8 illustrates this process of running a botsim simulation as a data flow diagram.

Setting up a simulation

A botsim simulation may be set up using configuration files. There are two main configuration files: the world configuration file and the entity file. In addition, each entity may have its own configuration file to configure the properties of the entity.

The world configuration file sets the properties of the world (such as the world dimensions and the number of time steps.)

The entity file lists all the entities and generators in the world. Each line of

the file specifies an entity, its Java class, its initial position in the world and an optional configuration file for the entity. Figure 2.9 shows an example entity file.

Simulation execution

A botsim simulation runs on a single threaded process. The simulation executes as a series of time steps. At each time step, entities may perform actions and their positions are updated. At each time step, the simulation attempts to resolve entity actions as if they were executed simultaneously.

The simulation event loop executes in the following order:

1. The generators are polled as to whether they want to add any new entities to the world.

2. The `act()` method of each active entity is called. At this point the active entities may request interactions with other entities. The interactions are placed in a list, they are resolved later. The active entities are also polled to obtain their intended movements. Again, the actual movements are resolved later in the loop.

3. The intended actions are resolved. For each action the simulation engine checks whether the requested action is possible. For an action to be possible the `validate()` method must succeed and the target entity and the initiating entity must pass within the required distance of each other during the time step. For example, if an entity tries to eat, the simulation engine will check the trajectories of the entity and the food during the course of the time step and calculate whether the entity passes within eating distance at any time. All actions are resolved, regardless of whether the initiating entity is eliminated from the simulation during the time step. For example, if two entities try

to kill each other during the same time step, then both kill actions will be enacted and the entities will be mutually annihilated. The order in which the actions are resolved is randomised rather than giving entities a priority. This is to give each entity an equal chance in the case where multiple entities try to simultaneously access a limited resource.

4. The positions of active entities are updated.

5. The positions of all the entities are recorded. Entities that implement the `Stated` interface also have their states recorded.

The result of a simulation execution is stored in a recording. The recording contains the positions of the entities at each time step and their states. The entities themselves may also record additional instrumentation.

Botsim interfaces

A botsim simulation may run using a graphical interface or by scripting a short program to setup the simulation.

The library interface

Botsim is a class library for programming simulations. A botsim program can run without writing a display to a graphical interface. The programmer may script batch jobs which output to text files. This gives botsim the facility to run large sequences of simulations. This capability is useful for:

- executing a simulation scenario with random influences multiple times to collect a large sample of evaluation data; and

- writing parameter optimisation algorithms.

To create a botsim simulation, the programmer needs to create an instance of the `World` class. The `World` constructor is passed, as arguments, the configuration settings of the world and a list of the entity initialiser objects. Entity initialiser objects are tuples which contain the name of the entity, its Java class, its position and its parameter settings. The `runSim()` method is called to run the simulation. Figure 2.10 lists a fragment of code that sets up and runs a botsim simulation.

The graphical user interface

A graphical user interface is an important tool for gaining an intuitive understanding of how entities interact in a simulation.

The botsim framework provide a graphical interface for playing back a simulation recording. The graphical interface is written as a Java applet so that it may be viewed in a web browser. Botsim uses only Java 1.1 classes so that it will run on the older versions of the Java Virtual Machine that are pre-installed with the most popular web browsers. A screen-shot of the graphical interface is shown in figure 2.11.

The graphical interface animates the simulation by showing the positions of all the entities in a series of time steps. Each entity is drawn in the world at the position it was recorded at that time step. How an entity is drawn is determined by the implementation of its `draw()` method.

The graphical interface has VCR like controls for playing back the simulation. The functions included are play, pause, step, step backwards, fast forward and rewind.

The graphical interface may be used to play the recording of a previously run simulation or to run a new simulation. New simulations are not animated in real-time. Instead, a recording is generated pre-time, by silently executing

the simulation in its entirety. This recording is then played to the user. Executing the simulation usually takes less than a second and the delay is not noticeable to the user. The simulation therefore *appears* to play in real-time.

If you wish to know about the crayfish escape circuit and all its adventures, you must read the next chapter.

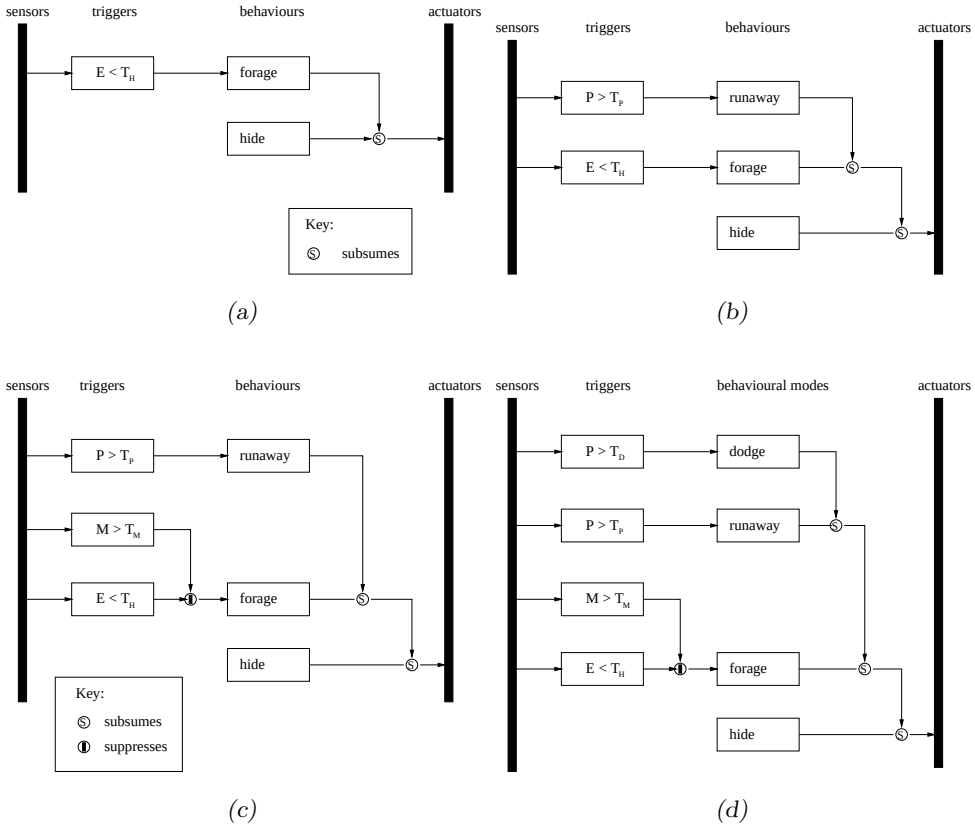

Figure 2.1: The architecture of the prey. (a) Hiding bot. By default the bot hides; if the bot is hungry, the forage behaviour takes control of the actuators. (b) The running away bot. Run away takes control of the actuators if the predator is seen. (c) The memory bot. If the predator was seen recently the forage behaviour is suppressed. (d) The dodging bot. Dodge takes precedence over all other behaviours.

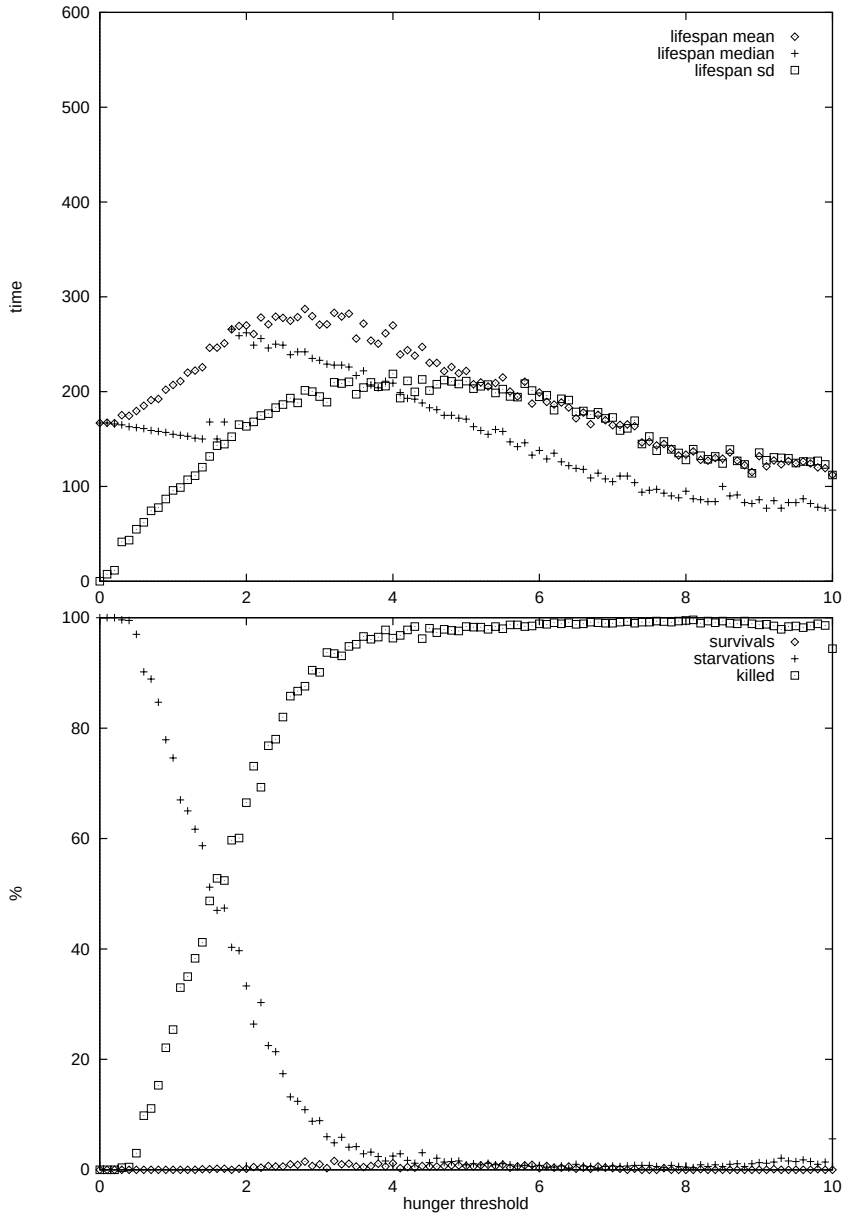

Figure 2.2: The survival times and cause of death the hiding bot. At low hunger thresholds the most common cause of death is starvation. At high hunger thresholds the most common cause of death is falling prey to the predator.

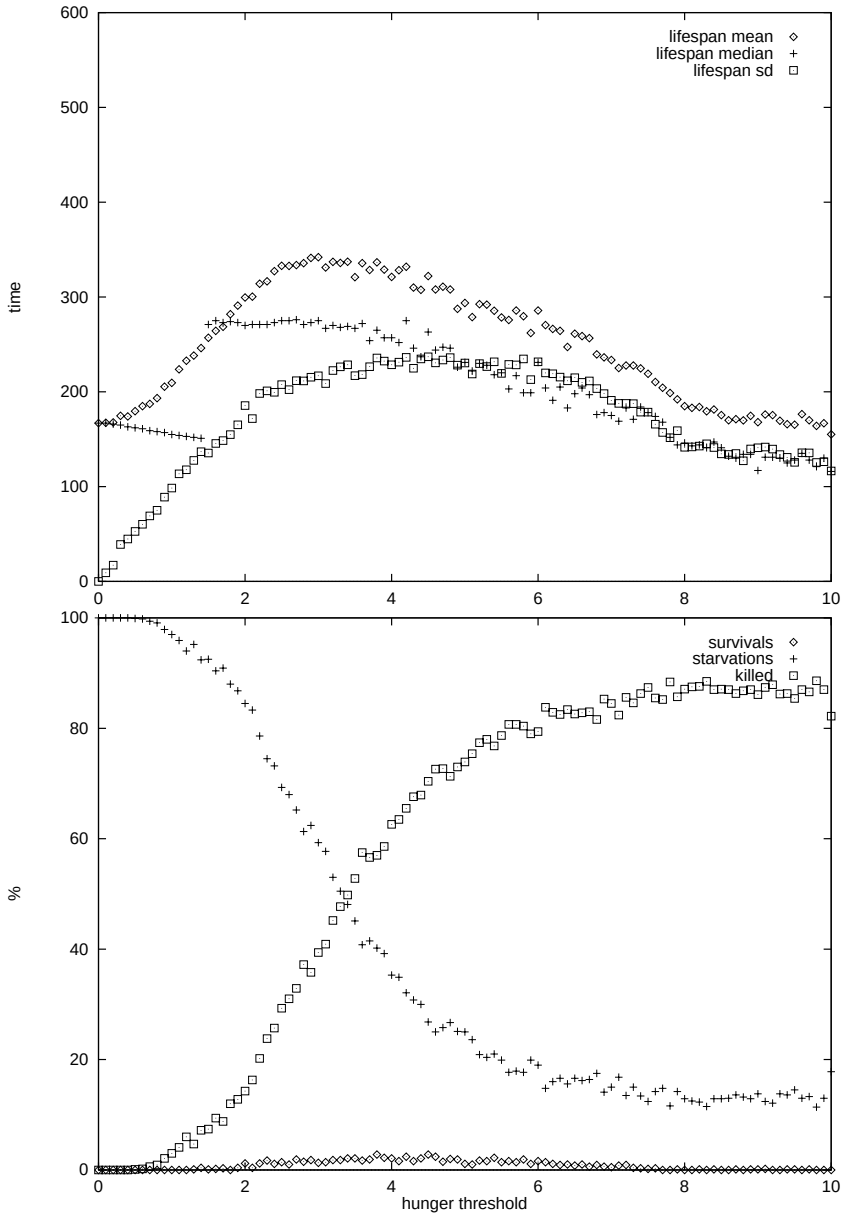

Figure 2.3: Running bot. The survival times and the cause of death versus the hunger threshold.

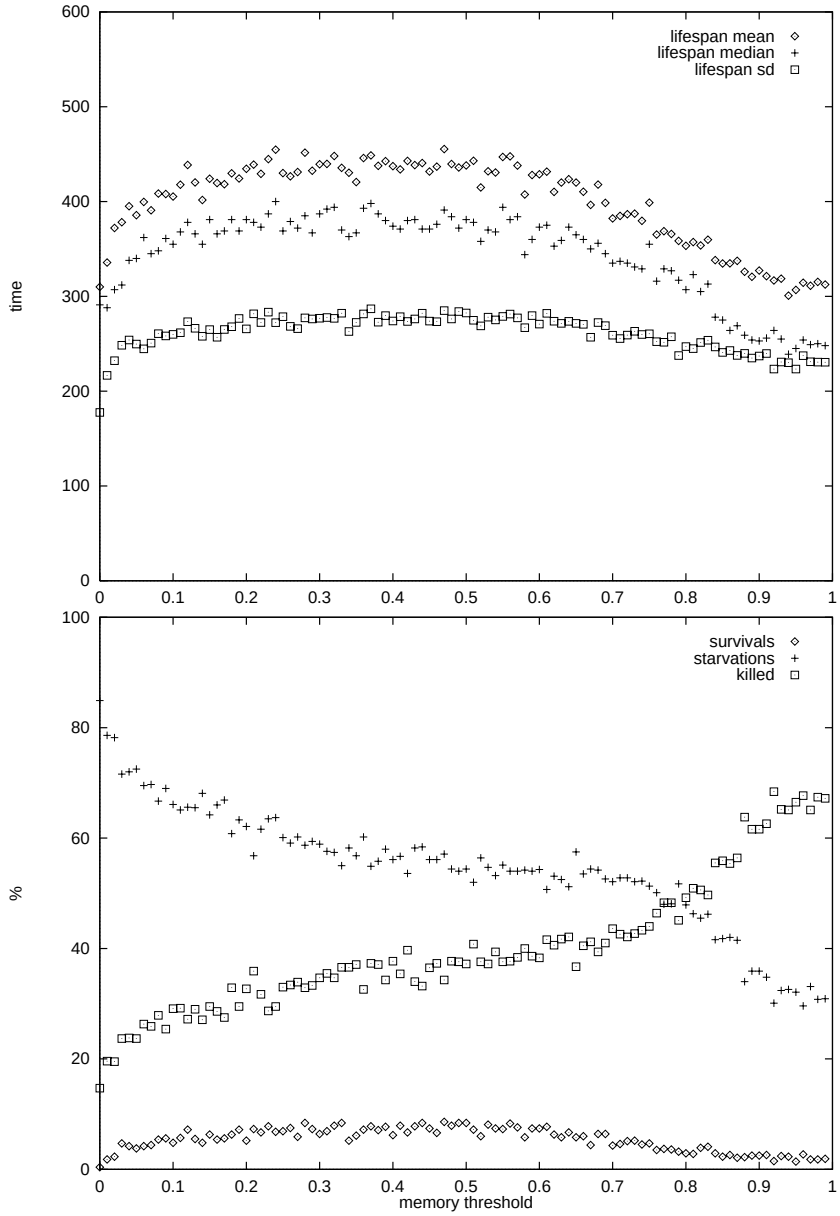

Figure 2.4: The survival times and cause of death versus memory threshold for the memory bot ($T_H = 6$).

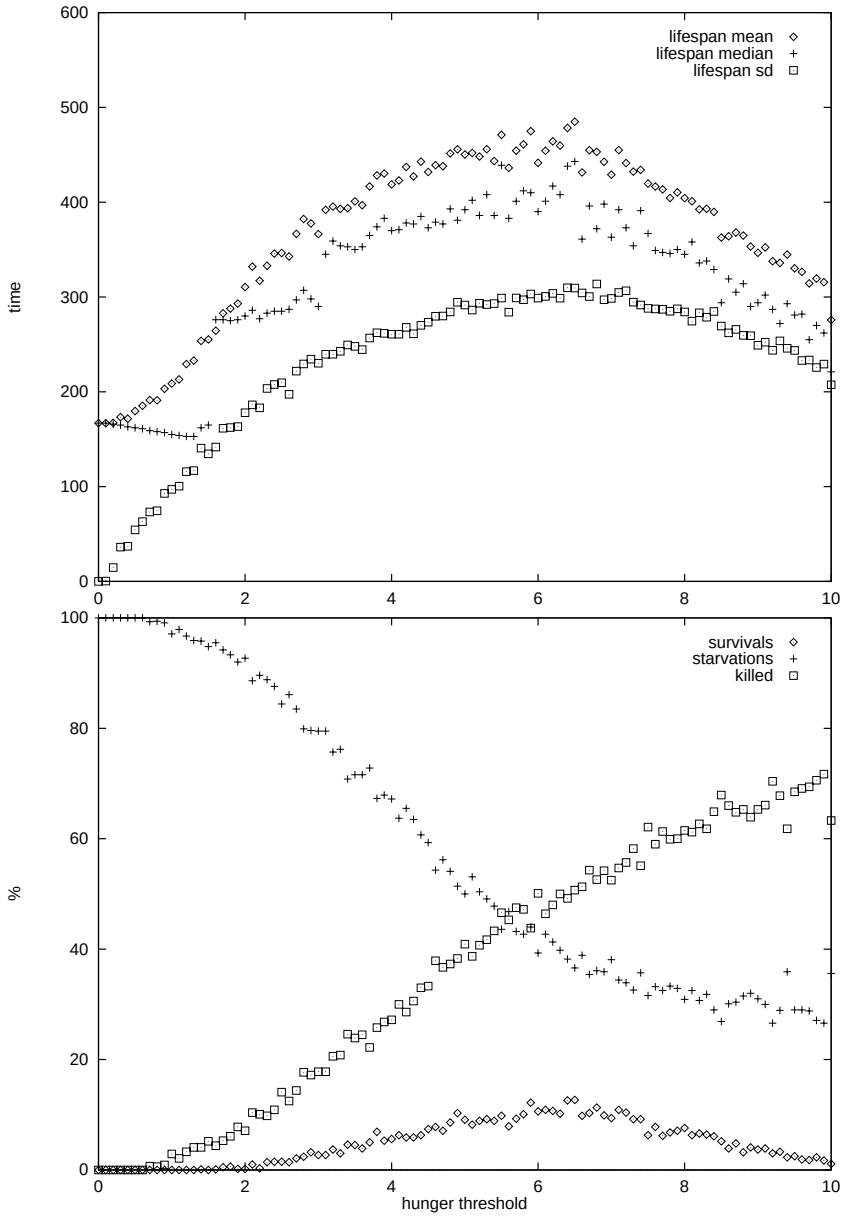

Figure 2.5: Memory bot ($T_M = 0.5$). How the survival times and cause of death vary with the hunger threshold.

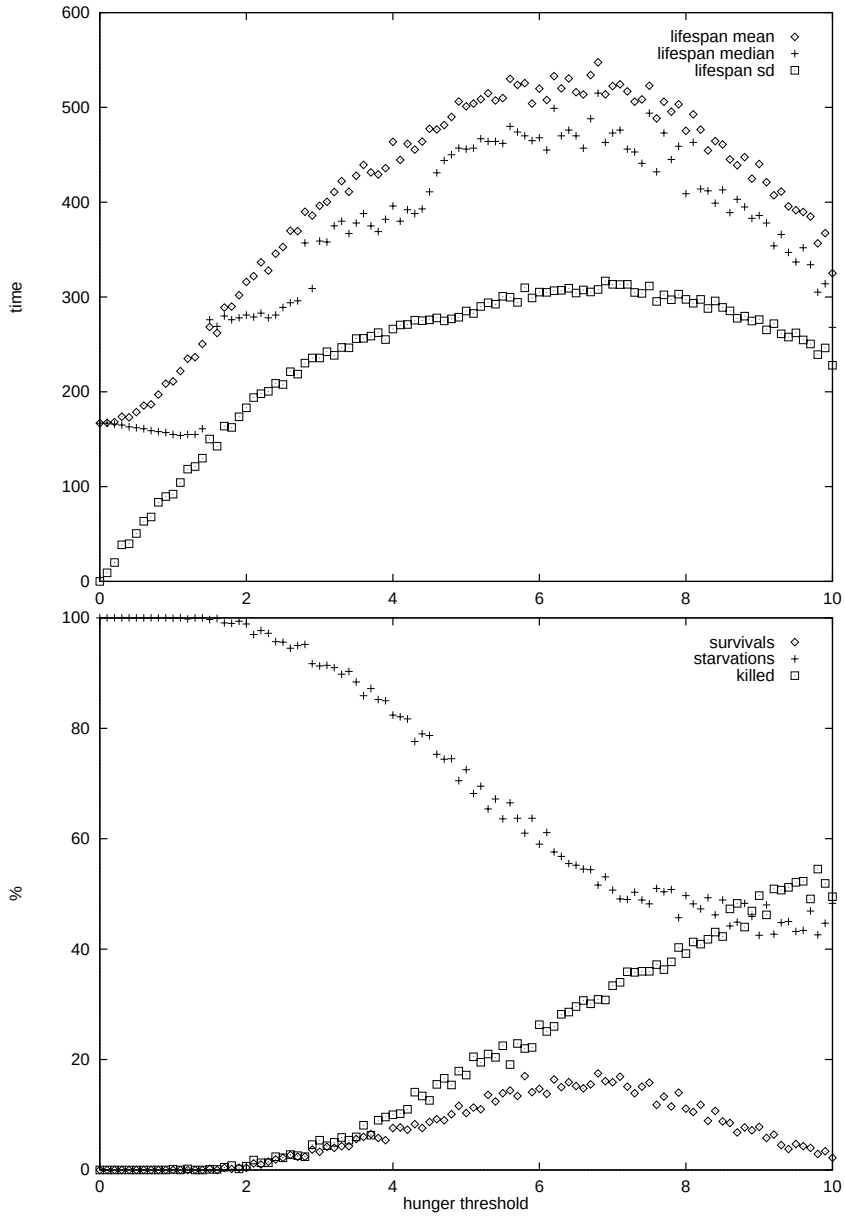

Figure 2.6: Dodging bot ($T_M = 0.5$). How the survival times and cause of death vary with hunger threshold.

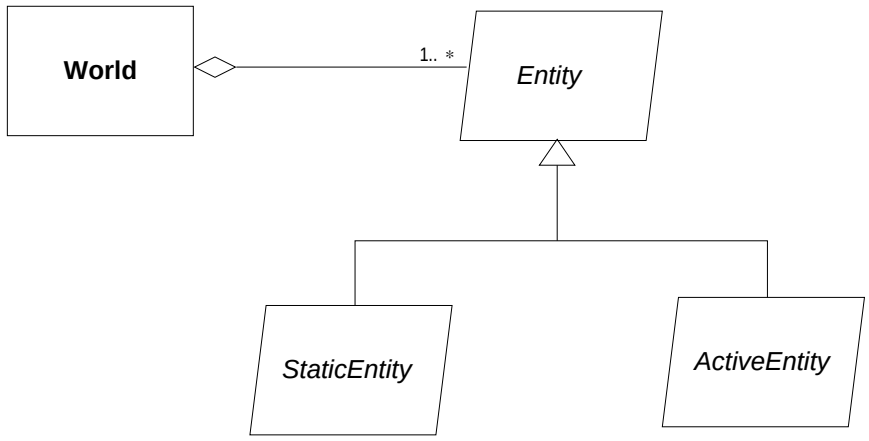

Figure 2.7: The botsim framework has a `World` *class, which defines an empty environment and contains the code for the simulation engine. In the world exist zero or more entities. Entities may be static or active.*

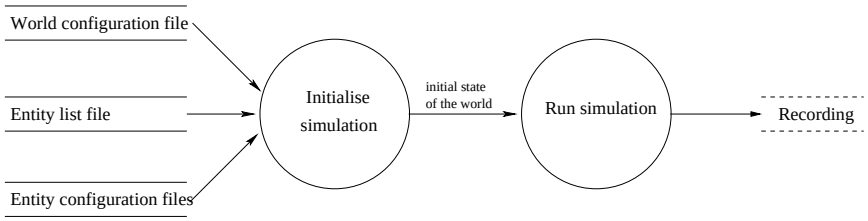

Figure 2.8: A high level data flow diagram of running a botsim simulation. Botsim reads configuration files to obtain the initial world settings and executes the simulation to produce a recording.

```
"generator", "sim.entities.FoodReplacer", null, "foodgen.dat"
"prey", "sim.simple.DodgingBot2", (250,250), "dodgingbot.dat"
"predator", "sim.entities.Predator", random, "predator2.dat"
"shelter", "sim.entities.Shelter", (250,250), "empty"
```

Figure 2.9: An example entity file. This file describes a simple world with a predator, a prey, a shelter and a food generator. Each line lists the entity name, the Java class, the initial position and the entity configuration file.

```
// Read the world configuration file
Parameters params =
     ConfigReader.readParameters("worldcfg.dat");
// Read the entity file
Vector ents = ConfigReader.readEntities("entlist.ent");
// Create the world
World world = new World(params, ents);
// Run the simulation
world.runSim(null);
```

Figure 2.10: An example Java code fragment which creates and runs a botsim simulation.

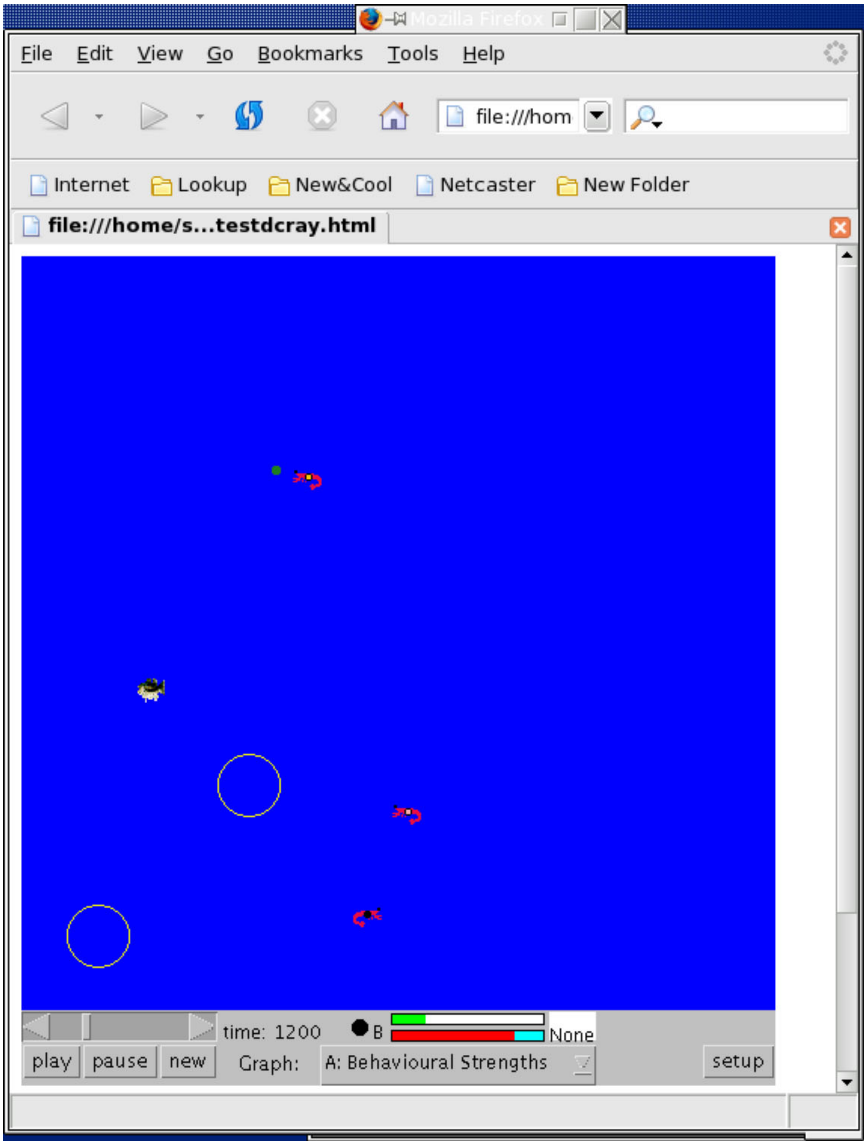

Figure 2.11: The graphical user interface to the botsim simulator.

CHAPTER 3

Neurobiology of the Crayfish Escape Circuit

The crayfish escape behaviour is a mechanism for the animal to escape from predators. The neural circuit controlling the escape behaviour has been the topic of extensive study in neuroscience. The circuit is very well understood, (at least compared with most other neural circuits.) This chapter reviews the current knowledge in neuroscience of how the crayfish escape circuit operates. The aim is to describe how the circuit functions from neurophysiological perspective. Before we can do that we first introduce some background of how neurons function. We also introduce some general information about crayfish biology and ecology, to put later discussions of the escape circuit in context.

Many biological terms are defined in this chapter. These terms are also listed in the glossary in appendix D.

3.1 An introduction to the neuron

Neurons are incredibly complex. This is a greatly simplified description of how they function, that focuses on the aspects that are relevant to understanding the rest of this thesis. Along the way, new terminology is introduced.

3.1.1 Overview of nervous systems

Complex animals have nervous systems. Nervous systems are responsible for processing sensory stimuli, making decisions and co-ordinating muscle movements.

From a functional perspective, one may think of the nervous system as an elaborate network for storing and processing information. Information is passed around as electrical signals. In addition to electrical signalling there is also chemical signalling. Chemical signals are able to trigger new electrical signals or change the structure of the network itself. Information is added to the nervous system when sensory organs convert stimuli into electrical signals. The nervous system is able to effect movement by transmitting electrical signals to the muscles.

Neurons are the basic building block of the nervous system. Neurons are specialised cells that can receive and transmit electrical signals. Neurons interconnect to form networks which together comprise the nervous system. In this thesis, the cells we are mainly concerned with are neurons; sometimes the words ‘cell’ and ‘neuron’ are used interchangeably.

Each neuron is responsible for integrating the electric currents it receives from other neurons. The integration process is complex. Different neurons integrate their inputs in different ways. The integration result is communi-

cated to the neurons it gives output to.

A neuron performs its integration asynchronously. Inputs may arrive during any slice of time. Unlike a micro-computer, the nervous system contains no central clock to synchronise the system.

There is an important subclass of neurons that are *spiking neurons*. Spiking neurons transmit electrical signals to their target cells, with which they communicate, in an 'all-or-nothing' manner. In a spiking neuron, if the sum of the neuron's inputs exceeds a threshold, a 'spike' is generated, which is communicated to the target neurons. The crayfish escape circuit is composed of spiking neurons.

Neurons communicate at their connections, called *synapses*. At synapses, electrical signals carry from one neuron to another. Synapses may be electrical or chemical and there are many subtypes. The variety of synapse types means that there is a huge selection of different mechanisms in the nervous system by which neurons can communicate. The current induced at synapses is a function of time. The currents induced at different types of synapses are characterised by different functions. The currents differ in their duration, peak amplitude and in the shape of their rise and fall. Synapses are not the only mechanism by which neurons can communicate. Neurons may also communicate through less direct mechanisms, such as through chemicals in the bloodstream.

The nervous system is not static. Neurons can grow or die. New neurons can form. The synapses between neurons can strengthen or weaken. New synapses can form between neurons. The way in which neurons integrate their input can change. Changes can be permanent or temporary. The dynamism of the nervous system is regulated by chemical signals between neurons and by chemical signals from other organs in the body.

In the coming sections we explain the mechanisms by which neurons operate.

3.1.2 Description of an individual neuron

A neuron consists of a *soma* (or cell body), *dendrites*, and an *axon* (or multiple axons.) These parts of the neuron are labelled in an example cell in figure 3.1. The soma contains the cell nucleus, which contains DNA. The DNA is read as part of the process of synthesising of proteins. The types of proteins assembled and where they are distributed determine the function of the cell. Extending from the soma are narrow tubes, called *neurites*. These tubes are filled with fluid (called cytosol) and conduct the electrical signals that pass through the neuron. There are two kinds of neurites: dendrites and axons. Input signals from other neurons are typically received in the dendrites. The dendrites subdivide into many branches to form a complex tree. The axon carries electric signals away from the cell onto the target neurons which the cell communicates with. Each cell typically has one axon, that may fork several times. Axons may project to cells that are nearby or in distant parts of the body.

Neurons differ greatly in their size and morphology (shape). Neurons serving different functions usually have vastly different morphologies. Even amongst neurons performing the same function there is considerable variation in the morphology. Figure 3.2 contains pictures of some different neurons.

The cell is enclosed by a barrier called the membrane. The membrane is studded with proteins, which perform various functions. Some of these proteins form channels. These are tiny holes in the membrane through which ions may pass. There are many different kinds of channels, that have different permeabilities to different ions.

The concentration of ions differs on the inside and outside of the cell. The

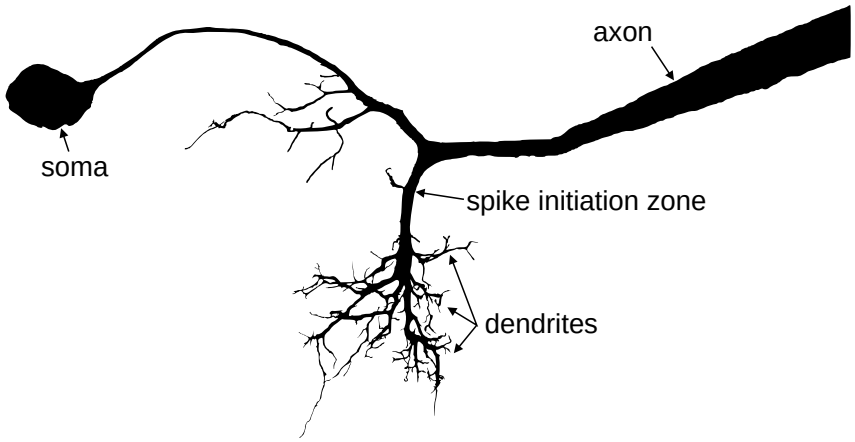

Figure 3.1: The anatomy of an invertebrate neuron. The neuron shown is the crayfish lateral giant. Picture: (Antonsen, 2003)

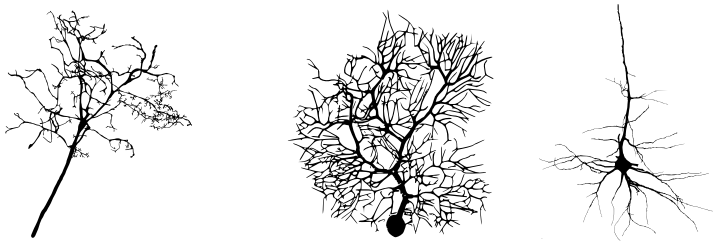

(a) *(b)* *(c)*

Figure 3.2: There is a great variation in the morphology of neurons. (a) A crayfish parasol cell. (b) A mammalian Purkinje cell. (c) A pyramidial cell from the mammalian cerebral cortex.

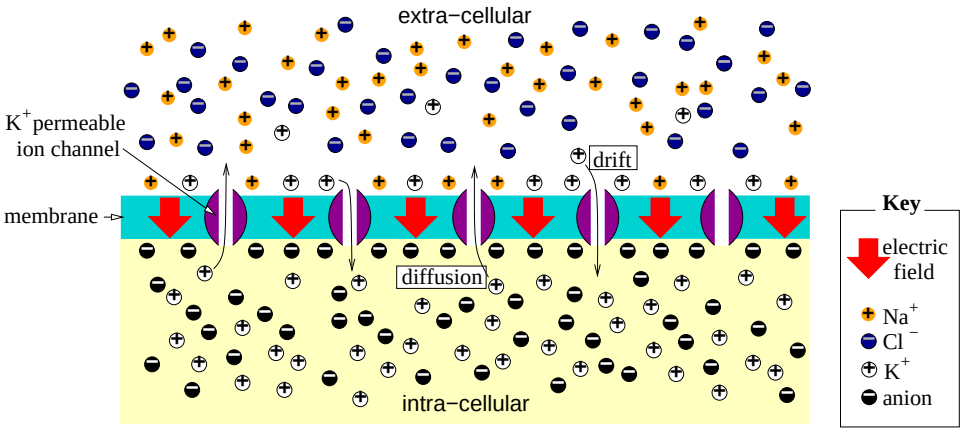

Figure 3.3: The neuronal membrane at rest. The membrane contains ion channels permeable (mostly) to potassium (K^+) ions. Due to the higher intracellular K^+ concentration, K^+ ions diffuse out of the cell. The extracellular build-up of positively charged ions creates an electric field, causing a drift current that counters the diffusion current.

concentration of sodium (Na^+) is higher outside the cell than inside the cell. Potassium (K^+) has a high concentration inside the cell and a low concentration outside. Other important ions include calcium (Ca^{2+}) and chloride (Cl^-). Due to the differences in ion concentrations, ions diffuse through the membrane. This diffusion of ions causes a buildup of charge, which creates an electric field across the cell membrane. This electric field drives ions in the opposite direction to the diffusion currents. When electric field generated current is in balance with the diffusion current the membrane is at its *equilibrium potential*, (or *resting potential*.) A typical cell membrane resting potential is around -65 to -70 millivolts. A conceptual diagram showing the cell membrane at equilibrium is shown in figure 3.3.

At rest, the membrane is said to be *polarised*, due to there being large numbers of oppositely charged ions on opposing sides of the membrane. Increasing the membrane potential above its resting potential is known as *depolarisation*. Causing the membrane potential to drop further below the resting potential is *hyperpolarisation*.

Relative to the total number of ions inside and around the cell, only a small number of ions need to cross the membrane to cause a large change in the membrane voltage potential. Importantly, large changes in the membrane voltage potential occur without significantly impacting the ion concentrations both inside and outside of the cell.

The differences in ion concentrations are created and maintained by ionic pumps. These are proteins in the membrane that transport ions across the membrane against their concentration gradients, at the expense of metabolic energy. For example, the sodium-potassium pump exchanges an intracellular sodium ion for an extra-cellular potassium ion. The ion pumps act as the 'quiet achievers', which work in the background, to maintain the ion concentration gradients, on which the working of the nervous system depends (Bear et al., 1996).

One of the main ways in which neurons transmit information is through spikes. In the dendrites, signals from other cells result in the generation of currents. These currents spread across the cell, depolarising the membrane. The start of the axon is called the spike initiation zone. If the membrane at the spike initiation zone is excited past a threshold then a spike is generated. This spike travels down the axon. Information is passed to other cells onto which the axon connects.

A spike is characterised by a sharp increase in the membrane potential, followed by a sharp decline. Spikes are caused and propagated by phenomena known as *action potentials*. Action potentials are generated in the following way. (Refer to figure 3.4 for a diagrammatic explanation.) A spike starts in the *spike initiation zone* (also known as the initial axon segment or the integrating segment.) This is usually located at the start of the axon. The membrane at the spike initiation zone is *active*. An active membrane contains voltage-gated sodium channels. These are channels than have open

and closed states. The channels are permeable to sodium ions when they are open and impermeable when they are closed. The state of channel depends on the membrane potential. When the membrane potential is at rest these channels are closed. Higher membrane potentials cause more channels to open. When an active membrane is depolarised, some of the voltage-gated sodium channels begin to open. The concentration of sodium is higher outside the cell, thus opening sodium channels causes an influx of sodium ions. If the membrane is depolarised past a threshold then a chain reaction is started. The opened sodium channels cause sodium currents which further depolarise the membrane which causes more sodium channels to open. The trickle of ions becomes a flood and the membrane potential is raised by over one hundred millivolts. A spike has been generated. (Refer to graphs (3) and (4) in figure 3.4 to see the shape of spike.)

We have explained how a spike rises, we now explain how it falls. Two mechanisms are involved. Firstly, the voltage-activated sodium channels only stay open for a short period. After about one millisecond they deactivate and return to a closed state. Secondly, the active membrane also contains voltage-gated potassium channels. These channels are also activated by a membrane depolarisation, but their opening is delayed. When the spike is at its peak, the potassium channels open, causing potassium ions to flow out of the cell. The outward movement of the positively charged ions causes the membrane potential to drop back to its resting potential. When the sodium and potassium channels return to their closed states, the cell is ready to generate the next spike.

A mathematical description of the sodium and potassium channel dynamics is given in section 4.1.1.

Once the action potential is initiated it propagates along the axon. The membrane along the axon is also active. Ions that flux into the cell at the

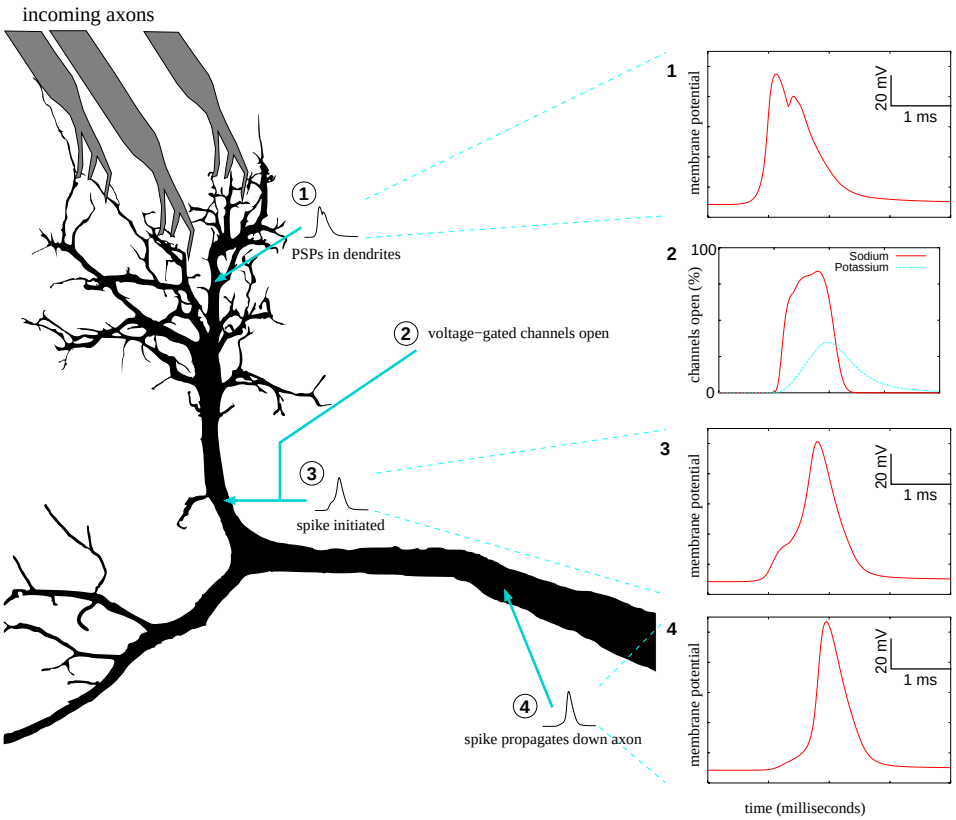

Figure 3.4: The generation of an action potential. (1) Spikes in the incoming axons cause depolarisations in the dendrites. (2) The depolarisation spreads to the spike initiation zone, causing the voltage-gated sodium channels to open. (3) The inflowing sodium ions cause a spike which declines as the voltage-gated potassium channels open. (4) The spike propagates down the axon.

spike initiation zone diffuse to the adjacent axon area, thereby depolarising the axon membrane to threshold. This activates the sodium currents causing a spike in the axon membrane potential. In this manner the spike is able to travel, as a wave, to the end of the axon. The active currents allow spikes to travel long distances without diminishing.

The dendrites affect the integration properties of the cell. Currents generated by the input from other neurons often need to travel through the dendrites before reaching the spike initiation zone. The morphology and

properties of the dendrites determine how the cell integrates its inputs. In the classical model of a neuron, the dendrite membrane is *passive*. Passive membrane contains no active channels. Thus, according to this model, currents initiating in the dendrites diffuse through the cell with attenuation. Ions leak out through the dendritic membrane as they diffuse towards the spike initiation zone. More recent research has shown that many dendrites do in fact contain different kinds of active channels. The classical model nevertheless serves as a useful reference point in the mental conception of how a neuron functions.

Functionally, neurons may be divided into three main subclasses: *afferents*, *interneurons* and *motorneuron*. Afferents carry sensory signals toward the central nervous system. Motorneurons (also called 'efferents') are neurons which connect onto muscle cells. Interneurons are all the neurons in between.

3.1.3 Synapses

Synapses are connections between neurons, through which electric currents in one neuron can generate electric currents in another. Synapses are the principal mechanism of neural communication.

Synapses may form at almost any part of a neuron. Most commonly, synapses connect an axon to the dendrites of another neuron. There also exist axon to axon synapses, axon to soma synapses and, in rare cases, dendrite to dendrite synapses.

The site on the membrane where a synapse forms is known as the *synaptic terminal*. Each of the two neurons, which are connected by the synapse, have a synaptic terminal. Transmission is usually uni-directional, (or at least more common in one direction.) Usually transmission is from an axon terminal to a target neuron. The axon terminal is said to be *pre-synaptic*

and the terminal at the target cell is *post-synaptic*.

Synapses can either be electrical or chemical.

An electrical synapse, or *gap junction*, is a connection between cells which allows electric current to flow directly from one neuron to another. At each terminal of the gap junction, there are pores, called *connexons*. These pores are wide enough to allow the passage of all the major ions and some small molecules. Electrical transmission is extremely fast and reliable.

Connexons are made out of proteins called *connexins*. There are different kinds of connexins. The type of connexins present determine the electrical conductance properties of the gap junction. Gap junctions can be ohmic (*i.e.* have a constant resistance) or be of variable conductance. In rectifying gap junctions, the conductance is dependent on the voltage difference between the pre and post synaptic terminals. The effect of this is that the conductance is higher for transmission in one direction than the other. Rectifying junctions play an important role in the crayfish escape circuit. This is explained in more detail in section 3.3.4.

Positive current flowing into the post-synaptic cell causes a depolarisation. This is called an *excitatory post synaptic potential* (EPSP).

In gap junctions current can flow in both directions. Nevertheless, there is still usually a direction in which the current most commonly flows, so there are still designated pre-synaptic and post-synaptic terminals. When current flows in the usual direction this is called *orthodromic* transmission. When current is transferred in the opposite direction (post-synaptic to pre-synaptic) this is *antidromic* transmission.

Chemical synapses are somewhat more complicated and there are many subtypes. In chemical transmission the pre-synaptic terminal releases molecules, called neurotransmitters, that cause a temporary change in the post-synaptic

membrane.

Neurotransmitter molecules are packed into capsules called 'synaptic vesicles'. The pre-synaptic terminal is arrayed with synaptic vesicles (see figure 3.5a.) The contents of a vesicle are released all at once, causing neurotransmitter release to occur in quanta.

The arrival of an action potential at the pre-synaptic terminal triggers neurotransmitter release. The membrane depolarisation opens voltage-gated calcium channels, causing an influx of calcium (Ca^{2+}) ions. Calcium is thought to catalyse the vesicles to release their contents. The vesicles fuse with the pre-synaptic membrane and neurotransmitter molecules spill into the synaptic cleft (the narrow space between the pre and post synaptic terminals.) Figure 3.5b shows a diagram of neurotransmitter release.

The neurotransmitter molecules bind to receptor sites at the post-synaptic terminal. This causes a change in the post-synaptic cell.

The response in the postsynaptic cell depends on the kind of neurotransmitter in combination with the kind of receptors at the post-synaptic terminal. The receptor effect falls into two categories:

transmitter-gated ion channels: the binding of the neurotransmitter molecules to the receptors cause ion channels to open. Synapses may be either excitatory or inhibitory, depending on the kind of ion channels activated. If a positive current is activated (such as when sodium channels are opened) a rise in the post-synaptic membrane potential is caused, *i.e.* a EPSP. It is also possible for the channels to cause negative currents (such as by opening chloride channels), resulting in an inhibitory post synaptic potential (IPSP). Figure 3.5c shows a diagram of transmitter-gated ion channels opening.

G-protein coupled receptors: the neurotransmitter binds with the re-

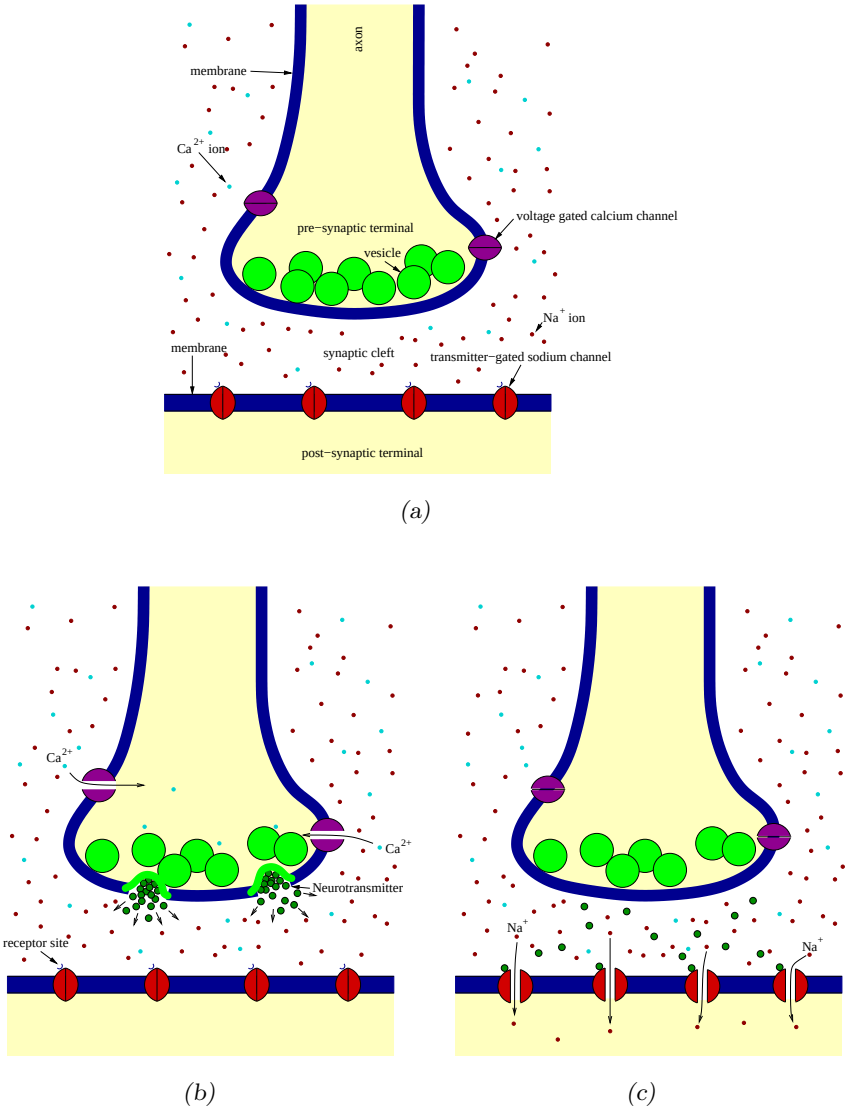

Figure 3.5: Chemical synaptic transmission. (a) The pre-synaptic terminal at an axon and the post-synaptic terminal at the target cell. (b) The arrival of an action potential at the pre-synaptic terminal opens voltage-gated calcium channels. The influx of calcium causes some synaptic vesicles to fuse with the pre-synaptic membrane and release neurotransmitter molecules into the synaptic cleft. (c) The neurotransmitters bind to receptors on transmitter-gated sodium channels, causing them to open and effecting an EPSP.

ceptor causing small molecules, called G-proteins, to be released in the post-synaptic interface. The G-proteins then interact with other proteins in the cell (called 'effectors') which causes a change in the cell. The effectors may activate an ion channel or cause 'second-messengers' to be released. Second messengers diffuse through the cytoplasm. They can regulate or modulate ion channels or change the cell in other ways, such as altering the cell's metabolism. In comparison to transmitter-gated ion channels, G-protein effects are generally slower and longer term.

There are many kinds of neurotransmitters. At present, over fifty molecules that have been identified as neurotransmitters and the number is growing with research. The neurotransmitters that are known to be involved in the crayfish escape circuit are: acetylcholine (ACh), gamma-aminobutyric acid (GABA) and serotonin (5-HT). Acetylcholine is usually an excitatory neurotransmitter and is used at the neuron to muscle synapses. GABA is an inhibitory transmitter. Serotonin has modulatory effects.

There may be many different receptor types corresponding to a neurotransmitter. The same neurotransmitter can have different kinds of effects at different synapses depending on the types of receptors.

From a functional perspective, electrical and chemical transmission each have their advantages and disadvantages. Electrical transmission is faster. Ions are transferred directly from one cell to another. There is no delay in this process. In transmitter-gated chemical transmission, the process of opening ion channels results in a delay (in the order of a millisecond.) Electrical synapses are common in circuits where time time is extremely critical, such as the crayfish escape circuit. Electrical synapses are limited in the kind of response they can effect in the post-synaptic cell. Electrical synapses are always excitatory and the duration and magnitude of the post-

synaptic response is dictated by the pre-synaptic spike.

Chemical synapses come in many types and subtypes, which vary greatly the response they produce in the post-synaptic cells. The characteristics of the response depend on the ion channels that are opened instead of the pre-synaptic spike. Chemical synapses may be excitatory or inhibitory. The duration of the post synaptic response may be much longer than the triggering pre-synaptic spike. If the chemical synapse has G-protein receptors, chemical transmission may not directly trigger an electrical signal but instead change the cell in another way. For example, the cell's spiking threshold may be changed. The characteristics of chemical transmission depend on what combination of neurotransmitter and receptors that are used in the synapse.

Neurohormones are another mechanism in nervous systems for communication. Whereas at a synapses one neuron communicates directly with one other neuron, neurohormones are used to communicate with populations of neurons. Neurohormones are neurotransmitters that are secreted by specialised neurons into the bloodstream causing them to be distributed widely and reach many neurons. Neurons that have receptors for the released neurohormones will be affected. Due to there being different receptor types, neurohormones can have many different affects on different cells. Sleep is an example of a behaviour that is affected by the release of neurohormones. Serotonin and octopamine are neurohormones that are thought to be involved in the regulation of the crayfish escape circuit.

3.2 An introduction to crayfish

Crayfish are freshwater dwelling decapods that are closely related to lobsters. Crayfish inhabit streams, rivers, swamps, ponds and dams. They are

scavengers and predators, and are themselves subject to heavy predation. To avoid predators, they hide, build burrows which they defend with their pincers, and also escape with a repertoire of rapid responses.

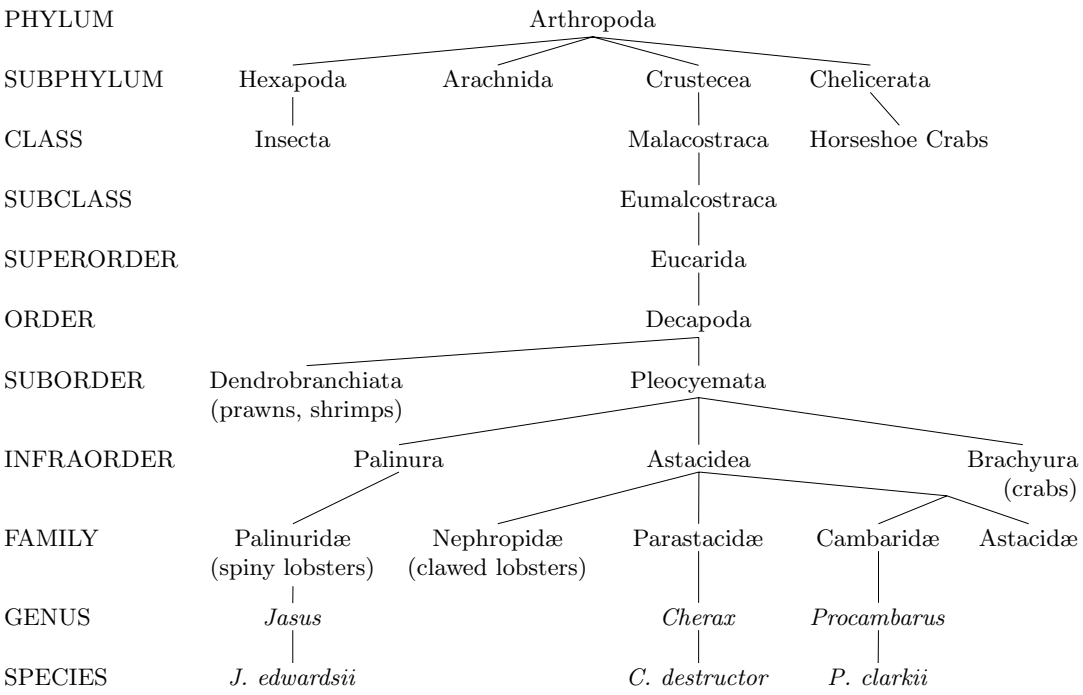

Figure 3.6: The scientific classification of freshwater crayfish and their relation to some other well known arthropods. Freshwater crayfish belong to the infraorder Astacidea, together with clawed lobsters. Cherax destructor and Procambarus clarkii are the species commonly used in experiments because of their abundance in their respective continents.

Across the world there are approximately 500 species of crayfish. Figure 3.6 shows the taxonomic classification of crayfish and their relation to some other well known arthropods. Crayfish belong to the order Decapoda, a large group of crustaceans with five pairs of thoracic limbs. Other decapods include prawns, shrimps, crabs and lobsters. Freshwater crayfish belong to the families Parastacidæ, Cambaridæ and Astacidæ. Crayfish from family Astacidæ are found in Europe and western North America. Crayfish from family Cambaridæ are distributed across eastern North America and east

Asia. The southern hemisphere is inhabited exclusively by species belonging to Parastacidæ. *Cherax destructor*, also known as the 'yabby', is most common species found in Australia. It is also the species usually used in neurobiological studies conducted in Australia. *Procambarus clarkii* is a species often used in North American studies. In Australia, the word 'crayfish' is also the common name for the ocean dwelling *Jasus edwardsii* (or 'rock lobster'), but that's not the animal we're referring to in the following discussions. (Huxley, 1896; Edgar, 2000; Merrick, 1993; Maddison, 2001)

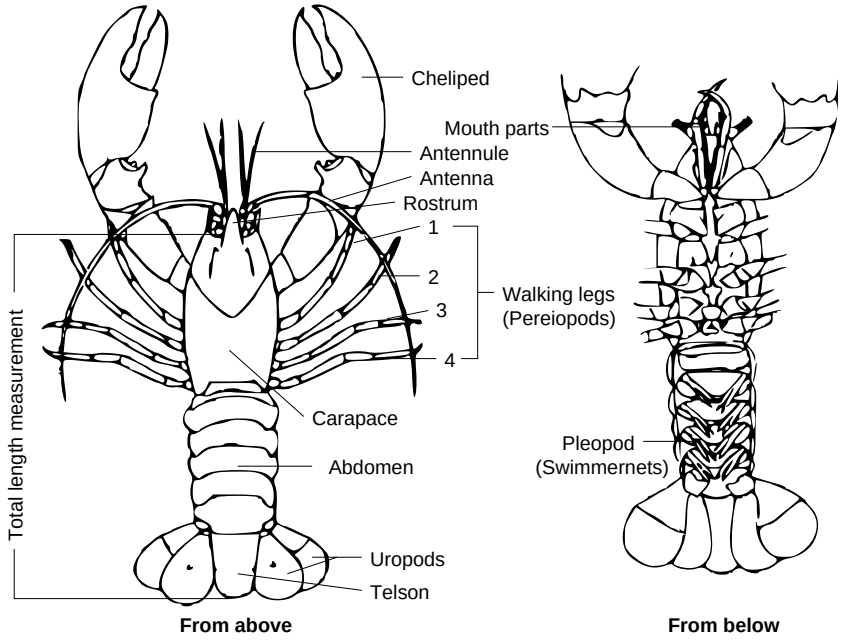

Figure 3.7: Anatomy of the crayfish. The species shown is Cherax destructor.

Figure 3.7 labels some of the main features of the crayfish anatomy. Crayfish have an exterior hard shell known as the exoskeleton that protects and supports the internal organs and structures. The exoskeleton is naturally divided into two distinct regions: a solid and continuous part at the front, called the *carapace*, that encases the head and thorax, and a jointed hind part, called the tail or *abdomen*. The crayfish has five pairs of thoracic

limbs. The rear (or posterior) four pairs of limbs are walking legs, known as the *pereiopods*. The front (or anterior) limbs are the *chelipeds* and end in claws, they form the crayfish's chief weapon in defence and offence. The front of the carapace ends in a sharp spine, extending like a beak, called the *rostrum*. Below the rostrum on each side is a pair of eyes that are mounted on moveable stalks. Below the eyes are two pairs of feelers, one short pair, the *anntenules* and a much longer pair, the *antennae*. The antennae are over half as long as the body, and can explore a considerable area in poor visibility. The abdomen is segmented into six segments. Attached under the middle four abdominal segments are a short pair of limbs called the *pleopods* or *swimmernets*. The pleopods assist the crayfish in propelling forward, and are used by a berried female crayfish to carry her eggs. The sixth segment has larger appendages, called the *uropods*, which together with the flaps on the terminal segment, called the *telson*, form the tail flap. The tail is used to rapidly thrust the crayfish through the water to escape sources of danger (Merrick, 1993; Holdich and Reeves, 1988; Withnall, 2000; Huxley, 1896).

The crayfish is equipped with sensory organs that allow it to see, hear, smell and touch. Mechanosensory hairs, or *setae*, are scattered over the crayfish's body and appendages. These hairs are able to detect disturbances. The antennules contain a complex array of olfactory receptors. The crayfish have a primitive ear. Crayfish have compound eyes. Crayfish compound eyes are similar to those of other arthropods (such as insects.) Although the details of how compound eyes exactly work is still the topic of research several general points have emerged: (i) they are able to sense light over a wide wavelength (greater than the visible spectrum, (ii), they provide a mosaic image with a wide field of view, and (iii) they are very sensitive to movement and changes in light intensity (Huxley, 1896; Merrick, 1993).

Four types of hairs found on telson. Some respond to low frequency (0-2 Hz)

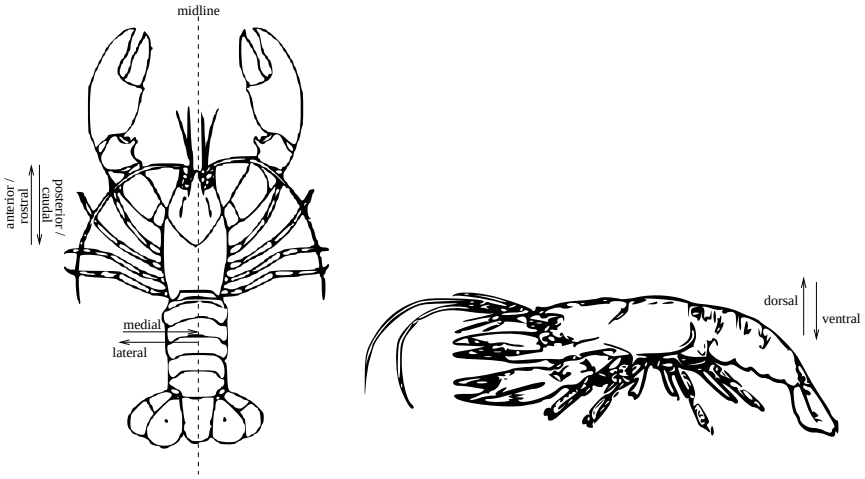

Figure 3.8: When referring to parts of the crayfish, directions are necessary. Biologists have given names to the directions along the three axes of an animal's body. The direction towards the rostrum is rostral or anterior. Towards the rear is said to be caudal or posterior. An imaginary line running down the rostrum to the rear of the crayfish, called the mid-line, divides the animal into symmetrical halves. The direction towards the mid-line is medial, away from the mid-line is lateral. The direction towards the top of the animal is dorsal, towards the bottom is ventral.

movement, others appear to be acceleration sensitive and have an optimal response at 80 Hz. Low frequency sensors are suppressed by high frequency stimulation (Wine, 1984).

The life-cycle of a crayfish progresses through a series of moults, during which the animal increases in size. During a moult a crayfish sheds its exoskeleton. Crayfish hatch from eggs in late spring and undergo seven to eight moults in their first year of life. The exact number of moults is governed by temperature. The growth cycle is repeated in subsequent years with the frequency of moults decreasing as size increases, to eventually one moult per year. After undergoing a moult a crayfish is vulnerable to attacks and cannibalism. It takes approximately three days for the new shell to harden. The nutrients in the shedded shell are recycled. Empty shells are eaten by crayfish, often by the crayfish that moulted it. (Lowery, 1988; Merrick, 1993)

Crayfish reach sexual maturity in their third or fourth year. The time of mating is dependent on the water temperature. Spawning usually peaks in late spring to early summer. During spawning, the male crayfish places a spermatophore between the female's fourth and fifth pair of walking legs. The female breaks open the spermatophore and mixes the sperm with the eggs she expels. The fertilised eggs become firmly attached to the pleopods and are thus carried under the abdomen. Depending on temperature, eggs take approximately 20-40 days to hatch. The young will remain attached to their mother for the first three moults of their life. (Withnall, 2000; Lowery, 1988)

Crayfish are omnivorous. They are scavengers and opportunistic predators. Their diet is diverse and adaptable. Three modes of feeding have been identified: browsing, hunting and bulldozing. When browsing they wander gathering food consisting mainly of *detritus* (decomposing plant and animal

material) but they will also graze on underwater plants if available. A study in Sweden, where the stomach contents of a large sample of crayfish were examined, revealed that, for the total population, plant material accounted for about two thirds of the total food consumed. The plant material was predominantly made up of decomposed leaves, roots and bark. Crayfish hunt a diverse range of smaller animals. Their most commonly cited prey include molluscs, aquatic insect larvae, worms, small crustaceans and amphibian tadpoles. Crayfish are also cannibalistic: large crayfish will kill and consume smaller ones. Other crayfish make up a substantial part of a crayfish's diet. When subjected to starvation a crayfish will bulldoze. This involves the crayfish scraping the material deposited on its exoskeleton and eating it. Crayfish prefer to feed at night or in low visibility conditions. There is considerable evidence that the diet of crayfish varies with age. Juvenile crayfish feed predominantly on aquatic invertebrate prey whereas adult crayfish feed predominantly on vegetation and detritus. It is hypothesised that this is because the latter's movements are less rapid and precise compared with juveniles. (Goddard, 1988; Merrick, 1993)

Crayfish are preyed upon by a great variety of animals, both aquatic and terrestrial. Crayfish are most vulnerable during juvenile stages and during the immediate post-moult period. Larger crayfish are one kind of commonly encountered dangerous predator. In their juvenile stages they often fall prey to invertebrates such as dragonfly nymphs or water bugs. Fish such as trout, pike, perch, catfish, bass and eels eat crayfish. Crayfish are also preyed upon from the air. Crayfish make up a substantial proportion of the diets of certain species of ibises, herons and egrets. Other birds such as crows and gulls will opportunistically feed upon crayfish. Other predators include reptiles, frogs and water associated mammals. (Hogger, 1988)

When crayfish live in a community they co-exist under a social dominance

hierarchy. An animal's social status governs its access to the limited resources, such as food and shelter. The dominance hierarchy is determined by fights between crayfish. An agonistic encounter between two animals proceeds according to strict rules of conduct. The typical scenario of an encounter follows a sequence of escalating behaviours starting with extensive threat displays upon first contact and terminating in a brief session of unrestrained combat (Huber and Kravitz, 1995). Following a series of fights, one of the crayfish emerges as the dominant. Subsequently, the fighting between the animals will diminish in frequency and duration (Issa et al., 1999; Herberholz et al., 2001). When faced with an encounter with the dominant animal, the subordinate animal will usually just back away.

The transition the losing crayfish makes being a dominant contender to a subordinate is swift. In observations of crayfish fighting, there is an identifiable moment when a crayfish gives up fighting and submits to the dominant (Herberholz et al., 2001). The dominance decision is followed by a change in the behaviour patterns of the animals. The subordinate animal ceases engaging in any aggressive behaviours such as approaching or making attacks and increases its frequency of submissive behaviours such as tail-flipping and retreating. The rapid change in behaviour is suggestive of a switch of neural state in the subordinate animal.

The internal dominance decision made in a crayfish also affects non-agonistic behaviours. The tendency of a crayfish to burrow, a behaviour not associated with fighting, changes in parallel with the dominance decision (Herberholz et al., 2003). Prior to the establishment of the dominance order, both the future dominant and future subordinate may engage in building burrows. After the dominance decision, the subordinate animal will cease burrowing. The dominant animal's burrowing activity will increase. The change in burrowing activity occurs simultaneously with the changes in agonistic be-

haviour that signals the dominance decision. The simultaneous transitions in agonistic and non-agonistic behaviour suggests that the changes may be mediated by the same neural mechanisms (Herberholz et al., 2003).

3.3 The neural mechanisms of the crayfish escape reflex

The crayfish produces a powerful tail-flip to avoid rapid strikes from predators. This behaviour is known as the escape response. The neural circuitry that controls this behaviour has been the subject of extensive research in invertebrate neuroscience.

In this section we give a review of the current scientific knowledge of the crayfish escape circuit. We begin by giving an overview of the circuit. We then describe the sensory part of the circuit followed by the motor part of the circuit. Rectifying gap junctions have been identified as playing a key role in the circuit's function and their role is discussed. The escape circuit is modulated by other parts of the nervous system. We describe the modulation it receives and their neurological mechanisms. Finally, we discuss the escape circuit as an example of decision making.

The description of the escape circuit in this section goes beyond what we model in this thesis. We give a complete description of the escape circuit as a platform for future work in modelling the crayfish escape circuit and to provide context for our model. The biological models in this thesis are of the sensory part of the LG circuit which is described in section 3.3.2. Our models rely heavily on rectifying junctions, which are described in detail in section 3.3.4.

The interested reader is referred to the reviews by (Wine and Krasne, 1982)

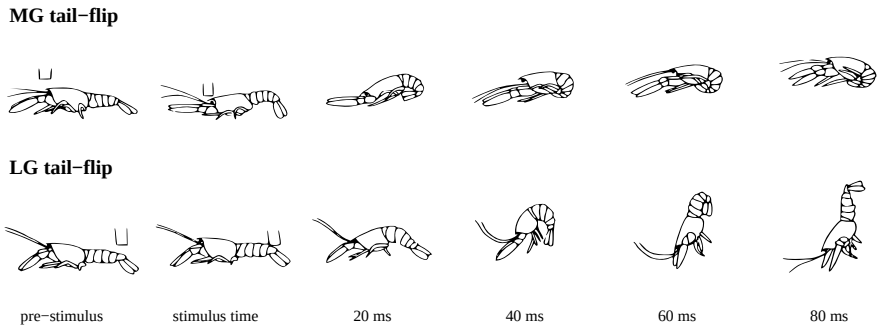

Figure 3.9: Crayfish tail-flips. The MG circuit responds to stimulus to the front of the animal and propels the crayfish backwards. The LG circuit responds to stimulus to the abdomen and lifts the crayfish forwards and upwards. Pictures from (Jackson and MacMillan, 2000).

and (Edwards et al., 1999) for further descriptions of the crayfish escape circuit.

3.3.1 Overview of the escape circuit

Within the scientific literature, three distinct crayfish escape circuits have been identified, each of which controls a variation of the escape response. They are known as the Lateral Giant (LG) circuit, the Medial Giant (MG) circuit and the non-giant circuit. The giant circuits produce a very fast, highly stereotyped responses. An LG tail-flip propels the crayfish forward, the MG tail-flip propels the crayfish backwards. The non-giant circuit produces a range of tail-flip movement patterns. Figure 3.10 shows an overview of the neural circuitry that controls a crayfish escape tail-flip. (Krasne and Wine, 1984; Wine, 1984)

The crayfish body is divided into segments. Each segment of the crayfish has a mini-brain called a ganglion. The crayfish central nervous system (CNS) is made up of a series of ganglia connected by the nerve cord. Each ganglion is divided into symmetrical halves called hemi-ganglia, one for the

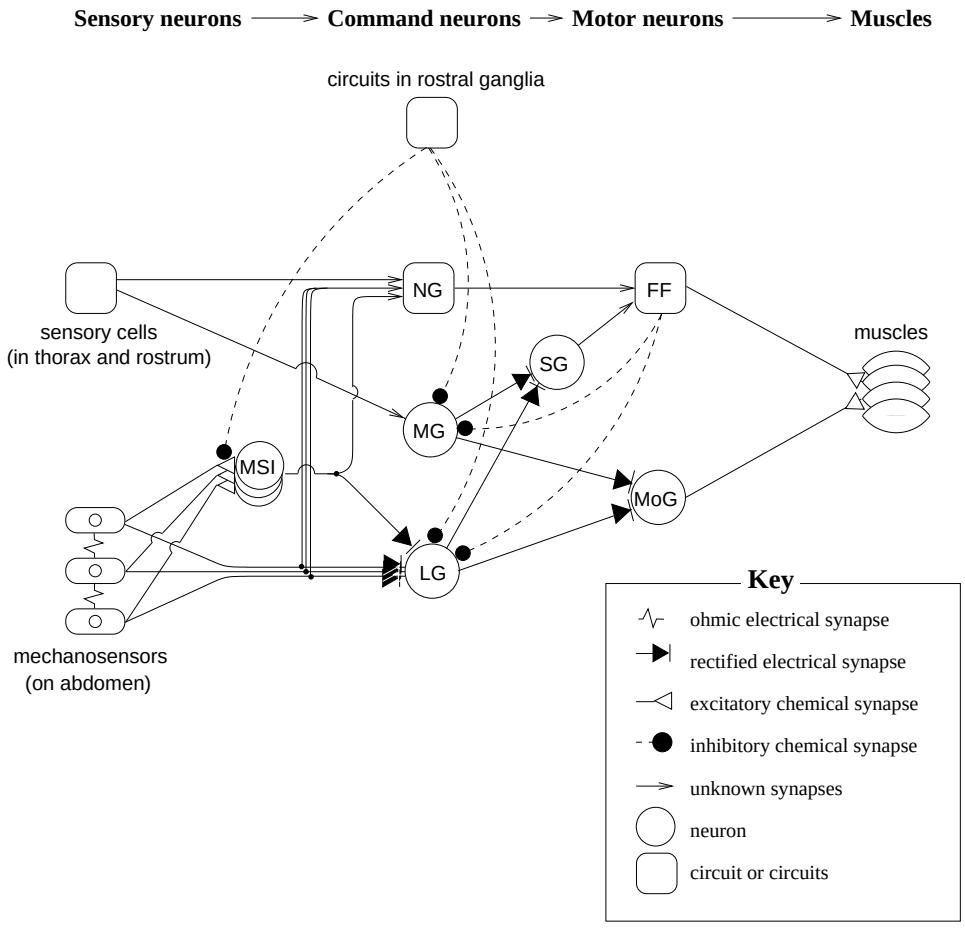

Figure 3.10: The neural circuits controlling the crayfish escape behaviour. The labelled cells are named as follows. MSI: mechanosensory interneurons. NG: non-giant. MG: medial giant. LG: lateral giants. SG: segmental giants. MoG: motor giant. FF: fast flexors.

left half of the segment and one for the right half. Neurons can be grouped in bilaterally homologous pairs, *i.e.* , each neuron is mirrored in the other hemi-ganglion. The giant circuits (the MG circuit and the LG circuit) are called such because each feature a pair of giant axons that run down the crayfish's spinal cord. These axons are large enough to be visible by the naked eye. The medial giants are a single pair of neurons. Their axons run through all the abdominal segments. The lateral giants are in fact a chain of neurons. Each abdominal segment has a pair of lateral giants, one in each hemi-segment. However the lateral giants are tightly connected by electrical junctions in a ladder-like structure causing the lateral giants to, in effect, act as a single neuron. If a spike is generated in one of the lateral giant axons, it will quickly generate spikes in the lateral giants in all the other abdominal segments. The lateral giant circuitry is better studied and better understood than that of the medial giant. However the two circuits are assumed to be similar in nature.(Wine and Krasne, 1982)

The LG and MG are the decision points for the circuits which they control (Krasne and Wine, 1984). For this reason they are are command neurons. Command neurons are single cells that can effect a complete complex behaviour. (Wiersma, 1961) first proposed the concept of command neurons when he discovered that electrical stimulation of single nerve cells elicited a recognisable behaviour. (Kupfermann and Weiss, 1978) extensively evaluated the command neuron concept and proposed the definition:

> a neuron that is both necessary and sufficient for the initiation of a given behaviour

The respective tail-flips that the LG and MG command are generated if and only if they fire. Sensory information feeds into these cells. The LG and MG command neurons integrate the sensory input and determine whether to evoke a short latency escape response. (Fraser, 1982; Krasne and Wine,

1984)

The lateral giant circuit and the medial giant circuit produce highly stereotyped responses of extremely short latency. The giant circuits make no evaluation of the location or the nature of the stimuli. The giants will generate a near identical response for all ranges of triggering stimuli. This lack of stimulus evaluation allows the circuits to have extremely fast responsiveness. For the lateral giant circuit as little as 5 milliseconds elapse from the time of stimuli to when the muscles are potentiated, and another 10 milliseconds elapse to the start of movement.(Krasne and Wine, 1984) The medial giant circuit responds to stimuli on the carapace and front of the crayfish and propels the crayfish backwards. The lateral giant circuit responds to stimuli on the abdomen and propels the crayfish upwards and forwards. (Wine and Krasne, 1982)

The non-giant circuitry supports a wide range of response patterns and their exact parameters of the response vary, probably largely continuously, as a function of stimulating events. The non-giant circuit also controls the escape swimming that follows any tailfan. The non-giant circuit's increased accuracy and greater flexibility comes at the cost of a slower reaction time. The delay between the triggering stimuli and the onset of movement is typically about 50-200 milliseconds, at least 10 times longer than the delay of the giant circuit tail-flips. The non-giant circuit is complex and comparatively little is known about its structure. The non-giant circuit is probably composed of a collection of different neural circuits. (Edwards et al., 1999; Wine and Krasne, 1982; Wine, 1984)

3.3.2 The sensory section of the escape circuit

A LG tail-flip is triggered by sharp taps to the abdomen. The LG receives its sensory input from mechanosensory hairs on the abdomen. The sensory

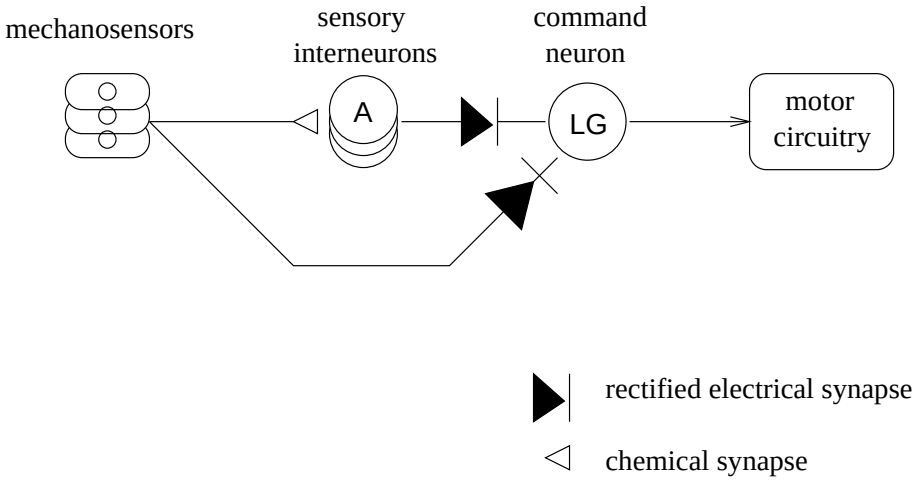

Figure 3.11: The sensory section of the crayfish LG escape circuit.

input is given to the LG through two main pathways: (1) mechanosensors synapse directly onto the LG, these neurons are called the *primary afferents*, and (2) mechanosensors (including some of those which synapse directly on the LG) synapse onto sensory interneuron cells which synapse onto the LG. Figure 3.11 shows a simplified diagram of the sensory portion of the LG circuit.

The direct connection from the mechanosensors is through rectifying electrical synapses (Edwards et al., 1991). Electron microscope imaging shows that there are also chemical synapses between the mechanosensors and the LG in parallel to the electrical synapses (Lee and Krasne, 1993). The role of these chemical synapses remains unclear. Electrophysiological recordings indicate that the electrical synapses play the dominant role (Edwards et al., 1991).

The synapses between the mechanosensors and the sensory interneurons are chemical. Acetylcholine (ACh) is the neurotransmitter used. These connections are depression prone, the strength of the connection diminishes with use. Axons from the interneurons form rectified electrical synapses onto

the LG. For an adult crayfish the sensory interneurons constitute the main synaptic drive onto the LG. The depolarisation needed in the LG at the spike initiation zone is about 8 millivolts. Each sensory interneuron generates an EPSP of up to 0.75 millivolts. Primary afferents monosynaptically connected to the LG produce an EPSP of up to 2 millivolts. It therefore takes a minimum of 10 to 12 simultaneous inputs to cause an LG spike (Wine and Krasne, 1982).

The shape of the LG's response reflects that there are two main sensory pathways onto the LG. The excitation in the LG in response to stimuli occurs in components called α, β and γ (see figure 3.12.) The α and β responses are caused by the primary afferents and the sensory interneurons, respectively. The γ component is a longer lasting smaller depolarisation. The γ component is caused by polysynaptic connections through interneurons. Some of the same sensory interneurons that contribute to the β component also contribute to the γ component. (Edwards et al., 1991)

There are about 1000 sensory hairs on the abdomen, connected to mechanosensory neurons, that provide sensory signals to the LG. Two main kinds of sensory hairs have been identified on the abdomen. There are feathered hairs that are directionally sensitive and respond to low frequency water movements. The most numerous type of hair is smooth. These respond to touch and high frequency water movement. Most afferents terminate in their ganglion, but some run anteriorly for several segments or posteriorly for one segment. (Wine and Krasne, 1982; Wine, 1984)

In the crayfish tailfan, the primary afferents are organised into groups called 'nerves'. Each nerve corresponds to a different sensory area on the tailfan. There are five of these nerves and each nerve is comprised of approximately fifty afferents. The lateral giant dendrites have five main branches. Each nerve connects to one of the main branches of the lateral giant, *i.e.* afferents

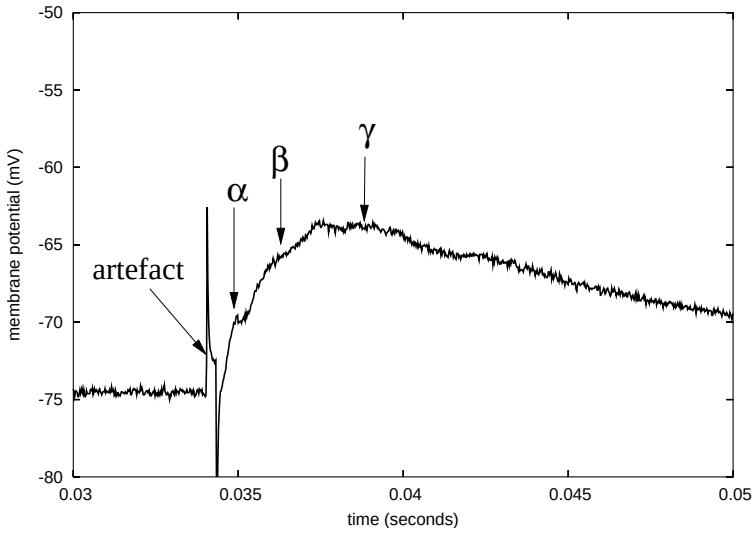

(a)

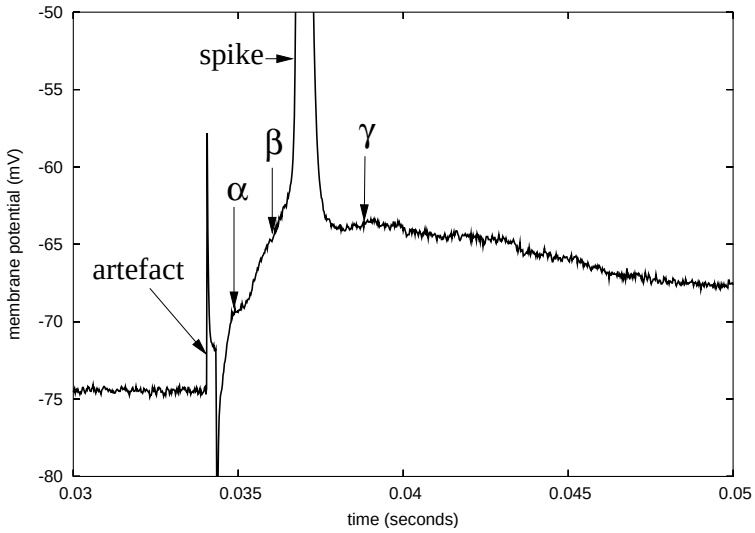

(b)

Figure 3.12: Recordings of the LG in response to (a) a sub-threshold stimulation and a (b) threshold stimulation which causes a spike. The LG response is rises in steps, with the first component caused by the primary afferents (α) and the second rise caused by the sensory interneurons (β). It may be noted that the LG spike rides off the β. The rest of the response is classified as the γ, which is caused by many different polysynaptic interneuron excitations. Recordings from (Herberholz, unpublished).

Afferent diameters (μm)			
	min	max	mean
Small afferents	*N/A*	*N/A*	4.7
Large afferents	*N/A*	*N/A*	12.3
Overall	2	18	5.6

Table 3.1: Summary of the measurements of the dimensions of the afferents, from (Antonsen and Edwards, 2003) (in 4 cm Procamburas Clarkii.) The afferents can be divided into two populations, one with a large diameter and one with a smaller diameter.

from the same nerve by and large synapse to the same main branch on the lateral giant. Electrophysiological experiments indicate that the afferent-to-afferent connections are mostly limited to afferents within the same nerve, *i.e.* afferents are likely to make connections to other afferents within the same nerve but are unlikely to form connections with afferents outside their nerve (Antonsen and Edwards, 2003).

There is a wide variability in the dimensions of the afferents in crayfish(Antonsen and Edwards, 2003). The distribution of the afferent diameters is bimodal. One population of afferents have wide diameters and the afferents in the other population have narrower diameters. Table 3.1 summarises the distribution of the dimensions of the afferents.

The sensory interneurons are not an homogeneous population, the cells respond differently to stimuli. The functional purpose of having some of the sensory excitation going through an interneuronal pathway is not well understood. It is assumed that the interneurons perform some kind of preprocessing such as combining sensory input or abstracting information from it. Two distinct types of sensory interneurons have been identified electrophysiologically. Neurons of one type have passively conducting dendrites and a single spike initiation zone. The other type of interneuron is 'pan-spiking', they have highly electrogenic dendrite membrane and multiple spike initiation zones. The functional significance of having these two different classes

of neurons is not understood. In general, selectivity of the primary afferents is retained by the interneurons and in some cases it is enhanced. Some interneurons respond to breaking water, some respond to touch, some respond to pinching. Some of the interneurons respond to high frequency water vibrations, some respond to low frequency vibrations. Not all sensory interneurons are discriminative, some are broadly tuned. The proportions of the different kinds of interneurons that connect to LG is not known. (Wine and Krasne, 1982; Wine, 1984)

The mechanosensors and the sensory interneurons are not the exclusive preneurons of the LG. Some of the interneurons also excite neurons in other circuits, and not all of the mechanosensors synapse onto the LG. In intracellular recordings of the LG membrane potential after a sensory nerve is stimulated, the arrival of the PSP from the primary afferents, and the PSP from the interneurons can be seen as separate rises in the potential. (Wine and Krasne, 1982)

The primary afferent axons connect together to form a lateral excitatory network (Herberholz et al., 2002). Afferents make direct connections to other afferents through ohmic gap junctions. Each afferent connects to between two and fifteen other afferents. The effect of the network is that stimulated afferents recruit non-stimulated afferents. The recruited afferents amplify the LG response by adding to both the α component and by adding the input to the sensory interneurons, thereby amplifying the β component.

In addition to there being direct connections between afferents, afferents are also indirectly coupled through the LG dendrites (Herberholz et al., 2002). Experiments show that depolarising the LG dendrites increases recruitment and hyperpolarising the LG dendrites decreases recruitment.

The lateral excitatory network shows that decision of whether to escape is not exclusively determined by the LG (Herberholz et al., 2002). The sensory

cells are themselves involved in processing the sensory signals. The escape decision is distributed across the sensory cells.

The LG is very selective in what it reacts to. Only an abrupt onset of a large group of highly synchronous sensory input will cause the LG to spike. Inputs that are a millisecond out of synchrony will fail to spike the LG. Inputs desynchronised by as little as 0.1 milliseconds will produce EPSPs about 25% lower (Edwards et al., 1998). The LG circuit acts as a high-pass filter. The LG is also somewhat sensitive to the direction of stimuli (Wine and Krasne, 1982).

Five potentially important features of the afferent circuitry have been identified that may contribute to its sensitivity to highly synchronous stimuli:

1. the chemical synapses between mechanoreceptors and sensory interneurons show rapid depression (Wine and Krasne, 1982; Wine, 1984).

2. the LG has a high spike initiation threshold and a short membrane time constant. This causes effective summation to require closely spaced input.

3. activity in afferents evokes recurrent presynaptic inhibition onto afferent terminals.

4. excitation of the lateral giant dendrites is followed by inhibition, caused by inhibitory interneurons (Vu et al., 1997). (See section 3.3.5 for description.)

5. the electrical synapses onto the LG are rectifying. (Edwards et al., 1998) showed that the rectifying electrical synapses contribute to coincidence detection in other cells in the crayfish. It is possible that the rectifying electrical synapses also contribute to the coincidence detection that is apparent in the LG.

The MG receives sensory signals from head and thorax. The stimuli that triggers a MG tail-flip is not entirely understood. It is known to be triggered by rapid moving visual stimuli, taps to the antennae and sharp taps to the carapace. However in recent studies of crayfish interaction (Herberholz et al., 2001), where a high speed video camera was used to record the interactions, for a large number of MG tail-flips observed the triggering stimulus was not identifiable. The MG neurons have a single input region and spike-initiation zone. Intracellular recordings of the MG response to stimulus show an early stable component and a later decrementing component. This is similar to the α and β components of the LG response, suggesting that the MG circuit has the same basic organisation at the LG circuit (Wine and Krasne, 1982).

3.3.3 The motor section of the escape circuit

Once triggered, escape has highest priority of any behaviour pattern in the animal's repertoire. All other competing patterns are inhibited or overridden. If spiked, the MG and the LG activate motorneurons in the abdomen to unleash a tail-flip. The MG and LG also send inhibitory signals to many other neurons across the crayfish's central nervous system.

With respect to the sensory part of the escape circuit, the LG and MG have completely separate pathways. On the motor side of the circuit the LG and MG share a lot of common circuitry.

A tail-flip involves the use of two groups of muscles: the flexion muscles and the extensors.

At the neurological level, a giant circuit tail-flip can be described as five steps (Wine, 1984):

1. sensory triggering of the flexion command axons

2. excitation of flexion muscles

3. short-lasting inhibition of extensor system

4. delayed, long-lasting inhibition of flexor system

5. delayed feedback excitation of the extensors.

A giant generated tail-flip is usually followed by escape swimming. This is a series of tail-flips that continue the escape motion. The escape swimming is controlled by the non-giant circuit (Krasne and Wine, 1984).

The MG and the LG command the excitation of the flexor muscles through two main pathways. The main path is through the motor giant (MoG). There is a secondary pathway through the fast flexor motorneurons (FF). Both pathways are innervated in parallel. The LG and MG each make a strong connection onto the motor giant through a rectified electrical junction, called the giant motor synapse (GMS). The motor giant is believed to be activated exclusively by the giants. No other circuits excite this neuron. The motor giant projects onto almost all of the phasic flexion muscles in its hemi-segment. A single impulse through this pathway is sufficient to produce strong contractions of the flexors, which causes a tail-flip. The medial giant excites the motor giants in every segment, whereas the LG only excites the motor giants in the anterior segments. This causes different tail motions for the LG and MG tail-flips. For an LG tail-flip the end of the tail remains straight, pushing the crayfish forward and up. For an MG tail-flip the tail curls completely, dragging the crayfish rapidly backwards. (Krasne and Wine, 1984)

The LG and MG also use the fast flexor (FF) motorneurons. The role the FF motorneurons play in a giant tail-flip is unclear. The LG and MG connect to the FF motorneurons through a premotor neuron called the segmental giant (SG) and also make a weak direct connection. The SG excites FF

motorneurons asymmetrically. Most of FF motorneurons have asymmetric dendritic domains. This possibly allows for direction sensitivity during non-giant escape. The role of the segmental giant is also unclear. Its possible roles include reducing the electrical load on the motor giant and giving a single point of control for activating the FF motorneurons. (Krasne and Wine, 1984)

Before a flexion is completed, all elements of flexion circuit is inhibited by long lasting IPSPs. These are triggered by the command signal and persist for duration of flexion re-extension (Wine, 1984).

The extensor muscles are not directly activated by the command neurons. The extensor muscles receive no excitation directed from the LG or MG during giant tail-flips. Instead, the extensor muscles are believed to be driven by sensory feedback. The sensory feedback is depression prone. During non-giant tail-flips, the sensory feedback is believed to be inhibited, to allow the non-giant circuitry to directly drives the extensor muscles (Wine and Krasne, 1982).

The LG and MG activate many other effects in addition to activating the motorneurons in the tail. The giants send massive inhibitory signals across the nervous system of the crayfish. The extensor-muscle stretch receptors are inhibited. The MG and the LG both inhibit themselves, each other and the motor giant. This ensures there will be only one or a few command neuron spikes and only a single motor giant spike. This prevents a continuous cycle of escape circuit reactivation (called recurrent inhibition, described further in section 3.4.2.) The primary afferents and mechanosensory interneurons are also inhibited. The afferents are inhibited presynaptically. This contributes to preventing re-excitation of the circuit and also prevents use-dependent habituation of mechanosensors (see section 3.4.3.) The LG and MG have also been shown to promote the limbs, they synapse directly

onto limb promoter neurons (Edwards et al., 1999).

Electrical synapses are used throughout the motor circuitry. Electrical transmission is fast and reliable making it suited for reflex decision making. Chemical transmission is only used by the motorneurons to activate the muscles.

It can be noted that the non-giant circuitry is recruited by the giant system but not vice versa. The motor giants are highly a specialised output of the giant system. The non-giant system outputs only to the fast flexors. The basic circuit for conducting tail-flips is the non-giant circuit. The non-giant circuit recruits the fast flexor (FF) motorneurons to excite the flexor muscles. The fast flexor system controls tail-flips and is completely separate to neurons that control slow movement of the tail. There are between 5 and 8 fast flexor motorneurons per hemi-ganglion (depending on the segment), each projects onto a separate section of the flexor muscles. The non-giant tail-flips are produced entirely by the FF motorneurons. The motor giants are strongly inhibited prior to and during non-giant tail-flips. (Krasne and Wine, 1984; Wine and Krasne, 1982)

3.3.4 Rectifying electrical gap junctions

Electrical synapses are used throughout the crayfish escape circuits. Experimental findings indicate that the majority of these are rectifying electrical synapses. Recent evidence suggests that the rectifying synapses may play a key role in the operation and even the decision making in the circuit. Rectifying electrical synapses have the obvious advantage of giving fast unidirectional transmission, but recent evidence indicates they may also play other more subtle roles including acting as a mechanism for coincidence detection, for synchronising inputs and as a modulation mechanism.

A rectifying gap junction is an electrical gap junction that has variable conductance that increases with the voltage difference between the pre-synaptic terminal and post-synaptic terminal. The conductivity is not dependent on the absolute voltage of either the pre-synaptic cell or the post-synaptic cell but the voltage difference between the cells. Electrical synapses that rectify were first discovered in the giant motor synapse of the crayfish (Furshpan and Potter, 1959). An experiment using voltage-clamp analysis found the conductance of the giant motor synapse to be governed by equation 3.1 (Giaume et al., 1987). The gap junctions between the mechanosensory interneurons and the lateral giants were found to be described by the same equation (Edwards et al., 1991).

$$g_{\mathrm{j}} = g_{\min} + \frac{g_{\max} - g_{\min}}{1 + e^{-A(V_{\mathrm{pre}} - V_{\mathrm{post}} - V_0)}} \tag{3.1}$$

Where g_{j} is the conductance of the junction, $g_{\min}$ is the minimal transynaptic conductance, $g_{\max}$ is the maximal transynaptic conductance, A and V_0 are constants.

Equation 3.1 describes the steady state conductance of a rectifying gap junction (Giaume et al., 1987). The conductance of rectifying gap junctions increases with a temperature sensitive delay. At temperatures typical of the crayfish's environment the conductance change occurs very rapidly. At 20 degrees Celsius, junctional currents across the giant motor synapse were already found to be constant 1 millisecond after step changes to the voltage across the junction was applied (Giaume et al., 1987; Edwards et al., 1999).

The physical mechanisms that operate rectifying gap junctions are poorly understood. There is no accepted theory of how they work.

Rectifying gap junctions are used throughout the LG circuit. Experimental evidence suggests that all the electrical junctions from sensory neurons onto

LG rectify (Edwards et al., 1991), *i.e.* , the monosynaptic primary afferent connections, the disynaptic connections through the sensory interneurons and the polysynaptic connections (the γ components.) Rectifying gap junctions are also extensively used throughout the motor part of the LG circuit. The giant motor synapse, that connects the LG to the motor giant, rectifies in the forward direction. The synapses along the secondary pathway of motor excitation also rectify. The synapses between the LG and the segmental giant and the segmental giant and the fast flexor motorneurons are forward rectifying (Heitler et al., 1991). Interestingly, the segmental giants on opposing sides of the mid-line may also be connected by rectifying junctions. Depolarising current injected into a segmental giant neuron spreads to its homologue more effectively than hyperpolarising current (Heitler et al., 1991). This suggests that the segmental giants may be connected by a pair of oppositely directed rectifying junctions, possibly via the fast flexor motorneurons. (This hypothesis was successfully tested in a computer simulation (Heitler et al., 1991).) It would seem that the general rule for the LG circuit is that the electrical synapses rectify. The major exceptions are the electrical synapses that connect an LG neuron to the LG neurons in the adjacent segments and to the LG neuron in the opposite hemi-ganglion in the same segment. These synapses do not rectify. (Edwards et al., 1991)

The synaptic connections between the neurons in the afferents and the LG are forward biased. The synaptic connections in the motor part of the circuit are reversed biased. The LG and MG have very hyperpolarised resting potentials: LG has resting potential been measured to be about -85 millivolts and MG has a similar resting potential. The primary afferents and sensory interneurons have resting potentials in the range -65 millivolts to -75millivolts (Herberholz, personal communication), causing the rectifying synapses onto the LG to rest in a state of forward bias. The neurons in the motor part of the LG circuit also have resting potentials above LG. The mo-

tor giant has a resting potential of about -73 millivolts, the segmental giant a resting potential of about -82 millivolts and the fast flexor motorneurons a resting potential of about -76 millivolts. This causes all the rectifying synapses in the motor part of the circuit to be in a state of reverse bias when at rest. The generality of the purpose of the reverse biasing is not known. One effect is to increase the safety factor by which the synapse prevents back-propagation. The most critical synapses in the motor part of the escape circuits, the LG and MG synapses to the motor giant, have the greatest transjunctional potential at rest. (Heitler et al., 1991)

The rectifying gap junctions are thought to perform a number of functions in the operation of the LG circuit. Firstly, being electrical synapses, they provide a mechanism for fast and reliable transmission.

When two neurons are connected by a non-rectifying electrical synapse the two cells are permanently coupled together. That is, any depolarisation or hyperpolarisation in one cell will quickly spread to the other cell, and vice versa, (as is the case with the LG neurons.) When the synapses are rectifying the cells are only coupled under certain conditions, that is when the transynaptic potential is at the right level. This allows for some potentially interesting effects.

Rectifying gap junctions provide unidirectional transmission. This prevents unwanted antidromic transmission. The rectifying junctions prevent spikes in the LG exciting the afferents, and help prevent spikes in the motor giant from re-exciting the LG (Edwards et al., 1999).

Rectifying gap junctions allow the LG to be excited by a small subset of its inputs. When the LG is being excited only the synapses from the active inputs are in a state of high conductance. The synapses to the non-active inputs are in a state of low conductance. Thus the inward currents to LG, in response to active inputs, are prevented from being antidromically shunted

down the low resistance contacts of the non-active inputs, as would happen if the synapses did not rectify (Edwards et al., 1991).

The rectifying gap junctions may also act as a mechanism for coincidence detection in the LG. Rectifying synapses select for phasic input, because slower synaptic currents, created by depolarising after potentials are more severely attenuated than fast currents. (Edwards et al., 1991) It has been shown with computer models that the intrinsic characteristics of rectifying junctions attenuate responses to stimuli that occur within as little of 0.1 milliseconds out of synchrony (Edwards et al., 1998). The model featured a target neuron receiving input from two afferent neurons through rectifying electrical synapses. If input is simultaneous both synapses will be highly conductive, causing the EPSP that each afferent neuron creates to sum. If inputs are slightly asynchronous, then the EPSP caused by the second synapse will be reduced because: (1) the first input will have partially depolarised the target neuron, therefore the transynaptic potential across the second synapse will be less, reducing the synaptic conductance (see equation 3.1), (2) the lower transynaptic potential reduces the driving force across the synapse, and (3) at the arrival of the second input the first synapse will still be in a state of higher conductance creating a shunt for the inward current of the second input. Greater delays between the first and second input reduce the inhibitory effect of the first input onto the second. It is hypothesised that this effect takes place in the summing of afferent input in the LG, where early inputs partially depolarise the dendritic tree, thereby reducing the effectiveness of later inputs. (Edwards et al., 1998; Edwards et al., 1999)

The impact of the rectifying junctions on coincidence detection is likely to be greater for inputs that are spatially closer. Depolarisation is greater at dendrites than at integrating segment. Out of phase inputs from the same source will be more severely attenuated than out of phase inputs synapsing

on different parts of the synaptic tree. The depolarisation in the dendrites at the synapse point of the late input will be greater in the former case. This would cause the rectifying junctions to have a greater impact in reducing the effectiveness of the late input. (Edwards et al., 1991)

The rectifying gap junctions provide a mechanism for modulating the response of a cell. Changing the depolarisation level of a cell modulates its response to inputs. Changing a cell's membrane potential changes the transynaptic potential across its rectifying junctions which changes their conductance. It has been experimentally shown that minor changes to the resting potential of the LG can have a major effect in increasing or reducing the PSPs in response to input. Depolarising the post-synaptic cell reduces transmission through the rectifying junctions by (1) reducing the conductance across the synapse, (2) reducing the amount of time the synapse is in a high conductance state, and (3) reducing the driving force across the synapse. This effect has been seen to be used in the LG and probably exists in other other parts of the escape circuits as well. The rectifying junctions cause depolarisation in the LG caused by the γ component to have a inhibitory effect on late afferent input. The effect can also be seen in regard to depolarising inhibitory post synaptic potentials (d-IPSPs), which are used widely in the escape circuits. d-IPSPs are generated through chemical synapses using GABA as the neurotransmitter. GABA's main effect is to open the chloride channels, which has a inhibitory effect on excitatory inputs. However chloride's resting potential is above the LG's resting potential, so opening these channels has the effect of raising LG's membrane potential by about 10 millivolts. This reduces the transmission through rectifying junctions. Hence with rectifying junctions, d-IPSPs have the duel effect of reducing transmission from the afferents and neutralising the depolarisation in the cell. d-IPSPs are transmitted to the LG after an initial excitation and during tonic inhibition. d-IPSPs seem to be widely used throughout the escape

circuit. (Edwards, 1991; Edwards et al., 1999)

3.3.5 Selecting for phasic stimuli

The lateral giants are highly sensitive to strong stimuli of abrupt onset (called phasic input) and insensitive to gradually increasing continuous inputs. A number of characteristics of the circuit contribute to this effect.

Firstly, charge is not allowed to gradually build up at the spike initiation zone. The LG has a short membrane time constant and a low input resistance causing incoming charge quickly redistributes throughout the cell's large structure and leak out. (Edwards et al., 1991)

Second, the sensory neurons also invoke inhibiting neurons, which send a delayed inhibition to the LG (Vu et al., 1993; Vu et al., 1997). This is called post excitatory inhibition (PEI). Post-excitatory inhibition is activated during falling phase of beta-component. The inhibiting neurons synapse onto the outer regions of the dendritic tree, close to where the excitatory input from the sensory neurons (which also excited the inhibiting neurons) synapse. This has the effect of making the inhibition location specific (Vu et al., 1993). The excitatory inputs from adjacent mechanosensors synapse at nearby locations on the dendritic tree of the LG. Hence if a mechanosensor at the back of the abdomen is activated, and a short time later mechanosensors at a nearby location is activated together with one at the front of the abdomen, then the signal from the back of the abdomen will be inhibited much more than the signal from the front. The major role ascribed to post excitatory inhibition is to discount gradually increasing stimuli.

The post-excitatory inhibition is mediated by GABA and activates chloride channels. The opening of the chloride channels causes excitatory input to be shunted out of the cell. Due to chloride's high reversal potential, the inhibi-

tion is also depolarising. The depolarisation contributes to favouring phasic input over gradually increasing tonic input. The depolarisation reverse biases the rectifying synapse, this has greater effect on slower currents than phasic currents. The depolarisation increases the resistance of synapses. The depolarisation may also deactivate the sodium channels, thereby increasing the LG threshold. This has been shown to occur in the MG (Edwards et al., 1991).

Finally, rectifying junctions have been shown to behave as coincidence detectors under certain conditions (as was described in section 3.3.4.)

3.4 Modulation of the escape circuit

The escape behaviour is highly modulated. The thresholds of the escape circuits are adaptive to the circumstances of the animal. The escape may be either inhibited or sensitised. The triggering of escape behaviour modulated by a wide range of influences. An escape tail-flip is an expensive action to perform and it interferes with the animal's other behaviours. The modulation of the escape thresholds provides the crayfish mechanisms for minimising the number of occurrences of an escape tail-flip in non-advantageous circumstances.

In this section we first describe the conditions which affect the escape behaviour threshold. Later sections give more details and describe what is known of the mechanisms that modulate these changes.

3.4.1 Conditions that modulate escape behaviour

The triggering of escape is adjusted by many factors and circumstances:

- directly after a tail-flip has been invoked many parts of the circuit are

massively inhibited to prevent repeated triggers of the tail-flip, this is called 'recurrent inhibition',

- the sensory input to the escape circuits habituate,
- when a crayfish is performing another behaviour the escape threshold may be modulated,
- the social status of the crayfish modulates the escape threshold, and
- the physical condition of the animal changes the threshold.

Escape is not the only defensive behaviour of the crayfish. When faced with a threat crayfish may also opt to walk backwards or go into a defence posture, in which it raises its claws and faces the direction of the source of danger. The LG tail-flip is inhibited when the crayfish is backward walking or in a defence posture (Beall et al., 1991). Preliminary evidence suggests LG escape is centrally inhibited during stimulation of interneurons that initiate backward walking motor program (Beall et al., 1991). Nothing is known of how tail-flip is inhibited during defense posture.

When the crayfish is moving its legs or is swimming the LG escape circuit is inhibited (Fricke et al., 1982). The inhibition is caused at least in part by proprioceptive hairs that pre-synaptically inhibit transmission at first central synapse.

Feeding causes the LG escape response to be suppressed. The suppression of escape appears to be related to the consumption of food rather than searching for food. Animals that are actively searching for inaccessible food do not show elevated thresholds. The non-giant escape is suppressed when the crayfish is feeding on a large piece of food, but is enhanced for when the crayfish is feeding on small pieces of food, with which the animal can escape. For the LG escape, size of food had no effect escape threshold. (Krasne and

Lee, 1988)

The rise in the sensory threshold required to trigger a tail-flip during feeding is reliable. The suppression appears to be result of reduced transmission to the LG from the primary afferents and sensory interneurons. Other parts of the circuit have been shown to be unaffected. There are changes in the thresholds of the sensory interneurons or in the transmission between the LG and the motor circuit neurons. It is unknown what causes the reduced transmission of sensory input to the LG. Control may be mediated by a serotonergic pathway. (Krasne and Lee, 1988)

When a crayfish is restrained the LG escape response is inhibited.

The escape threshold is modulated in response to the circumstances of the animal. When the animal is required to be more bold, the escape threshold may be raised. When the crayfish is hungry and searching for food the escape thresholds are raised (Wine and Krasne, 1982). Females in berry (carrying eggs) also have a higher escape threshold.

In other circumstances, where the crayfish is vulnerable, the escape threshold is lowered. Freshly molted animals have very low escape thresholds (Wine and Krasne, 1982).

The escape response changes as the animal grows. Young crayfish invoke the escape very frequently. The escape threshold is very fickle for small crayfish. As the animal grows the escape threshold rises.

There are many different mechanisms that facilitate the modulation of the escape circuits. In comparison to the rest of the circuit the modulatory mechanisms are less understood and the subject of ongoing research. Most of the modulating interneurons have not been identified (nor is it likely that they soon will be.) Some modulation is mediated by pathways coming from anterior ganglia (*i.e.* the brain), other kinds of modulation are more local in

origin. (Wine and Krasne, 1982)

3.4.2 Inhibition

When a tail-flip is invoked, the command circuits are massively inhibited (Roberts, 1968), this is called recurrent inhibition. Recurrent inhibition prevents a new tail-flip movement starting while a tail-flip is in progress. Recurrent inhibition prevents the massive sensory stimuli the animal experiences while performing a tail-flip from perpetually re-exciting the command neurons. Such continual re-initiation would obviously interfere with performing a tail-flip. Recurrent inhibition directly affects LG but also other sites.

Following an LG spike, the LG and the motor giant are inhibited. This ensures there are only one or a few LG spikes and only ever one motor giant spike. The inhibition at the LG takes place close to the spike initiation zone (Vu and Krasne, 1993b). The proxal placement causes the inhibition to always be effective. Even very strong stimuli are unable to cause the LG to spike. Large depolarisations at the dendrites are ineffective at the spike initiation zone as they are shunted out of the cell. The inhibition is mediated by GABA and causes chloride channels to open (Roberts, 1968). Due to chloride's reversal potential being higher than the LG's resting potential, the inhibition is depolarising.

An LG spike also causes inhibition at the primary afferents. The inhibition is again GABA-ergic and activates chloride channels, which causes a depolarisation. It is therefore referred to as the primary afferent depolarisation (PAD). The PAD inhibition is long lasting (in the order of tens of milliseconds) and typically causes a 3-5 millivolt depolarisation (Kennedy et al., 1980). The PAD inhibition descends from inhibitory interneurons which are excited through poly-synaptic pathways which must be ultimately be driven

by the giant neurons (Kirk and Wine, 1984; Kirk, 1985).

3.4.3 Habituation

The mechanosensory neurons that excite the LG reduce their input onto LG in response to repeated stimulation. This is known as habituation.

Repeated tapping on the abdomen reduces the probability of an LG escape response. In experiments it has been shown that tapping the abdomen once per minute can drive the responsiveness to near zero. Three hours after the tapping has ceased there is little recovery in the response. Habituation occurs even if the nerve connection to the thorax is cut. This points to the source of habituation as being local rather than descending from higher centres in the nervous system. (Wine and Krasne, 1982)

Habituation is understood to be caused, at least in part, by a diminished release of neurotransmitters from the pre-synaptic terminal between the mechanosensors and the sensory interneurons (Zucker, 1972b). The chemical synapses between the sensory afferents and the interneurons are pre-synaptically depressed.

Habituation is specific to the sensory afferents that are repeatedly stimulated. Other sensory afferents are not affected.

Habituation helps filters out sustained stimuli. This is a sensory adjustment so that continuous sensation, which is unlikely to be caused by an approaching predator, is ignored in making the escape decision. Due to the point of depression being at the synapses, the output of the sensory afferents to other systems is unaffected (Wine and Krasne, 1982).

3.4.4 Neural mechanisms of modulation

Three neurotransmitters have been identified to exhibit a modulating effect on the LG escape circuit: GABA, serotonin and octopamine.

The escape response is known to be reliably inhibited during feeding and restraint. This modulation is called 'tonic inhibition' (Krasne and Wine, 1975). Tonic inhibition originates from rostral ganglia (Krasne and Wine, 1975; Vu and Krasne, 1993b). A GABA channel from the rostral ganglia to the abdominal segments mediates tonic inhibition (Vu and Krasne, 1993a).

Tonic inhibition takes effect at the LG dendrites, far from the spike initiation zone (Vu and Krasne, 1993b). GABA receptors open chloride channels in the LG dendrites which increases the membrane conductance at the dendrites (Vu and Krasne, 1993a). The effect of this is that currents, caused by excitatory input, are shunted out of the cell, greatly reducing their effect. Both the α and β inputs onto the LG are depressed by tonic inhibition (Vu and Krasne, 1993a). Tonic inhibition does not directly affect the sensory afferents or the motor neurons (Krasne and Wine, 1975).

The placement of tonic inhibition in the LG dendrites, as opposed to close to the spike initiation zone, may serve two functions. The distal placement of tonic inhibition allows it to be overridden if there is a large enough stimulus. Intuitively, this is a good property, as during feeding false escapes are especially undesirable, but the animal should still be able to escape when seriously endangered. In contrast, recurrent inhibition, which should never be overridden, occurs close to the spike initiation zone. The distalness of tonic inhibition could also be used to selectively suppress parts of the sensory system (Vu and Krasne, 1993b). Different sensory neurons synapse at different places in the LG dendritic tree. Sensory neurons that synapse close to the point of inhibition would have their input severely inhibited. Input

from other sensory neurons would be less affected.

Animals can be provoked into tonically inhibiting by giving them small object to hold in their chelipeds (Vu and Krasne, 1993b). This supports the hypothesis that the closing of the claws is the trigger for tonic inhibition (Wine and Krasne, 1982).

Tonic inhibition has also been shown to affect learning. Habituation is thought to take place at the synapses between the primary afferents and the sensory interneurons. It was supposed that this was mediated entirely by local synaptic plasticity. However in experiments, animals that had their GABA-ergic channel cut, did not habituate as much as intact animals (Krasne and Teshiba, 1995). Tonic inhibition was found to be responsible for a large part of habituation. It is possible that tonic inhibition is used as a transient step, while the synapses are being modified. This experiment showed that learning is caused by the events at higher centres of the nervous system rather than exclusively by local plasticity.

Serotonin has also been supposed as the neurotransmitter effecting tonic inhibition (Glanzman and Krasne, 1983). Serotonin receptors are known to occur in the LG and applying serotonin to LG also almost mimics tonic inhibition exactly, however other experiments do not support serotonin as the mediator (Vu and Krasne, 1993a). It is likely that the serotonin receptors in the LG are used for something. The hypotheses for their purpose include: being used for long term regulation of the LG excitability, or perhaps in some way linked with crayfish's ability to inhibit more strongly, more rapidly over repeated restraint sessions (Vu and Krasne, 1993a).

Serotonin has been shown to have a number of effects on the LG escape circuit. The effects also depend on the animal's social status, as is described in section 3.4.5.

Serotonin inhibits the LG escape response. When serotonin was injected into the crayfish's ventral artery, the EPSPs recorded in the LG were reduced (Glanzman and Krasne, 1983). Serotonin only affected the sensory interneuron (β) component of the LG input. EPSPs caused by the direct mechanosensory inputs were not affected. Serotonin raises the thresholds of the sensory interneurons. This suggests that serotonin causes synaptic transmission to be depressed between sensory interneurons and LG or between afferents and sensory neurons. Serotonin does not alter the rate of habituation.

Octopamine has been shown to enhance the excitability of LG escape response. In experiments, injecting octopamine into the crayfish's ventral artery, caused the EPSPs recorded in the LG to increase (Glanzman and Krasne, 1983). Octopamine mainly effects interneuron (β) component of LG EPSP. The direct mechanosensory (α) component is only slightly raised (Glanzman and Krasne, 1983). Octopamine reduces the thresholds of the sensory interneurons.

Serotonin and octopamine may affect other parts of the escape system. Serotonin may enhance tonic flexion. Octopamine may depress tonic flexion and enhance tonic extension. (Glanzman and Krasne, 1983)

The modulatory effect of serotonin differs depending on the animals social status (see section 3.4.5) and on the way it is applied. Serotonin can be either inhibitory or excitatory depending on the concentration and duration of the injection (Teshiba et al., 2001).

Serotonin and octopamine have been shown to modulate the threshold of the escape system. We know that the animal does modulate its escape threshold. The hypothesis is that serotonin is used by the animal to mediate suppression of LG escape and octopamine may mediate sensitisation of LG escape (Glanzman and Krasne, 1983). There is no direct evidence however

that the crayfish releases octopamine to mediate this effect.

3.4.5 Dominance and serotonin

Serotonin has been linked with the aggression and dominance related behaviours in crayfish. Injections of serotonin in crayfish and lobsters evokes a dominant posture and promotes aggression (Livingstone et al., 1980; Kravitz, 1988).

The effect of serotonin varies according to social status of animal. Serotonin causes a rise in the LG excitability in subordinates but inhibits the LG response in dominant or isolated animals (Yeh et al., 1997). In freely behaving animals, subordinates show a substantial fall in LG excitability while they are engaged in fighting (Krasne et al., 1997). In contrast, dominant animals experience only a slight fall in LG excitability (Krasne et al., 1997). It is hypothesised that these changes in the LG excitability are mediated by serotonin, although there is no direct evidence to this effect. Outside of agonistic encounters, the LG threshold appears to be independent of its social status (Krasne et al., 1997). While a subordinate animal's LG escape is repressed during fights, the excitability of the non-giant escape circuit is enhanced. The subordinate animal frequently invokes non-giant tail-flips (Yeh et al., 1997; Herberholz et al., 2001).

The different effects caused by serotonin is due to there being different populations of receptors. Experiments suggest that there are at least two different kinds of serotonin receptors acting on the LG circuit (Yeh et al., 1996; Yeh et al., 1997). One receptor type inhibits the LG response, and the other enhances it. The different responses to serotonin in dominants and subordinates may be due to there being different compositions of the quantities of the different receptor types.

The story gets more complicated. The prior social experiences of a crayfish affects its response to serotonin (Yeh et al., 1996; Yeh et al., 1997). The inhibitory effect of serotonin experienced by subordinates is reversible. Subordinates who later become dominants experience enhancement to the LG response when serotonin is applied. The enhancement affect of serotonin in dominants is irreversible. Dominants who later become subordinates still respond with enhancement when serotonin is applied.

That the LG escape response is inhibited in subordinates and sensitised in dominants seems counter-intuitive. A possible explanation is that it is a consequence of competition between giant escape and non-giant escape. In fighting, subordinates evoke tail-flips far more often than dominants. The majority of these are non-giant tail-flips (Herberholz et al., 2001). The greater flexibility and precision of the non-giant escape circuit makes it the seemingly preferable choice for an animal that knows it will be doing a lot of tail-flipping. The LG escape could be inhibited to promote the use of non-giant escape (Krasne et al., 1997). In contrast, dominants are able to go about their business, anticipating that they will not be disturbed, and rely on the giant circuits to respond to unexpected danger (Krasne et al., 1997).

3.4.6 The development of the escape circuit

The crayfish escape circuit changes as the animal develops. During the larval stage, lobsters do not elicit LG or MG tail-flips (Jackson and MacMillan, 2000), they use only the non-giant circuitry. At the next stage of development however, the post-larval stage, lobsters and crayfish show the same ranges of responses as adults. Lobster larvae do also not appear to habituate their escape response (Jackson and MacMillan, 2000).

In observations of crayfish and lobsters, juvenile crayfish tail-flip very fre-

quently (Lang et al., 1977; Jackson and MacMillan, 2000). As the animal grows tail-flips are used less frequently. Adult crayfish invoke tail-flips much more rarely. Adults are more inclined to defend themselves from threats (Lang et al., 1977).

The reduced frequency of tail-flipping that occurs as a crayfish grows, coincides with an increase in the LG threshold. The rise in threshold can be accounted for by the growth of the LG (Hill et al., 1994; Edwards et al., 1994b). A small cell requires less current to charge it up to firing threshold. The growth of the LG is nonuniform. The dimensions of the axon increase at a greater rate than the dendrites. The effect is, the input resistance of the LG is much lower an adult crayfish than in a juvenile. The strength of the input onto the LG does not increase at a rate that compensates for the cell's larger electrical load. As a consequence, the stimulus strength required to invoke an LG tail-flip increases as the crayfish gets older. In juvenile crayfish, the α input alone, caused by the direct connections from the mechanosensors, is sufficient to fire the LG. In adult crayfish, the α input alone will never fire the LG. An LG spike always rides off the β component, caused by the sensory interneurons (Hill et al., 1994; Edwards et al., 1994b).

3.4.7 Priming the non-giant circuit

It appears that the non-giant circuit is able to be primed to respond faster. In experiments, the response of the non-giant circuit was almost as fast as the giant circuits if the crayfish was kept in the presence of a predator before it attacked (Herberholz et al., 2004). This may be explained by the crayfish sensing that it is in the presence of a predator and that an attack is therefore likely. Accordingly, it primes the non-giant circuit so that it is able to respond faster. This may involve pre-calculating the direction of the escape. Priming the non-giant circuit would allow the combination of the

non-giant circuit's accuracy with giant speed. However, in the context of tests conducted in the experiments, non-giant tail-flips were less effective in avoiding capture than either LG or MG tail-flips.

3.4.8 LG potentiation

It has also been shown that the LG can become sensitised in response to stimulation.

(Krasne and Glanzman, 1986) showed that strong AC shocks to the abdomen can cause sensitisation in the LG escape response. A single traumatic shock to the abdomen, reliably causes a fall in the LG threshold. The threshold typically recovers in about one hour. Repeatedly applied traumatic shocks caused greater drops in the LG threshold and the sensitisation was long lasting. A full threshold recovery took up to 24 hours. Crayfish with initially higher thresholds tended to show greater threshold drops. The sensitisation of the LG appears to be caused by an enhancement of the chemical transmission between the mechanosensors and the sensory interneurons (Miller et al., 1987). Traumatic shocks can also cause the sensory interneurons to be sensitised. When the nerve cord between to the thorax and abdomen was cut, there was a reduced amount of sensitisation in LG and the sensory interneurons (Krasne and Glanzman, 1986). This suggests that sensitisation is mediated through a neural pathway descending from more rostral ganglia. Octopamine may be the neurotransmitter responsible for mediating sensitisation. The functional purpose of sensitisation may be to compensate for habituation.

The LG escape response can also undergo long term potentiation (LTP) (Yeh et al., 2002). The synaptic strength at the rectifying junctions between the primary afferents and the LG increases in response to rhythmic sub-threshold stimulation. Over the course of an hour, such stimulation can

cause the EPSPs in LG to increase by 60%. The potentiation persists for six hours. LTP requires coincident pre-synaptic and post-synaptic depolarisation.

3.5 Discussion of the crayfish escape circuit architecture

Six key features are apparent of the architecture of the crayfish escape circuit.

Firstly, the crayfish escape circuit is optimised for a short latency response. The physical characteristics of components of which the circuit is constructed cause the response to be as fast as possible:

- The chain of neurons between the stimuli and response is very short. There are only four synapses between the mechanosensors and the muscles.
- Electrical synapses are used throughout the escape circuit. Electrical synapsing is extremely fast and reliable. Electrical circuits are used wherever possible. Chemical synapses are used only between the mechanosensors and the sensory interneurons and at the neuromuscular junctions.
- Giant axons carry the signals of the escape command neurons. Wide axons facilitate fast, reliable conduction.

Secondly, in the MG and LG circuits, the decision of whether to escape is made at a single point of control. Within their respective circuits, the LG and MG command neurons exclusively determine whether or not an escape tail-flip will be invoked.

Thirdly, the test criteria that the giant escape circuits use to determine whether to invoke a tail-flip is relatively simple. The LG circuit continuously filters its sensory input for an abrupt massive stimulus. There may also be a spatial component to the criteria (Wine and Krasne, 1982). The test is calculated within chains of three neurons and two synapses.

Fourthly, there are multiple escape circuits. The different circuits serve different roles. Two circuits, the LG and the MG, evoke highly stereotyped responses but they are extremely fast in their response. The non-giant circuit is much more versatile in the kind of responses it commands, but it is slower. The escape responses exist within the context of the crayfish' overall defensive strategies. Other defensive behaviours within the crayfish' behavioural repertoire include backward walking, swimming and defence posturing. We can think of these behaviours existing within a hierarchy. The more complex defensive behaviours offer greater versatility and are more intelligent in their response. It is preferable to identify a source of danger early, and tactically withdraw to the burrow than to rapidly shoot backwards and end up who knows where. The giant escape circuits fulfill the role of the last resort. When ambushed by a predator and faced with death it is better to respond quickly than smartly. An escape tail-flip is the final role of the dice before submitting to capture.

Fifthly, the escape response overrides any other behaviour that the crayfish is engaged in. Once invoked, escape takes the highest precedent and completely takes over the behaviour of the animal.

Finally, the escape circuits are heavily modulated. The test of whether to escape is relatively simple but it is not static. The crayfish' nervous system's evaluation of the animal's overall state is used to modulate the escape threshold. Whilst performing an escape tail-flip, starting another one is completely blocked. The inhibition or sensitisation may be general or

specific. The circuit may make the escape response less excitable to some specific inputs or to all inputs.

(Bennett, 1984) observed that startle responses are present in many different kinds of animals (including humans) from different parts of the 'evolutionary tree'. He compared the startle responses in different animals to identify common attributes. Some of the key features we have identified in the crayfish escape circuit are common to the organisation of startle responses in general. In most animals, the axons carrying the startle response signal are usually of large diameter axons to facilitate rapid conduction. Electrical synapsing is frequently used for its faster transmission and reliability. Startle responses often use command neurons: which function as a single decision point. Not all startle responses are stereotyped: mammalian startle response is highly dependent on the state of the organism.

In the LG and MG circuits, the decision to enact a tail-flip behaviour is made by command neurons. The behaviours can be selectively controlled by its trigger(s), and occurrence of behaviour are signalled uniquely by the firing of the trigger(s) (Krasne and Lee, 1988). Are behaviours in general driven by command neurons? Consideration of vector coding predicts that more than one command neuron should be involved in most behaviours (Fraser, 1982). (Edwards et al., 1999) identifies that animals have a spectrum of premotor organisations, which can be grouped into three categories:

- parallel distributed networks, in which shifts in the pattern of activity are caused by corresponding shifts in a pattern of activity of a population of premotor interneurons,
- command systems: categorically different movements are selected and guided by patterns of activity in distinct groups of neurons, and
- command neurons.

(Edwards et al., 1999) hypothesises that it may emerge that localist architectures are used when discrete numbers of categorically different responses are to be produced and intermediates would be maladaptive. Distributed architectures are used when response properties vary continuously with stimulus.

The crayfish escape circuits also offer themselves as a case study of how an animal chooses between different conflicting behaviours. Different behaviours are incompatible because they move the same body parts and because they are adaptive responses to different sensory stimuli (Beall et al., 1991). In the crayfish it can be observed that the escape circuits inhibit each other and other parts of the nervous system. Conversely, the LG has been shown to be inhibited during feeding and restraint. (Edwards, 1991) postulates that the crayfish and other animals may use mutual inhibition among behavioural command systems as the mechanism for switching between incompatible behaviours.

In relative terms, the crayfish escape circuit is one of the simpler neural circuits found in biology. It is however by no means trivial, as is evident from its description. The interactions between the components of the system and other parts of the crayfish nervous system are complex. The escape circuit exhibits many interesting properties, such as learning and modulation. Its interaction and involvement with other behaviours, such as fighting and dominance, make it an interesting system to study combinations of behaviours. The crayfish escape circuit serves as a good model system to study many of the interesting properties exhibited by biological systems.

This description of the crayfish escape circuit highlights two aspects that are true of neuroscience in general. One should keep in mind that the crayfish escape system is one of the simplest and best understood neural systems. Firstly, the interactions of the nervous system are extremely complex. There are interactions between many cells in different ways. Secondly, what we

understand is much less than what we do not understand. There are many unanswered questions in regard to the function of the crayfish escape circuit. Our understanding of how the circuit functions is continuously evolving with research. Inevitably, the descriptions and models become more complicated.

3.6 Reflex behaviours in artificial systems

Almost every complex animal, from disparate parts of the evolutionary tree, has some form of startle response (Bennett, 1984). Reflex behaviour is clearly important in animal behaviour. They are probably important in artificial systems as well.

Escape and reflex response shave obvious and analogous applications in robots. One can easily imagine that robots, for example a mine-sweeping robot, being threatened with physical danger. (Hoy, 1993) argues that the equivalent of an escape behaviour has a role in other artificial systems as well. Hoy argues that artificial systems should have an escape system that is activated when the integrity of the agent is threatened. The escape system would cause the agent to diverge from its normal operations and take evasive action. As an abstract example, imagine a web agent whose purpose is to collect information about rental rooms from the web. The agent may start at real estate portals and follow links to real estate pages where it parses the pages and collects the information. The equivalent of an escape system may be to have a component of the system that tests whether the agent is still processing real-estate pages, or has it followed links that have landed the agent in a completely unrelated portion of the web. The response may be to backtrack to a place where the agent is sure it is processing real-estate pages.

The crayfish escape circuit offers insights into how to design reflex behaviours

for artificial systems. Drawing on what we have identified as the key points of the crayfish escape circuit's architecture (see section 3.5) and generalisations (Bennett, 1984) of other escape circuits, we propose these guidelines of how to construct reflex behaviour in robots and agents.

1. The escape functions should not interfere with the general function of the agent while an escape manœuvre is not actually in play.

2. An escape response should be fast. This is relevant to real-time systems, such as robots. In a real-time system, an escape should be fast in its triggering and its execution. In the crayfish escape system, the circuit is highly optimised for the maximum possible speed to give it an edge in avoiding predators.

3. The escape system needs to continuously monitor the agents environment to determine whether the system is at risk an an escape should be triggered.

4. The trigger for an escape should be relatively simple to calculate. In a real-time system this is a necessity to facilitate a fast response. In any system, a simple test for escape is desirable so that it does consume all the processing resources. In the crayfish escape circuit, the test for escape is calculated within a network of neurons three layers deep.

5. An agent may be presented with multiple sources of danger. Or different dangerous scenarios from the same source. The crayfish escape system suggests that one should not try to handle all dangerous scenarios with one mechanism, but rather have multiple mechanisms for each scenario. Such a decomposition is conducive to keeping the test to trigger escape simple. The crayfish escape system has at least three different escape circuits to respond to different kinds of predator attacks.

6. When an escape is triggered, it needs to be able to take over the entire system, and do so near instantaneously. In the crayfish an escape tail-flip causes the crayfish to stop whatever it was previously doing. A massive inhibitory signal is sent across the nervous system of the animal to prevent other muscle movements that might interfere with a tail-flip.

7. An escape reflex needs to be modulated.

We have seen in the simulations in the scenario in chapter 2 that when one adds effective escape behaviour to a system, the agent is able to use a near optimal strategy in how it approaches its primary functions. From an architectural design perspective, escape and reflex behaviours are a mechanism for decoupling the agent's primary functions from how to respond when the system's integrity is threatened. Escape and reflex behaviours are analagous to exception handling in computer programs.

If you wish to learn about the endeavours of Artificial Intelligence researchers, in their quest to build artificial systems that recreate natural intelligence, you must read the next chapter.

CHAPTER 4

Previous Models

In this chapter we introduce computational neuroscience, describe previous models of the crayfish escape circuit and review some previous studies in adaptive behaviour.

4.1 Computational Neuroscience

The field of computational neuroscience is concerned with simulating neurons and neural circuits on a computer.

4.1.1 Modelling techniques

A dendrite conducts ions. Electricity conducts along a dendrite in a similar manner to which it conducts along a cable. If we represent a dendrite as a cylinder then cable theory (Thomson, 1855) can be applied to calculate an exact analytical solution to the electrical potential along the dendrite as

a function of space and time in response to an arbitrary current injection (L. Davis and Lorente de Nó, 1947; Rall, 1957; Rall, 1989; Rall and Agmon-Snir, 1998).

Cable theory only allows us to find analytical solutions to dendrites shaped as a cylinder. To be able to model a branching structure, it is necessary to calculate the equivalent cylinder (Rall and Rinzel, 1973; Rall, 1989; Rall and Agmon-Snir, 1998). This method can only be applied if the dendritic branching structure adheres to strict geometric constraints. Unfortunately, real dendritic structures rarely conform to these rules. It is still possible to find analytical solutions to arbitrary structures using recursive methods, however the solutions become extremely complex and unwieldy. Analytical solutions cannot be found when there are any voltage dependent channels in the membrane, so these methods cannot be applied to an active membrane or at synapses. To be able to model complex dendritic trees with active channels other methods are necessary.

The idea of compartmental modelling is to represent the electrical properties of a neuron as a series of connected compartments (Segev and Burke, 1999). The compartments are made small enough so that the portion of the neuron represented by a compartment is considered to be at isopotential. Current injected into the compartment may conduct along the dendrite to adjacent compartments, or leak out through the cell membrane. Along each path the current faces an impedance. The impedance of the path along the dendrite is modelled by a resistor, called *axial resistance*. The impedance of current leaking out of the cell is modelled by a resistor, called *membrane resistance*, in parallel with a capacitor, called *membrane capacitance*. Figure 4.1 shows a three compartment model of a section of dendrite.

The impedance values of the compartment depend on the electrical properties of the cell and dimensions of the dendrite. The axial resistance is given

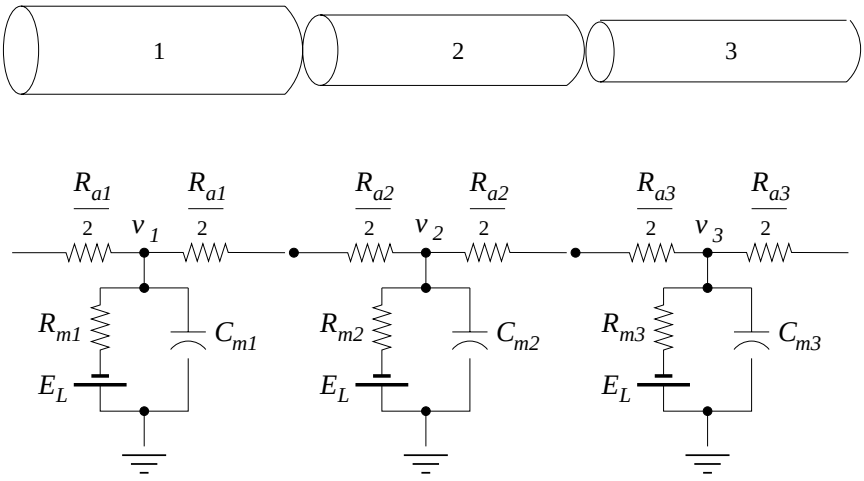

Figure 4.1: A three compartment model of a section of dendrite. The dendrites morphology is modelled as three cylinders. Each cylinder is represented as an electrical compartment.

by:

$$R_a = \frac{L/2}{\pi(D/2)^2} r_a = \frac{2L}{\pi D^2} r_a \tag{4.1}$$

where R_a is the axial resistance in MΩ, r_a is the axial resistivity in Ωcm, L is the compartment length in micrometres and D is the dendrite diameter in micrometres.

The membrane resistance is calculated by dividing the cell's membrane resistivity by the compartment's area:

$$R_m = \frac{r_m}{\pi DL} \tag{4.2}$$

where R_m is the membrane resistance in MΩ and r_m is membrane resistivity in $\Omega.\text{cm}^2$.

If we invert equation 4.2, we have an expression for conductance, called the leak conductance g_L.

$$g_L = \frac{1}{R_m} = \frac{\pi DL}{r_m} \tag{4.3}$$

The membrane capacitance is calculated by:

$$C_m = \pi D L c_m \tag{4.4}$$

where C_m is the compartment's membrane capacitance in μF, and c_m is the membrane capacitance per unit of area in $\mu\text{F.cm}^{-2}$.

At each compartment, the currents and membrane potentials can be described by a set of differential equations. These equations can be integrated numerically. On a computer, the response to arbitrary input in complex dendritic structures can be calculated in this way.

We now consider active membrane. In a compartment representing active membrane, we add voltage dependent channels in parallel to the leak resistance.

Hodgkin and Huxley (1952) proposed a mathematical model to describe the voltage dependent sodium and potassium channels that are responsible for generating action potentials. They used data from the recordings in the squid giant axon as the basis of their model. Hodgkin and Huxley proposed that each sodium channel had three activation components and one inactivation component. For the channel to be active all three activation components had to be in the 'on' state and the inactivation component had to be in the 'off' state. A potassium channel was made up of four activation components and no inactivation components. The probability of an activation component changing state is related to the voltage of the membrane.

Let us consider a section of the membrane with millions of sodium and potassium channels. Let us use the variable m to represent the proportion of sodium activation components in the on state and the variable h to represent the proportion of sodium inactivation variables in the off state. At any

given time the sodium conductance can be calculated as the product of the maximum conductance (where all the sodium channels are open) and the proportion of channels that are open, which is given by the product m^3h. Hodgkin and Huxley proposed the following set of differential equations to describe the membrane potential of an active membrane.

$$C_m \frac{dV}{dt} = -\overline{g}_L(V - V_L) - \overline{g}_{\mathrm{Na}} m^3 h(V - V_{\mathrm{Na}}) - \overline{g}_{\mathrm{K}} n^4 (V - V_{\mathrm{K}}) \tag{4.5}$$

$$\frac{dm}{dt} = \alpha_m(V)(1 - m) - \beta_m(V)m \tag{4.6}$$

$$\frac{dh}{dt} = \alpha_h(V)(1 - h) - \beta_h(V)h \tag{4.7}$$

$$\frac{dn}{dt} = \alpha_n(V)(1 - n) - \beta_n(V)n \tag{4.8}$$

Where V is the membrane potential, t is time, C_m is the membrane capacitance, $\overline{g}_L$ is the membrane leak conductance, V_L is the membrane leak resting potential, $\overline{g}_{\mathrm{Na}}$ is the maximum sodium conductance, V_{Na} is the sodium reversal potential, $\overline{g}_{\mathrm{K}}$ is the maximum potassium conductance, V_{K} is the potassium reversal potential, n is the potassium activation variable and $\alpha_m(V)$, $\beta_m(V)$, $\alpha_h(V)$, $\beta_h(V)$, $\alpha_n(V)$ and $\beta_n(V)$ are voltage dependent functions governing the rate of change of the sodium and potassium activation and inactivation variables.

For the squid axon the α and β functions, the activation variables m, h and n, can be described as follows:

$$\alpha_m(V) = \frac{A_{\alpha m}(V - V_{0\alpha m})}{e^{(V - V_{0\alpha m})/B_{\alpha m}} - 1} \tag{4.9}$$

$$\beta_m(V) = A_{\beta m} e^{(V - V_{0\beta m})/B_{\beta m}} \tag{4.10}$$

$$\alpha_h(V) = A_{\alpha h} e^{(V - V_{0\alpha h})/B_{\alpha h}} \tag{4.11}$$

$$\beta_h(V) = \frac{A_{\beta h}}{e^{(V - V_{0\beta h})/B_{\beta h}} + 1} \tag{4.12}$$

$$\alpha_n = \frac{A_{\alpha n}(V - V_{0\alpha n})}{e^{(V - V_{0\alpha n})/B_{\alpha n}} - 1} \tag{4.13}$$

$$\beta_n = A_{\beta n} e^{(V - V_{0\beta n})/B_{\beta n}} \tag{4.14}$$

4.1.2 Neuron modelling software

The Hodgkin-Huxley equations have no analytical solution. To calculate the time course of the membrane potential it is necessary to numerically integrate the equations. It is thus necessary to write a computer program to numerically integrate the membrane equations.

Programs to simulate neuron membrane properties have been written many times. A number of libraries and tools have emerged that save one from having to write the numerical integration engine from scratch. The two most popular general purpose neural simulators are GENESIS (Bower and Beeman, 1998) and NEURON (Hines, 1984; Hines, 1993; Hines, 1994; Hines, 1998; Carnevale and Hines, 2003).

GENESIS and NEURON have equivalent functionalities. The programmer defines the compartmental structure of the neurons being simulated, the

external components, what values need to be recorded and the integration method. The simulation engine converts the description into a simulation of an RC circuit with variable conductances and numerically integrates the currents to compute the simulation results. The programmer defines the simulation components and setup in scripts. NEURON and GENESIS each come with their own scripting languages.

GENESIS and NEURON each come with a library of built-in components that are commonly used in neuronal models. The libraries include frequently used membrane channel types. NEURON and GENESIS have predefined passive membrane channels and Hodgkin-Huxley sodium and potassium channels. NEURON and GENESIS are both extensible. New channel types can be written in C and integrated into the simulation. The simulation libraries also provide other components that can be used in the simulation such as synaptic mechanisms, voltage clamps and current clamps. Both NEURON and GENESIS allow the programmer to write new components in C.

4.1.3 Model parameter searching

The data used as the basis of neural models is often incomplete. There are usually many parameters whose biological values are not precisely known. Their values are typically known to be constrained within a range. The modeller must choose parameters that cause the model to act in a way that is consistent with experimental observations. There are often multiple sets of parameters which satisfy the constraints. The process of finding suitable values is called parameter fitting. There are many techniques for doing this, including:

1. tuning the parameters by hand,

2. using automated techniques to map out the parameter space or subsections of it, and

3. using an automated search algorithm (for example gradient descent or genetic algorithms) to find the 'best fit' parameter values.

4.1.4 Previous models of the crayfish escape circuit

There have been previous computational models of portions of the crayfish escape circuit.

There have been a number of models examining the role of rectifying junctions in the escape circuit. (Edwards et al., 1991) examined the function of rectifying electrical synapses onto LG. The function of the rectifying electrical synapses between LG and the motor giant (MoG) motor neurons has also been simulated (Edwards, 1990a; Edwards, 1990b; Heitler and H.Edwards, 1998). (Heitler et al., 1991) examines the role of the rectifying junctions between the segmental giant (SG) and the fast flexor motor neurons (FF).

(Edwards et al., 1998) showed, with single compartment models, that rectifying junctions can enhance coincidence detection.

(Edwards and Mulloney, 1987) simulated synaptic integration between the SG and the FF motor neurons.

(Vu and Krasne, 1993b) simulated some aspects of synaptic integration by LG to show the effects of proxal versus distal inhibition.

The effects of growth in the lateral giant has been studied in computational models (Edwards et al., 1994a; Edwards et al., 1994c; Hill et al., 1994). The models predict that the lateral giant's spiking threshold rises in larger animals as the electric load increases disproportionately to the size of its inputs.

MacMillan and Patullo (2001) created a simulation of the control of the position and movement of the abdomen, with the goal of obtaining insights for robotic design. MacMillan and Patullo observe that the crayfish obtains very fine control with very few control elements. They argue that understanding this has obvious benefits for robotic design. They conclude that further studies in both biology and neural simulations are necessary in order to obtain the necessary understanding to exploit the full potential of the crayfish's circuit design.

4.1.5 Models of coincidence in other animals

Coincidence detection has been studied in other neural circuits. Mammals and birds use the small time difference at which a sound arrives at its opposite ears to estimate the location of the source. Barn owls are often used as a case study of this phenomenon due to their keen ability to use sound to localise their prey at night. The calculation relies on coincidence detector neurons in the auditory system that respond maximally to simultaneous inputs from both ears. The coincidence detection neurons in the barn owl have been modelled. The coincidence detector cells should fire when inputs from two independent sources coincide but not fire when inputs from the same source coincide.

In models, the coincidence detection is accounted for by having the inputs from different sources synapse at different dendritic branches (Carr and Boudreau, 1993; Agmon-Snir et al., 1998; Simon et al., 1999; Peña et al., 2001). Alternative dendritic branches causes synchronous inputs from different sources to sum more linearly and the other branch acts as a current sink when inputs are asynchronous or from the same source. Inputs synapsing at the same location sum nonlinearly. This is because the depolarisation is largest at the synapse point. The driving force is proportional to the volt-

age difference across the membrane. Thus if two inputs synapse at the same locations, they will reduce each others driving force, causing them to sum nonlinearly. When inputs synapse at different branches, they are electrotonically separated, and the depolarisation that one input causes at the other branch is not very great. Hence, inputs on alternative dendritic branches do not significantly reduce each other's driving force. Simultaneous inputs from both ears therefore sum effectively. Simultaneous inputs from the same ear sum ineffectively.

More detailed simulations show that active potassium currents may be involved in discriminating between asynchronous and synchronous inputs (Grau-Serrat et al., 2003).

4.1.6 Simulations of other biological systems

There have been many previous simulations of parts of nervous systems in different animals with the goal of obtaining insights into the general design of animal behaviour.

(Laurent, 1993) examined integration in spiking and non-spiking local neurons in the locust central nervous system. Their studies highlighted the importance of cellular and synaptic properties for the function of the network. The main design principles that emerged from (Laurent, 1993)'s studies included:

- Tactile and proprioceptive signals are sparsely distributed. The signals are sparse and specific. Each interneuron may share sensory input, but each interneuron has a different receptive field. The same receptors may act differently on different interneurons.

- The network lacks internal feedback pathways.

- The coexistence of spiking and non-spiking integrative modes in the same networks must be important.

(Lockery, 1993) created an artificial neural network model of an escape reflex in the leech.

(Ritzmann, 1993) examined the neural organisation of cockroach escape and its role in context-dependent orientation.

(Morse et al., 1998) created a neural network model of part of the nervous system of the nematode *C. elegans* to study robust navigation based on smell and chemical sensory information.

4.2 Adaptive behaviour models

Artificial Intelligence has existed as a discipline since the 1950's.

The aim of artificial intelligence is to build computer systems that display the kind of intelligence we expect in humans.

Central to artificial intelligence is the idea of an agent. An agent is an entity situated in an environment (Wooldridge and Jennings, 1995; Wooldridge, 1999; Russell and Norvig, 2003). An agent is able to sense and act on its environment. An agent is a broad term which can be applied to include robots and also software agents, such as web agents.

The predominant approach to artificial intelligence has been to create a symbolic representation of the state of the world and apply logical reasoning to the symbolic model.

A major change in the approach to artificial intelligence was instigated by Rodney Brooks during the late 1980s (Brooks, 1986). Brooks conjectured that the intelligence displayed by 'simple' animals, such as insects, was vastly

superior to the state-of-the-art in artificial systems. Brooks argued that rather than try to emulate human intelligence we ought to first try to replicate the levels of intelligence displayed by 'simple' animals. Brooks argued that rather than having a complex symbolic model of the world, a robot should use the world as its own model. Brooks proposed a subsumption architecture. Brooks' subsumption architecture attempted to decompose behaviour into layers. The lower layers were responsible for the most elementary functions, for example obstacle avoidance. The higher layers were responsible for higher level reasoning, for example navigating to the goal. Brooks work inspired a new approach to robotics, called behaviour based robotics, where an agent's the problem is decomposed by behaviours.

The subsumption architecture was an example of behaviour based robotics. A behaviour based robot has a set of behaviours. The environment determines which behaviour should have control at any given time. Behaviour based robots differ from purely reactive systems because 'they can use different forms of internal representations and perform computations on them in order to decide what effector action to take.' (Mataric, 1992)

(Nolfi, 1997) further explored the concept of behaviour based robotics. Nolfi hypothesised that to obtain robust solutions behaviours should be broken down 'proximally' rather than 'distally', *i.e.* in terms of sensorimotor loops instead of high level observed behaviours. Mobile robots are able to affect which kind of stimuli they are exposed to, *i.e.* sensory state, consequence agents are able to produce behavioural sequences without using memory. Some sensory states are more important than others. (Nolfi, 1997) proposed an 'emergent modular architecture', *i.e.* an architecture where a neural network is constrained such that each motor neuron is associated with a separate small group of control neurons. Their experiments supported their claim that emergent modular architectures have improved performance.

Beer (Beer, 1990; Beer et al., 1989) argued for studying 'simple animals'. He argued to examine an existing intelligent system in nature to look for clues in how to build an artificially intelligent system. Beer (Beer, 1990) in his PhD thesis successfully modelled insect walking, arguing that it is a step to understanding intelligence as adaptive behaviour.

There is now a sub-discipline of artificial intelligence that is concerned with the simulation of adaptive behaviour in animals (Meyer and Wilson, 1990; Meyer et al., 1992; Cliff et al., 1994; Maes et al., 1996; Pfeifer et al., 1998; Berthoz et al., 2000; Floreano et al., 2002).

4.3 A comparison of biological and artificial neurons

Looking at neurons for inspiration in regard to performing computation is not new. (Rosenblatt, 1962) presented the idea of a perceptron: an abstraction of a neuron. Networks of perceptrons were shown to be capable of solving certain problems. (Minsky and Papert, 1969) mathematically analysed the capabilities of the perceptrons and proved which classes of problems can and cannot be solved with perceptrons. The field remained largely dormant until it was revived by Rumelhart and McClelland (1986a, 1986b) who extended the theory of perceptrons and demonstrated it could be used to solve complex problems.

There are some fundamental differences between the computations of a neural network based on classical perceptrons versus how a biological neural networks functions. The perceptron model, and variations of it, has many useful applications. It is however insufficient to capture the entire functionality of biological neural networks. We highlight nine differences between artificial and biological networks.

1. Artificial neural networks process their data synchronously. Biological neural networks are asynchronous. In an artificial neural network the information propagates between layers synchronously. In biological neural networks cells communicate completely asynchronously. Although examples exist in biological circuits where computation is synchronised (e.g. coincidence detectors and central pattern generators), there is no central clock whereby all neurons communicate at the same time.

2. A biological neuron integrates its input in a highly nonlinear way. Active channels in the dendrites and variable conductance at transmitter gated channels cause PSPs to sum in highly nonlinear ways. In biological neurons, integration is complex even without any active channels. With only passive membrane the dendritic tree is a linear system, however it is a very complicated linear system. An accurate model of the integration of input onto a passive dendritic structure is a complicated procedure. It is doubtful whether a neuron's integration process can be modelled accurately with a simple to compute function.

3. A biological neural network has many different kinds of synapses. In artificial neural networks synapses differ only by their weight (see figure 4.2.) Since artificial networks are synchronous, the function of excitation versus time varies only by its magnitude. In biological neurons different synapse types have vastly varying excitatory responses. There are many different kinds of neurotransmitters and there are different types of receptors for each neurotransmitter. Each receptor type responds differently. These combinations of neurotransmitters and receptors cause different synapse types to vary greatly in the duration and the shape of the post-synaptic excitation.

4. In biological neural networks, communication at the synapses is not

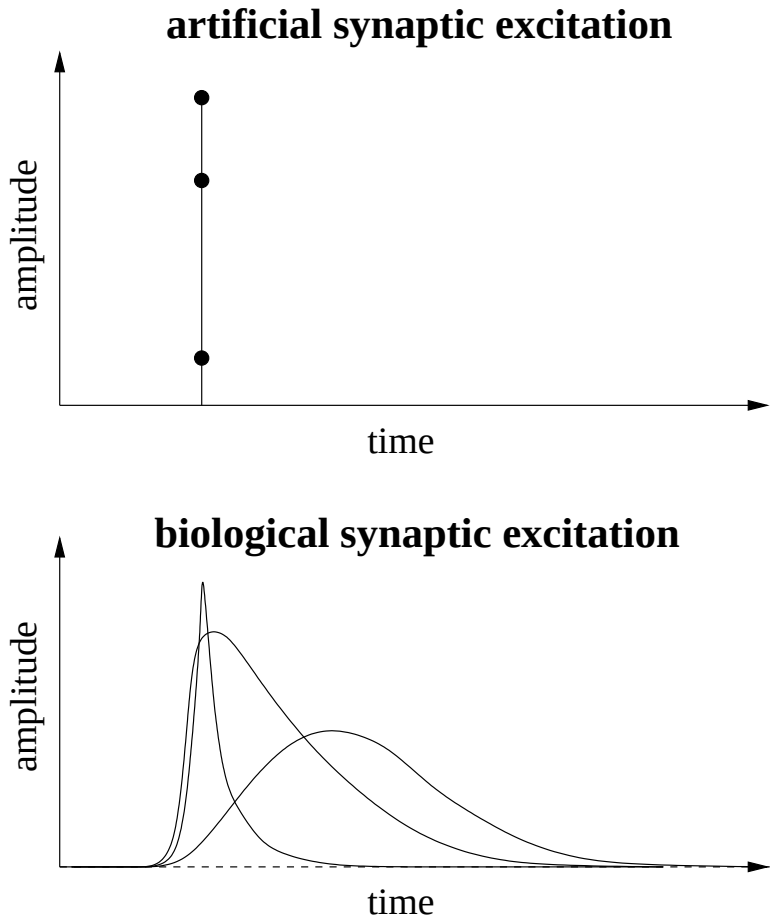

Figure 4.2: In artificial neural nets, synapses only differ by their weight (see graph above). In biological neural networks, the input at a synapse is a continuous function of time (see graph below).

limited to excitatory or inhibitory currents. Neurotransmission may effect a second-messenger response. In a second messenger response no PSP is generated directly but another process is activated which may change the intrinsic properties of the target neuron.

5. There are neurohormones which are able to affect many different cells across the entire brain in different ways. Neurohormones usually cause modulation. Different cells are usually modulated in different ways by the same neurohormone because the response in a cell depends the receptor types and the density of these receptors at the target cell. There is equivalent to neurohormonal transmission in artificial neural networks.

6. In biological neuron PSPs sum over different time courses, due to their asynchronous computation. In an artificial neural network weights are added in a single time step and all inputs are effectively given the same duration. In a biological neuron the duration of the PSP is different for different synapses depending on the channels involved. Furthermore, whereas in an artificial neuron weights are effectively a step function, in a biological neuron the magnitude of the PSP varies over a time course, although it can often be modelled with exponential functions.

7. Biological neurons are be modulated. The properties of a cell can be changed in many alternative ways by neuromodulators. Neuromodulation can affect synapses or the integration properties of the neuron. For example, a cell may be made more responsive (potentiation) or less responsive (depressed.) The duration of modulation is depends on what neuromodulators and receptors are involved.

8. Related to modulation is habituation. In biology, a neuron may become less (or more) responsive to input if it is contagiously activated.

9. Biological neural networks are dynamic. A biological neural network is continuously changing. In addition to modulation and habituation the actual structure of the network and cells is changing. Dendrites grow. New synapses form. Synapses are removed. Synapses are strengthened, synapses are weakened. The structure of a biological neural network is not static.

Each of these described biological mechanisms has been shown to have functional significance. The classical perceptron model does not capture the properties described in the afore mentioned points. One could still argue that biological neural networks and perceptron-based artificial neural networks differ in their mechanisms but ultimately achieve the same results. Although a matter of conjecture, it is difficult to computationally reconcile these fundamentally different systems. These differences between biological nervous systems and classical artificial neural networks offer clues as to what may be need to be added to artificial systems to achieve more robustness.

Much recent work by artificial neural network researchers works towards addressing this gulf that exists between biological and artificial neural networks. Some comtemporary models attempt to incorporate the basical biogological mechanisms within their assumptions. (Elliot and Shadbolt, 2003a; Elliot and Shadbolt, 2003b) include neurotrophic processes in their neural network models. Activity dependent feedback dynamically adjusts synaptic strength. The inclusion of neurotropic mechanisms was shown to vastly improve robustness. In experiments with robots, the neural networks evolve to be well-tuned to the specifics of the hardware and the environment.

If you wish to learn how the lateral giant can be modelled with resistors and capacitors then you must read the next chapter.

CHAPTER 5

Compartmental models of the Lateral Giants

The biological models in this thesis are of the sensory part of the lateral giant (LG) circuit (see figure 3.11.) We focus on the sensory part of the circuit to consider how the crayfish arrives at its decision to invoke an LG tail-flip.

The key cell in the escape circuit of the crayfish is the LG. It is the decision point to trigger escape behaviour. Much information is known about the neuroanatomy and physiological properties of the LG. Accordingly, the key step in this research is to build compartmental models of the LG. The compartmental models are described in this chapter.

Much is known about the structure and physiology of the LG. Much more is unknown. It is necessary to make assumptions and simplifications. The models attempt to replicate the cell's observed behaviour. The model designs are constrained by what is known. The gaps of the unknown are filled

with plausible choices that produce a model that behaves in a way that is similar to the experimentally observed cell.

As we discussed in Chapter 4, computer programs are appropriate for solving the differential equations governing the neuronal membrane potentials in the case that there is not an analytical solutions. Two models were built. The first attempts to capture the essential features of the LG morphology. It was created before a complete set of neuroanatomical measurements was available. The model was created using GENESIS and is described in section 5.1. This model is included because it is simpler than the complete neuroanatomical model yet it exhibits many of the same properties. The second model has its morphology taken directly from biophysical measurements. It is written using the NEURON simulator and described in section 5.2.

The models described in this chapter are used as the basis for simulations addressing specific questions in chapters 6 and 7.

5.1 Basic compartmental model of the lateral giants

The lateral giants are large complex cells. A picture of one of the cells dye-filled may be seen in figure 5.1. The cell's structure can be qualitatively described as follows:

- The cell has a large wide axon.
- The axon is fed by 5 large dendritic branches. These branches converge on the main dendritic branch at roughly the same point.
- At the end of the dendritic tree are a large number of very narrow dendritic tips.

Figure 5.1: A picture of the Lateral Giant cell filled with dye. Picture from (Antonsen, unpublished.)

Two approaches were taken to building compartmental models of the lateral giants. The first was to create a simplified model that includes all the features of its qualitative description. The second was to create a model of the cell's structure directly from the pictures taken of the dye-filled cells using a software tool (refer to section 5.2.)

There are of course many different variations of the lateral giant cell. Each crayfish has a pair of lateral giants in each abdominal segment. There are small variations in the lateral giant structure between segments. The dimensions and number of branches in the lateral giants change as a crayfish grows. The dimensions of the lateral giant modelled were based on the lat-

eral giants in the final abdominal segment of a 6 centimetre *Procambarus Clarkii*. This was chosen because there was the most data available.

5.1.1 Abstracted morphology model description

The LG is modelled as a structure built of connected cylinders and these cylinders are modelled with electrical compartments. A diagram of the simplified model adopted of the lateral giant cell is shown in figure 5.2. The simplified structure has the qualitative attributes of the lateral giant cells. The following constraints were put on the structure and dimensions of the cell:

- The axon width used matched an actual cell's width.
- The diameter of the base of the dendrites matched that of an actual cell.
- There are five main dendritic branches.
- The number of dendritic tips approximately matched that of an actual cell.
- The length and diameter of the dendritic tips approximately matched the average dimensions of these attributes for actual cells.

In the model the dendritic base branches out into five symmetric branches. Each of these branches out into 10 symmetric sub-branches. Each of these sub-branches branch out to 10 narrow dendritic tips. The lengths and diameters used for the segments of the dendritic tree at various branching levels were chosen by measuring the lengths of segments in dye-filled photos of the lateral giants and by consulting with crayfish biologists. Table 5.1 shows the dimensions used in the model for the various stages in the branching

Branching level	Number	Length (μm)	Diameter (μm)
Base	1	100	20
Major branches	5	20	10
Minor branches	50	8	1
Tips	500	6	0.4

Table 5.1: The dimensions used for the branches in the model of the lateral giant dendritic tree.

structure of the model lateral giant. These values are based on the measured observations from the lateral giant anatomy (Antonsen 2000, personal communication).

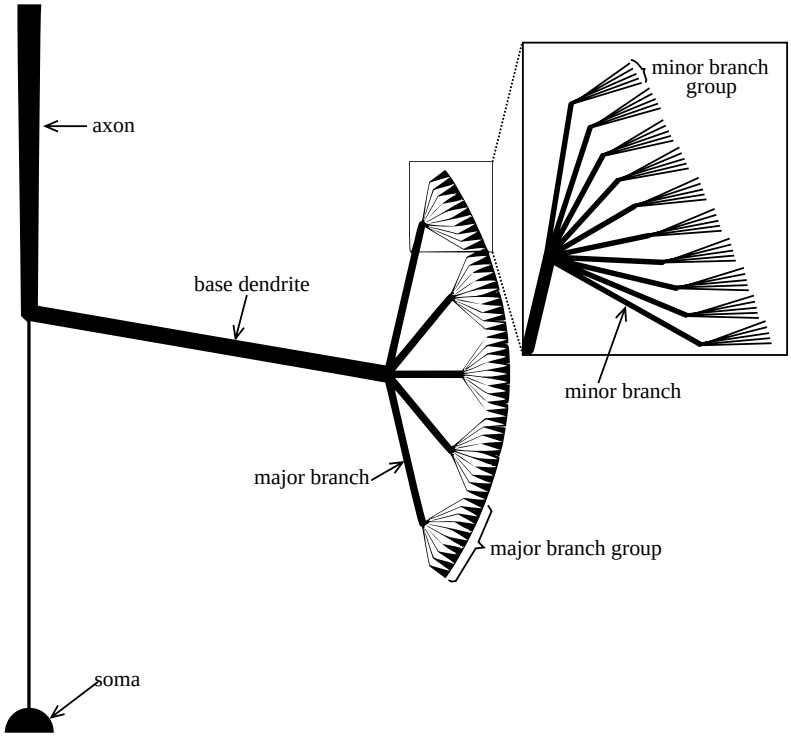

Figure 5.2: The symmetrically branching structure used to model the LG.

The dendrites were assumed to have a passive membrane, *i.e.* they passively conduct current. There is recent evidence to suggest that, contrary to classical models, most dendrites have some active channels. Active channels were not included in the modelling of the dendrites. This simplification is both

necessary and justified. First, nothing is known about any active channels in the LG dendrites. There is no evidence available to show their existence, let alone details of their properties. Second, electrophysiological experiments show that the response of the lateral giants to current injection is highly linear (Heitler et al., 1991), suggesting that active channels do not play a major role in the conduction of the sensory signals. Finally, understanding a passive membrane model of the dendrites is a prerequisite to any active membrane models (Segev and Burke, 1999).

The cytoplasmic resistivity of the dendrites was assumed to be 60 Ωcm, an experimentally measured value for crustaceans (Katz, 1966). This is the same value that has been used in previous models of the lateral giants (Edwards et al., 1994b; Hill et al., 1994).

The number of compartments used to model any segment of the dendritic tree is such that for all compartments, the electrotonic length of the dendritic segment being modelled did not exceed 0.05λ. The electrotonic length (L) is defined as the physical length (l) divided by the space constant (λ), as shown in equation 5.1. The space constant is the distance over which the voltage level of a signal will decay by a factor of $\frac{1}{e}$. It may be calculated from membrane resistance (R_{M}), the axial resistance (R_{A}) and the dendrite diameter using equation 5.2. Using segments of electrotonic length 0.05λ gives high accuracy results whilst keeping the number of compartments to a level that is computationally possible (Segev and Burke, 1999).

$$L = \frac{l}{\lambda} \tag{5.1}$$

$$\lambda = \sqrt{\frac{(d/4)R_{\mathrm{M}}}{R_{\mathrm{A}}}} \tag{5.2}$$

The lateral giant is known to have a very short membrane time constant.

The axon is modelled as a series of compartments of wide diameter. The final model had a very long axon, with a diameter initially matching that of the base of the dendrites, widening out over the initial section of the axon to 100 μm. The continuation of the axon maintains this diameter. The total length of the model axon was 5mm. This is the approximate length of a lateral giant axon in a real crayfish of the size being modelled. A case can be made that this very large charge sink should be greater still: the axon terminates at electrical gap junctions onto other lateral giants, allowing the ions to freely flow into the next cell. However (as is discussed in a following paragraph) the effect on the response at the dendrites caused by an axon approaching infinite length compared to one of the length in the model is not significant. Since the additional charge sinking capacity of synapses onto the adjacent lateral giants would be less than an infinite axon, this factor can be safely ignored.

The axon has a big effect on the response properties of the cell. Due to the low resistance of the axon compared to the dendrites, it acts as a charge sink. Ions from presynaptic cells enter at the dendrites. Due to the narrowness of the dendrites, high membrane potentials are created. However the charge quickly dissipates down the low resistance axon. Due to the axon's extremely large size in comparison to the dendrites, as the charge distributes evenly through the cell the membrane potential very quickly drops to a minuscule fraction of what it was at the dendrites. The effect of the large axon is to thus cause any significant membrane potentials in the dendrites to be extremely short lasting.

The effect of the axon and the sensitivities of the choices made were examined in more detail. The cell's response to input was made with three other variations of the lateral giant cell:

Axon model	EPSP (mV)
Axon model adopted	0.89
No axon	5.20
Infinite axon	0.89
Axon with uniform diameter	0.36

Table 5.2: The effects of the axon on cell's response to input. Table compares the EPSP caused by a single spiking afferent at the base of the dendrites for the model used to alternative modelling choices for the axon.

1. A cell with no axon,

2. A cell with an infinite axon, and

3. A cell with an axon of uniform width.

Table 5.2 shows the EPSP generated at the base of the dendrites in response to a single firing afferent for these different axon models. It can be observed that the existence of the axon has a large effect on the magnitude of the EPSP. The EPSP in a model cell with an axon is less than a fifth of one without one. The difference between an EPSP in the model cell and a cell with an infinite axon is not very great. This is because the length of the axon in the model is comparable to the space constant. Modelling the axon as initially narrow and broadening out also has a significant effect. The higher resistance of the initial segment of the axon causes the EPSPs recorded at the base of the dendrites to be slightly higher.

Figure 5.3 shows how the shape of the EPSP changes depending on how the axon is modelled. A model cell without the axon has sustained slowly decaying EPSPs. Model cells that include the axon have short duration rapidly decaying EPSPs. This is because in model cells with an axon charge can dissipate down the axon, where as in a cell without an axon it can only leak out through the membrane. The length and shape of the model of the axon do not appear to have a significant effect on the shape of the EPSP.

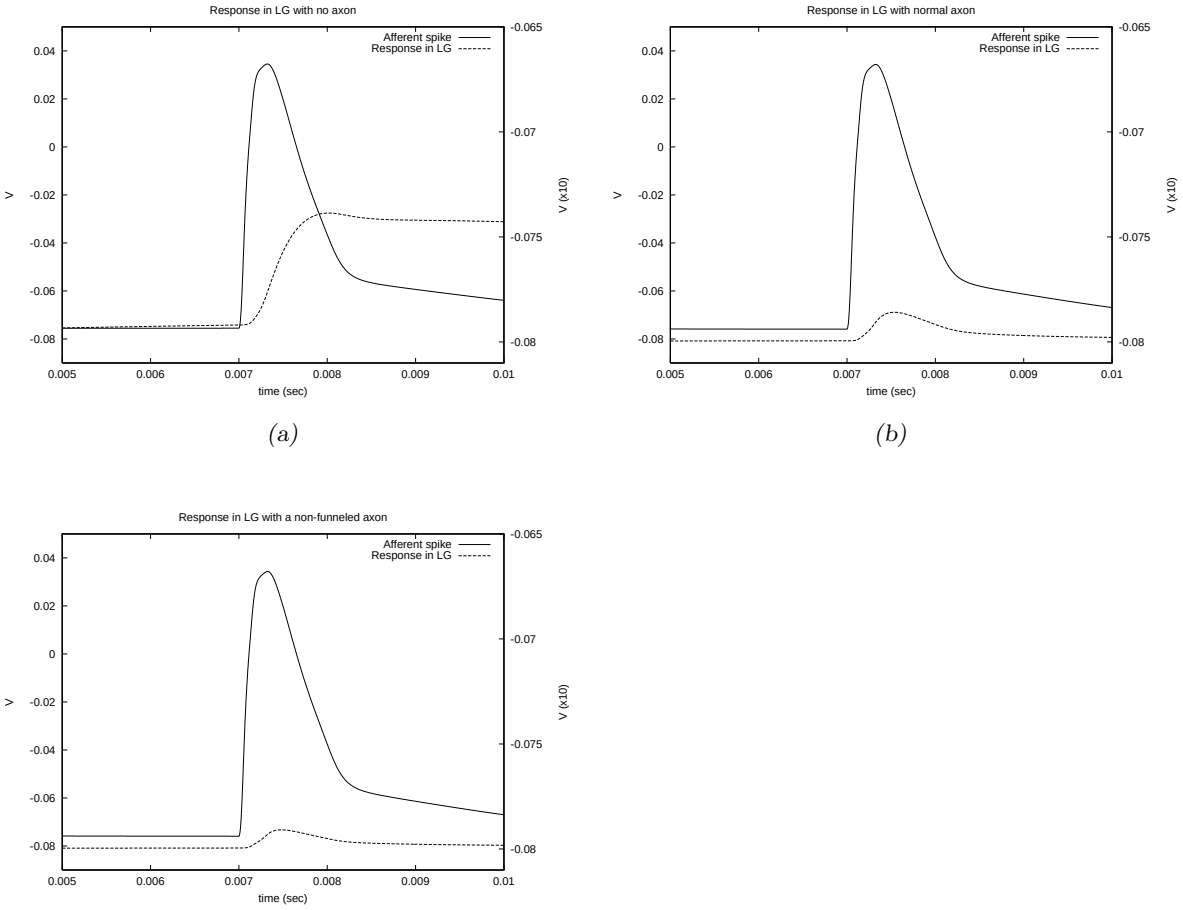

Figure 5.3: The shape of EPSPs in cells with different axon models. The graphs plot the response at the base of the dendrites to a single spiking afferent.

The resting potential for the model lateral giant cell was set at -80mV. This is a typical resting potential recorded for the lateral giants in neurophysiology experiments (Heitler et al., 1991; Zucker, 1972a). This resting potential was achieved in the model by setting the membrane leakage voltage for each compartment in the lateral giant to -80mV. Because there are no other channels in the model this is enough to achieve the correct resting potential.

Very simple models were used for the afferents. The question being explored in the model was how do the lateral giants combine the input given to them by the afferents and interneurons. The afferents were assumed to be sensory messengers, conveying stimulus to the interneurons and lateral giants. Hence the complexities of the afferents' morphology were ignored. being explored was how do the lateral giants As described in section 3.3.2 the afferents can be broadly divided into two categories. In the model, all afferents were made to be identical. This decision was made to keep the model simple. An extension of the work would explore what effects the afferents of different sizes have on the behaviour of the circuit. Only the axon part afferents were modelled, because this is the part that connects with LG. The afferent axons were modelled as a cylinder of uniform diameter spanning 300 μm. Six similar linearly connected compartments were used to model each afferent. Figure 5.4 shows the compartmental model used for afferents.

Dye-filled pictures of the afferents show that they do not terminate at the lateral giants. The afferents are not exclusively used by the lateral giants: their axons synapse onto the lateral giants, but most continue after this point to make synapses with other neurons. The model afferents make their synapses onto the lateral giant in the fourth compartment. The two following compartments provide an alternative sink for the current travelling down the axon.

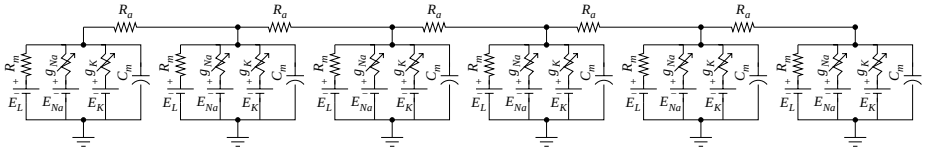

Figure 5.4: The compartmental model of the afferent axons. Each afferent is represented by six identical compartments connected together in a chain. The compartments contain Hodgkin-Huxley channels.

All the compartments in the model of the afferent axons are active, to allow them to propagate a spike. Hodgkin-Huxley sodium and potassium channels are used.

The shape of the spikes seen in crayfish tail afferents are atypical compared to spikes exclusively caused by the fast sodium and delayed potassium channels. Figure 5.5 shows the shape of a crayfish afferent spike compared to a spike caused by Hodgkin and Huxley sodium and potassium channels. The spikes are sharper, have a shorter duration and do not undershoot. This suggests that there are additional currents involved. Calcium currents are likely to be involved (Edwards, 2000, personal communication). There are probably other currents acting as well.

To get accurate simulations of the reaction to afferent spikes in the LG it is necessary to make the shape of voltage-time curves of the simulated afferent spikes to those of actual ones. The ideal way to do this is to correctly model all the currents and channels involved in an afferent spike. However since little is known about the currents acting in the afferents this would be a difficult and lengthy task. Since we are most interested in the response of the LG, a compromise solution was adopted. The constants, used in the Hodgkin-Huxley equations for the sodium and potassium channels, were modified to give the simulated afferent spikes a shape similar to actual ones. An alternative approach would have been use a real spike recording to directly set the voltage of an afferent. However the load on the afferent cell

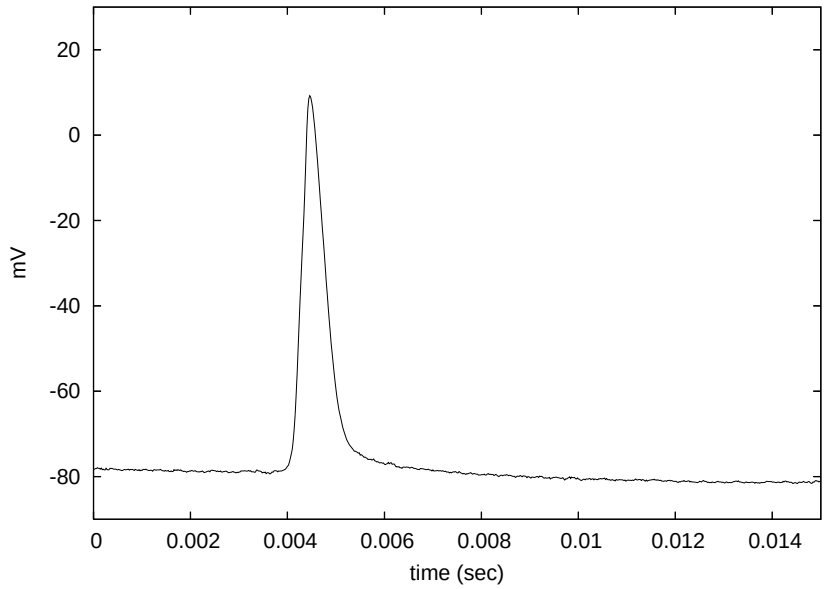

(a)

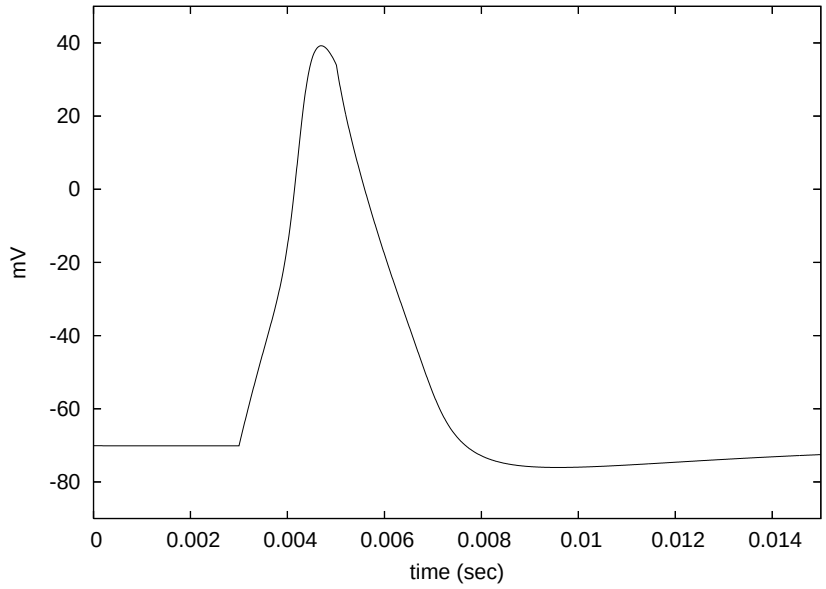

(b)

Figure 5.5: Recordings of the membrane potential during (a) a crayfish afferent spike, and (b) a spike in a simulated cell with standard squid Hodgkin-Huxley channels. Note that the crayfish afferent spike is of very short duration and does not undershoot.

changes as the resistance of channels changes. This effect would not be captured if the voltage was set directly. The compromise solution approximates this effect.

The Hodgkin-Huxley equations (see section 4.1.1), which describe action potentials in a squid neuron, were used as the basis for modelling the action potentials of crayfish afferents. The constants were adjusted so that a simulated spike be of a similar shape to that of spikes real crayfish afferent cells. Figure 5.6 compares a simulated spike using modified Hodgkin-Huxley channels to a real crayfish afferent spike recording. Appendix B.2.1 lists the values of the constants used.

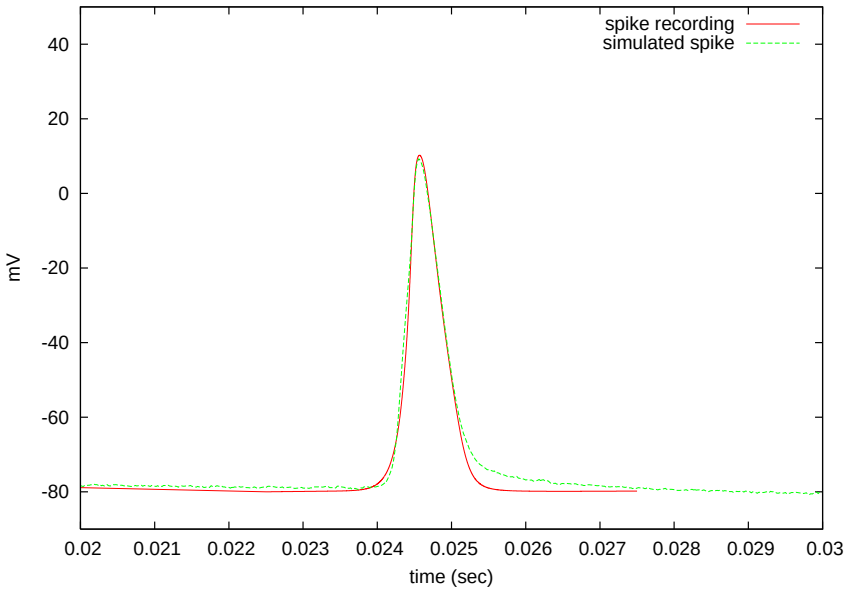

Figure 5.6: A comparison of a simulated crayfish afferent spike versus a real laboratory recording of a crayfish afferent spike. The shape and duration of the spikes are closely matched. The simulated spike was obtained by using Hodgkin-Huxley channels with modified values for the constants in equations 4.9 to 4.14.

The membrane constants of the lateral giant were chosen so that the model matches the experimental membrane time constant.

5.2 An accurate morphology compartmental LG model

The model described in section 5.1 generalises the essential elements of the lateral giant's morphology. In this section we describe a more accurate model, composed directly from the morphological data. In section 5.2.1 we describe how our model can be linked with other cells.

The compartmental model was developed in two stages. The first stage is to produce a morphological model, without electrical properties, using biophysical measurements. A set of 17 confocal stacks of a dye-filled lateral giant neuron, taken from a 4 cm animal, was imported into the Neurolucida (Microbrightfield Inc., 2003) software package. A projection of the confocal stacks is shown in figure 5.7. Neurolucida was used to manually trace around the dendritic structure of the lateral giant to measure the dendrite diameters and lengths. Neurolucida compiles these measurements to construct a morphological model of the cell as a composition of cylinders. A three dimensional rendering of the morphological model of the lateral giant can be seen in figure 5.8.

A schematic view of the morphological model is shown in figure 5.9. The morphological model is made up of a total of 1114 branch segments.

Th second stage of the development of the model was importing the morphological model into NEURON to produce a compartmental model. A script written in Java was used to convert from the Neurolucida format into a NEURON compartmental model. The number of compartments for each branch segment, was expressed as a function of the axial resistance and the membrane resistance. No compartment is allowed to be longer than 0.05 of the steady-state electrotonic length constant, λ. The number of compartments changes dynamically with the values set for axial resistance and

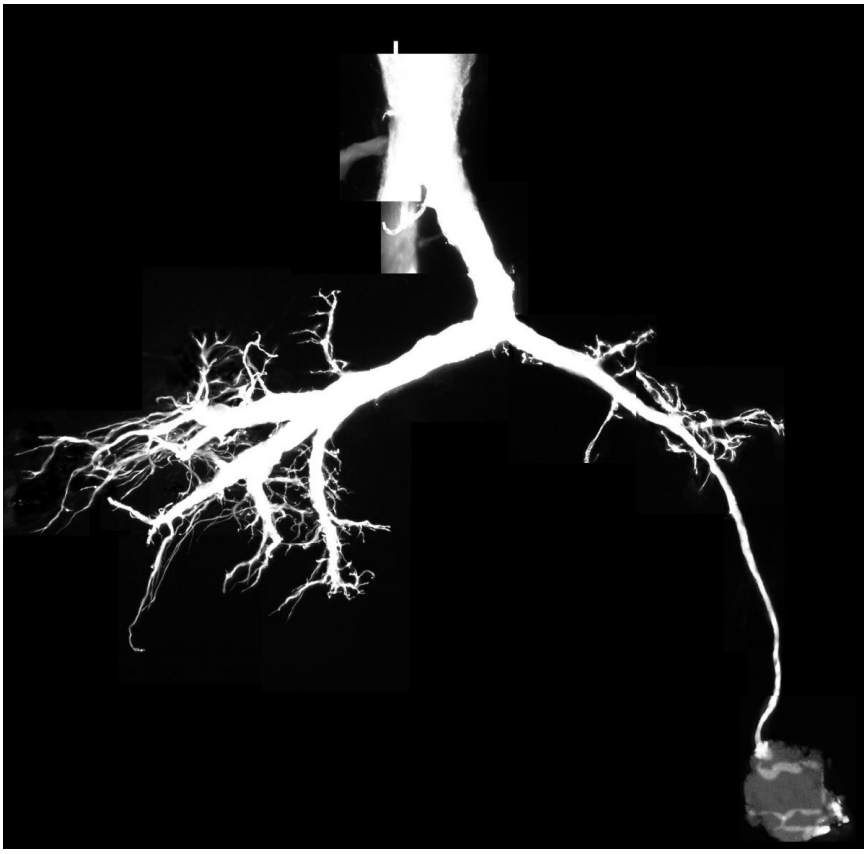

Figure 5.7: A projection of 17 confocal stacks showing a dye filled lateral giant neuron.

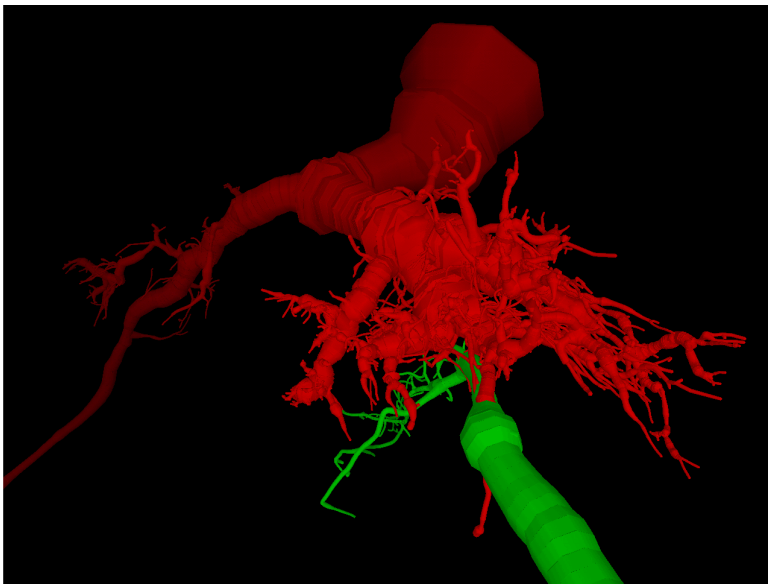

(a)

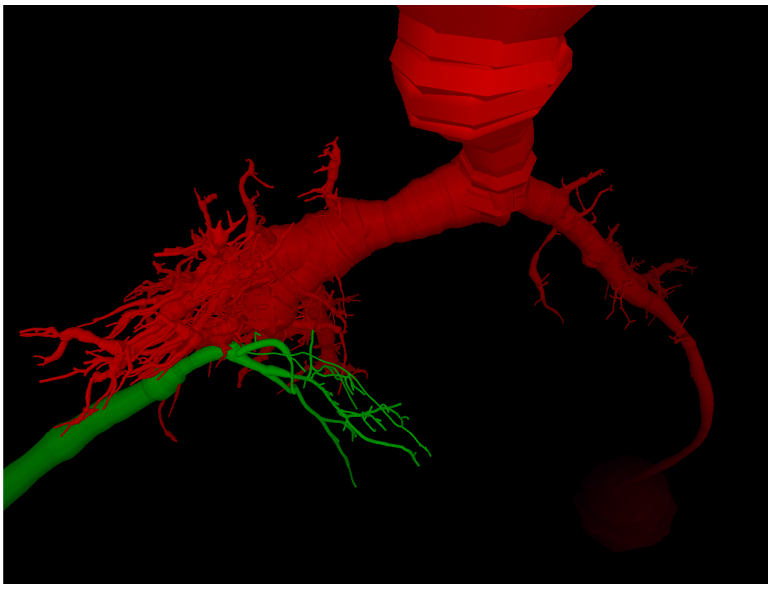

(b)

Figure 5.8: A 3D rendering of the morphological model of the lateral giant (large neuron in red). Also shown in the picture is a model of a primary afferent axon (in green.)

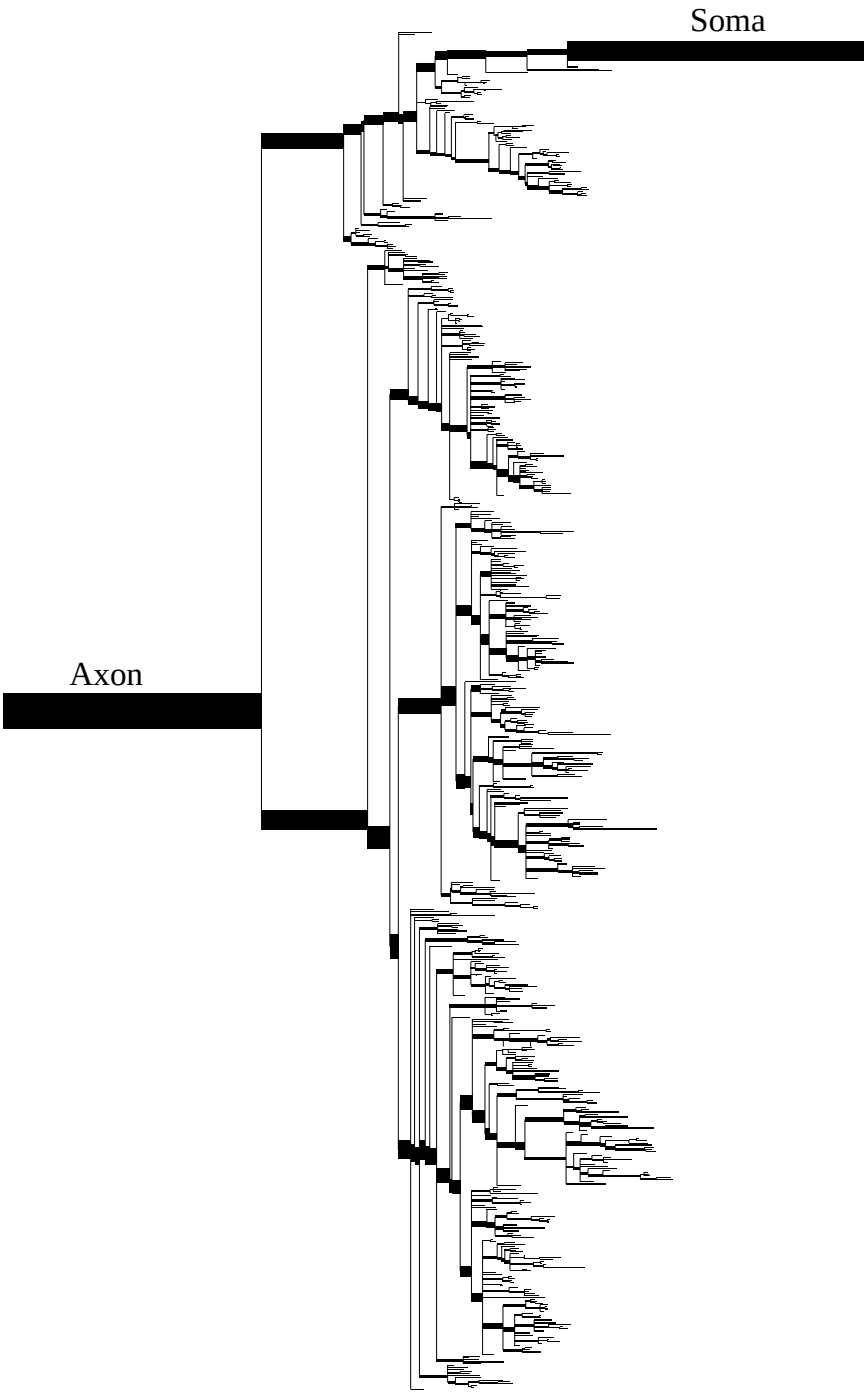

Figure 5.9: A schematic diagram of the lateral giant. Branch lengths and diameters are drawn to scale.

membrane resistance. In practice most branch segments were represented with a single compartment. Multiple diameter measurements within branch segments were handled using NEURON's `pt3dadd()` function. NEURON combines the subsections to calculate the total capacitance, membrane resistance and axial resistance for the compartment.

The following choices were made in the lateral giant model generated by the script:

- A passive membrane was used in the lateral giant dendrites.
- Hodgkin-Huxley channels were used in the lateral giant axon. The parameters were adjusted so that the shape of the action potential conforms with experimental recordings using the same procedure as described in section 5.1.1.
- The value used for axial resistivity was 60 Ωcm. This is the measured cytoplasmic resistivity in the crayfish medial giant (Glantz and Viancour, 1983) and consistent with other measures of cytoplasmic resistivity in crustaceans (Katz, 1966).
- The membrane resistivity was adjusted so that the response of the simulated cell to a current clamp was consistent with experimental recordings.

5.2.1 Selecting LG synapse points

When we want to use our model of the lateral giant in a multi-cell simulation, it is necessary to connect the lateral giant to other cells. To do this we need to decide the points on the lateral giant dendrites where the other cells connect to, *i.e.* the synapse points. In our multi-cellular models the lateral giant receives its input from primary afferents. We therefore need to identify

the points at which the primary afferents synapse onto the lateral giant.

The exact points on the dendritic tree where afferents make synapses is not included in the morphological data. Precise data of all the synapse locations is difficult to obtain. Using anatomical observations (Antonsen and Edwards, 2003) of where synapses occur in the lateral giant, a method was formulated to identify synapse location candidates in the model. This method was then applied to the model to pick a set of synapse locations. The method used to pick the candidate set is described in the following algorithm:

let d_{synthr} be a threshold diameter value
let the set B contain all the dendritic branches within the five main branches
foreach $b_i \in B$ **do**
  let d_i be the diameter of the branch b_i
  let the set C contain the diameters of b_i's child dendrites
  if $\exists d_j \in C$ *such that* $d_j < d_{\text{synthr}}$ **or**
    b_i *is a terminal branch* **and** $d_i \geq d_{\text{synthr}}$
  then
    the branch b_i is a synapse point candidate
  end
end

Figure 5.10 marks the locations of synapse point candidates after applying the algorithm.

5.2.2 Simplifying the model

For many simulations we are only interested in what happens in part of the lateral giant. For example we may be providing all of our input onto only one major branch. In that case we are probably interested in how PSPs and current spread throughout the stimulated branch and how it is integrated at the integrating segment and the axon, but we are probably not very interested in what happens in detail at the non-stimulated branches.

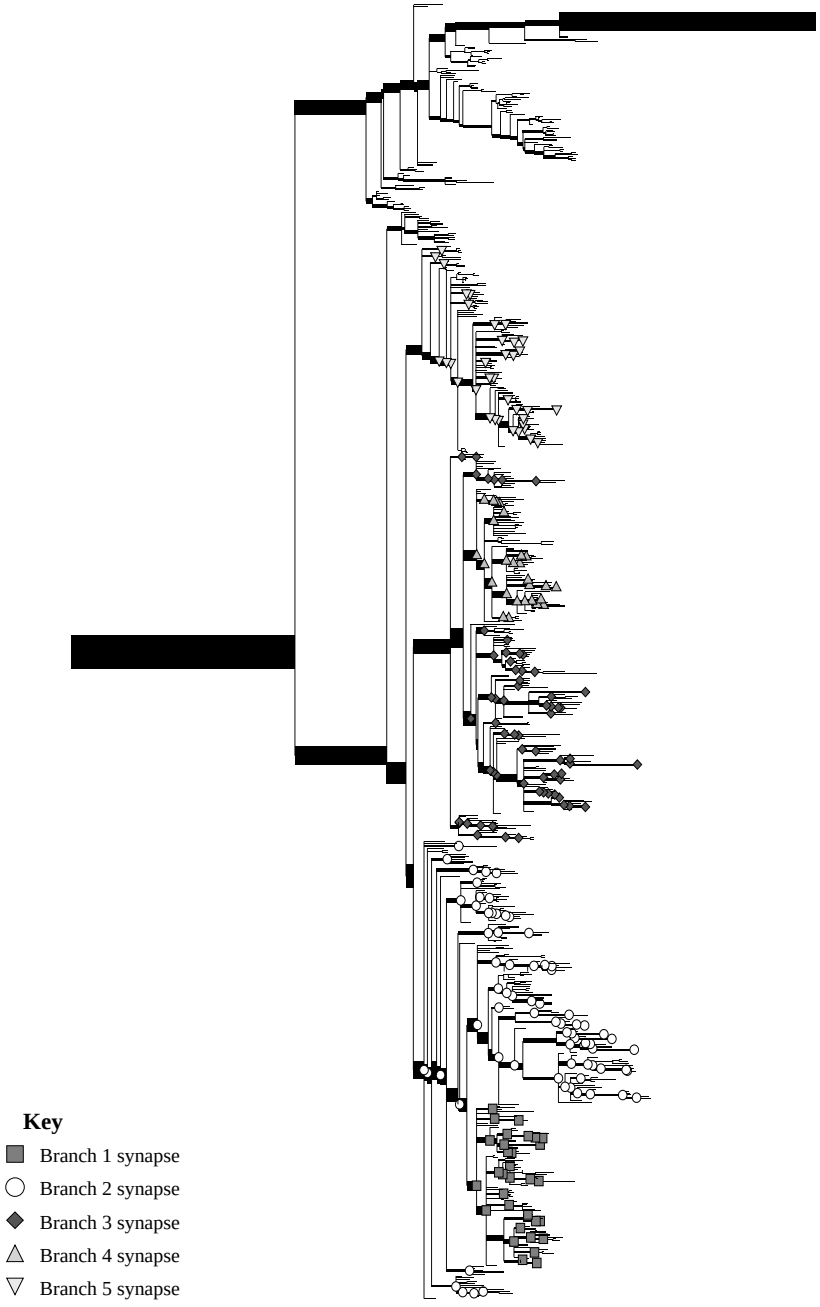

Figure 5.10: The candidate synapse location points on the lateral giant, identified with the algorithm described in section 5.2.1.

Calculating how current spreads over the entire tree is very computationally expensive.

There are various methods for reducing the number of compartments in a model neuron (Bush and Sejnowski, 1993). The approach we use is to approximate the complex structures of the parts of tree that we are not interested in with two or three compartments. We ensure that the electrical load of the approximating compartments matches that of the original structure.

For most of our simulations we were only interested in what happens in one of the lateral giant's major branches. For each other major branch we approximated it with a three compartment model using the following steps:

1. The transfer function of the detailed compartmental model of the branch was found. This was determined by applying a step current and recording the response. Since our model of the lateral giant dendrites is passive, the electrical circuit is composed entirely of resistors and capacitors. It is therefore a linear time invariant system, albeit a very complicated one. The response of a linear time invariant system to a step function gives the transfer function.

2. A three compartment model that has a transfer function which closely matches that of the original branch model was found. Each compartment has three free variables: the axial resistance, the membrane resistance and the capacitance. This gives a total of nine free variables. The Principle Axis method (Brent, 1973) was used to find suitable values for these parameters. The Principle Axis method is able to find the minimum of a multi-variable function, $f(\mathbf{x})$, without requiring the derivative. The NEURON simulator provides an implementation of the Principle Axis method, from the SCoP Math Library (National Biomedical Simulation Resource, Duke University), in its toolkit. This

tool was used to estimate the parameter values. The function given to the optimiser to solve was the sum of the mean squared difference between the step response of the original branch morphology and the three compartment approximation.

This method was used to simplify four of the major branches and the soma branch. This left one major branch still modelled in detailed morphology. Figure 5.11 shows a diagram of the simplified model. The simplified model contains 281 compartments versus the original's 1113. Running a simulation with the simplified model is approximately four times as fast as with the full lateral giant morphology.

Tests were conducted to verify that the simplified model performed similarly to the original. Both variants of the model were current clamped and voltage clamped and the responses in the dendrites in branch 2 were recorded. For each voltage clamp and current clamps a step signal and 500 Hz sinusoidal signal were given. This was to test the cell response to both direct current and high frequency alternating current. The recorded output of the two models was almost inseparable. Figure 5.12 shows a comparative plot of the response of the two models to sinusoidal current. The response of the original model to the 500 Hz sinusoidal current was very slightly faster: the original model's response led the simplified model by approximately 0.01 milliseconds.

5.3 Simulating the rectifying junctions

The rectifying junctions have been identified as an important functional component of the lateral giant escape circuit. The behaviour of the rectifying junctions are described in section 3.3.4. Equation 3.1 defines the variable conductance across the rectifying gap synapses.

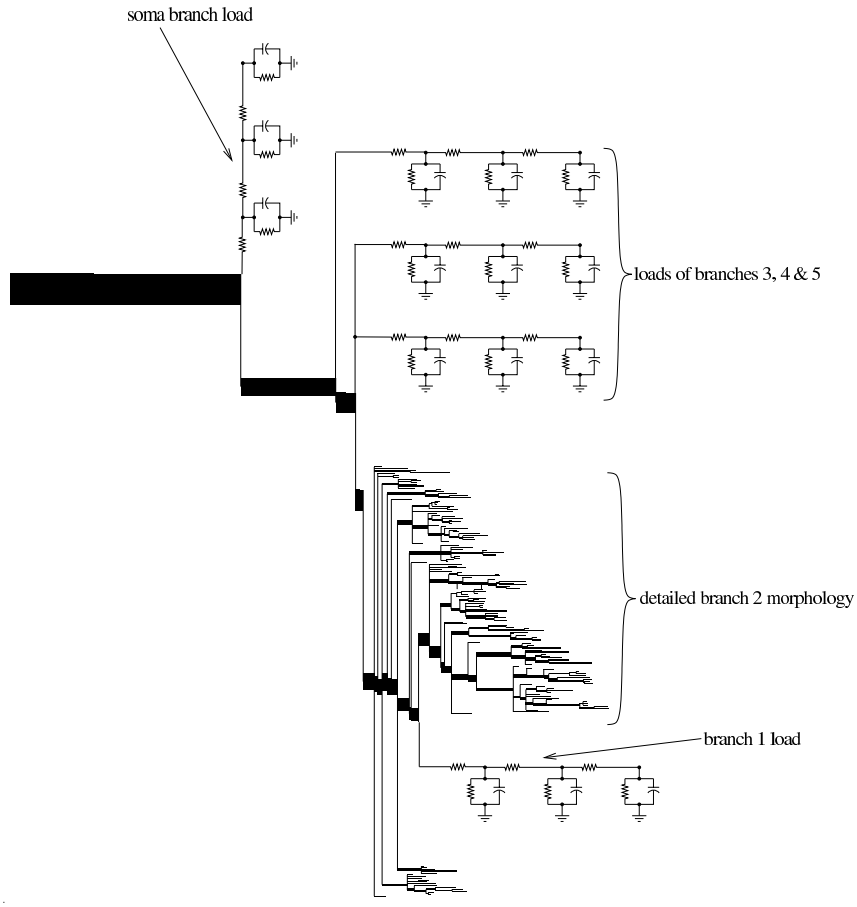

Figure 5.11: The morphological model of the lateral giant is simplified by keeping the detailed morphology of only one of the major branches (branch 2) and replacing each of the other major with three compartments that closely approximate their electrical load.

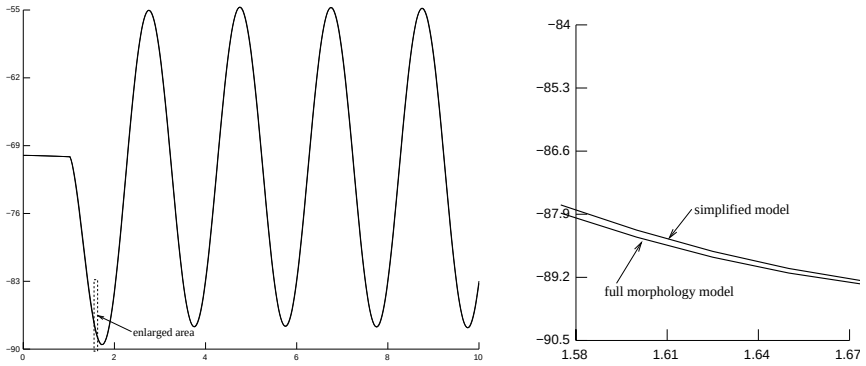

(a) *(b)*

Figure 5.12: The responses of the full morphology model and the simplified model to a 500 Hz sinusoidal current injection are plotted on the same axes. Only by enlarging the graph to a very short time scale (see (b)) is one able to distinguish the results.

Neither GENESIS or NEURON has built-in support for rectifying junctions. The rectifying junctions were implemented as extension objects. A rectifying synapse was implemented as a pair of customised current sources. One current source was located at the pre-synaptic terminal and another at the post-synaptic terminal. The rectifying junction objects were given pointers to the voltages at the pre-synaptic and post-synaptic terminal. At each simulation time step the synapse conductance was calculated. This was done by calculating the steady state conductance as defined by equation 3.1 and then updating the instantaneous conductance. The instantaneous conductance converges exponentially towards the steady state conductance with a time constant τ. Equation 5.3 shows how the instantaneous conductance is calculated:

$$g[t] = g[t-1] + \Delta g \frac{\Delta t}{\tau} \tag{5.3}$$

$$\Delta g = g_{\text{steady}} - g[t-1]$$

$$\tau = \begin{cases} \tau_{\text{open}} & \text{if } g_{\text{steady}} > g[t-1] \\ \tau_{\text{close}} & \text{otherwise} \end{cases}$$

Where $g[t]$ is the conductance of the synapse at time t, g_{steady} is the steady state conductance calculated according to equation 3.1 and Δt is the numerical time step of the simulation.

The current injected into each cell was calculated according to Ohm's Law as the product of the synapse conductance and the voltage difference between the pre and post synaptic compartments. At the post-synaptic terminal the injected current is thus:

$$i[t] = (V_{\text{pre}} - V_{\text{post}})g[t] \tag{5.4}$$

Where V_{pre} is the pre-synaptic membrane potential and V_{post} is the post-synaptic membrane potential.

According to Kirchoff's Current Law, the net current going into the post-synaptic terminal is necessarily equal to the net current leaving the pre-synaptic terminal. For computational efficiency the conductance was only calculated at the post-synaptic terminal. The pre-synaptic rectifying junction object was given a pointer to the conductance at the pre-synaptic terminal and this value was used to calculate current.

The source code of the NEURON implementation of rectifying gap junction objects is listed in appendix C.1.

5.4 Dendritic tips

One striking feature of the lateral giant's anatomy is that there are very narrow (less than 1 μm in diameter) dendrites at the ends of the tree. This is in stark contrast to the wide dendritic base of the tree and giant axon the cell projects. The afferents make their synapses near these very narrow dendrites (Herberholz et al., 2002; Atwood and Pomeranz, 1977). Figure 5.13 shows

some narrow dendrites wrapping around an afferent. At the point of contact the dendrite forks into very narrow branches and these branches wrap around the afferent. There is some ambiguity over the actual synapse location. There are effectively three possibilities:

1. there may be a single large synapse at the fork;
2. there may be smaller synapses along the fine dendrites; or
3. synapses may exist at the fork and along the fine dendrites.

We investigate what effect the location in the dendrites of the afferent to lateral giant synapses has on the cell's integration of input.

5.4.1 Simulation setup

The model consisted of a single afferent connecting onto a single branch of the lateral giant. A branch from the morphologically realistic model of the lateral giant was selected to be included in the model. The branch selected was a likely synapse point as identified by the method described in section 5.2.1. The rest of lateral giant structure modelled with three compartments to model its electrical load. The simplification method used is described in section 5.2.2. The setup of the model components is drawn in figure 5.14.

The afferent was connected to the lateral giant branch with rectifying junctions (as described in section 5.3.) The parameters used for the rectifying junctions were $\tau_{\text{open}} = 0.5\text{ms}$, $\tau_{\text{close}} = 0.5\text{ms}$, $V_0 = 0$, $A = 0.15\text{mV}^{-1}$ and $\frac{g_{\max}}{g_{\min}} = 100$. Alternative values were used for $g_{\max}$, varying from 0.01 to 0.15 μS.

The synapses were located at the three different possible configurations. This is illustrated in figure 5.14. Let us call a single synapse at A 'config-

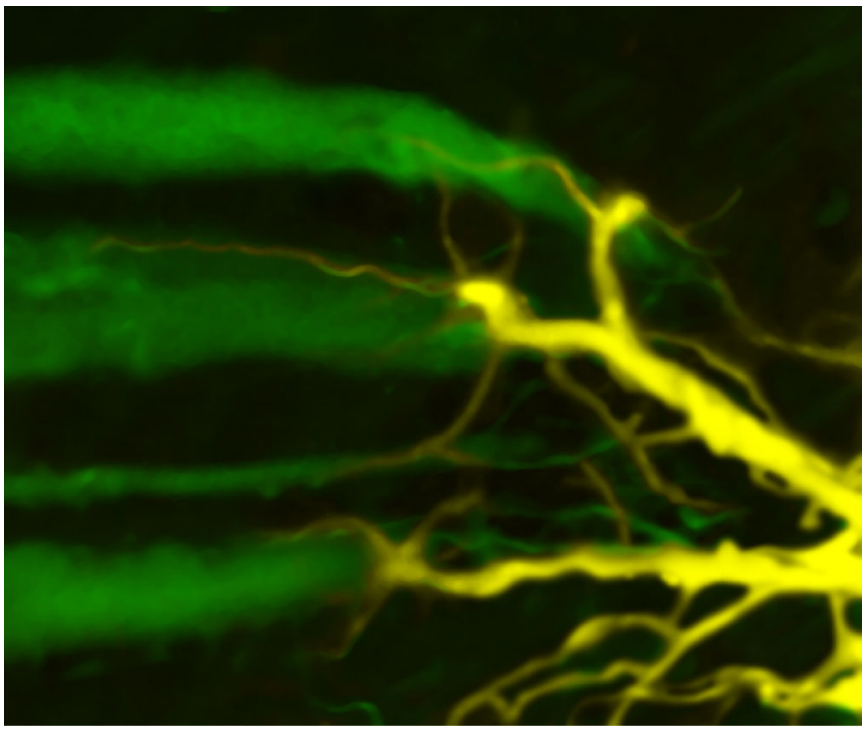

Figure 5.13: An image of dye filled afferent axons and the lateral giant dendrites showing the contact point between an afferent and the lateral giant. At the contact point the lateral giant dendrite forks into very narrow branches which wrap around the afferent axon. Picture: (Antonsen and Edwards, 2003).

uration 1', multiple synapses at locations B 'configuration 2', and synapses at both locations A and B 'configuration 3'.

The sum of the synapse conductances used in the three alternative configurations was made to be equal. Let us use the notation g_{A1} to denote the maximum synapse conductance at point A for configuration 1. To ensure that the total conductance for configuration 2 is equal to that of configuration 1:

$$g_{B_2} = \frac{g_{A_1}}{3} \tag{5.5}$$

For configuration 3, the conductance at A was made to be double that of the conductance of a synapse at point B. Hence:

$$g_{A_3} = \frac{2g_{A_1}}{5} \tag{5.6}$$

$$g_{B_3} = \frac{g_{A_1}}{5} \tag{5.7}$$

For each configuration the afferent was fired and the simulated EPSPs in the lateral giant dendrite were recorded so that the effects of the synapse location could be compared.

To evaluate the effects of synapse location on antidromic transmission the simulated lateral giant was spiked. This was repeated for all three possible synapse location configurations. The PSPs in the afferent were recorded so they could be compared.

5.4.2 Simulation results

Figure 5.15a compares the EPSPs produced in the lateral giant for the three synapse configurations considered. The EPSPs produced in LG when the contact point is at the wider dendrite (contact point A) are larger than those obtained when the contact points are out in the fine dendrites (contact point

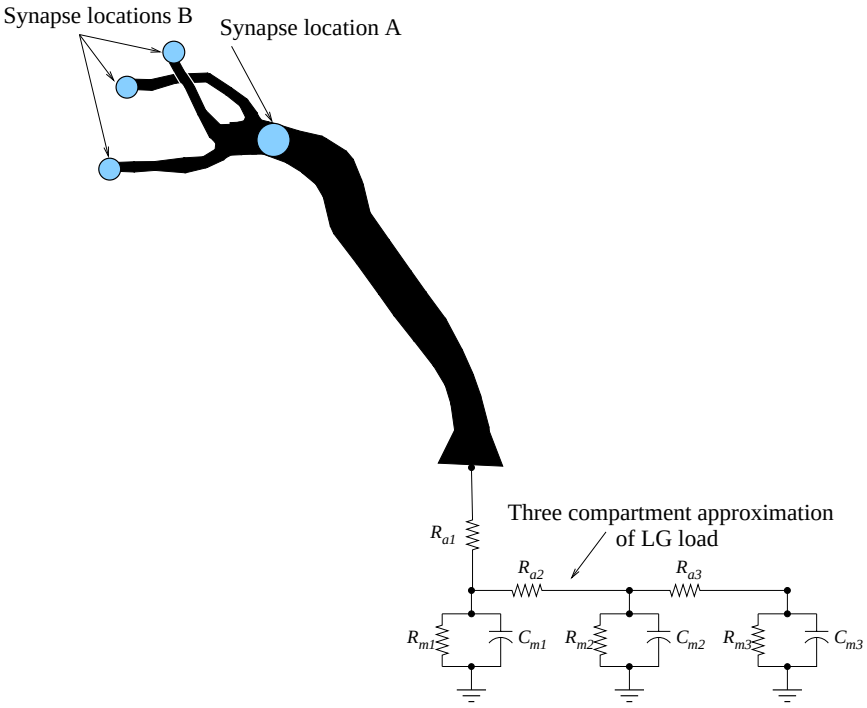

Figure 5.14: Three alternative configurations were tried for the locations of synapse points: (1) at location A: single synapse located on the larger branch, (2) at locations B: three smaller synapse located on the ends of the fine dendrites and (3) a combination of locations A and B.

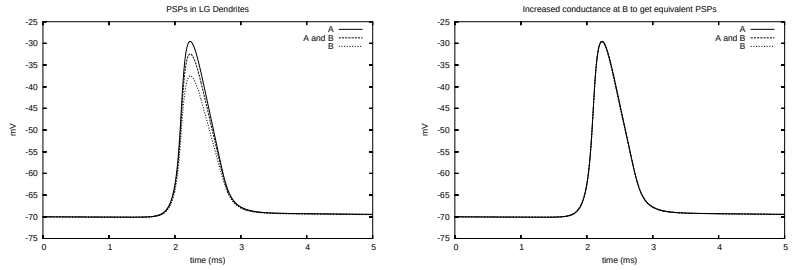

(a) *(b)*

Figure 5.15: A comparison of the orthodromic EPSPs produced in the lateral giant dendrites for the three alternative synapse location configurations. (a) Shows the EPSPs when the total synapse conductance is kept equal for all three configurations. (b) Compares the shape of the EPSPs when the synapse conductance of alternative locations is unequal but adjusted so that all configurations produce a EPSP of equal magnitude.

B.) This is true if for each synapse location setup we use an equal total conductance. Although the size of the EPSP is changed, the EPSP shape appears similar for all three synapse location schemes.

To test whether the EPSP shape is in fact unchanged and whether equivalent EPSPs can be obtained using contact points at B if higher conductance values are used, we increased the conductances for the fine dendrite synapses. Figure 5.15b compares the EPSPs when the conductances at B have been increased. One can see that there are no major differences in the shape of the EPSP regardless of the synapse location.

This result demonstrates that fine dendrite synapses can produce similar EPSPs as a synapse located at the fork provided the conductance at the fine dendrite synapse has been upwardly adjusted. The factor by which the fine dendrite conductance needs to be increased (fine dendrite synapse conductance divided by large dendrite synapse conductance) increases with the large synapse conductance. There exists a threshold beyond which even an infinitely conducting fine dendrite synapse is unable to equal the EPSP

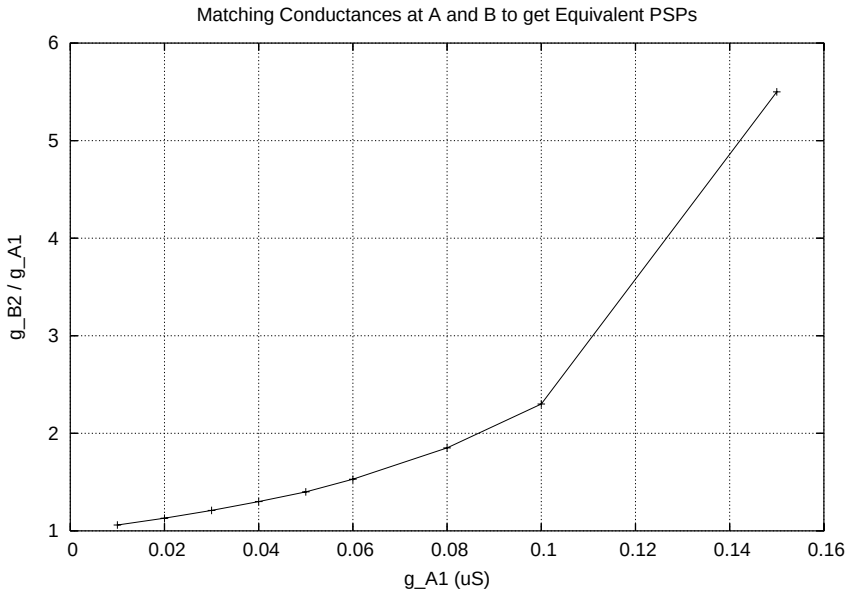

Figure 5.16: The increased conductance required at the fine dendrite synapses (B) in order to obtain the same sized EPSPs as when the dendrites are located at the branch fork (A).

of a large dendrite synapse. This point is reached when the resistance of the through the fine dendrites exceeds the resistance through the large dendrite synapse. Figure 5.16 plots the relationship between large dendrite synapse conductance and the equivalent required fine dendrite synapse conductance.

The antidromic EPSPs in the afferent, caused by spiking LG, are similar in size regardless of the synapse location (see figure 5.17a.) Although the EPSPs are slightly larger when the synapses are located at the larger dendrite, the difference is only very minor. However if one uses higher conductances at the fine dendrites, such that the orthodromic PSPs are of equivalent size, then the antidromic EPSPs will be higher for the fine dendrite pathways (see figure 5.17b.)

The results for configuration 3 sit neatly between those of configurations 1 and 2. Having synapses at both the larger dendrite and on the fine dendrites is not as efficient as having only a large synapse at the large dendrite, but

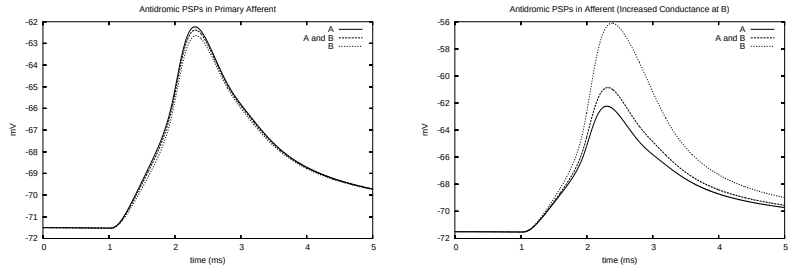

(a) *(b)*

Figure 5.17: A comparison of the antidromic EPSPs produced in the lateral giant dendrites for the three alternative synapse location configurations. (a) Shows the EPSPs when the total synapse conductance is kept equal for all three configurations. (b) Compares the size of the antidromic EPSPs if the synapse conductances have been adjusted to produce equivalent orthodromic PSPs.

more efficient than having synapses only at the fine dendrites.

5.4.3 Discussion

The location of the synapse in the dendrites does not have a big effect on the shape of the EPSP. However synapses located in the fine dendrites will produce smaller orthodromic PSPs than synapses located at a larger dendrite, if the total synapse conductance is the same for both cases. The synapse location has very little effect on antidromic transmission.

By using a higher conductance at a fine dendrite synapse it is possible to produce equivalent sized PSPs as a lower conductance large dendrite synapse. This is possible so long as the axial resistance through the fine dendrite pathway does not exceed the resistance of the large branch synapse. However a higher conductance at the fine dendrites causes the antidromic PSPs to be larger. Thus, it is not possible to choose conductance values such that fine dendrite synapses operate equivalently to large dendrite synapses if one

considers transmission in both directions.

The cause of the differing results for antidromic and orthodromic transmissions is the rectifying junctions. The rectifying junctions' complex dynamics and varying conductance cause radical changes in the circuit for differing transmission directions. In tests where the rectifying junctions were replaced with purely ohmic gap junctions, changing the synapse locations caused proportionally equal changes to the PSP magnitude for both transmission directions.

These results have implications for the models of the lateral giant circuit described in the following chapters. If what we are only considering transmission in one direction in our model then the location of the synapse is not important. In effect, the synapse location is merely a multiplying factor applied to the synapse conductance. If the model results are dependent on both orthodromic and antidromic transmission then the location of the synapse needs to be considered as an extra variable that will affect the results.

If you would like to know how the model lateral giants are able to detect simultaneous excitation you will have to read the next chapter.

CHAPTER 6

Coincidence detection in the Lateral Giant

We now focus our model on simulating coincidence detection in our computational models. In particular we will examine the role that the rectifying junctions play in coincidence detection. The lateral giant has been shown to be an effective coincidence detector (Edwards et al., 1998). This is a critical capability in the circuit in being able to robustly distinguish a predator attack. Being able to reproduce this capability in our models would be a useful result in our quest of understanding robust biological systems.

We now test the effect that the rectifying junctions can have on coincidence detection in simulations of our morphological and anatomically realistic model of the lateral giant. This chapter describes simulations which test hypotheses regarding the role of rectifying junctions in coincidence detection. Previous simulations (Edwards et al., 1998), using single compartment models of neurons, have shown that rectifying junctions can affect coinci-

dence detection. We now test what role the rectifying junctions play in the lateral giant using realistic models of the morphology. We will attempt to determine what parameter settings of the rectifying junctions give the best results for effective coincidence detection.

We examine coincidence detection in the two alternative models of the lateral giant described in chapter 5. We simulate coincidence detection in the simplified model in section 6.3 and in the morphologically accurate model in section 6.4. [1]

The setup of the simulations in the two models differ, to allow us to ask slightly different questions. A consequence of this is that not all of the results are directly comparable between the two models.

The questions we want to address are:

- Can we demonstrate that we can reproduce the functionality of coincidence detection, that is observed in a crayfish, in an artificial system based on simulating the crayfish neurons?
- What effects do rectifying junctions have on the summation of inputs to the lateral giant? How does the response obtained when rectifying junctions are used compare to that obtained when ohmic gap junctions are used?
- Can rectifying junctions enhance coincidence detection in the lateral giant?
- What are the necessary parameter settings to obtain good coincidence detection?

[1] The coincidence detection simulations done using the simplified model in section 6.3 were done prior to the direct neuroanatomical data being available. There is some overlap between the results in sections 6.3 and 6.4. Both sections are included to show that the conclusions drawn in this chapter hold for both models.

- Do the rectifying junctions provide a sufficient explanation of how the lateral giant achieves coincidence detection?

- Are the rectifying junctions necessary to explain coincidence detection?

- How do the simulation results obtained from the abstracted morphology model compare to those from the morphologically realistic model?

The rest of this chapter is concerned with answering these questions. We first present some definitions so that we may objectively measure coincidence detection. We then explain the parameters that govern the behaviour of the rectifying junctions. We then set up a series of simulations using both our abstracted morphology model and our realistic morphology model to address the questions we have just mentioned. We then draw on the results of these simulations to answer these questions.

6.1 Quantifying Coincidence Detection

In order to determine under what conditions the lateral giant acts best as a coincidence detector, it is necessary to define some quantitative measures. To this end we now introduce two metrics.

The efficiency ratio is a measure of how effective the lateral giant is at summing multiple inputs. Imagine that we have two groups of afferents: group 1 and group 2. If we stimulate the afferents in group 1, then we effect a PSP in the lateral giant, let v_1 be the peak PSP. Let v_2 designate the PSP that results when group 2 is stimulated. Finally, let us call v_{peak} the PSP obtained when both groups are stimulated at some time, Δt, apart. Let us

define the *efficiency ratio*, E, as:

$$E(\Delta t) = \frac{v_{\text{peak}}(\Delta t)}{v_1 + v_2} \tag{6.1}$$

An efficiency value of 1 signifies a perfect linear summation of the two groups of inputs. We would always expect the efficiency to be less than 1 (barring active currents or weird nonlinear effects.) An efficiency value of 0.5 is the lower bound of inefficiency. If the efficiency is 0.5, then the PSP obtained from stimulating both groups is no greater than the response of stimulating just one of the groups.

A good coincidence detector has a high efficiency when the two groups are fired simultaneously and a rapidly declining efficiency as the interval of time between firing the two groups is increased.

The *coincidence ratio* is a measurement of the comparative decline in PSP magnitude that occurs when the two groups of afferents are fired out of coincidence. The coincidence ratio is defined as the ratio of the peak PSP obtained when the two inputs are separated by a time Δt to the PSP obtained when the two inputs are applied synchronously as shown in equation 6.2.

$$C(\Delta t) = \frac{v_{\text{peak}}(\Delta t)}{v_{\text{peak}}(0)} \tag{6.2}$$

where C is the coincidence ratio and v_{peak} is the peak PSP. Figure 6.1 graphically illustrates the measures from which the coincidence ratio is calculated.

If we combine equations 6.1 and 6.2, the coincidence ratio may also be expressed in terms of efficiency:

$$C(\Delta t) = \frac{E(\Delta t)}{E(0)} \tag{6.3}$$

By definition, the coincidence ratio is always 1 at $\Delta t = 0$. In a coincidence

detector, the coincidence ratio will decline as Δt is increased. Assuming our two inputs are equal, then 0.5 is the lower bound on the value of coincidence ratio. This value indicates that if we desynchronise the input by Δt the second input causes no increase to the peak response. Conversely, a coincidence ratio of 1, for some Δt, indicates that the two inputs sum together just as efficiently as when they were synchronised. Coincidence ratios greater than 1 signify that the responses sum more effectively when the input is desynchronised to when it is synchronous, *i.e.* the circuit is acting as an 'anti-coincidence' detector. A good coincidence detector has a coincidence ratio approaching 0.5 as we increase Δt.

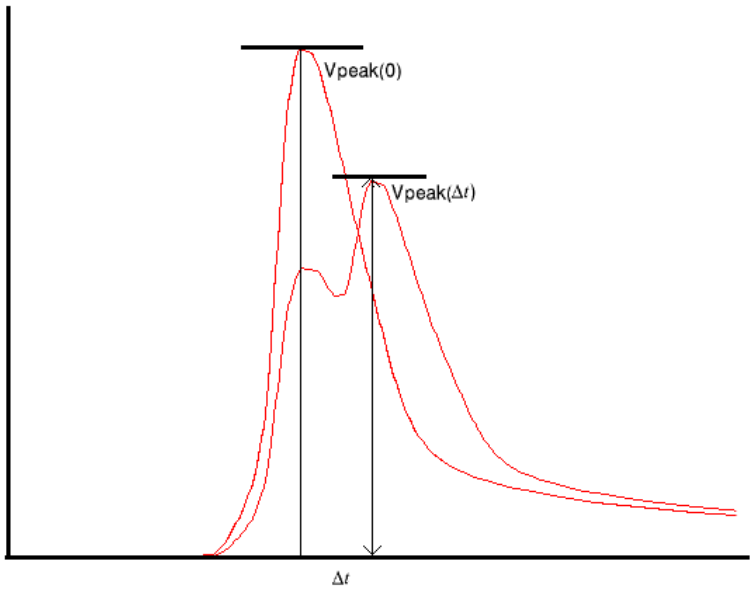

Figure 6.1: Measuring coincidence detection. The coincidence ratio is the ratio of the asynchronous input PSP ($v_{\text{peak}}(\Delta t)$) to the synchronous input PSP ($v_{\text{peak}}(0)$).

6.2 The rectifying junction parameters

The sigmoidal equation (see equation 3.1) that describes the conductance across a rectifying synapse is governed by six constants:

- $g_{\min}$
- $g_{\max}$
- A
- V_0
- τ_{open}
- τ_{close}

We now explain how each of these constants affect the operation of the synapse. $g_{\min}$ and $g_{\max}$ determine the minimum and maximum conductance of the synapse, respectively. V_0 is the *voltage activation threshold*. When the voltage difference between the pre and post synaptic terminal is equal to V_0, then the synapse operates at its mean conductance (*i.e.* $\frac{g_{\min}+g_{\max}}{2}$.) Increasing V_0 causes the synapse to require a greater voltage difference to move into a high conductance state. The constant, A, determines the slope of the transition between low conductance and high conductance. A higher value of A narrows voltage range over which the synapse changes from low conductance to high conductance. These four constants govern the steady state conductance of the synapse. Figure 6.2 illustrates how these constants affect the steady state conductance response of the synapse. τ_{open} and τ_{close} are the time constants. They determine how quickly the junction moves from its current conductance to the steady state value. There are different time constants for when the synapses move to a more open state or a more closed state.

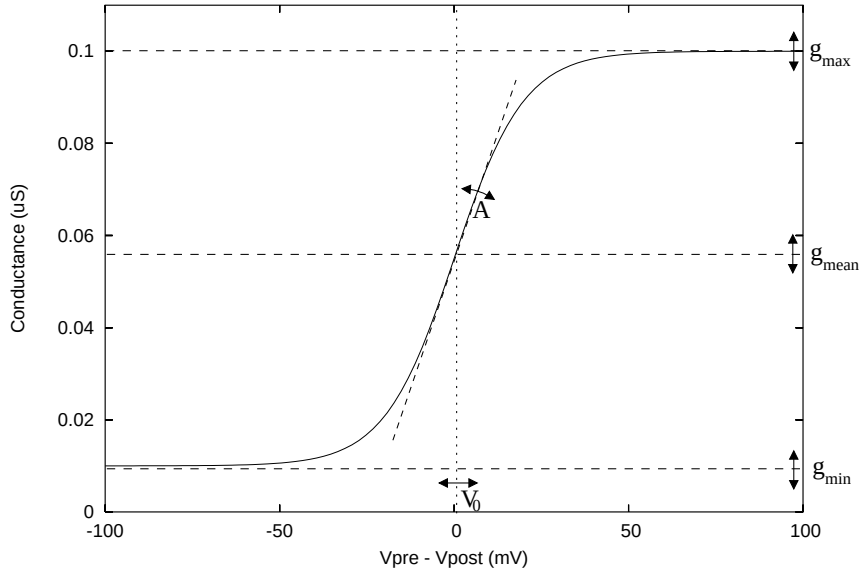

Figure 6.2: The steady state conductance of a rectifying junction versus the voltage difference between the pre and post synaptic terminals. The steady state conductance curve is governed by four constants: $g_{\min}$, $g_{\max}$, V_0 and A. $g_{\min}$ and $g_{\max}$ determine the minimum and maximum conductances, respectively. Changing V_0 moves the curve to the left or right. A affects the slope of the transition from minimum to maximum conductance. The curve plotted has $g_{\min} = 0.01\mu S$, $g_{\max} = 0.1\mu S$, $V_0 = 0$ and $A = 0.1\mathrm{mV}^{-1}$.

It is sometimes useful to re-factor the conductance variables in terms of g_{mean} and $\frac{g_{\max}}{g_{\min}}$, where g_{mean} is defined as:

$$g_{\text{mean}} = \frac{g_{\max} + g_{\min}}{2} \tag{6.4}$$

This allows one to change the amount of rectification of the synapse while keeping the mean conductance constant, thus allowing a fairer comparison of two synapses of different rectification.

It is not possible to experimentally directly measure the parameters of the rectifying junctions onto the lateral giant. This is due to them being located at the ends of the lateral giant dendrites. These dendrites are too small and narrow to probe with an electrode. It is therefore not possible to directly measure the voltage characteristics across the junctions and thereby infer their parameters. In our simulations we therefore use the experimentally measured parameters of the giant motor synapse (Furshpan and Potter, 1959) as a starting point for the lateral giant junction parameters. Parameter searching and estimation is then used to optimise the rectifying junction parameters to obtain agreement with the biological observations.

6.3 Coincidence detection in the abstracted morphology model

We now test coincidence detection in the abstracted morphology model, that is described in section 5.1.1. Our aim is to determine what contribution the rectifying junctions make to coincidence detection. We measure our model lateral giant's response to different input stimulations. To determine the necessary parameter settings for effective coincidence detection, different rectifying junction parameter settings are tried.

To test this question, a simulation was setup, comprising of the lateral giant and a group of afferent axons providing input. The afferent axons were modelled as described in section 5.1.1. The afferents formed rectifying synapses (as described in section 5.3) onto the terminal ends of the finest lateral giant dendrites. The constants used to model the membrane properties are listed in appendix B.1. The simulation was run by spiking a subset of the afferents and recording the response in the lateral giant.

As a starting point, we recorded the response of the lateral giant to afferent input using rectifying junctions with parameters similar to that measured in the giant motor synapse (GMS). The coincidence ratio and efficiency was calculated for these settings.

We then searched for parameter values that would give better coincidence detection by conducting a parameter search. Parameter exploration was undertaken in two stages:

1. Different combinations of the rectifying junction parameters were tried to determine which settings give the best coincidence detection.

2. For some specific rectifying junction settings, as determined by the previous step, the integration response of the lateral giant was evaluated for different input combinations.

In our search for the best coincidence detector, we calculated the coincidence ratios in response to a specific kind of input to the lateral giant: two groups of ten afferents. In order to be able to calculate the coincidence ratio, the two groups were fired synchronously and 0.5 milliseconds apart. All twenty afferents connected to dendrites on the same major branch. Terminal dendrites from different minor branches were equally represented in choosing the synapse points. Figure 6.3d illustrates where the afferents make their synapses. This input represents a medium level of input to a branch of the

Parameter	Values	Units	Number of data points
$g_{\max}$	1, 2.15, 4.64, 10, ... , 1000	nS	10
$g_{\max}/g_{\min}$	1, 10, 100, 1000	1	4
V_0	0, 5, 10, 15, ... , 70	mV	15
A	0.02, 0.04, 0.1, 0.2, 0.5, 1, 2, 4	mV^{-1}	7
τ_{open}	0, 0.1, 0.2, 0.3, 0.4, 0.5	ms	6
$\tau_{\mathrm{open}}/\tau_{\mathrm{close}}$	0.5, 1, 2, 4, 8	1	5
Total number of combinations			105000

Table 6.1: The tested rectifying junction parameter values.

lateral giant, separated by an interval where strong coincidence detection is expected to occur. The limited number of inputs reduces the computational expense, which is desirable in a large parameter search. This choice of input was therefore deemed a suitable input for repeated testing, in search for effective coincidence detection.

For this specific kind of input, different combinations of the rectifying junction parameter settings, were tried. Each parameter was varied along a range, using its experimentally measured value at the GMS as a mid-point (Giaume et al., 1987). The range by which each parameter was varied, extended beyond the bounds of biologically plausibility. This was done in order to gain an understanding of how the parameter can affect the results at the extremities. Table 6.1 tabulates the tested range for each rectifying junction parameter. In each simulation run, the parameter settings were applied to all of the rectifying junctions. All synapses were assumed to be identical.

For each parameter combination the PSPs caused in the lateral giant dendrites were recorded and the coincidence factor was calculated using the method described in section 6.1. We then examined the results to search for which parameters give the lowest coincidence ratios. We are thus able to ascertain, which parameter settings give the best coincidence detection

under the described input conditions.

For parameter setting identified as giving good coincidence detection, a second stage of parameter exploration, concerning the input, was undertaken. The aim was to determine how the circuit acts as a coincidence detector to different kinds of input. Different combinations of input were given to the lateral giant to test how it integrates these inputs and how different inputs affect coincidence detection.

The different combinations of input that can be given to the lateral giant is of course infinite. This holds true even within the subset of inputs consisting two groups of afferents, where within each group all afferents are fired simultaneously. A scheme for parameterising the input space is thus warranted. The input was parameterised according to the following variables:

- The time interval, Δt, between stimulating the first input and the second input.
- The number of afferents, n_1, involved in the first input.
- The number of afferents, n_2, involved in the second input.
- How widely the distributed the first input is. Input may be distributed: (a) to be within the same minor branch group (see figure 5.2), (b) within the same major branch group (see figure 5.2), or (c) across the entire tree.
- The distribution of the second input.
- The separation of the two inputs, *i.e.* the location of the second input, relative to the first input. The second input may be: (a) on the same minor branch as the first input, (b) on the same major branch as the first input, or (c) be on a different major branch.

Input Parameter	Values	Number of data points
Δt (ms)	0, 0.1, 0.2, 0.3, ... , 1.0	11
n_1	0, 1, 2, 5, 10, 20, 50	7
n_2	1, 2, 5, 10, 20, 50	6
Input 1 distribution	minor branch, major branch, whole tree	3
Input 2 distribution	minor branch, major branch, whole tree	3
Input separation	same minor branch, same major branch, tree	3
Total number of combinations		12474

Table 6.2: The input combinations for which the model lateral giant was tested.

Figure 6.3 show diagrams of differently distributed and separated input combinations. Table 6.2 lists the input combinations for which the response of the lateral giant was tested.

6.3.1 Simulation Results

When the parameter values of the giant motor synapse (GMS) are used for the rectifying junctions there is little coincidence detection. For $\Delta t = 0.5$ milliseconds, coincidence detection ratio is around 0.78-0.8 for most of the input patterns tried. Coincidence detection is less effective for input patterns where a large number of afferents are concentrated on a small area of the dendrites. Figure 6.4 shows how the coincidence detection ratio varies for patterns of input. The EPSPs caused by the two inputs at various parts of the LG tree are plotted in figure 6.5a. There is no significant decline in conductance across the synapse for the second input. This can be seen in the conductance plot in figure 6.5b.

The coincidence detection performance of rectifying junctions, with GMS parameter settings, is similar to that of ohmic junctions. To be able to compare the rectifying junction to an ohmic junction, the mean conductance (as calculated by equation 6.4) of the rectifying junction was used for the

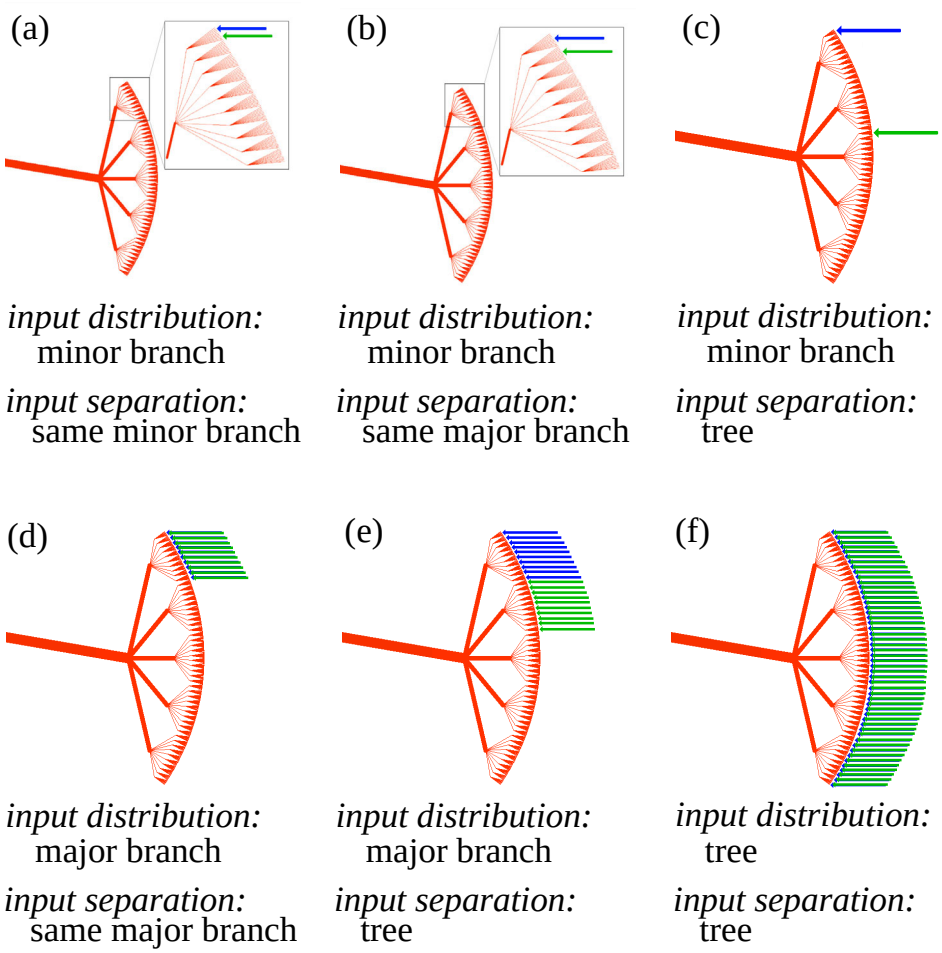

Figure 6.3: Two inputs can be combined differently by varying their spatial distribution and spatial separation.

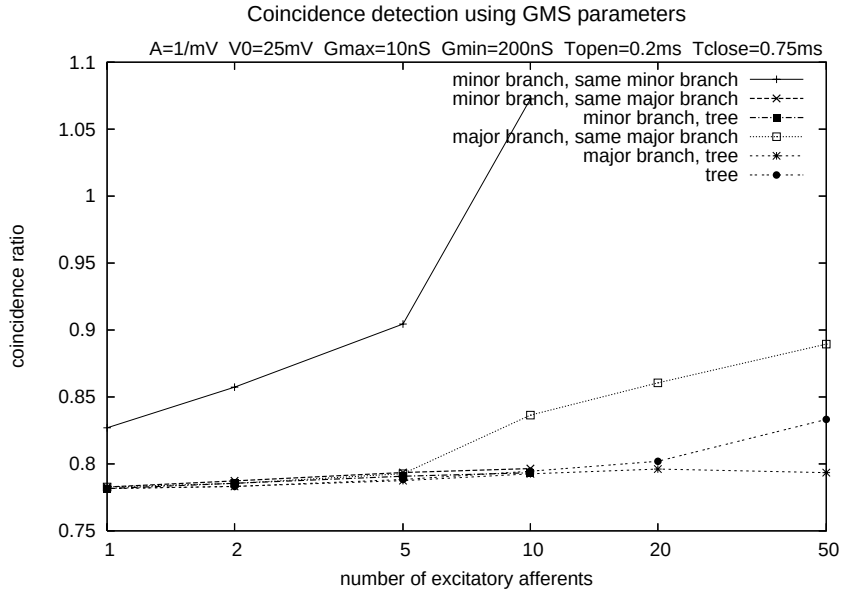

Figure 6.4: The coincidence detection ratio for different patterns of input (see figure 6.3) using rectifying junction parameters similar to those in the GMS.

ohmic junction. Figure 6.6 plots the coincidence detection ratio for different input patterns. The curves look very similar to those for rectifying junctions, with the coincidence detection ratios being only marginally higher.

The parameter search revealed that rectifying junctions can cause very effective coincidence detection, with appropriately set parameters. For certain combinations of parameters, coincidence detection ratios as low as 0.61 were achieved. Such effective coincidence ratios were only attainable using very low conductance values that caused low EPSPs and unrealistically steep values for A. By constraining the solutions to only include solutions that cause EPSPs of at least 50mV in the dendrites and not allowing A to exceed 0.2mV^{-1} coincidence ratios of approximately 0.70 were achieved. The parameters of the best solution that satisfies these constraints are listed in table 6.3. The noteworthy features of this solution include: (1) a very high value for V_0, (2) the low synapse conductance and (3) the fast opening and

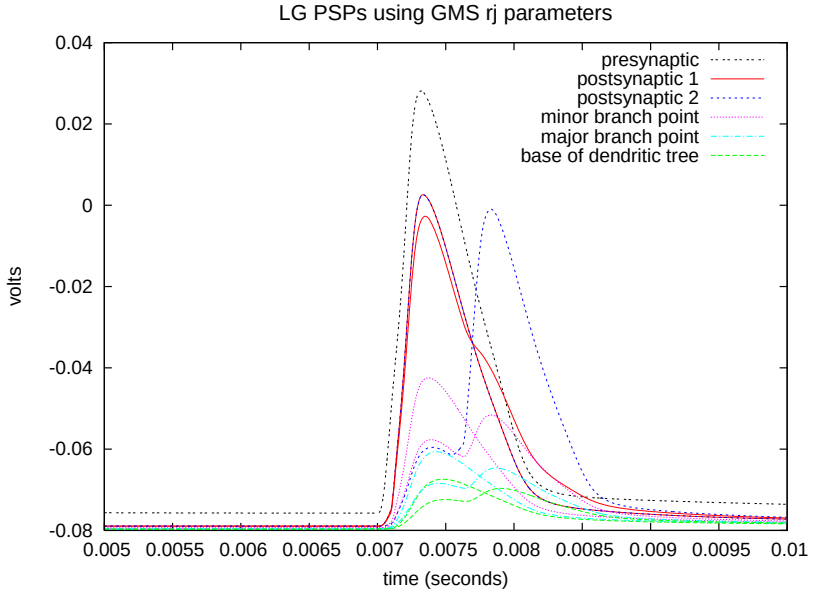

(a)

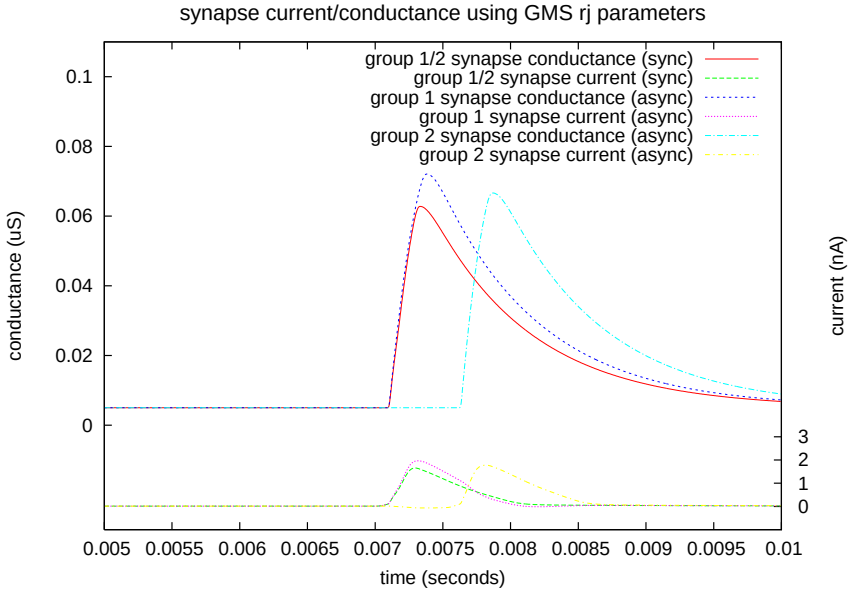

(b)

Figure 6.5: The response of the model lateral giant dendrites to two groups of input (10 afferents per group, each group is distributed onto a separate minor branch as drawn in figure 6.3b.) The values used for the rectifying synapses are those measured at the giant motor synapse. The PSPs at various points in the LG dendrites is plotted in (a), (b) plots the conductance change at the rectifying synapses. (Parameter settings used: $A = 1\mathrm{mV}^{-1}$, $V_0 = 25mV$, $g_{\max} = 100nS$, $g_{\min} = 5nS$, $\tau_{\mathrm{open}} = 0.2ms$, $\tau_{\mathrm{close}} = 0.75ms$.)

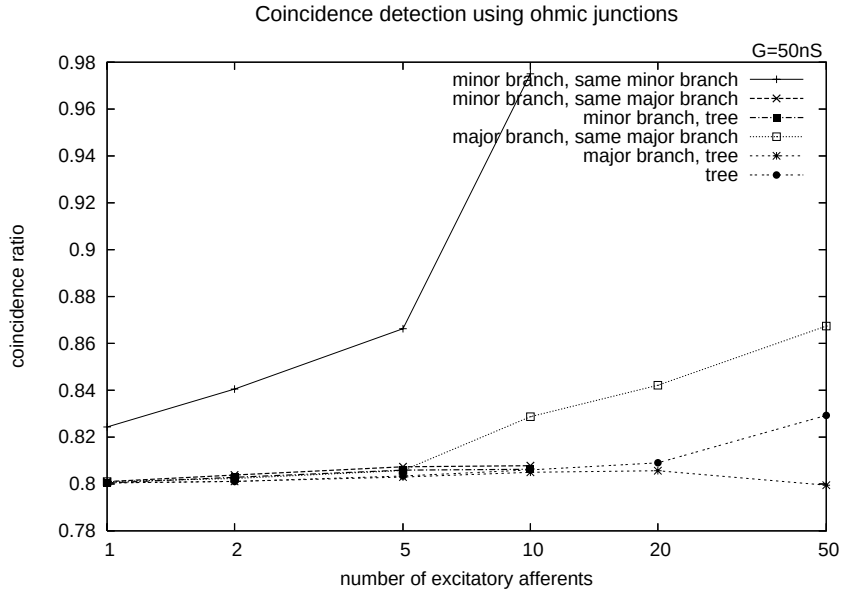

Figure 6.6: The coincidence detection ratio for different patterns of input using ohmic gap junctions. ($g = 50nS$)

A	V_0	$g_{\max}$	$g_{\max}$	τ_{open}	τ_{close}
0.2mV^{-1}	70 mV	46 nS	0.46 nS	0.2 ms	0.4 ms

Table 6.3: Parameter settings that gave the best coincidence detection

closing time constants. The effect of these attributes is that a very high differential between the pre and post synaptic cells is needed before the rectifying junctions start to open. This causes the junctions to open quite late. This combined with the very low resting conductance means that significant currents only start moving into the LG dendrites late into the rising phase of the pre-synaptic spike. The fast closing time causes the junction to start closing again even before the spike is completely over (see figure 6.8b.) The combined effect of this is to shorten the duration of the EPSP, as can be seen in figure 6.8a. Short duration EPSPs need to be highly synchronised to sum effectively. Asynchronous inputs sum poorly, as can be seen in figure 6.8a.

The optimised rectifying junctions were compared with high resistance ohmic gap junctions to assess whether the improved coincidence detection is caused

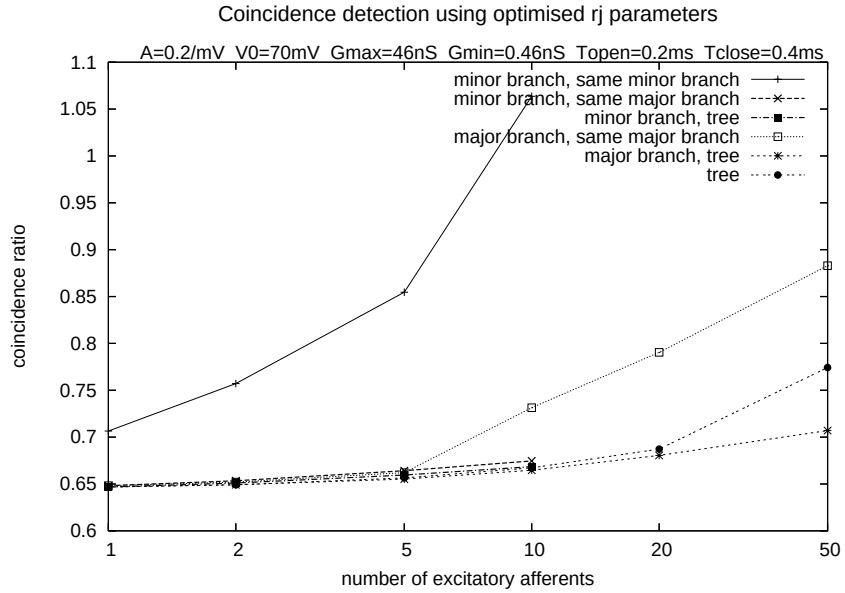

Figure 6.7: The coincidence detection ratio for different patterns of input (see figure 6.3) using optimised rectifying junction parameters.

by the rectification or the low conductance. Figure 6.9 plots the coincidence ratios calculated for ohmic junctions. With low conductance ohmic junctions are able to achieve coincidence ratios of about 0.8, which is of no improvement on high conductance ohmic gap junctions. The rectifying junctions thus perform substantially better at coincidence detection than ohmic junctions of comparable conductance.

The parameter search also revealed an alternative rectifying junction configuration that produced effective coincidence detection. In the alternative solution the rectifying junctions are half open at rest. This is achieved by having a value of V_0 close to 0. The junctions have a large value for τ_{open} and a small value for τ_{close} making them also slow to open and fast to close. Early stimulation is able to enter the LG because the junctions are open by default. However the early current causes the LG dendrites to depolarise closing the gap junctions, making it difficult for late stimulation to enter the LG. The graph in figure 6.10a shows the EPSPs at various stages in the LG

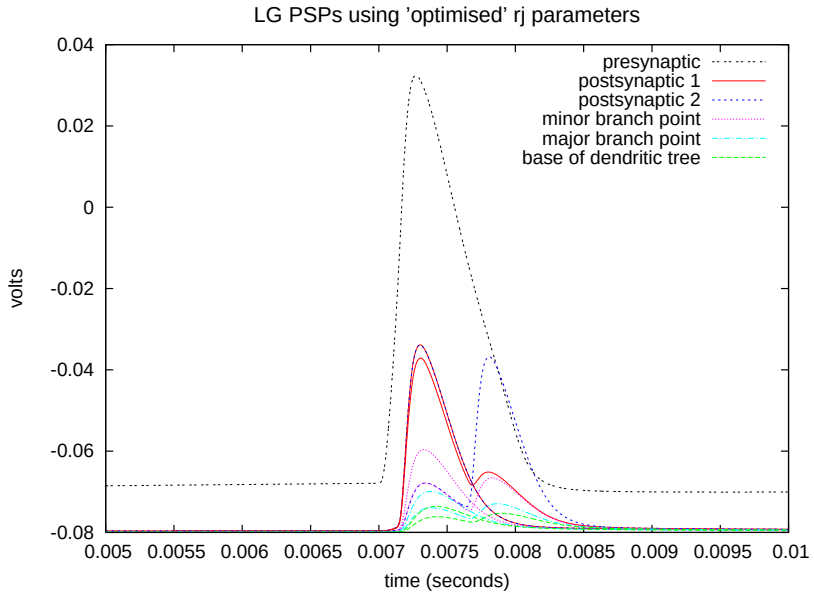

(a)

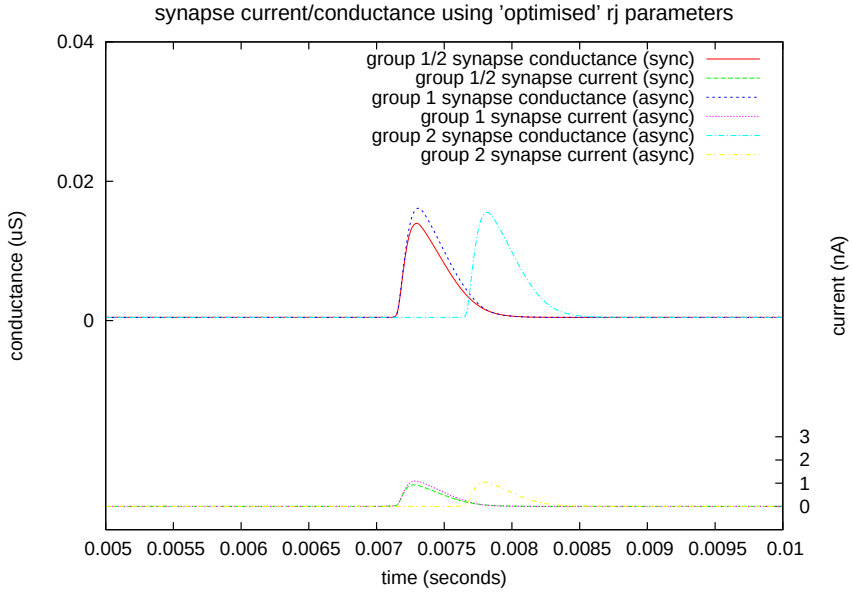

(b)

Figure 6.8: The response of the model lateral giant dendrites to two groups of input when the rectifying synapse parameters are optimised for effective coincidence detection. The PSPs at various points in the LG dendrites is plotted in (a), (b) plots the conductance change at the rectifying synapses. (Parameter values: $A = 0.2\text{mV}^{-1}$, $V_0 = 70\ mV$, $g_{\max} = 46.4nS$, $g_{\min} = 0.464nS$, $\tau_{\text{open}} = 0.2ms$, $\tau_{\text{close}} = 0.5ms$. *Each input group contains 10 afferents which synapse onto separate minor branches of the LG as drawn in figure 6.3b.)*

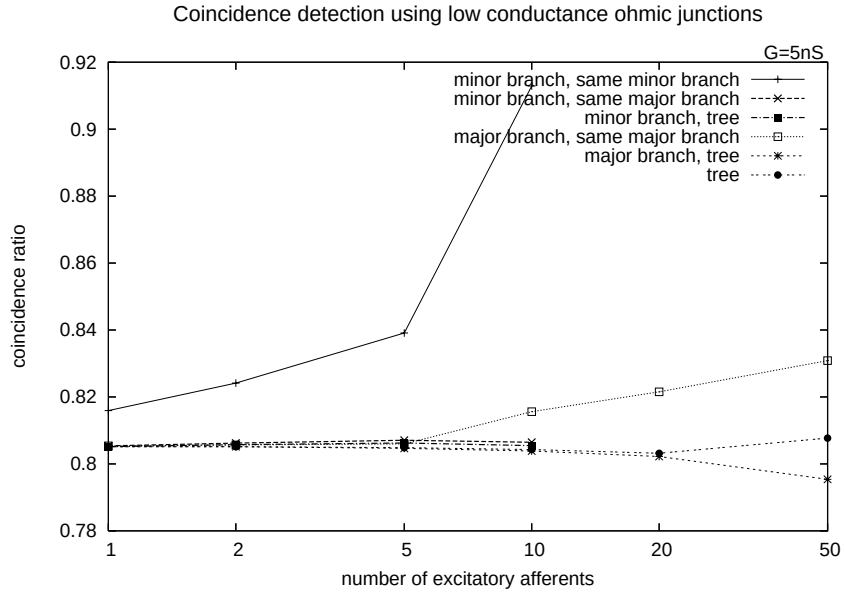

Figure 6.9: The coincidence detection ratio for different patterns of input (see figure 6.3) using low conductance ohmic gap junctions.

dendritic tree when these open by default junctions are used. Figure 6.10b contains a graph that shows the how the conductance across the junctions is affected. This is an interesting result because it identifies a new possible dynamics of how rectifying junctions can enhance coincidence detection.

The simulation findings did not corroborate a previous hypothesis of how coincidence detection may occur in the LG. The hypothesis being the post synaptic depolarisation caused by the early inputs reduces the cross synaptic voltage difference for the late inputs, thereby reducing the maximum conductance of the synapse, thereby reducing the current that is transmitted. The lateral giant's large axon acts as a current sink which very quickly drains charge from the dendrites. For inputs separated by 0.2 milliseconds or more, the depolarisation in the dendrites caused by the first input is no longer significant when the second input arrives.

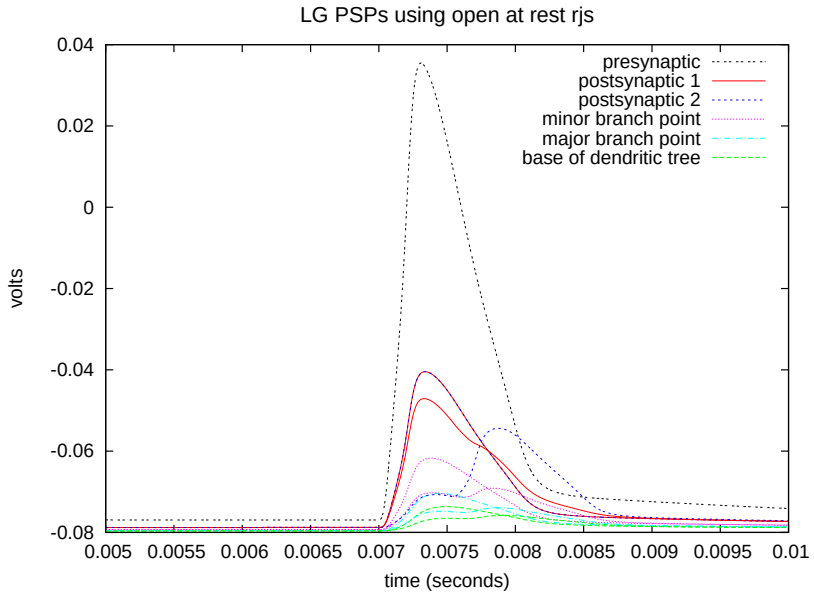

(a)

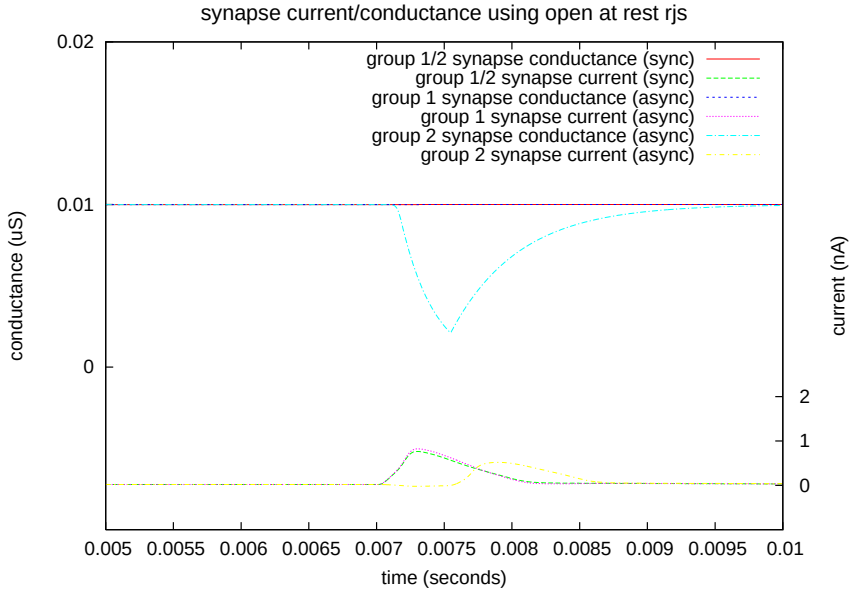

(b)

Figure 6.10: The response of the model lateral giant dendrites to two groups of input. The rectifying junctions are configured such that they are open at rest and are closed when the synapse is reverse biased. The PSPs at various points in the LG dendrites is plotted in (a), (b) plots the conductance change at the rectifying synapses. (Parameter values: $A = 4\text{mV}^{-1}$, $V_0 = 0$, $g_{\min} = 10nS$, $g_{\min} = 0.01nS$, $\tau_{\text{open}} = 0.5ms$, $\tau_{\text{close}} = 0.25ms$. Each input group contains 10 afferents which synapse onto separate minor branches of the LG as drawn in figure 6.3b.)

6.3.2 Discussion

These simulations show that rectifying junctions can contribute to coincidence detection in a simulated cell with a lateral giant like structure. Rectifying junctions with well selected parameter settings perform significantly better at coincidence detection than ohmic gap junctions.

Coincidence detection in the simulation was not achieved in the same way as demonstrated by (Edwards et al., 1998). The large axon combined with the LG's fast membrane time constant causes charge to rapidly disperse from the LG dendrites. Depolarisations in the LG dendrites are very temporary. There is only a slight effect of early depolarisation preventing the gap junctions opening for out of synchronisation input.

The parameter search revealed two alternative tactics for achieving coincidence detection using rectifying junctions in the lateral giant. In the first, the rectifying junctions are given a low resting conductance and a high voltage activation threshold. This has the effect of shortening the duration of post-synaptic potentials in the lateral giant dendrites, thereby requiring inputs to be highly synchronised to sum effectively. In the alternative solution, the rectifying junctions are open at rest. Depolarisation of the lateral giant dendrites by early input causes the gap junctions to close, blocking the path of later input. The principles are similar in concept to the effect simulated by (Edwards et al., 1998), the subtle difference is that the rectifying junction's 'door' is closed instead of being prevented from being opened.

The biological evidence (Herberholz et al., 2002) suggests it is unlikely that the rectifying junctions in the lateral giant are configured in a manner consistent with the first solution. The extremely high voltage activation threshold would make the conduction of antidromic currents virtually impossible. This would contradict the experimental observations, which suggest that

these synapses conduct significant currents bi-directionally. In this regard, the second solution is consistent with the biological information.

An implication of the rectifying junctions being open at rest is that this would cause a continuous current between the afferents and the lateral giant. The 10 mV difference between the resting potentials of the two cells in combination with open connecting synapses would cause there to be a constant current from the afferents to the lateral giant. If this is true, the fine dendrites at the edge of the LG tree, where the afferents are thought to synapse play an important functional role. These narrow dendrites greatly increase the resistance between the afferents and the lateral giant integrating segment. The higher resistance would reduce the current flow at rest and prevent a short circuit.

6.4 Coincidence detection in the morphologically realistic model

We now examine coincidence detection in our morphologically realistic model of the lateral giant (described in section 5.2.) In the previous section we have demonstrated that rectifying junctions can enhance coincidence detection in the abstracted morphology model. We now conduct a set of similar simulations to test whether we can reproduce these results in the morphologically realistic model.

The specific questions we wish to address with our simulations are:

- Can rectifying junctions enhance coincidence detection in the context of using the morphologically realistic model of the lateral giant?
- What are the necessary parameter values to give to the rectifying junctions to cause them to enhance coincidence detection?

- Can the rectifying junctions cause the lateral giant to act in other ways, i.e can the rectifying junctions affect other things other than just coincidence detection?
- How does each parameter affect the results?
- How do the results obtained from simulations to the morphologically realistic model compare with those obtained from the abstracted morphological model?

6.4.1 Simulation setup

Our simulation is set up with a model nerve comprising of fifty afferent axons connecting onto our model of the lateral giant (as described in section 5.2.) The afferent model is described in section 7.1.1. Each afferent makes a rectifying synapse (described in section 5.3) onto the lateral giant. The afferents form synapses at locations identified by the criteria described in section 5.2.1.

To address our aim of determining what parameter settings are necessary to obtain good coincidence detection, we set up a parameter search. The six parameters (see section 6.2), that govern the rectifying junctions, were varied with the values listed in table 6.4. Trials were conducted in which the rectifying junctions were parameterised with each combination of the values listed. The parameter values were tested over a range that extended beyond the limits of biologically plausible values. For example, it may be noted that we include testing with instantaneously opening and closing junctions. Testing to extreme limits increases our confidence that we have not missed the best solution and allows us to assess the impact of the parameters at their extreme limits.

In our tests the lateral giant received input from two groups of spiking

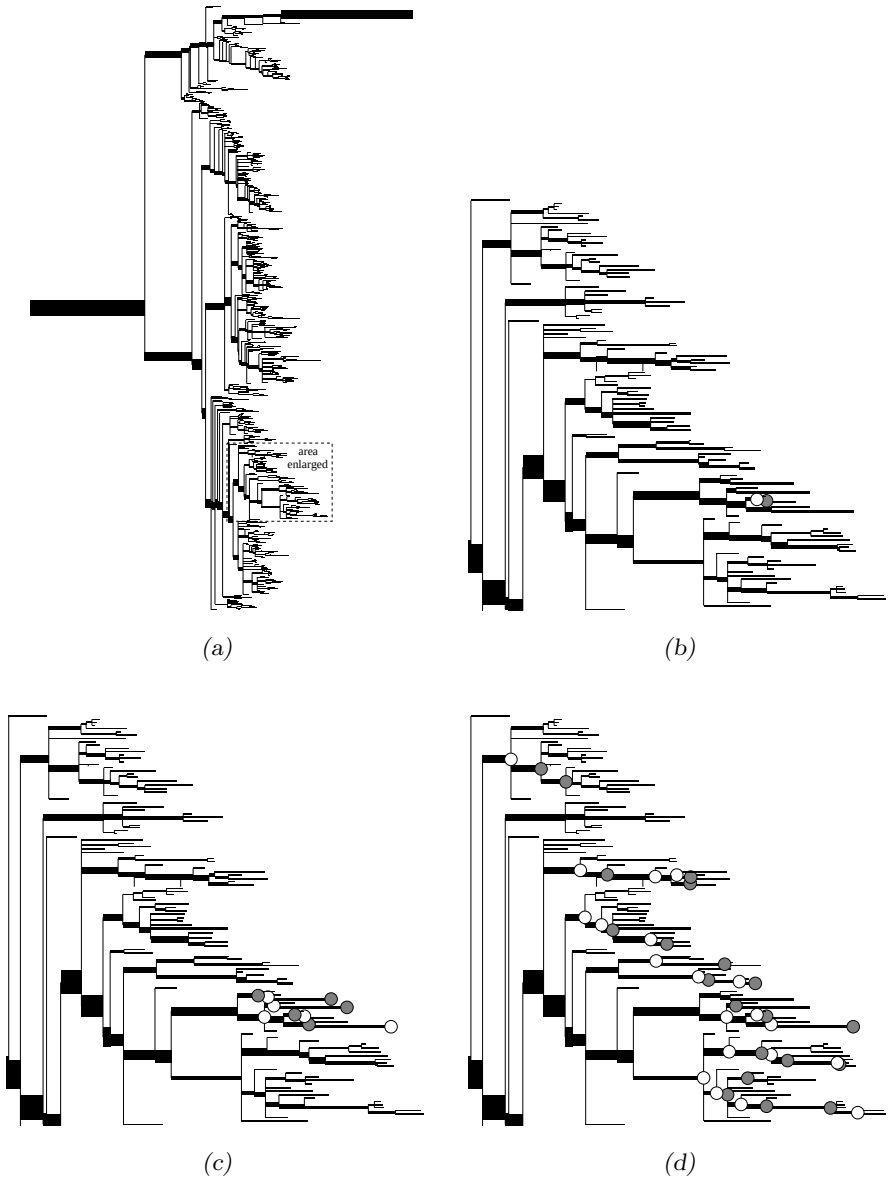

Figure 6.11: The synapse points on the lateral giant at which afferent input is received. (a) shows the entire schematic layout of the lateral giant tree. The area where the synapse points are concentrated is marked out with a box drawn in dashed lines. This area is enlarged in the subsequent sub-figures. (b), (c) and (d) show the specific locations of the synapse points for two groups of one afferent, two groups of five afferents and two groups of twenty afferents, respectively. The locations of the first input group synapses are marked with white circles. The locations of the second input group synapses are marked with grey circles.

Parameter	Values	Units
g_{mean}	0.05, 0.1, 0.2, 0.4	μS
$g_{\text{max}}/g_{\text{min}}$	1, 2, 5, 10, 20, 50, 100, 200, 500, 1000	1
V_0	0, 10, 20, 30, 40, 50, 60, 70	mV
A	0.05, 0.1, 0.125, 0.25, 0.5, 1	mV^{-1}
τ_{open}	0, 0.125, 0.25, 0.5, 1	ms
τ_{close}	0, 0.125, 0.25, 0.5, 1	ms

Table 6.4: The tested rectifying junction parameter values.

Input parameter	Values	Units
Afferent group size	1, 5, 20	1
Temporal separation	0, 0.025, 0.05, 0.1, 0.2, 0.4, 0.8	ms

Table 6.5: The parameters of the input stimulus given to the circuit to test coincidence detection.

afferents. Each afferent group contained an identical number of afferents. All afferents within a group were fired simultaneously but the two groups were stimulated at different times. The input to the circuit was varied by changing the number of afferents in each group and the interval by which their firing is separated. Table 6.5 lists the precise input parameters tested. We tested the response of the lateral giant to a pair of afferents, to two medium sized groups of afferents and to two large groups of afferents. In all three cases, the afferent synapses were chosen to be closely located to each other. We did not test the response to inputs synapsing at locations spread far apart on the lateral giant tree (as we did with the tests on abstracted morphology model in section 6.3.) The exact afferent synapse locations of the stimulated afferents are marked in figure 6.11.

In all simulations, all fifty afferents and their connecting synapses to the lateral giant were included in the model. Afferents were included regardless of whether they were stimulated or not. This was necessary because unstimulated afferent synapses provide an exit for depolarising current and add significantly to the electrical load in the lateral giant dendrites. To only include the stimulated afferents in the model would provide inaccurate and

skewed results. Such results could not be used to fairly compare the effects of increasing the number of afferents stimulated.

As may be noted from table 6.5, the separation time between input groups was varied from being synchronous to the inputs being 0.8 milliseconds apart. The longest separation time approximately corresponds to the length of an afferent spike. Longer input separation intervals were considered uninteresting as it exceeds the window in which we are looking for coincidence detection effects and by far exceeds the time frame in which coincidence detection occurs in the animal. More points were chosen at closely spaced separation times to determine how narrow a separation time is sufficient for coincidence detection effects to emerge.

A simulation was run to test each combination of the parameters listed in tables 6.5 and 6.4. In total this comprises 1,008,000 simulation runs. For each simulation run the peak EPSP was recorded at three different locations in the lateral giant, including at the based of the stimulated major branch, at the spike initiation zone and at the base of the axon. Figure 6.12 marks the recording points. The coincidence and efficiency ratios (see section 6.12) were calculated at each of these recording points.

6.4.2 Simulation results

Overall, the results reflect the findings of the abstracted morphology model simulations (see section 6.3), which is somewhat comforting. Due to the sheer quantity of the results produced by the parameter search, we describe the subset that addresses the questions we are interested in. We present the results by:

- describing the results that gave the best coincidence detection,
- describing the effects of each rectifying junction parameter on the re-

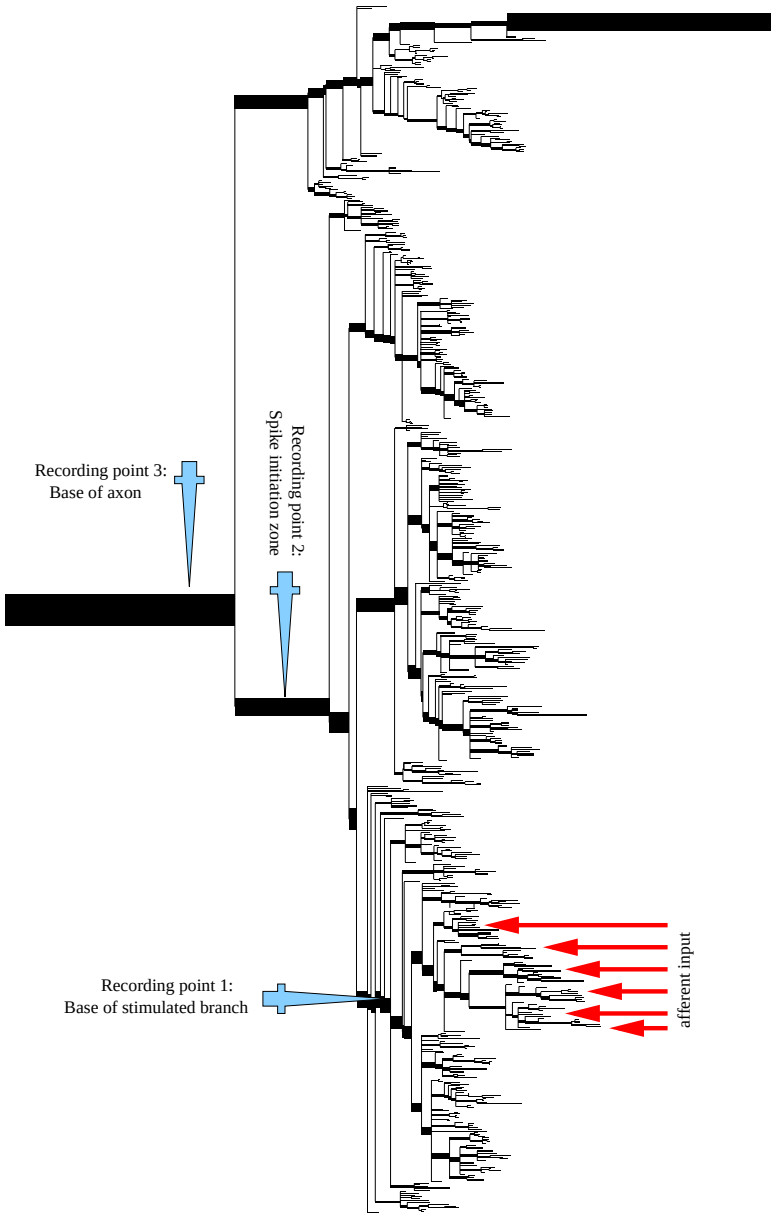

Figure 6.12: The points on the lateral giant where the maximum EPSPs were recorded.

sults, and

- reporting some of the unusual results observed in the parameter space.

As described in our simulation setup, we recorded the EPSPs at three locations in the lateral giant (refer to diagram in figure 6.12.) When we compare the recorded PSPs at these points, the EPSPs are higher at the more distal points in the LG dendrites and decline as we move proxally towards the axon. The efficiency ratios are greater at the recording points closer to the axon. However, this decline in EPSP magnitudes and the increase in calculated efficiency ratios are uniform. There is a high correlation between the values corresponding to the different points (the Pearson correlation coefficient of the efficiency ratios calculated at recording points 1 and 2 is equal to 0.99.) The coincidence ratios stay relatively constant across the recording points. In the following reporting of results, we refer to the efficiency and coincidence ratios calculated at the base of the major branch (recording point 1 in figure 6.12) unless otherwise stated.

The best coincidence detection

The parameters settings that give the best coincidence detection, (as defined by equation 6.2 for $\Delta t = 0.4$ ms), for each of the three different numbers of stimulated afferents tested, are listed in tables 6.6 and 6.7. Table 6.6 includes results from instantaneously opening and closing rectifying junctions in the rankings. This is of course biophysically impossible. Table 6.7 excludes these results. The most effective coincidence detection occurs when the rectifying junctions are open at rest, have a fast closing time and a slow opening time. The best coincidence detection of all is caused by instantaneously closing rectifying junctions (see table 6.6.)

We now describe how the rectifying junctions cause coincidence detection,

Afferents fired	g_{mean} (μS)	$\frac{g_{\min}}{g_{\max}}$	A (/mV)	V_0 (mV)	τ_{open} (ms)	τ_{close} (ms)	C ($\Delta t = 0.4$ ms)	$E(0)$
2	0.05	1000	0.15	0	1	0	0.5652	0.9735
10	0.05	10	0.25	0	1	0	0.6971	0.7665
40	0.05	20	1	0	1	0.25	0.7111	0.7782

Table 6.6: The parameter settings that gave the best coincidence detection for the three tested scenarios of afferents fired.

Afferents fired	g_{mean} (μS)	$\frac{g_{\min}}{g_{\max}}$	A (/mV)	V_0 (mV)	τ_{open} (ms)	τ_{close} (ms)	C ($\Delta t = 0.4$ ms)	$E(0)$
2	0.05	1000	0.25	0	1	0.125	0.6435	0.9113
10	0.05	200	0.15	0	1	0.125	0.6977	0.7539
40	0.05	20	1	0	1	0.25	0.7111	0.7782

Table 6.7: The parameter settings that gave the best coincidence detection, when instantaneously opening and closing rectifying junctions are excluded.

when configured with the parameters listed in tables 6.6 and 6.7. When the input is synchronised, all the junctions are open and a large PSP is caused in the lateral giant (see figure 6.13a.) When the inputs are desynchronised, open junctions receive the early input causing the lateral giant dendrites to depolarise. The reverse bias across the junctions from non-stimulated afferents causes them to close (including the junctions that will deliver the late input.) When the late input arrives the junctions are still closed and the input is ineffective (see figure 6.13b.)

We defined a good coincidence detector to have an efficient response to synchronised input, with a rapidly declining efficiency as we desynchronise our input. In our results we were able to fit the lateral giant with parameters that cause it to meet this objective. When we choose parameters that give good coincidence detection, then the summed PSP recorded in the lateral giant declines along a bell shaped curve (see figure 6.14) as one increases the time interval, Δt, by which the two input groups are separated.

One can optimise the parameter settings to deliver the best coincidence de-

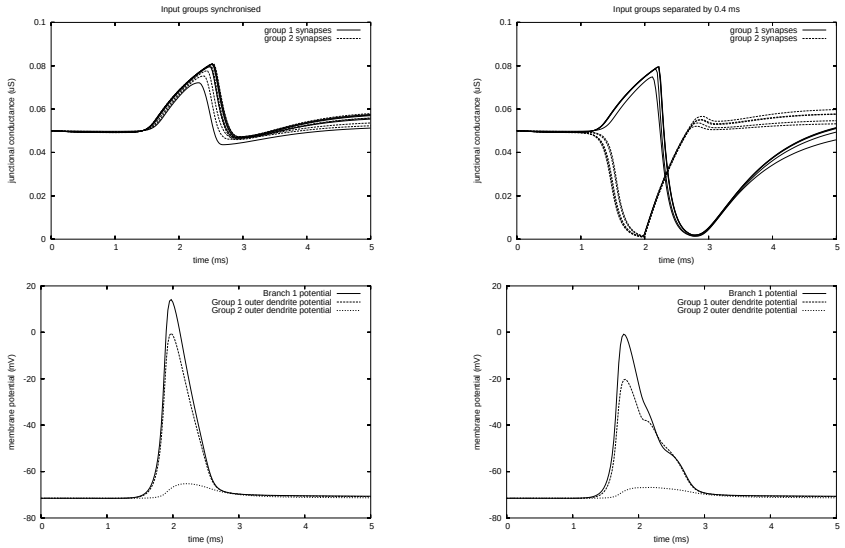

(a) *(b)*

Figure 6.13: The junctional conductance versus time (top graphs) and the resulting PSPs (bottom graphs) in response to (a) ten synchronised afferents (b) two groups of five afferents (synapsing at the points shown in figure 6.11c) desynchronised by 0.4 milliseconds. The plots of the conductance, at all ten junctions receiving afferent input, are overlayed. For each input group (Group 1 and Group 2) the PSP recorded at one of the receiving dendrite branches is plotted. We also plot the PSP at the base of main branch 1. (Parameter settings: $\frac{g_{\max}}{g_{\min}} = 200$, $A = 0.15\mathrm{mV}^{-1}$, $V_0 = 0$, $\tau_{\mathrm{open}} = 1ms$, $\tau_{\mathrm{close}} = 0.125ms$.)

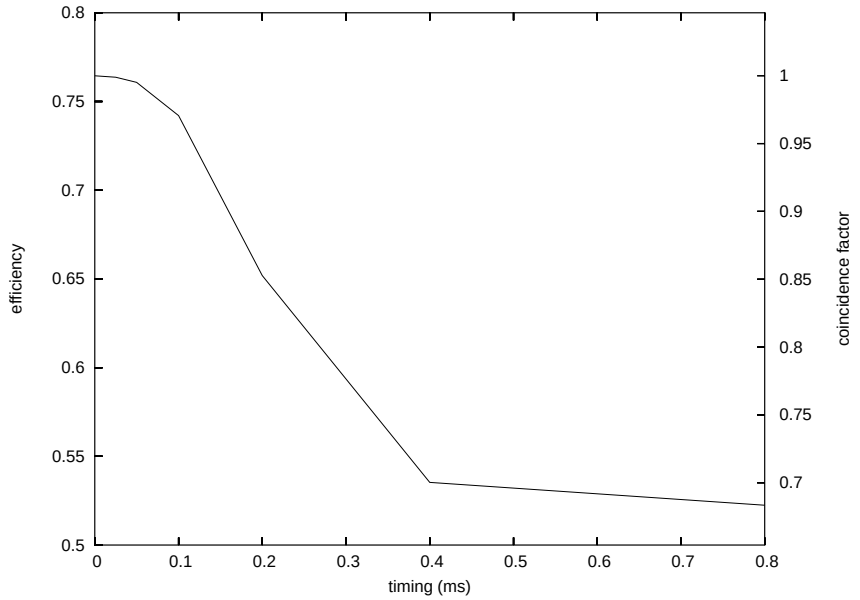

Figure 6.14: The coincidence detection performance of the lateral giant to two groups of five inputs when the rectifying junction parameters have been optimised. (Parameter settings: $\frac{g_{\max}}{g_{\min}} = 200$, $A = 0.15\text{mV}^{-1}$, $V_0 = 0$, $\tau_{\text{open}} = 1ms$, $\tau_{\text{close}} = 0.125ms$.*)*

tection for firing small numbers of afferents or large numbers of afferents. The best parameters settings are different for the the alternative scenarios of the number of afferents stimulated. Tables 6.6 and 6.7 show the parameter choices that gave the best coincidence detection (for $\Delta t = 0.4$ milliseconds) for two stimulated afferents, ten stimulated afferents and forty stimulated afferents. Figure 6.15a graphically illustrates how the best parameter settings vary for the different numbers of stimulated afferents. The best values for three of the parameters stay constant. For all numbers of stimulated afferents, the mean conductance is low ($g_{\text{mean}} = 0.05\mu\text{S}$), the voltage activation (V_0) is zero and there is a slow junction opening time constant ($\tau_{\text{open}} = 1$ ms.) The best value for the other three parameters, however, changes with the number of afferents stimulated. A larger value of A is required as the population of stimulated afferents is increased. The best rectification value ($\frac{g_{\max}}{g_{\min}}$) varies widely as we change the number of afferents stimulated. This

needs qualification however, as exact amount of rectification, beyond a factor of ten has little impact on the operation of the circuit. Any amount of rectification, beyond a factor of ten, will produce similar coincidence ratios, (assuming the other parameters are kept constant.) One surprising result is that for the larger group of stimulated afferents, a junction closing time constant of 0.25 milliseconds does better than instantaneous closing. It should be noted that the improvement in coincidence detection, obtained by using non-instantaneous closing, is ever so slight.

The highest obtainable coincidence ratio is generally higher for small groups of stimulated afferents than for large groups. Tables 6.6 and 6.7 and figure 6.15b show that better coincidence ratios are obtainable for small groups of stimulated afferents than for large groups. Similarly, the efficiency declines as the number of afferents stimulated is increased.

Instantaneously closing junctions deliver the best coincidence detection. Using instantaneously closing junctions gives a dramatic advantage when only two afferents are stimulated, as may be observed by comparing the coincidence ratios in tables 6.6 and 6.7. The advantage is not as significant when there are a greater number of afferents stimulated. An explanation for this may be that in the two afferent scenario, the synapse points are closely spaced, whereas in the forty afferent scenario the synapse points are spread across an entire major branch of the lateral giant.

The coincidence detection ratio, C, is a function of the input stimulus separation time, Δt, as defined in equation 6.2. Figure 6.14 plots the coincidence ratio as a function of the input separation time, for the lateral giant using data obtained from our results. We observe that the coincidence ratio declines with Δt, along a bell shaped curve. The parameter values that give the best coincidence detection, differ depending on which the separation time we wish to optimise for. For example, the parameter settings that give the

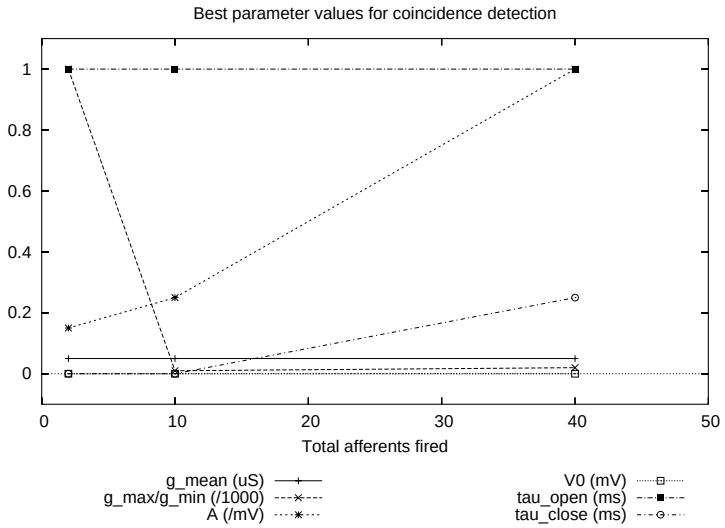

(a)

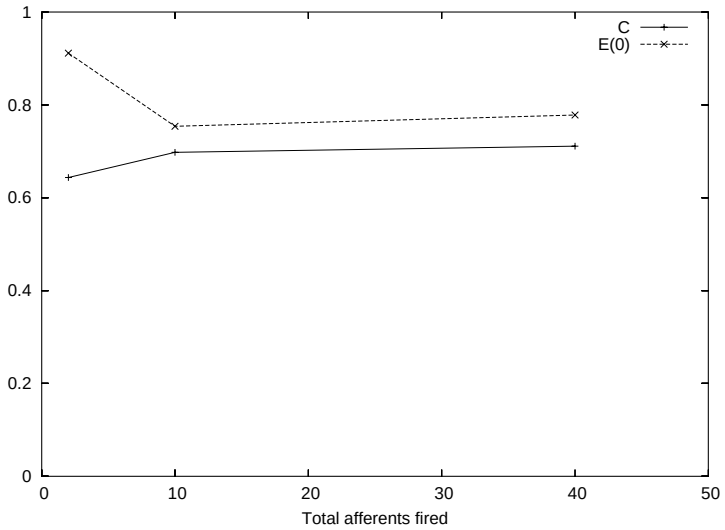

(b)

Figure 6.15: (a) The parameter settings that gave the best coincidence detection for the three tested scenarios of afferents fired. (b) The best coincidence ratios obtained (C), for the three tested scenarios of afferents fired, and their corresponding efficiency ratio for synchronised input ($E(0)$).

best coincidence detection for $\Delta t = 0.4$ ms may not give the best possible coincidence detection for $\Delta t = 0.2$ ms. Table 6.8 tabulates the parameter settings that gave the best coincidence detection for each Δt tested. An inspection of the table shows that there are two kinds of solutions: one that does well for small Δts and one that does well for large Δts. The threshold at which the latter outperforms the former lies somewhere in the space $0.1\text{ms} < \Delta t < 0.2\text{ms}$. The solutions that give good coincidence detection for small input delays have:

- a high maximum conductance,
- a very large value for A,
- a high voltage activation threshold, V_0,
- a very fast junction opening time constant, and
- a long junction closing time constant.

This is in contrast to parameters that give the best coincidence detection for larger input delays, which have:

- low conductance,
- a low voltage activation threshold,
- a slow junction opening time constant, and
- a fast junction closing time constant.

The transition between the two kinds of solutions is clearly delimited as can be seen in the graph in figure 6.16a.

That it is possible to configure the rectifying junctions such that they discriminate between perfectly synchronised input and only slightly out of

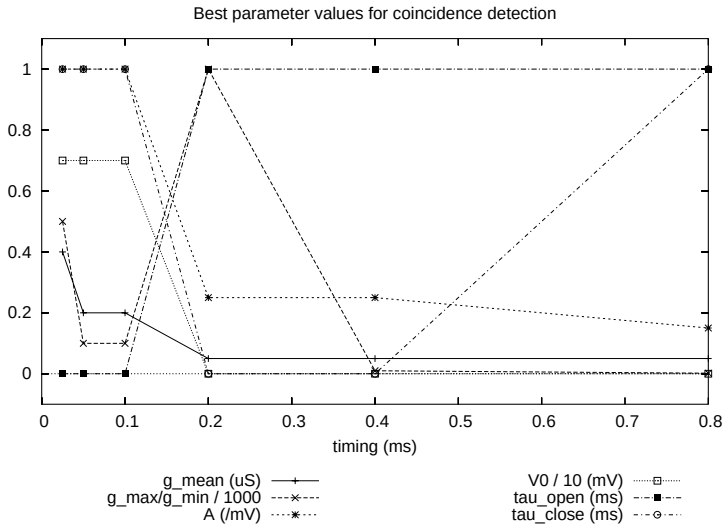

(a)

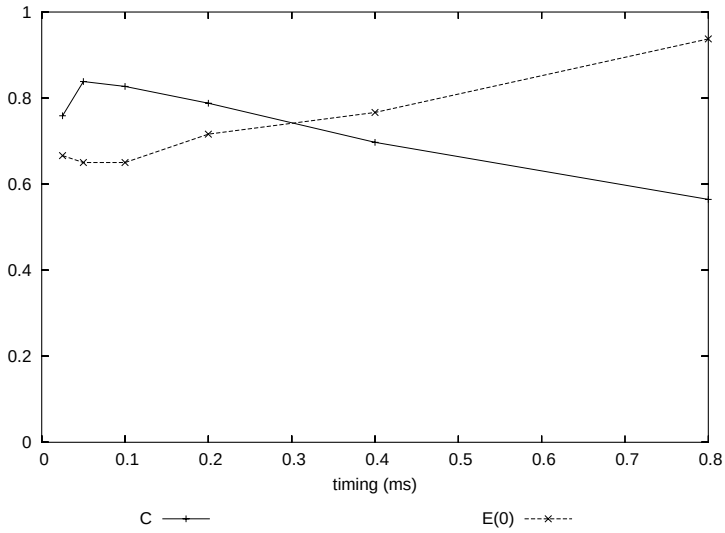

(b)

Figure 6.16: (a) The parameter settings that gave the best coincidence detection for the different input separation timing, Δt. (b) The best coincidence ratios (C) obtained, for the input separation timings tested, and their corresponding efficiency ratio for synchronised input ($E(0)$).

Δt (ms)	g_{mean} (μS)	$\frac{g_{\min}}{g_{\max}}$	A (/mV)	V_0 (mV)	τ_{open} (ms)	τ_{close} (ms)	$C(\Delta t)$	$E(0)$
0.025	0.4	500	1	70	0	1	0.7590	0.6663
0.05	0.2	100	1	70	0	1	0.8384	0.6503
0.1	0.2	100	1	70	0	1	0.8271	0.6503
0.2	0.05	1000	0.25	0	1	0	0.7881	0.7162
0.4	0.05	10	0.25	0	1	0	0.6971	0.7665
0.8	0.05	2	0.15	0	1	1	0.5645	0.9376

Table 6.8: The parameter settings that gave the best coincidence detection for different input separation times, Δt. if instantaneous opening and closing of the rectifying junctions is excluded.

phase input is a remarkable result. In effect we have a kind of 'super-coincidence' detector. As is tabulated in table 6.8, desynchronising the input by as little as 0.025 milliseconds can cause the response to be almost a quarter lower, for some parameter settings. To understand how this can occur, let us examine what happens to the conductance at the synapses in this scenario. Figure 6.17 compares the conductance changes at the rectifying synapses in response to synchronised and slightly desynchronised input. When the input is desynchronised, the junctions from the latter group of stimulated afferents fail to open. Due to the depolarisation of the lateral giant dendrites, the cross junctional voltage difference never reaches the high activation threshold required. In this configuration, the junctions operate as hypothesised by (Edwards et al., 1998). However, this result is achieved only through unrealistic parameter settings, namely: instantaneously opening gap junctions, a high value for A, a large rectification factor and a very high activation threshold. These factors combine to cause the junctions to instantaneously switch from being completely closed to completely open. Such kinetics are biologically implausible.

Although it is theoretically possible to achieve a high coincidence ratio for very small values of Δt, the efficiency of that circuit in summing synchronised input is very low. Firing all ten afferents causes the PSP to be just is

just 1.3 times higher than if only five afferents are fired. This is an undesirable property of a good coincidence detector.

The effects of the rectifying junction parameters

We now investigate how each of the six parameters affects the results. We wish to determine how sensitive the results are to each of the parameters. We use the parameter settings that give the best coincidence detection as our starting point (for $\Delta t = 0.4$ milliseconds, see table 6.7 for the parameter values.) For each parameter, we vary the value, while keeping the values of the other parameters constant. It should be noted that we are only observing the trends that occur around the best found coincidence detection point. Our observations need to be qualified: trends caused by changing a parameter value around best settings point may not occur, or even be reversed, in other regions of the parameter space.

Figure 6.18 shows how the rectification of the gap junctions affects the coincidence detection of the lateral giant. The first observation from this graph is that rectification has a big impact. Circuits with rectifying junctions have far better coincidence detection than occurs with non-rectifying junctions (*i.e.* $\frac{g_{\max}}{g_{\min}} = 1$). A number of other specific observations can be made. Firstly, rectification decreases the efficiency of the circuit. Synchronised inputs sum almost linearly when the junctions are non-rectifying. When the junctions have a high rectification the efficiency of summing synchronised input drops to about 0.75. Secondly, rectification improves the coincidence detection of the circuit. When the junctions are non-rectifying, the summed PSP gradually declines as the separation between the input timing, Δt, is increased. For rectifying junctions, the PSPs drops off much more sharply as Δt is increased. The summed PSP starts to drop off sharply as Δt exceeds 0.1 ms. At $\Delta t = 0.4$ ms the high rectification curves bottom out. In

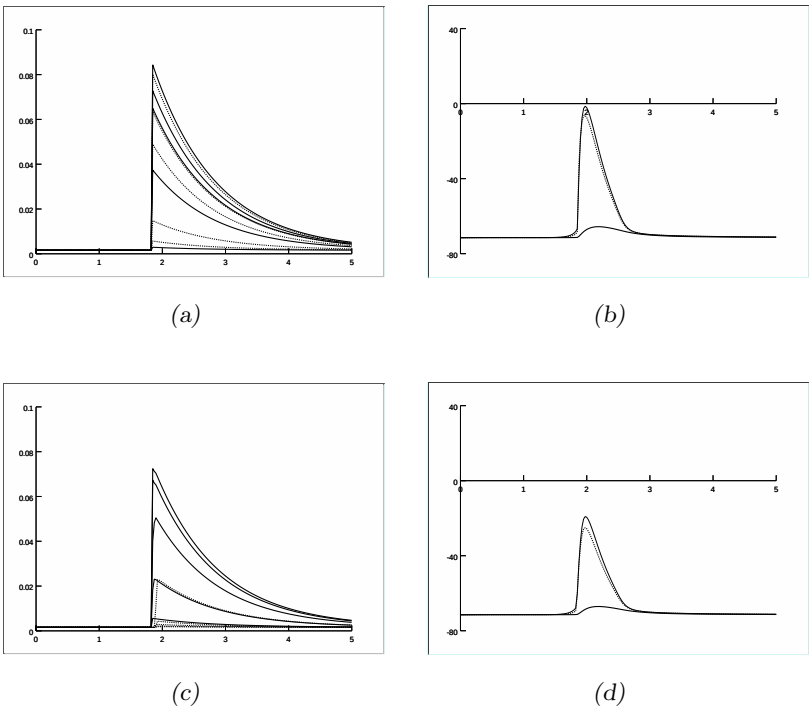

(a) *(b)*

(c) *(d)*

Figure 6.17: The conductance changes in the rectifying synapses and PSPs in the lateral giant when the rectifying junctions are configured to cause coincidence detection for small input separation timing (Δt). These graphs compare the responses to the firing of ten afferents when the are fired simultaneously or as two groups separated by $\Delta t = 0.025$ ms. (a) The conductance change at the synapses for simultaneous input, (b) The PSPs in the lateral giant for simultaneous input, (c) The conductance change at the synapses when the two afferent groups are fired 0.025 milliseconds apart, (d) The PSPs in LG when the afferent groups are fired apart. For (a) and (c) the conductances of synapses from group 1 afferents are plotted with solid lines and group 2 synapse with dotted lines. For (b) and (d) the PSP at the post synaptic site of a group 1 afferent is plotted with a thin solid line, the PSP from a group 2 afferent with a dotted line and the PSP at the base of the main branch is plotted with a thick solid line. (Parameters used: $g_{\text{mean}} = 0.4\mu S$, $\frac{g_{\max}}{g_{\min}} = 500$, $A = 1\text{mV}^{-1}$, $V_0 = 70\ mV$, $\tau_{\text{open}} = 0$, $\tau_{\text{close}} = 1\ ms$.)

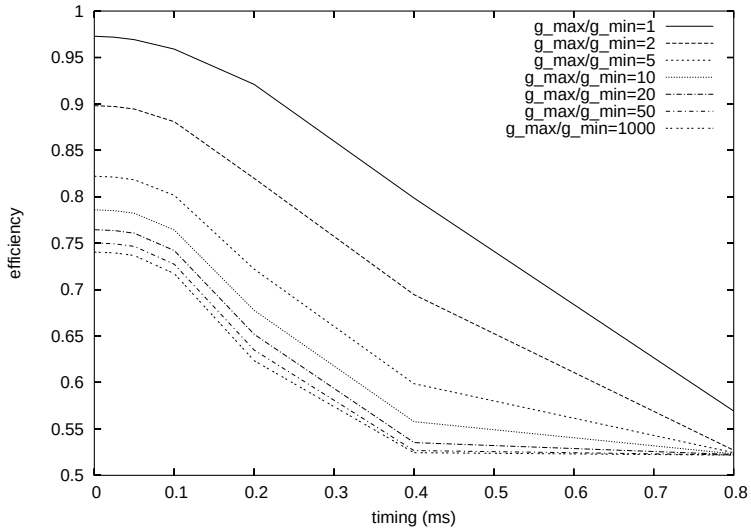

Figure 6.18: The effect of rectification, $\frac{g_{\max}}{g_{\min}}$, on coincidence detection.

excess of this input timing separation the latter input causes only a marginal increase to the PSP in the lateral giant dendrites. This is in stark contrast to the non-rectifying curve where the bottom still has not been reached at $\Delta t = 0.8$ ms. Finally, it can be observed that increasing rectification has diminishing returns. Increasing rectification beyond a factor of about ten does not have a very great effect. The differences between the rectification of a factor of fifty and a thousand is marginal. The curves for $\frac{g_{\max}}{g_{\min}} = 50$ and $\frac{g_{\max}}{g_{\min}} = 1000$ are fairly similar (and as are all the non-plotted curves in between.)

Figure 6.19 plots the efficiency of the lateral giant summing two groups of input for different junctional conductances. In general two groups of input sum more efficiently at lower conductance values, this is especially true for synchronised summation. When the inputs are separated by 0.4 milliseconds or more the conductance has little impact on the efficiency. Because a lower junctional conductance causes increased efficiency for synchronised input ($\Delta t = 0$) but has little effect on asynchronous input ($\Delta t \geq 0.4$ ms) lower conductances give better coincidence detection.

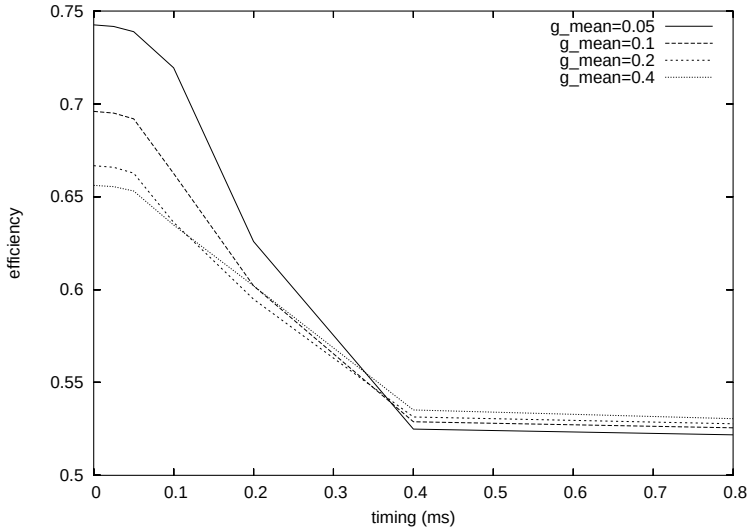

Figure 6.19: The effect of the mean junctional conductance, g_{mean}, on coincidence detection.

Increasing the rectifying junctions' voltage activation threshold, V_0, reduces the coincidence detection of the lateral giant. Increasing V_0 reduces the PSPs in the lateral giant, as can be seen in figure 6.20. A high voltage activation threshold causes the junction to be closed at rest. This combined with the slow junction opening time constant that we are using, causes the junctions to be too slow to open in response to a pre-synaptic spike. The junctions never reach a state of high conductance because the spike has passed before the conductance can significantly increase. Thus at high values of V_0 (combined with our other parameter values) the junctions act like a high resistance ohmic junctions. As can be seen in figure 6.20 this causes the PSPs to be lower and the response versus Δt curve to be shaped like that of a non-rectifying junction (compare with $\frac{g_{\max}}{g_{\min}} = 1$ in figure 6.18.)

Tampering with the value of the parameter, A, has mixed consequences on the coincidence detection properties of the circuit. In the context of the other parameter settings, higher values of A result in larger PSPs in the lateral giant dendrites, as is recorded in figure 6.21. Higher PSPs occur in response

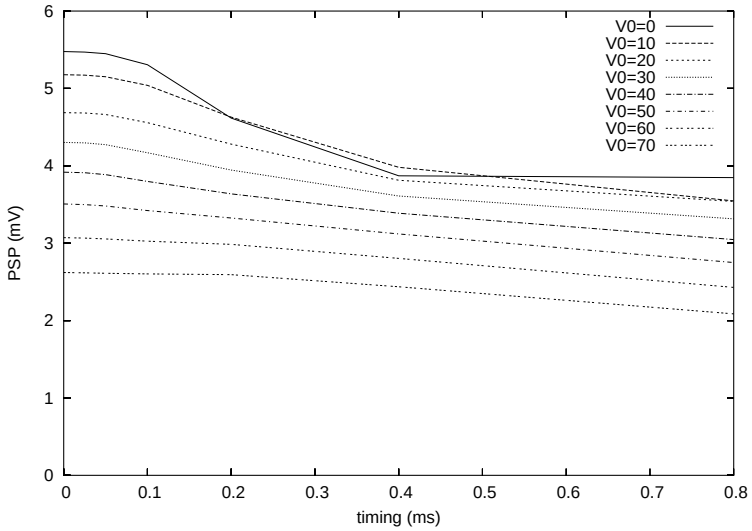

Figure 6.20: The effect of the activation threshold, V_0, on the PSPs in the lateral giant caused by two groups of five afferents coincidence detection.

to both synchronised and asynchronous input. By our definition, the best coincidence detection occurs when the ratio of the synchronised PSP to the asynchronous PSP is at its highest: this occurs when the value of A lies in the mid-range of the values tested. Higher values of A give larger PSPs because the junctions require a smaller voltage differential between the pre and post synaptic cells to switch into a state of high conductance. Lower values of A cause the junction conductance to be less affected by changes in the voltage differential. As the value of A tends to zero, the junction behaves more like an ohmic gap junction. The consequences of this effect can be observed in figure 6.21: for the lowest value of A plotted ($A = 0.05\text{mV}^{-1}$)), the shape of the response curve resembles the response of an ohmic junction (compare with ohmic curve in figure 6.18.)

The junction opening and closing time constants also greatly affect the response of the lateral giant. In the context of our other parameter settings, the slower the the junction is to open, the better the coincidence detection

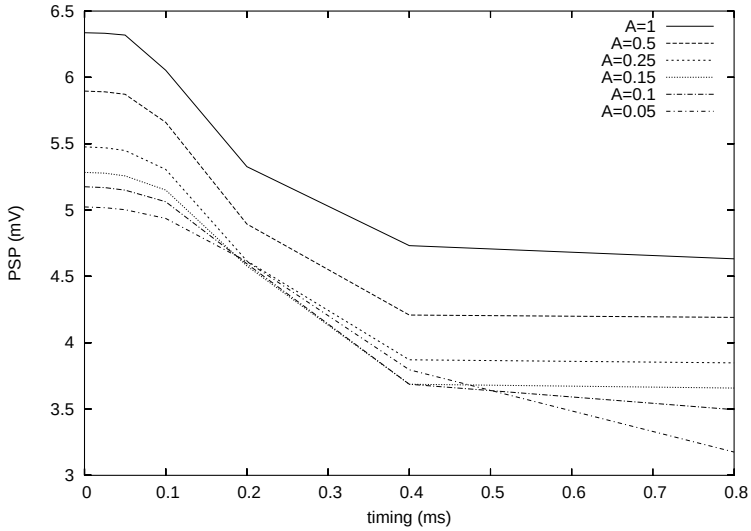

Figure 6.21: The effect of the parameter, A, on the PSPs generated in the lateral giant dendrites for different input separation timings.

(see figure 6.22a.) A slower opening time constant causes the lateral giant to be more efficient at summing synchronised input and less efficient at summing asynchronous input. This is a desired property in a good coincidence detector. A slow opening time constant is crucial for the performance of the circuit. Input from the first group of afferents depolarises the lateral giant dendrites, which reverse biases the junctions gating the late input, causing them to close. A slow opening time constant prevents these closed junctions from reopening too quickly when the late inputs arrive, thus blunting the impact of the late inputs.

The effect of the closing time constant is more ambiguous. (see figure 6.22b.) A slower closing time constant causes synchronised input to sum more efficiently, but it also causes asynchronous input to sum more efficiently. A good coincidence detector minimises ratio of the efficiency of asynchronous input and synchronised input. Around this point in the parameter space, the closing time constant is correlated with both. The best value occurs when a balance is reached. In the context of the other parameter values this

is at $\tau_{\text{open}} = 0.125$ milliseconds.

Unusual effects

With carefully selected parameter values we have seen that rectifying junctions can enhance coincidence detection. It is a false impression to think rectifying junctions enhance coincidence detection under any parameter settings. The results show that the circuit can be made to behave in a wide range of ways, depending on the parameter settings.

Rectifying junctions, with alternatively chosen parameter settings, can also cause the lateral giant to behave as a very poor coincidence detector, or an 'anti-coincidence' detector. Within the parameter space, there exist combinations where the rectifying junctions cause the lateral giant to respond better to asynchronous input than synchronised input. Figure 6.23 graphs the response of the lateral giant dendrites under such conditions. In this example, the lateral giant responds extremely inefficiently to synchronised input. The circuit's efficiency to synchronised input is 0.55 (*i.e.* firing ten afferents together, instead of five, only increases the PSP by 10%.) However the lateral giant responds very efficiently to asynchronous input ($E(\Delta t = 0.4\text{ms}) = 0.8$.) We have thus turned the circuit into a good anti-coincidence detector.

To understand how rectifying junctions can cause the lateral giant to respond better to asynchronous input it is helpful to look at the behaviour of the junctions themselves. Figure 6.24 plots the conductances of the rectifying junctions in response to synchronous and asynchronous inputs, and the resulting PSPs in the lateral giant dendrites. From these graphs it can be seen that when all ten afferents are fired simultaneously, the junctions do not reach the same level of conductance as they do when the two groups of afferents are fired at different times. While synchronising the input causes

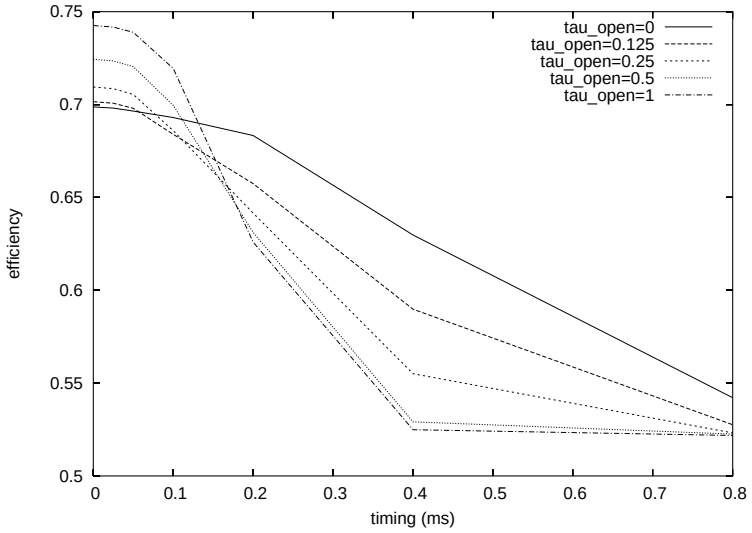

(a)

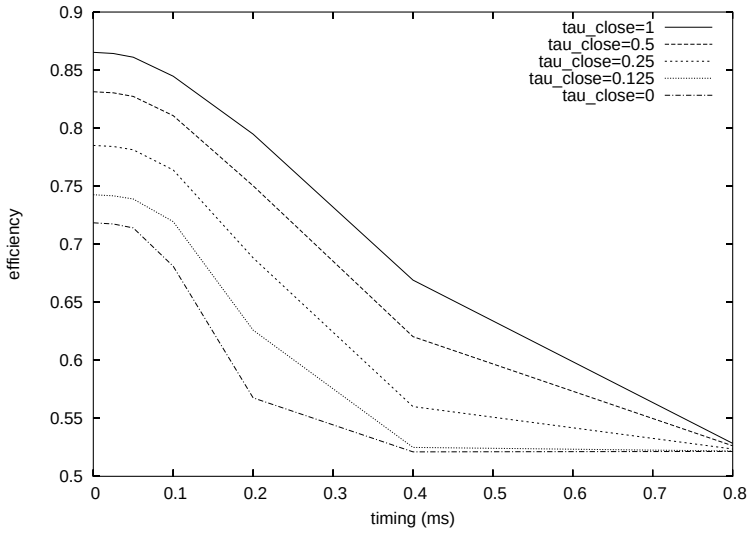

(b)

Figure 6.22: The effect of (a) the junction opening time constant, τ_{open}, and (b) the junction closing time constant, τ_{close}, on coincidence detection.

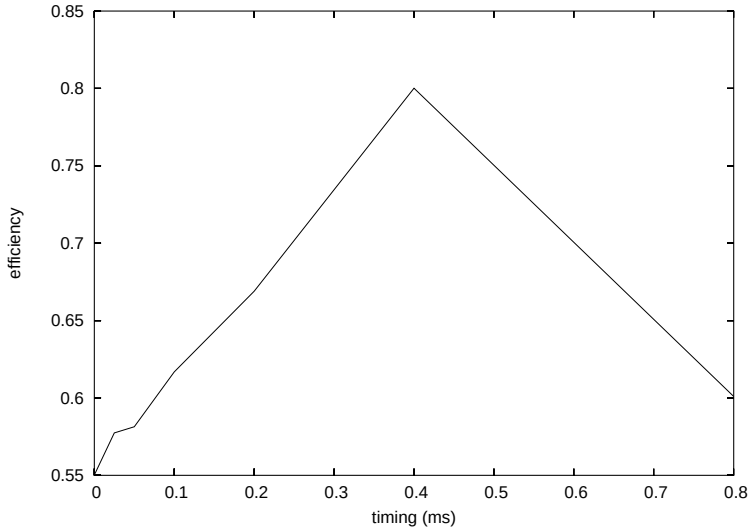

Figure 6.23: Under some conditions, the rectifying junctions can cause the lateral giant to respond better to asynchronous input than synchronised input. (Parameters used: $g_{\text{mean}} = 0.2\mu S$, $\frac{g_{\max}}{g_{\min}} = 1000$, $A = 1\text{mV}^{-1}$, $V_0 = 70\ mV$, $\tau_{\text{open}} = 0$, $\tau_{\text{close}} = 0.25\ ms$.*)*

the lateral giant dendrites to receive the currents from ten afferents simultaneously, the lower conductance reduces the amount of current transferred from each afferent. Thus, when one fires ten afferents simultaneously, instead of five, most of the gain obtained by firing more afferents is offset by lower conductance at the synapses.

This raises the question, how does synchronised input cause lower synaptic conductances? The steady-state synaptic conductance is a function of the the trans-synaptic voltage difference (see equation 3.1.) In figure 6.25a, it can be seen, that the trans-synaptic voltage difference is marginally lower in response to synchronised input than to asynchronous input. Ten simultaneously firing afferents causes a greater depolarisation in the lateral giant dendrites, which reduces the voltage difference between the pre and post synaptic cells. The difference between the peak voltage difference in response to synchronous and asynchronous input is only slight: 71.9 versus 74.8 millivolts. To understand how such a small difference can affect the

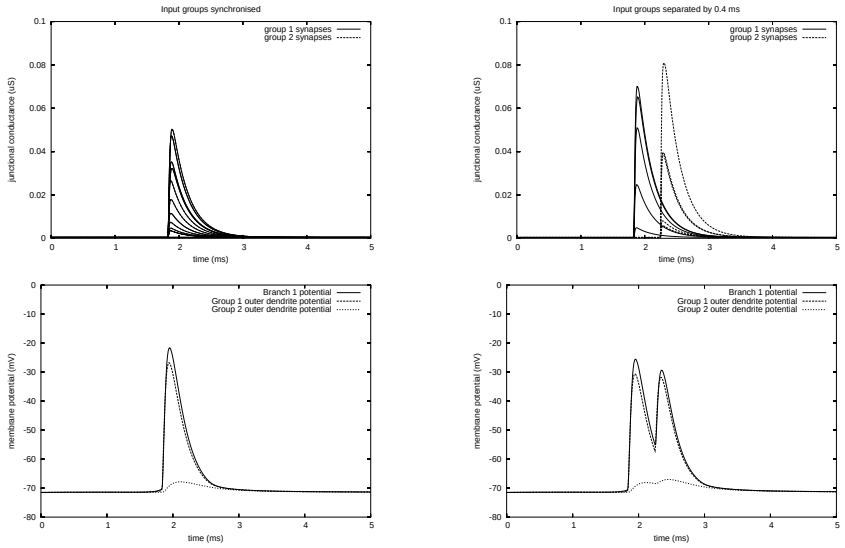

(a) *(b)*

Figure 6.24: For parameter settings that cause the lateral giant to respond better to asynchronous input, the junctional conductance versus time and the resulting lateral giant dendrite PSPs are plotted. (a) Plots the junctional conductance and PSP in response to synchronised input. (b) Plots the junctional conductance and PSP in response to input desynchronised by 0.4 milliseconds. (Parameter settings: $g_{\text{mean}} = 0.2\mu S$, $\frac{g_{\max}}{g_{\min}} = 1000$, $A = 1\text{mV}^{-1}$, $V_0 = 70\ mV$, $\tau_{\text{open}} = 0$, $\tau_{\text{close}} = 0.25\ ms$.*)

junctional conductance it is necessary to view a plot of the conductance equation, when it is given, the currently discussed, anti-coincidence detection parameter values. Figure 6.25b plots the relationship between synaptic conductance and the voltage difference. It may be observed that the conductance makes a very sharp transition from low conductance to high conductance when the trans-synaptic voltage difference surpasses 70 millivolts. Thus, although synchronising the input causes only a slight change in the trans-synaptic voltage difference, the change occurs within the range where the conductance function is at its most sensitive. A small change in the trans-synaptic voltage difference causes a large change in the synaptic conductance.

Taken all together, when the afferents are synchronously fired, there is a slightly larger depolarisation in the lateral giant dendrites, which reduces the magnitude of the peak trans-synaptic voltage difference, which reduces the conductance of the synapses, which causes each afferent to contribute less current to the lateral giant. When the afferents are fired in two groups, the synapses operate at a higher conductance, generating larger currents. Each group of five afferents fired separately contributes almost as much current as ten afferents are fired simultaneously. Thus, firing the afferents in two groups, causes a larger current transfer overall.

It should be restated at this point that this effect is entirely theoretical, and highly unlikely to occur naturally as it relies on a combination of parameter values outside of the biologically plausible range. Anti-coincidence detection is an artifact of the extreme parameter values chosen.

We have shown that depending on how we choose our parameters, rectifying junctions can facilitate coincidence detection or anti-coincidence detection. We now show that with another set of parameters they can cause the lateral giant to respond to inputs of different synchronisation in a manner

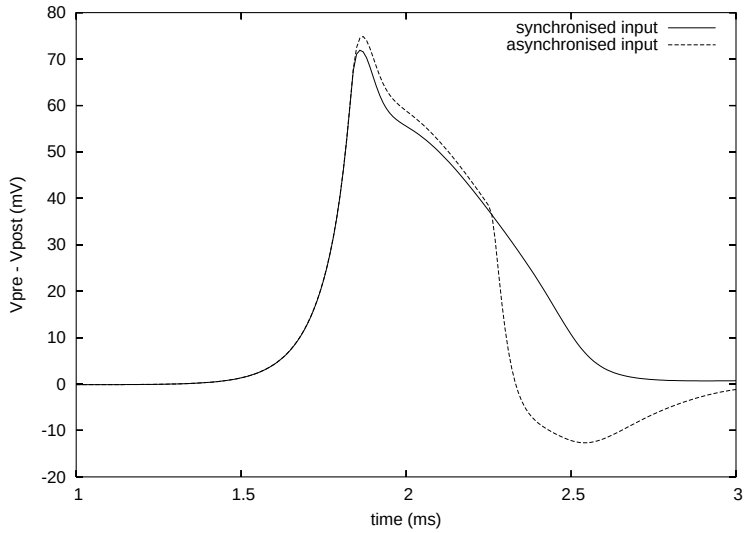

(a)

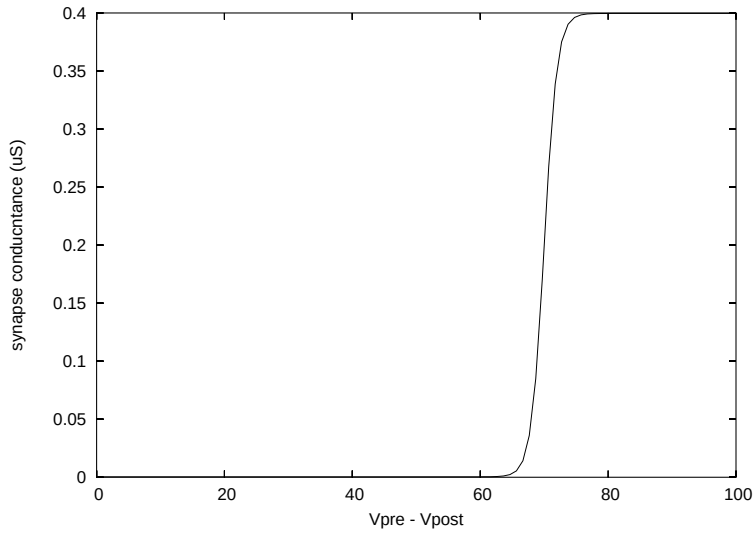

(b)

Figure 6.25: (a) The trans-synaptic voltage difference plotted over time in response to ten afferents fired synchronously, and in response to two groups of five afferents fired 0.4 milliseconds apart. (b) The steady-state conductance relationship to the trans-synaptic voltage difference, (refer to equation 3.1.) The constants used cause the lateral giant to act as a 'anti-coincidence' detector: $g_{\text{mean}} = 0.2\mu S$, $\frac{g_{\max}}{g_{\min}} = 1000$, $A = 1\text{mV}^{-1}$, $V_0 = 70$ *mV*, $\tau_{\text{open}} = 0$, $\tau_{\text{close}} = 0.25$ *ms.*

than can only be described as 'weird'. Figure 6.26 plots a scenario where the lateral giant efficiency fluctuates with the input synchronisation. The explanation for this rather unusual looking graph is that the governing parameter settings give rise to two different previously discussed phenomenon. The rectifying junctions act as hybrids of 'anti-coincidence' detectors and 'super-coincidence' detectors. When the input groups are slightly out of synchronisation ($\Delta t = 0.05$ milliseconds) the depolarisation caused by the early input reduces the junctional conductance of the late input, in a manner that is similar to the previously discussed results shown in figure 6.17. When the inputs are broadly separated, the depolarisation in the lateral giant dendrites has passed, and thus the late input can be effective. In these circumstances, the circuit acts like the just discussed 'anti-coincidence' detector. Like anti-coincidence detection and super-coincidence detection, the hybrid behaviour relies on a conductance curve with a high activation threshold ($V_0 = 70$ millivolts) and a very steep transition ($A = 1\text{mV}^{-1}$). Again, it is an interesting phenomenon of the model, but unlikely to occur biologically.

6.4.3 Conclusions

The simulation results answer the questions we set out at the start of this section.

Rectifying junctions can significantly add to coincidence detection in our morphologically realistic model of the lateral giant.

In general, the best coincidence detection occurs when the parameters are set so that the junctions are partially open at rest, and close quickly when reverse biased. Early input depolarises the lateral giant dendrites causing the junctions to close, thus blocking the pathway for late inputs. This result is in agreement with one of the solutions obtained with the abstract morphology model.

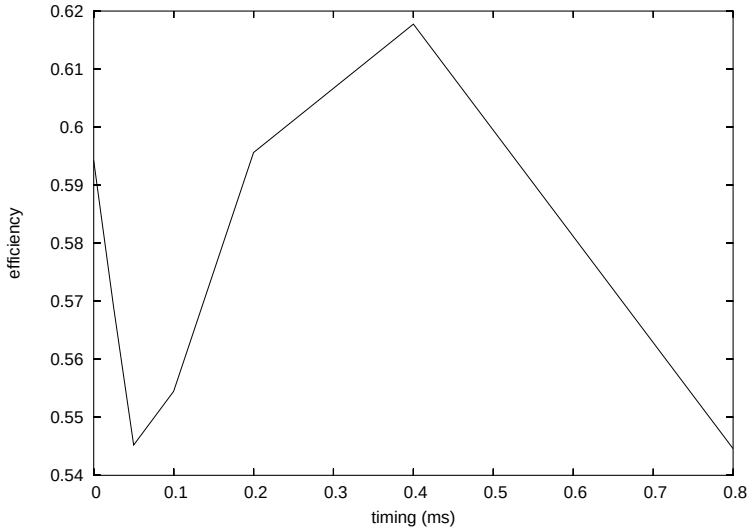

Figure 6.26: The efficiency of summing two groups of afferent inputs as a function of the time interval separating (Δt) their firing. Under the chosen rectifying junction parameter values the efficiency fluctuates with Δt. (Parameter setting: $g_{\text{mean}} = 0.4\mu S$, $\frac{g_{\max}}{g_{\min}} = 1000$, $A = 1\text{mV}^{-1}$, $V_0 = 70\ mV$, $\tau_{\text{open}} = 1\ ms$, $\tau_{\text{close}} = 1\ ms$.)

Finally, our results show that the effect of rectifying junctions on coincidence detection can vary widely depending on what parameter values are chosen to govern their conductance-voltage relationship. With the appropriately chosen parameters, they can greatly enhance coincidence detection. With other parameter settings, the rectifying junctions can cause the lateral giant to sum non-simultaneous inputs better than simultaneous ones, and act like a 'anti-coincidence' detector. The effects of this give rise to interesting but difficult to intuitively predict phenomena. One thing can be stated with absoluteness: their effect is not neutral.

6.5 General discussion and conclusions

In this chapter we have explored the role that rectifying junctions can play in our model in regard to coincidence detection.

In the results of our simulations, it is clear that rectifying junctions have a great impact on the way the lateral giant receives input from the afferents. The exact nature of this change depends on the parameter settings of the junctions. With rectifying junctions we were able to achieve outcomes that were impossible using ohmic gap junctions.

We have shown in, two different multi-compartment models of the lateral giant tree, that rectifying junctions can cause coincidence detection. This complements previous simulations (Edwards et al., 1998), that demonstrated rectifying junctions can causing coincidence detection in single compartment models.

Our simulations show that the rectifying junctions play a critical role in coincidence detection. If we use ohmic gap junctions in place of rectifying junctions, to connect the afferents to the lateral giant, then we are unable to get as effective coincidence detection. Indeed, it is only through the use of rectifying junctions that we are able to achieve the degree of coincidence detection that matches the biological observations. We have demonstrated that rectifying junctions are a necessary condition of coincidence detection.

In our simulations we have identified that there are alternative ways of achieving coincidence detection with rectifying junctions. The method that gave both good results and is biologically feasible, has the rectifying junctions partially open at rest. The depolarisation caused by early input reduces the conductance across the junctions making later input less effective. Achieving coincidence detection in this manner was demonstrated in both the abstracted morphology model and the morphologically accurate model. The idea of having the junctions open at rest is a new hypothesis in regard to the operation of the lateral giant circuit. The electrophysiological data available (Edwards and Herberholz, personal communication) is consistent with this hypothesis.

The rectifying junctions must not only act as coincidence detectors but they must also be conducive to admitting enough current to fire the lateral giant. Excellent coincidence detection that produces EPSPs that are too small to reach firing threshold is not helpful to the animal. A general trend that was observed in our results was that lower junctional conductance gave better coincidence detection but lower EPSPs. A trade-off is necessary in choosing a junctional conductance that delivers good coincidence detection but still allows the lateral giant to get to firing threshold.

The conceptual model performs comparably to the morphologically accurate one in the scenarios tested in this chapter. The best junction parameter values for coincidence detection have similar values when we compare the results of the parameter searches of the two models. Hence we can say, at least in the context of coincidence detection, the conceptual model successfully captures the essential features of the lateral giant morphology.

Rectifying junctions are simple devices that have very useful applications. The time and voltage variance of the rectifying junctions gives rise to complex dynamics. In our simulations, we have demonstrated that rectifying junctions are very effective as coincidence detectors. We have also demonstrated that rectifying junctions can serve other functions. Rectifying junctions are very versatile devices. Other functions we have demonstrated include: selecting for inputs that are out of coincidence or selecting for multiple kinds of input phasicness, e.g. responds well to highly phasic input and widely desynchronised input but poorly to slightly desynchronised input.

We have demonstrated that rectifying junctions are highly configurable in their stringency requirements of the synchronisation of their input. Our simulated rectifying junctions can be set up to discriminate for inputs to be within 0.05 milliseconds within each other or to have less stringent requirements as to how synchronised the input must be.

The amount of coincidence detection we achieve with our model compares favourably with the biological data. In lateral giant, inputs desynchronised by as little as 0.1 milliseconds, produce EPSPs that are about 25% lower (Edwards et al., 1998). In our simulations we are able to replicate the narrowness of the lateral giant's response to phasic input and the drop in EPSP that occurs when the input is out of phase. We have demonstrated that rectifying junctions provide a sufficient explanation of how coincidence detection occurs in the lateral giant.

Our model and simulation results make a number of predictions that have implications for the neurobiological understanding of the lateral giant escape circuit. Firstly, our simulations predict that the rectifying junctions are at least partially open at rest. Our first prediction has implications that give rise to our second prediction. If the gap junctions are open at rest then this combined with the fact that the afferents and lateral giant have different membrane resting potentials, then there would be a continuous current flow between the two cells.

Finally, although we have shown that rectifying junctions are very effective at coincidence detection, we have also demonstrated that they can be configured to serve a wide range of functions. Gap junctions have been shown to widely exist in the nervous systems of many different animals (Bear et al., 1996). There have been shown to be many different kinds of gap junctions. Their voltage dependent conductive properties vary widely depending on the makeup of their connexins (Ebihara, 2003). We have shown that rectifying junctions can play different functional roles depending on the kinetic properties of their trans-junctional conductance. We predict that rectifying junctions play different important functional roles across many neural circuits. Experiments where connexins are replaced with other kinds of connexins cause functional deficits in animals (Plum et al., 2000; White and

Bruzzone, 2000; White, 2003). The results of these experiments are consistent with our hypothesis.

We have demonstrated that we can achieve coincidence detection in our artificial scenario. We have shown, in our computational model, that the rectifying junctions can make a large contribution to coincidence detection. A highly coincidental sensory stimulus is a key signature of a predator assault. Detecting this efficiently and robustly is an important property of the crayfish escape system. We have demonstrated that we can explain and reproduce this important property in an artificial system.

If you wish to learn how a network of afferents can collaborate to influence the escape decision you will have to read the next chapter.

CHAPTER 7

Modelling the afferent network

This chapter describes a computational model of the afferent network.

The primary afferents in the lateral giant circuit (Herberholz et al., 2002) are connected to form a lateral excitatory network, whereby stimulated afferents recruit non-stimulated afferents, thereby amplifying the response. Afferents make direct connections to other afferents and indirect connections through the gap junctions to the lateral giant dendrites (Herberholz et al., 2002; Antonsen and Edwards, 2003).

The model described in this chapter is used to consider these questions:

- What are the necessary model conditions and parameter settings in order for the model to produce recruitment as observed experimentally? This is addressed in section 7.2.2.
- The afferents are connected both directly and indirectly through through the lateral giant dendrites, what are the relative contributions of these two couplings to recruitment? Sections 7.2.3 and 7.2.4 address this

question.

- Could afferents in one nerve recruit afferents in another? The coupling between nerves is examined in section 7.2.5.

- How does the response of afferent network change with different phasic input patterns? This is addressed in section 7.3.

- How does the response of afferent network change with different spatial input patterns? This question is answered in section 7.4.

- What effects do different connection patterns between afferents have on recruitment in afferent network? This is explored in section 7.4.1.

- What is the function of the interneurons in the circuit? This is addressed in section 7.5.

7.1 Model description

As described in section 3.3.2, the afferents in the crayfish tailfan are organised into five 'nerves', each of which contains about fifty afferents. In each nerve, the afferents are organised into a network. Each afferent network functions mostly independently of the networks in the other nerves. We made a model of an afferent network in one of the nerves. The networks in the other nerves are assumed to operate in a similar fashion. Our model ignores the occasional direct connections between afferents in different nerves (Antonsen and Edwards, 2003), whose impact on the operation of the circuit we assume to be insignificant. In our model, as is thought to be the case in the biological circuit (Antonsen and Edwards, 2003), nerves interact with each other by antidromic currents via the lateral giant. This interaction is modelled in section 7.2.5.

The afferent network was modelled using the NEURON simulator (Hines, 1993; Hines, 1994; Hines, 1998). Figure 7.1 depicts an overview of the afferent network model components and how they are connected. The model is comprised of a group of fifty afferents and a model of the lateral giant.

A morphologically accurate model of the lateral giant was used in the simulations. The LG model used is described in section 5.2. The afferent model is described in section 7.1.1.

The model afferents are connected together in a network. In most simulations afferents were connected together in a regular pattern, whereby each afferent was connected to n of its nearest neighbours. Alternative connection patterns and their consequences on the afferent network response are simulated in section 7.4.1.

The model afferents also make weakly rectifying synapses onto the lateral giant. These connections cause the afferents to be coupled indirectly via the lateral giant, as has been shown in biological experiments (Herberholz et al., 2002).

The interneurons are not included in simulations until their role is examined in section 7.5. Excluding the interneurons makes the model simpler and allows us to examine the role of the afferent network in isolation. By studying the afferent network in isolation we can determine what kind of input the afferents provide to the interneurons. The communication between the afferents and the sensory interneurons is unidirectional. The synapses are chemical which means there are no antidromic currents to consider. It is therefore a reasonable to assume that the operation of the afferent network can be studied independently of the interneurons.

The interneurons are a vital part of the lateral giant circuit. They generate the β component of the lateral giant input, which is the most significant

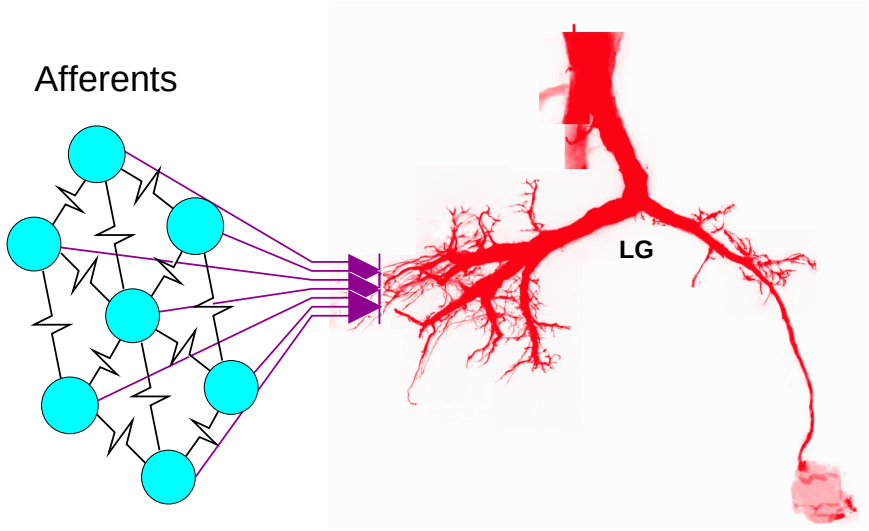

Figure 7.1: The schematic layout of the circuit used to simulate the afferent network. A group of fifty afferents (only seven shown in light blue) are connected together and form weakly rectifying synapses onto one of the major branches of the lateral giant (LG.) A morphologically accurate LG model is used.

component. We treat the interneurons as a 'black box'. The interneurons are considered as a complex function that integrates the input provided to it by the afferents and communicates the result to the lateral giant. Our simplified treatment of the interneurons is not only reasonable but also necessary. The level of knowledge needed to make an accurate quantitative model of the interneuron integration is lacking.

Formally, the simulation can be thought of as a function. Let $\mathbf{x}(t)$ be the vector of input functions, representing the external stimuli functions given to the afferents. Let $\mathbf{y}(t)$ represent the outputs of the simulation that we choose to measure. Let the vector θ describe the parameters characterising the setup of the simulation. The output of the simulation can be thought of as a function of the input stimuli, $\mathbf{x}(t)$, and the simulation parameters, θ. Equation 7.1 formalises this relationship. This is depicted in a diagram in figure 7.2.

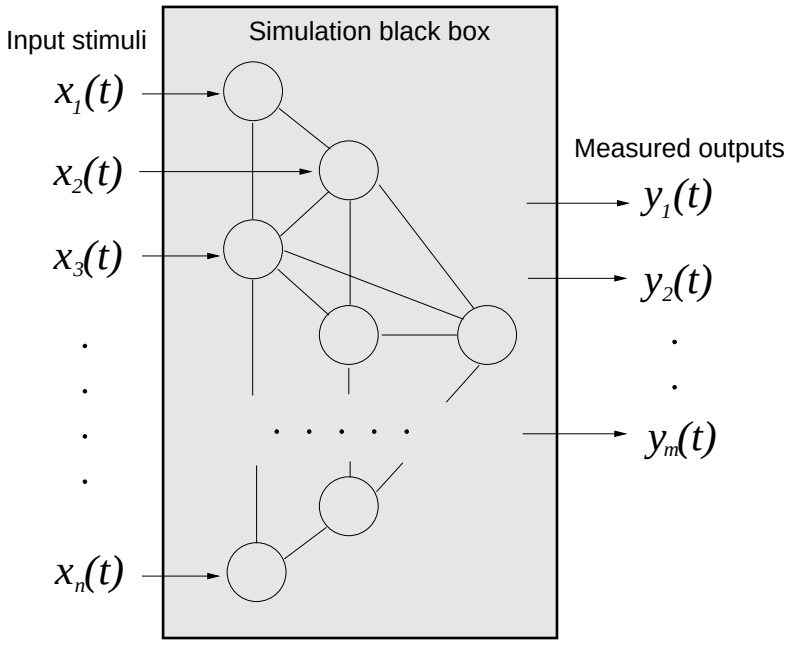

Figure 7.2: The relationship between the vector of input stimuli $\mathbf{x}(t)$ *and the vector of measured outputs* $\mathbf{y}(t)$*. The simulation is considered a black box.*

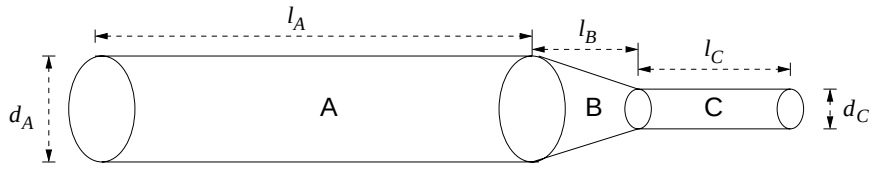

Figure 7.3: A schematic diagram of how the afferents are modelled. The model is divided into three sections: (A) a wide axon projecting from the sensory hairs, (B) an intermediate zone, where the axon narrows, and (C) a narrow section modelling the terminal end of the afferent.

$$y_i(t) = f_i(\theta, h_1(t) * x_1(t), \ldots, h_n(t) * x_n(t)) \tag{7.1}$$

Where $y_i(t)$ represents the ith simulation output, x_j(t) represents the jth simulation input, θ is the parameter vector, f is a function that produces the output of the simulation and $h_j(t)$ is a transfer function.

7.1.1 Afferent model

The afferents are modelled as tubes with a narrowing at their terminal ends. The afferent model is similar to the model used in Chapter 5. In both cases, the afferents are modelled as tubes with active Hodgkin-Huxley sodium and potassium channels. A higher density of compartments is used at the terminal ends. This is done because the afferents connect to each other at the terminal ends, so the morphology at this part of the cell may play a critical role in the activation of spikes. The narrowing of the afferent's diameter is included in the model. Two extra sections are added: a narrow section at the end, and a transitory section, where the diameter narrows from the wide diameter to the narrow diameter (see figure 7.3.) It is not known whether the narrow end of the afferent is passive or active. Both possibilities were experimented with and this choice did not appear to have a large impact on the simulation results.

The variability in the afferent dimensions (Antonsen and Edwards, 2003) was ignored in the afferent network simulations. To simplify the simulation all afferents were given identical dimensions. All afferents were given a diameter of 5 micrometres, and a length of 3 millimetres. These chosen dimensions are consistent with the overall mean dimensions of the afferent populations in actual crayfish (see section 3.3.2). Giving all afferents consistent dimensions reduces the complexity of the network which lessens the variability of the afferent network output. The simplification makes it easier to study the effects of the other parameters, for example the effect of the connection pattern between afferents.

7.2 The necessary conditions for recruitment

One of the key questions being addressed by the simulation is to identify the subset of the parameter space that causes the recruitment response of the model afferent network agrees with the experimental observations (Herberholz et al., 2002). In biology, recruitment does not occur unless the nerve is stimulated above a threshold. If the stimulus magnitude is increased beyond that point, non-stimulated afferents are recruited, thereby enhancing the PSPs in the lateral giant. When all available afferents have been recruited, the saturation point is reached. Beyond this point, further increases in the magnitude of nerve stimulation do not recruit any additional afferents (Herberholz et al., 2002). We expect our simulation to produce a stimulation-recruitment response with similar characteristics. As we increase the fraction of afferents stimulated to be above threshold, there should be a steep transition from virtually no recruitment to maximum recruitment. Our aim is to determine what parameter settings are necessary for a recruitment response that matches this description. We also wish to learn how each parameter affects recruitment.

In order to be able to formulate objective criteria by with which we can evaluate the model results it is necessary to define a quantitative measure of recruitment. Let recruitment gain, G, be defined as the total number of afferents that fire divided by the number of afferents initially stimulated:

$$G = \frac{n_{\text{stimulated}} + n_{\text{recruited}}}{n_{\text{stimulated}}} \tag{7.2}$$

Where $n_{\text{stimulated}}$ is the number of afferents initially stimulated and $n_{\text{recruited}}$ is the number of afferents recruited.

7.2.1 The parameter space

The setup of the afferent network simulation is parameterised by many variables.

Some of the simulation parameters, θ, include:

Rectifying junction parameters: The parameters governing the voltage dependent conductance across the rectifying junctions that connect the primary afferents to the lateral giant (see section 5.3.)

Afferent-afferent conductance: the ohmic resistance connecting afferents.

Number of connections: the number of connections between afferents.

The membrane properties of the LG: including the capacitance, the axial resistance, the passive membrane resistance, the leakage current and the properties of the active membrane channel.

The membrane properties of the afferents: each of the membrane parameters described for the LG also applies to each afferent.

The dimensions of the afferents: the size and morphology of the model afferents.

Connection organisation: the organisation of the connections between afferents. Some possible organisations include:

- a regular pattern, where afferents are arranged in a hexagonal grid and connect with n of the adjacent afferents. Such an organisation is unrealistic but easier to understand.
- completely random connection organisation, where each afferent has an equal chance of connecting to any other afferent, regardless of proximity.
- a probabilistic connection organisation, where the probability of two afferents being connected is proportional to the distance separating them. Such a scheme is biologically plausible.

We setup simulation with the following simplifications and assumptions:

- The properties of all rectifying junctions were made to be identical
- All afferents were given identical properties, *i.e.* all afferents were of the same length and diameter and had the same membrane properties.
- The ohmic connections between afferents were given identical properties, *i.e.* the same conductance was used for all afferent-afferent synapses.

Even with these simplifications, the parameter space is too large to comprehensively search. There are two alternatives: (1) to search for 'optimal' parameters using a parameter optimisation algorithm such as a genetic algorithm (section 4.1.3 discusses parameter searching methods), or (2) to map out a subset of the parameters. We opted for the latter approach. The

parameter space was mapped out rather than optimised for two reasons. Firstly, mapping out the parameters is more conducive to analysing how each parameter affects the results. Secondly, there are multiple solutions which give results that are compatible with the experimental observations. We desire to know the set of possible parameter settings rather than have just one or a few solutions.

Four simulation parameters were chosen to be mapped out:

- the number of connections between afferents,
- the conductance between afferents (g_{aff}),
- the maximum conductance of the LG-afferent synapses ($g_{\max}$ - see section 5.3), and
- the rectification across the LG-afferent synapse ($g_{\max}/g_{\min}$ - see section 5.3.)

The these parameters are the most relevant to the questions presented in the introduction. Table 7.1 summarises the range over which the selected parameters were searched. The parameter search ranges were determined by choosing as a 'reasonable' value as a midpoint. 'Reasonable', as determined by the biological constraints and the results of tests in preliminary simulations. Values greater and less than this value were then searched to the bounds of plausibility. Parameters were varied on either a linear or log scale. The steps between values were quite large. It was necessary to balance on the one hand, choosing enough values to gain an understanding of how the parameter affects the network, and on the other hand, limiting the size of the search space.

All other parameters of the simulation was kept constant. Refer to appendix B.3 for the values used.

Table 7.1: The parameter space mapped out in the afferent network simulations

Parameter	Min	Max	Distribution	Number of points
Stimulated afferents (n)	0	50	linear	11
Number of connections (c)	0	12	linear	13
Afferent-afferent conductance (g_{aff})	0.05 μS	0.8 μS	log	5
Afferent-LG max. conductance ($g_{\max}$)	0.0125 μS	0.8 μS	log	7
Rectification factor ($\frac{g_{\max}}{g_{\min}}$)	1	1000	log	10
Total number of simulations				50050

The input to the circuit was a group of simultaneously spiking afferents. The input was changed only by increasing the number of afferents spiked. The afferents which were initially stimulated were chosen randomly. So that the response of the networks with different parameter settings could be directly compared, the same random seed was used to choose the initially spiked afferents for all parameter settings. For example, when five afferents were spiked the same five were spiked for all simulation runs.

For each simulation run we record the following:

- the number of afferents recruited (r), and
- the PSP at the base of the lateral giant dendrites (v_{lg}). Since the interneurons are not included in the simulation this is only a measure of the α-PSP.

A simulation was run for each of 50050 parameter possibilities, and the above outputs were recorded for each simulation.

The number of connections between afferents was one of the varied parameters. For all numbers of connections, a uniform connection pattern was used. This connection organisation is described in section 7.4.1 and drawn in figure 7.21.

7.2.2 Overall results

The recruitment response varied greatly with the different parameter values tested in the parameter space. A general trend that occurred very frequently throughout the parameter space was that the onset of recruitment was very sharp. In the simulated network as the number of stimulated afferents was increased past a certain threshold, the recruitment gain made a very steep transition from from unity (that is no recruitment) to its maximal value. The band at which recruitment occurs is very narrow. We refer to this phenomenon, as 'hair trigger' recruitment. Hair trigger recruitment was a robust property of the network that occurred for a wide range of parameters. Hair trigger recruitment is also what has been shown to occur experimentally (Herberholz et al., 2002). We can therefore say that for many different parameter settings, the response of the simulated afferent network was consistent with the experimental results. Figure 7.4 shows an example recruitment gain curve. Figure 7.5 shows the recruitment gain curves for all setups of the simulation within the subspace of having six afferent connections. Let us call the threshold of the number of stimulated afferents required for recruitment to start occurring the 'trigger threshold' (shown graphically in figure 7.4.) Changing the parameters of the network changed the slope of the curve and the trigger threshold. In this and the following sections we will talk more of the characteristics of these relationships.

The shape of the recruitment gain curves is well expressed with the equations:

$$G(x) = \frac{x + r(x)}{x} \tag{7.3}$$

$$r(x) = p(x)(1 - x) \tag{7.4}$$

Where x is the fraction of afferents stimulated, $r(x)$ is the fraction of re-

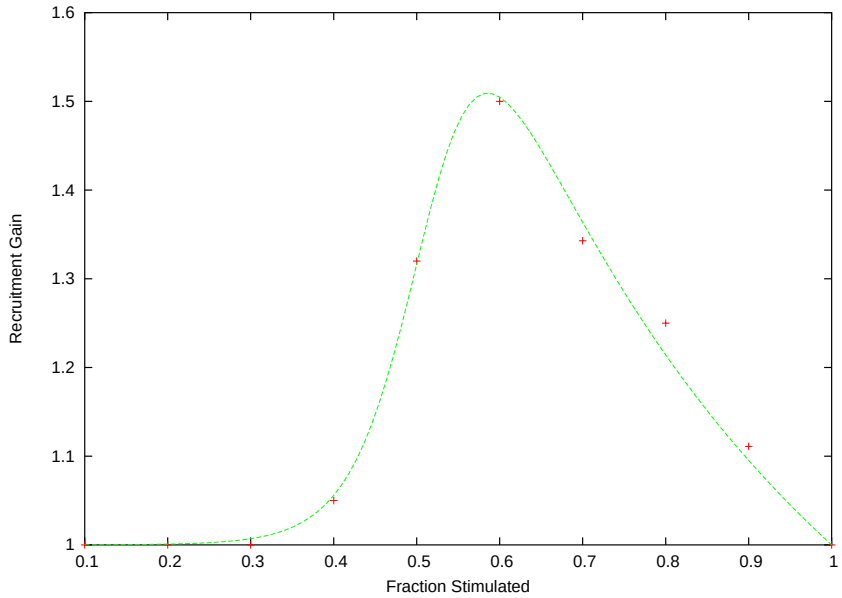

Figure 7.4: An example recruitment gain curve of the simulated afferent network. Along the X-axis is the number of neurons initially stimulated. Along the Y-axis is the recruitment gain. It can be seen that the recruitment gain makes a steep transition from no recruitment to peak recruitment as the number of stimulated afferents passes the trigger threshold.

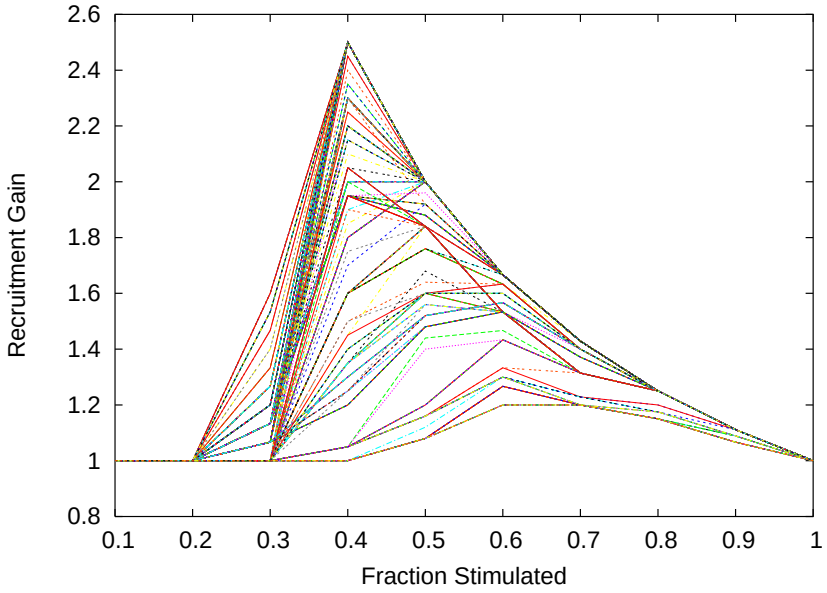

Figure 7.5: The good, the bad and the ugly. The recruitment gain curves for 6 afferent connection networks for all parameters searched.

cruited afferents and $G(x)$ is the recruitment gain. The function $p(x)$ is the probability of a non-stimulated afferent being recruited. These equations model recruitment as the product of the fraction of non-stimulated afferents and the probability of such afferents being recruited.

We found that the probability of recruitment was best modelled with a sigmoidal function:

$$p(x) = \frac{p_{\max}}{1 + e^{-k(x - x_R)}} \tag{7.5}$$

Where $p_{\max}$ is a constant which represents the maximum recruitment probability, k is a constant which determines the steepness of recruitment and x_R is proportional to the trigger threshold.

Figure 7.4 shows a recruitment gain curve obtained from the computational model fitted with equations 7.3, 7.4 and 7.5.

A quantitative method is needed to determine whether a simulated network's gain response meets the criteria for hair trigger recruitment. This method

can then be applied to the simulation results to determine which parameter subspace produces recruitment curves consistent with the experimental observations.

One simple method is to measure the steepness of the change from minimum to maximum recruitment. This may be achieved by calculating the gradient of the graph as it moves from some arbitrary small recruitment gain to an arbitrary large value.

Our aim is to quantify the rapidness of the change from low recruitment to high recruitment. In order to compare the rapidness of this change in curves with different trigger thresholds it is necessary to first normalise the gain. Curves with an earlier trigger threshold are likely have a higher maximum gain. As can be seen in equation 7.2 the maximum possible gain at any point is $\frac{n_{\text{total}}}{n_{\text{stimulated}}}$. One method of normalisation is for the data points of each curve to be divided by the maximum gain for that curve:

$$z(x) = \frac{G(x) - 1}{G_{\max} - 1} \tag{7.6}$$

Where z is the normalised gain, $G_{\max}$ is the maximum gain of the curve and x is the fraction of afferents stimulated ($x = \frac{n_{\text{stimulated}}}{n_{\text{total}}}$.) Since the smallest possible gain for all curves is unity, this can be subtracted from the value. This method normalises the gain to have a value of between 0 and 1 for all curves.

We can then quantify the steepness by calculating the gradient of the curve as it moves from a chosen small value, z_{small} to a chosen large value, z_{large}. In practice the values for z_{small} and z_{large} that were used were 0.1 and 0.9, respectively.

Due to the sparseness of the data points, the exact point at which the graph reaches z_{small} and z_{large} is usually not exactly known. However their

values can be approximated through interpolation, (refer to illustration in figure 7.6.) Hence, let steepness, S, be defined as:

$$S = \frac{\Delta z}{\Delta x} = \frac{z_{\text{large}} - z_{\text{small}}}{x_{\text{large}} - x_{\text{small}}} \tag{7.7}$$

Where x_{small} and x_{large} are defined such that $z(x_{\text{small}}) = z_{\text{small}}$ and $z(x_{\text{large}}) = z_{\text{large}}$.

Interpolation has limited accuracy, especially for a very steep curves, where the recruitment transition may take place entirely between data points. However, due to the sigmoidal shape of recruitment curves the method underestimates the steepness. It is thus useful for obtaining a lower bound of a curve's steepness.

Table 7.2 tabulates the steepness range of the recruitment curves for all the parameters tested. These results show that by this measure, the majority of the curves have a steep transition from minimal to maximum recruitment. Let us consider a gradient of 4 as the threshold for a steep curve, (this corresponds to 80% of the recruitment gain rise occurring within a Δx of 0.2 or less.) By this criteria, we can say that 76% of the curves made a steep recruitment transition.

For some of the curves the recruitment gain never exceeded 1, *i.e.* no extra neurons were ever recruited. No steepness value can be calculated for these curves. These curves are shown in the data as having a value on 'NA' and represent approximately 7% of the curves. Most of these curves come from afferent networks with zero connections between afferents.

One limitation of the previously described steepness measure is that there is no discernment between curves that have a high maximum gain and curves that have low maximum gain. An alternative method for normalisation is to divide all values by the maximum achievable gain at the maximum

Steepness range	number of curves	fraction of curves
NA	318	0.0699
$0 < S < 1$	0	0.0000
$1 \leq S < 2$	19	0.0042
$2 \leq S < 3$	159	0.0349
$3 \leq S < 4$	584	0.1284
$4 \leq S < 5$	254	0.0558
$5 \leq S < 6$	907	0.1993
$6 \leq S < 7$	286	0.0629
$7 \leq S < 8$	205	0.0451
$8 \leq S < 9$	36	0.0079
$9 \leq S < 10$	612	0.1345
$S > 10$	1170	0.2571

Table 7.2: Steepness range of recruitment curves

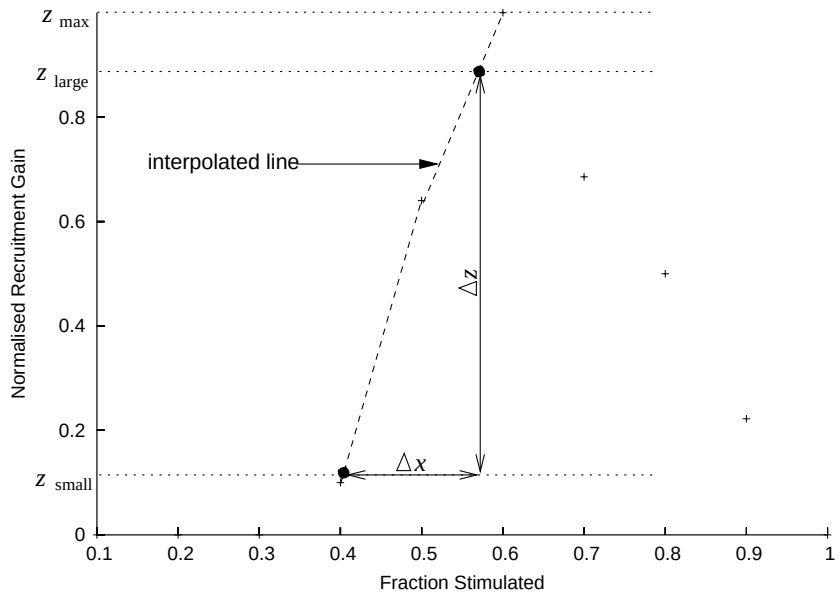

Figure 7.6: Method for measuring the steepness of the graph. Where steepness, S is defined as the gradient $\frac{\Delta z}{\Delta x}$ as the gain rises from the values G_{small} to G_{large}.

gain point rather than by the maximum gain of the curve. The maximum achievable gain is:

$$G^{\star}_{\max} = \frac{n_{\text{total}}}{n_{\max}} = \frac{1}{x_{\max}} \tag{7.8}$$

Where $n_{\max}$ is the number of afferents stimulated when the maximum gain of the network is achieved.

Normalised gain then becomes:

$$z^{\star}(x) = \frac{G(x) - 1}{G^{\star}_{\max} - 1} \tag{7.9}$$

Where $z^{\star}$ is the normalised gain and $G^{\star}_{\max}$ is the maximum achievable gain at the maximum point.

The steepness value then becomes:

$$S^{\star} = \frac{\Delta z}{\Delta x} = \frac{z^{\star}_{\text{large}} - z^{\star}_{\text{small}}}{x_{\text{large}} - x_{\text{small}}} \tag{7.10}$$

Using this method 45% of curves had a steepness value of 4 or greater. This is a much smaller proportion than the previous method. The steepness was unmeasurable for 54% of curves because the gain never reached the 0.9 threshold of maximum achievable gain. Table 7.3 tabulates the distribution of curve steepnesses using the second method.

How recruitment affects the α-PSP

Recruitment causes additional primary afferents to fire, which amplifies the LG PSP. We now use simulations to examine the relationship between recruitment and the LG PSP.

The recruited afferents may synapse onto interneurons or the directly onto lateral giant itself (or do both.) In the former case, the β input to LG

Steepness range	number of curves	fraction of curves
NA	2439	0.5360
$0 \leq S^\star < 1$	0	0.0000
$1 \leq S^\star < 2$	10	0.0022
$2 \leq S^\star < 3$	27	0.0059
$3 \leq S^\star < 4$	48	0.0105
$4 \leq S^\star < 5$	74	0.0163
$5 \leq S^\star < 6$	409	0.0899
$6 \leq S^\star < 7$	173	0.0380
$7 \leq S^\star < 8$	120	0.0264
$8 \leq S^\star < 9$	45	0.0099
$9 \leq S^\star < 10$	486	0.1068
$S^\star \geq 10$	719	0.1580

Table 7.3: Distribution of steepness using method 2

is enhanced, in the latter, the α is enhanced. The effect of recruitment on the α response is easily simulated. How the β is affected is difficult to accurately simulate because it depends on how the interneurons integrate their inputs. The interneuron integration is not currently understood at a sufficient level of detail to be able to predict the effect on the β response. We therefore concentrate our simulations on predicting recruitment's effect on the α response.

Figure 7.7 graphs the peak α-PSP in the LG dendrites versus the fraction of afferents stimulated. The total number of afferents spiked is plotted against the same axes. There is a steep increase in the number of afferents spiked when recruitment takes hold. However, there is no corresponding steep increase in the peak α-PSP. The α-PSP continues to increase, more or less linearly, as a function of the fraction of afferents initially stimulated. The α PSP continues to increase with the number of afferents initially stimulated even after all non-stimulated afferents have been recruited. This is in contrast to recruitment where once all available afferents are recruited the network is saturated, thus increasing the fraction of afferents stimulated has no further effect on the total number of afferents spiked. The recruited

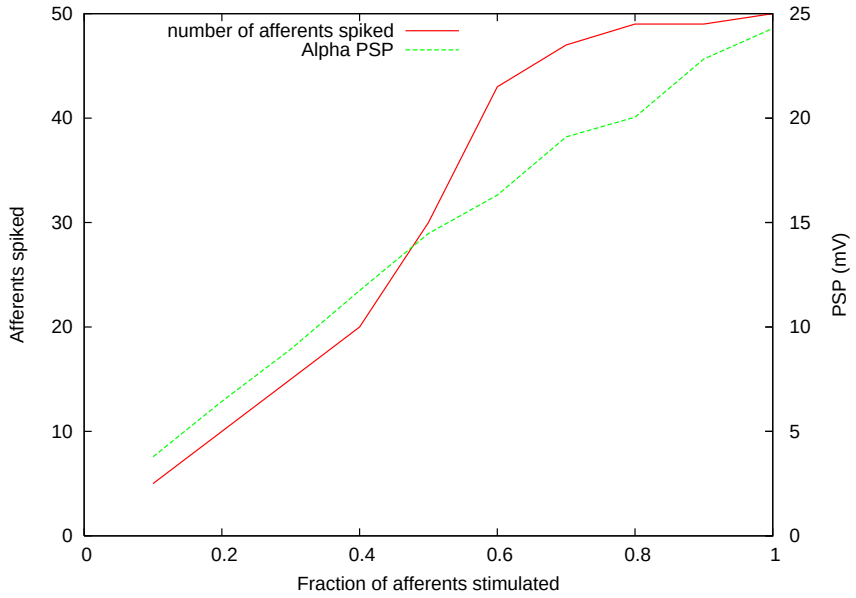

Figure 7.7: A comparison of how the peak α-PSP (in the LG dendrites) and the number of afferents spiked increase with the fraction stimulated. (Parameter settings used: 8 afferent-afferent connections, $g_{\text{aff}} = 0.05\mu S$, LG $g_{\max} = 0.025\mu S$, $\frac{g_{\max}}{g_{\min}} = 4.64$.)

afferents do not appear to have a significant impact on the α-PSP. There are some cases where recruitment causes a slight increase in the α-PSP (see figure 7.8 for an example.) The sizes of the increases are not proportional to the amount of recruitment.

The lack of an increase in the α-PSP corresponding to recruitment is caused by the recruited afferents spiking later than the stimulated afferents. This causes directly stimulated spikes and recruited spikes to be out of phase. The coincidence detection properties of the lateral giant (as seen in section 5.3) causes the out of phase input from the recruited afferents to have little extra effect on the peak PSP.

Even though recruitment has little effect on changing magnitude of the peak α-PSP, it does change the shape of the PSP by prolonging it. A sustained α PSP may have an important effect on the α's summation with the interneu-

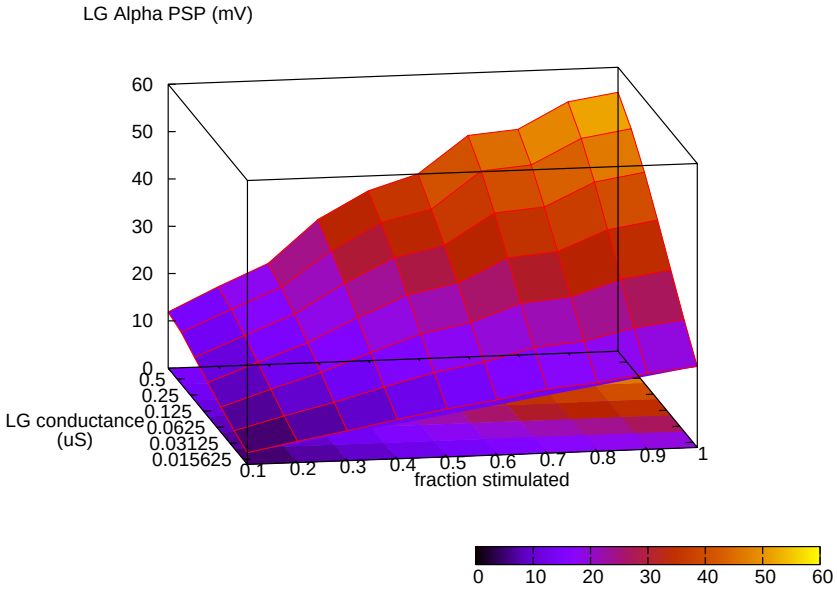

Figure 7.8: In some cases the recruited afferents increase the α PSP in the LG dendrites. This graph plots the fraction of afferents stimulated versus afferent-LG synapse conductance versus the α-PSP. Jumps in the α-PSP attributed to recruitment can be seen for high values of LG synapse conductance. (Parameter settings used: 6 afferent-afferent connections, $g_{\text{aff}} = 0.1\mu S$, LG $\frac{g_{\max}}{g_{\min}} = 46.4$.)

ron induced β-PSP. (This is discussed in section 7.5.)

7.2.3 Effect of afferent-afferent coupling

The effect of afferent-afferent coupling on the recruitment response of the circuit was examined by varying the number of afferent connections and the conductance of these connections.

The simulation results show that in most cases recruitment increases as the afferent-afferent conductance is increased. Similarly, increasing the number of afferent to afferent connections, generally, increases the recruitment. This is an expected result. Figure 7.10 plots the recruitment gain versus these two variables for a slice of the parameter space.

There are some circumstances where adding more afferent connections and/or increasing the conductance of the connections actually reduces the recruitment of the network. If there are a large number of afferent connections and each connection is of a very high conductance then the electrical load on a stimulated afferent may become too high. Adding further connection between afferents may actually reduce recruitment for two reasons. Firstly, the amount of current a spiked afferent is able to feed to its neighbours is limited by the conductance of its sodium channels. If the electrical load placed on the afferent by the connections to its neighbours starts to approach this limit, then adding additional connections will reduce the amount of current that each neighbour receives. Secondly, extra afferent connections reduces the resistance of each cell. Therefore each cell will require more current to bring it to spiking threshold. Recruitment is reduced because non-stimulated cells, connected to a spiked cell, receive a smaller current when they require a higher current.

As an example, consider two spiked afferents which make high conductance

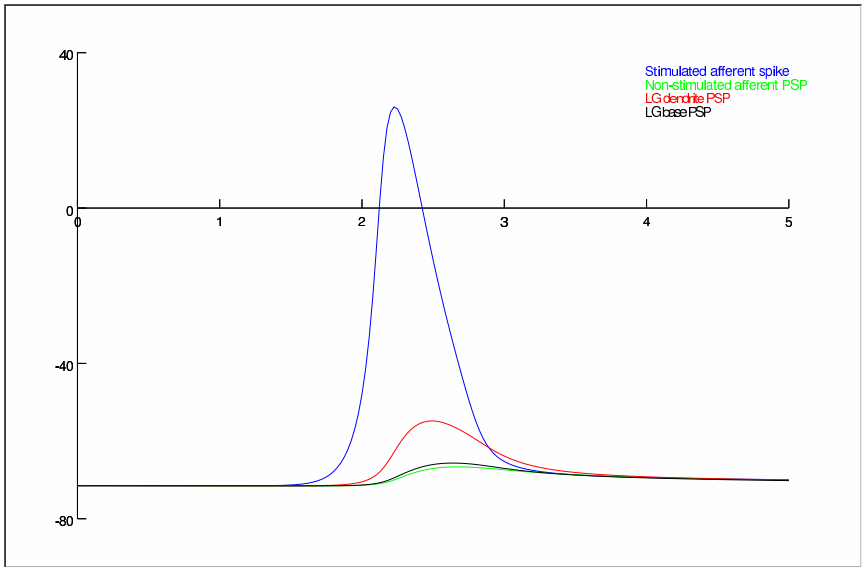

(a)

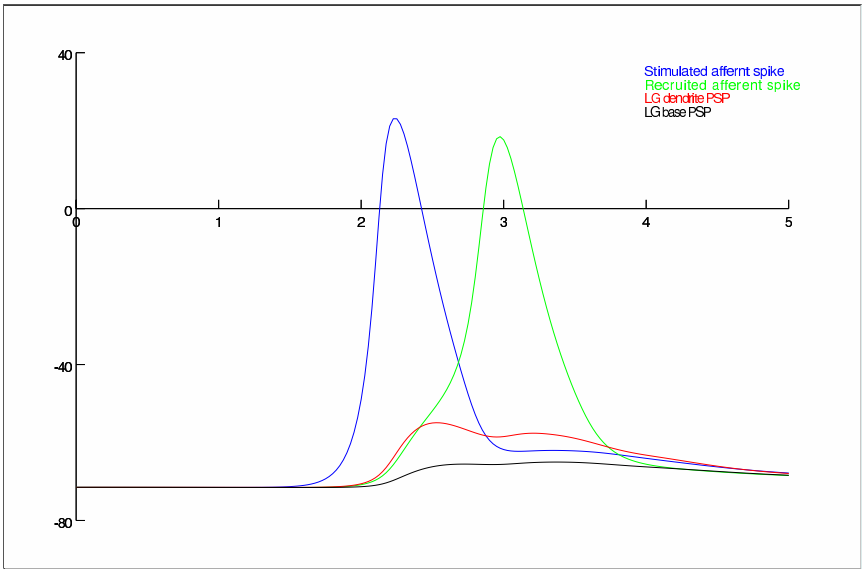

(b)

Figure 7.9: The effects of the recruitment on the α-PSP in the LG and the LG dendrites. ***(a)*** *Shows the PSP at the base of one of the LG branches and base when the afferents are unconnected and hence there is no recruitment.* ***(b)*** *Shows the PSP for a connected afferent network where there is recruitment.*

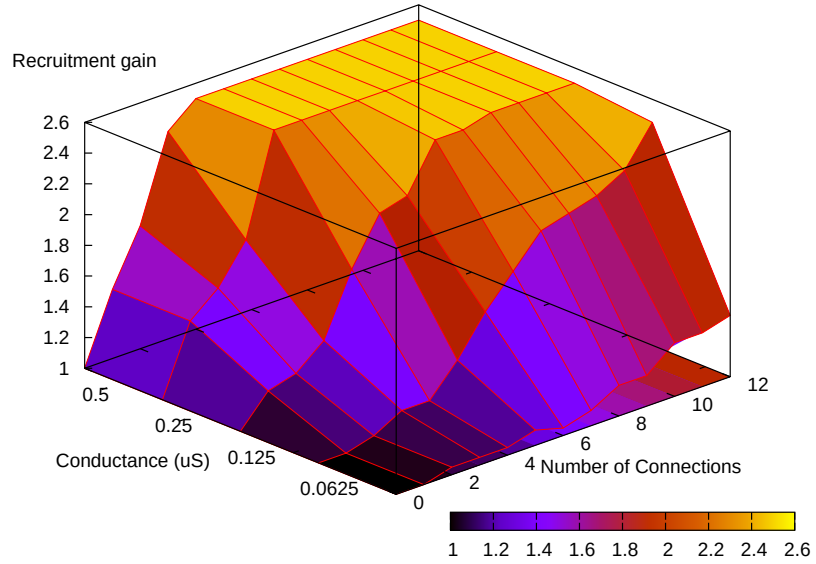

Figure 7.10: The recruitment gain increases with afferent-afferent conductance and with the number of afferent-afferent connections. A slice of the parameter space through these two axes is plotted here. The values of the other parameters are: fraction stimulated (x) = 0.4, LG $g_{\max} = 0.1\mu S$, $g_{\max}/g_{\min} = 46.4$

connections to the same non-spiked afferent. If the spiked cells connect to only a few other non-spiked afferents then they may be able to recruit this cell which they both connect to. However, if they are connected to a large number of afferents, each by a high conductance connection, then they may not be able to pass on as much current and may not recruit any other afferents.

Figure 7.11 plots two examples of where adding extra connections to the network actually reduces the recruitment performance.

It should be stressed that the effect of adding extra connections and getting diminished recruitment only happens in special cases. For most parameter settings, afferents are not overloaded. The general rule is that adding more connections enhances recruitment.

7.2.4 Effect of afferent-LG coupling

The simulation results show that the synapses between the lateral giant and the afferents are an important pathway for recruitment. Antidromic currents through the lateral giant synapses help bring unstimulated afferents to threshold. Recruitment increases significantly as the conductances of the afferent to lateral giant synapses are increased. Figure 7.12 plots recruitment curves with different values of conductance at the LG synapses.

The afferent to lateral giant synapses significantly contributed to recruitment for all sections of the parameter space. Increasing the conductance of these synapses causes recruitment peak to be earlier and larger. The contribution of the lateral giant synapses to recruitment becomes less important when the afferents are strongly coupled, *i.e.* when the afferent synapses have high conductance.

The coupling between afferents, through the lateral giant synapses, is insuf-

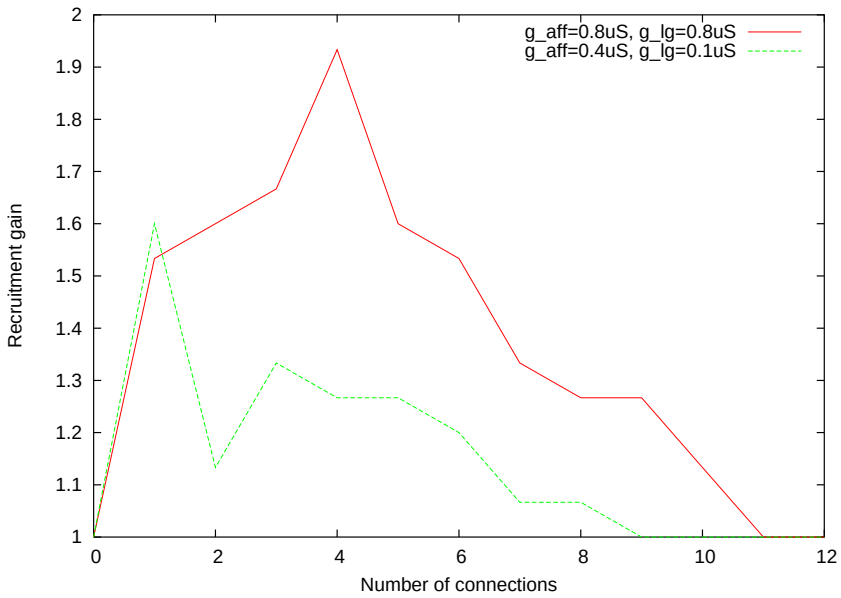

Figure 7.11: In some circumstances adding extra connections actually reduces the gain of the network (for a certain fraction stimulated.) This graph plots two examples. The red line plots a network where both the afferent-afferent synapses and the afferent-LG synapses have very high conductance ($g_{\text{aff}} = 0.8\mu S$, LG $g_{\max} = 0.8\mu S$, $\frac{g_{\max}}{g_{\min}} = 10$.) When 0.3 of the afferents are stimulated, the maximum gain is achieved with 4 connections. The recruitment performance declines as more afferents are added. The green line plots another example, where the maximum gain is reached when there is only one connection (0.3 of afferents stimulated, $g_{\text{aff}} = 0.4\mu S$, LG $g_{\max} = 0.1\mu S$, $\frac{g_{\max}}{g_{\min}} = 100$.)

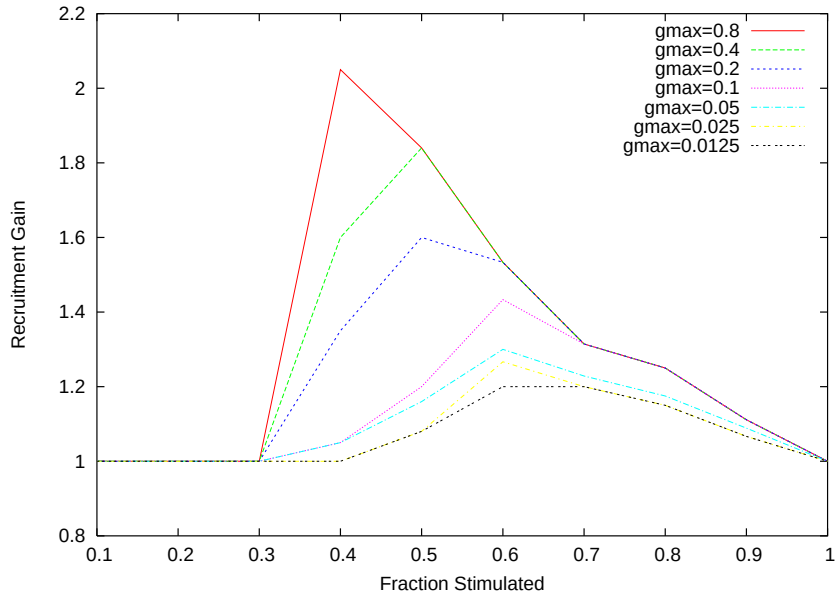

Figure 7.12: Antidromic currents through the LG synapses facilitates recruitment. As the conductance through LG synapses is increased recruitment increases. The parameters of the recruitment curves graphed vary only by their value of the LG synapse conductance $g_{\max}$. The values of the other parameters are: $g_{\max}/g_{\min} = 46.4$, Afferent synapse conductance $g_{\text{aff}} = 0.05\mu S$, Number of afferent connections = 6.

ficient to achieve recruitment by itself. In simulations where afferents made no direct connections to each other, and were coupled only through the lateral giant, no recruitment usually was achieved. Only when extremely high conductance values were used for the lateral giant synapses ($g_{\max} \geq 0.4\mu$S) was there a small amount recruitment. Even in these cases, relatively few afferents were recruited. Full recruitment was never reached. At the maximum gain, there were still more afferents available for recruitment. Figure 7.13 plots the recruitment gain for networks where afferents are coupled only through the lateral giant.

The rectification factor ($\frac{g_{\max}}{g_{\min}}$) of the afferent to lateral giant synapses significantly impacts on the performance of the network. The highest recruitment results occurred when there was no rectification. Networks without any

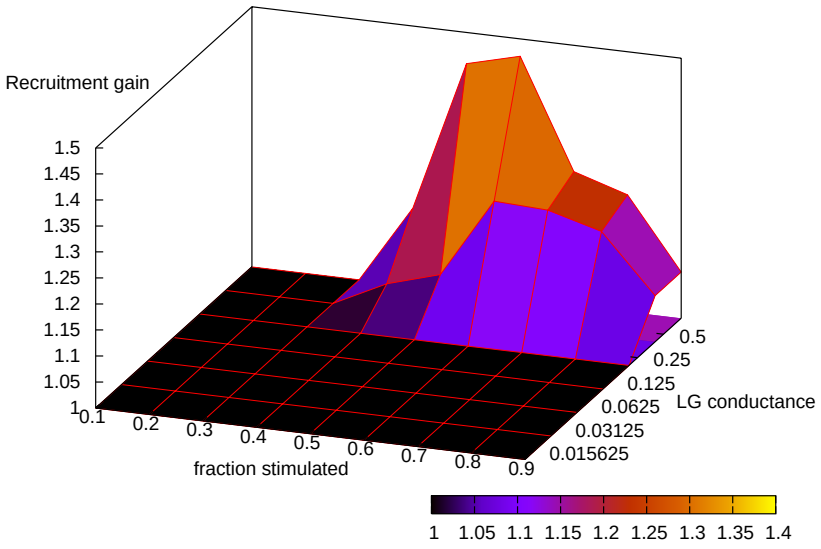

Figure 7.13: Recruitment caused through the lateral giant synapses alone. This graph shows recruitment gain versus fraction stimulated and afferent-LG synapse conductance for networks with no direct afferent to afferent connections. Only if extreme values are used for LG conductance is some recruitment achieved.

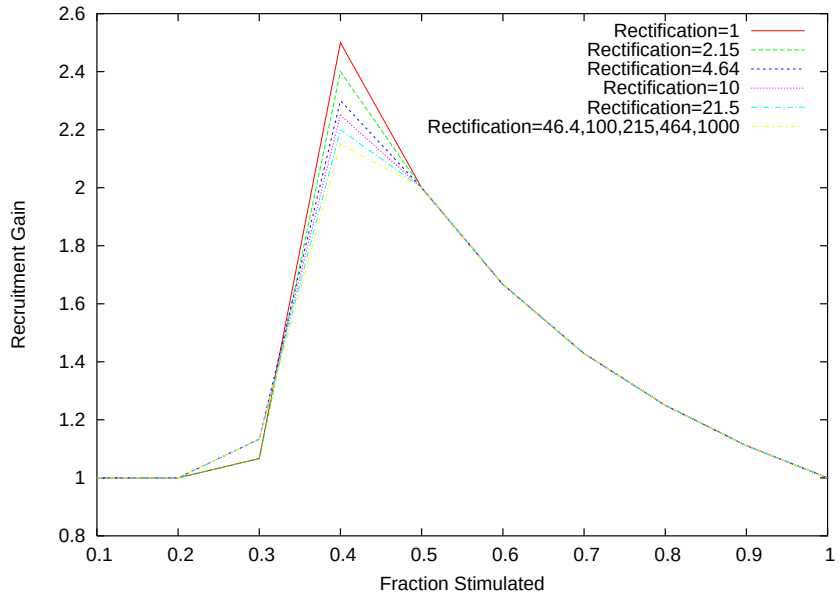

Figure 7.14: Lower rectification gives better recruitment. The recruitment curves for different rectification factors are plotted. The values of the other parameters are: LG $g_{\max} = 0.1\mu S = 46.4$*, Afferent synapse conductance* $g_{\text{aff}} = 0.2\mu S$*, Number of afferent connections = 6.*

rectification at the LG synapses had higher gain and the peak was reached earlier. This result is expected since no rectification allows higher antidromic currents. There is a quite sharp drop-off as one introduces some rectification. The point at which adding extra rectification has no additional effect on reducing recruitment occurs quite early. In some simulations it was as early as $\frac{g_{\max}}{g_{\min}} = 4.64$. That is for some simulations, choosing a rectification factor of 4.64 produced an identical recruitment response as choosing 1000. For all the parameter space tested, increasing the rectification factor beyond 46.4 had no effect on recruitment. These results should be qualified within the context of the values chosen for the other rectifying junction parameters. Figure 7.14 plots some recruitment curves with differing rectification factors.

The synapses between the afferents and LG are a potential mechanism for

modulating the escape response. We have shown that changing the conductance of the synapses has a dramatic effect on the size of PSPs in the LG. Serotonin is known to depress the response of LG. The synapses between the afferents and LG is a possible site where serotonin may be acting. The experimental changes in the response of LG can be accounted for by changing the conductance of the synapses. Simulations that decrease the conductance are consistent with the experimental application of serotonin.

7.2.5 Inter-nerve interactions

Afferents have been shown to recruit other afferents within the same nerve but what is the nature of interaction of afferents between nerves? In the model, the only source of such interaction is through the lateral giant synapses. If the afferents in one nerve are stimulated, it is useful to know what size PSPs are produced in the dendrites of the major branches of the other nerves and what size PSPs antidromically propagate into the afferents themselves.

This may be tested with a simple simulation. By stimulating one of the branches with afferent input and recording the potentials at synapse positions in the other branches. Only the PSPs by direct primary afferent connections are simulated, *i.e.* the α-PSPs. Figure 7.15 shows the locations of the recording points.

The results of this simulation show that a significant portion of the PSP in the stimulated branch spreads to the other branches. Figure 7.16 plots the PSPs at all the recording points. The PSP at B, (*i.e.* the branch that is closely connected with the stimulated branch), is almost equal in magnitude to the stimulated branch.

Figure 7.17 summarises the attenuation of the PSP at A to the other recording sites. The attenuation at the four other branches is averaged. Because

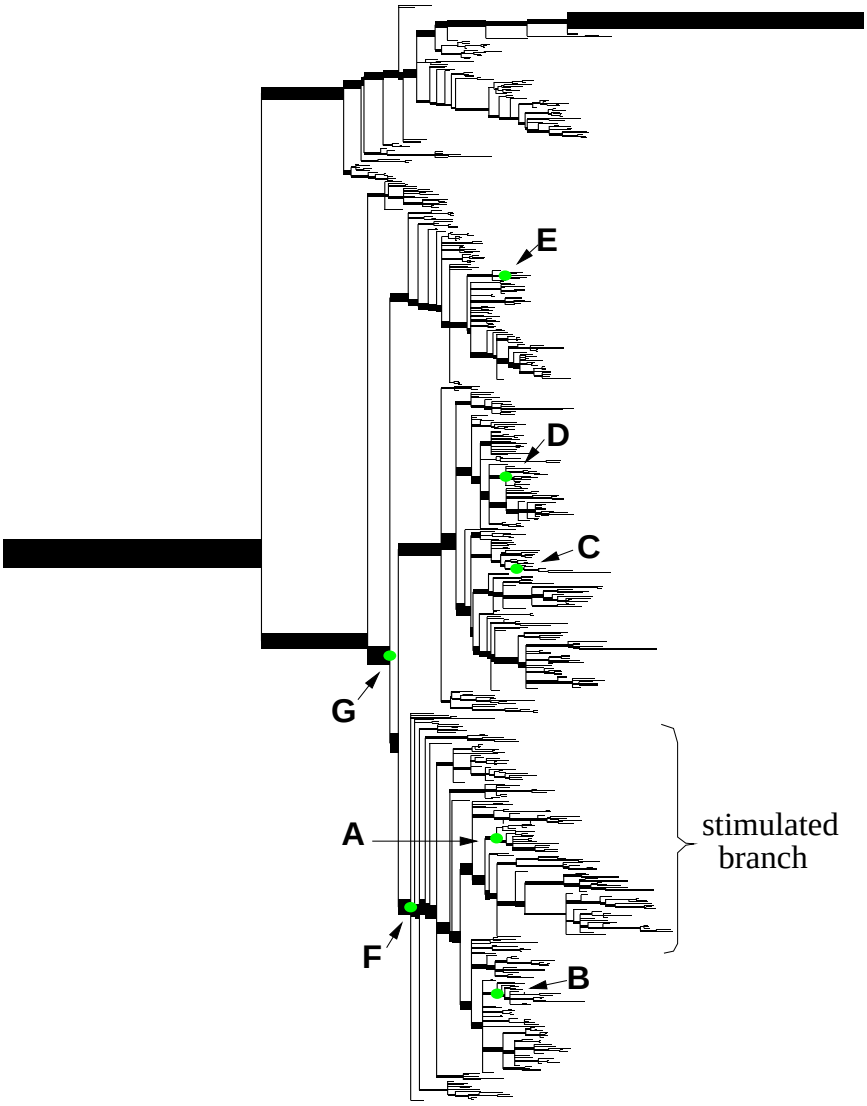

Figure 7.15: A simple simulation to test how the PSPs generated in one branch spread to other branches. Seven recording points record the PSPs at locations on the LG tree. Point A is at a synapse position in the stimulated branch. Points B, C, D and E are at selected synapse points on the other branches. Points F and G are at lower order branches from which A is descended.

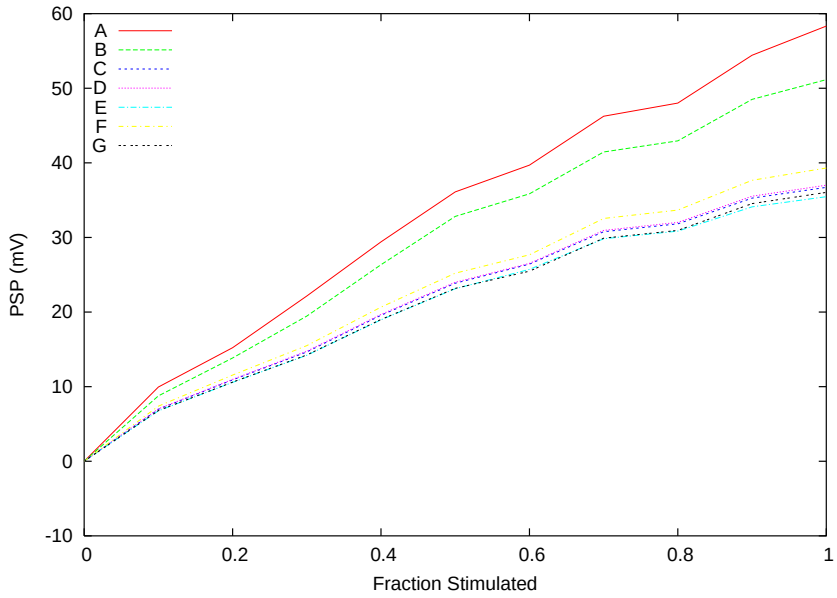

Figure 7.16: The α-PSPs at all the recording points as a function of what fraction of the afferents on the stimulated branch were stimulated.

the attenuation at B is so much less than for the other branches, it is also plotted separately. One can see that for all the branches excluding B, the PSP attenuates to between 0.6 to 0.7 of the PSP at A. This means that a significant fraction of PSP generated at a stimulated branch makes it into other branches. The amount of attenuation does not appear to change significantly as the stimulation at A increases. For point B, the attenuation is to approximately 0.9 of the PSP. Hence, in this particular simulation, the adjacent branch is depolarised almost as much as the stimulated branch itself.

These simulations show that if one branch of the LG is depolarised, a significant fraction of the PSP spreads to the other branches. Indeed, adjacent branches may be excited almost as much as the stimulated branch. It was shown (see section 7.2.4) that the antidromic currents through the lateral giant play a significant role in enhancing the recruitment of non-stimulated

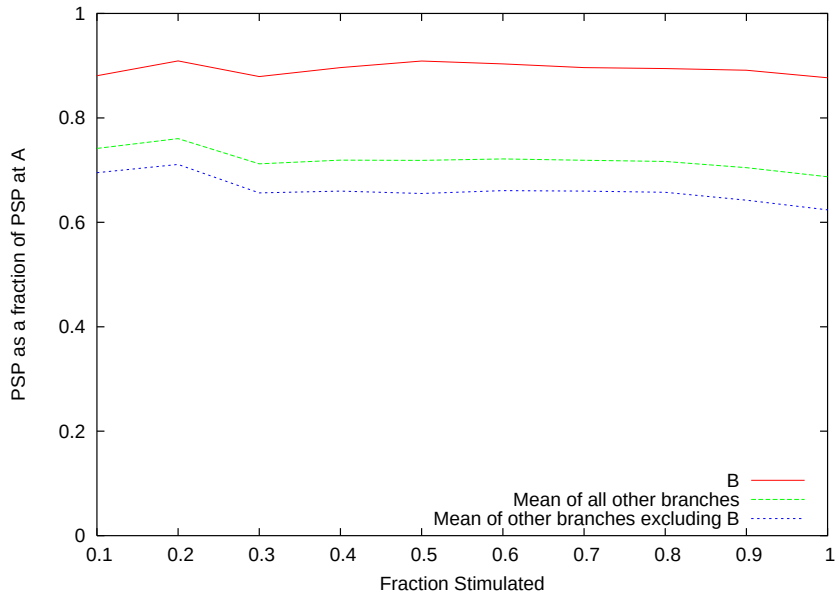

Figure 7.17: The PSPs recorded at the other branches as a fraction of the PSP recorded at the stimulated branch. The top line plots the PSP at B, which is in the branch closest to the stimulated branch. The middle line plots the mean of all the recording sites of the other branches. The lower line excludes point B.

afferents. Given that it appears that most of the PSP generated at the stimulated branch spreads to the other branches, one can expect afferents in the non-stimulated nerves to receive antidromic currents as well. These currents would assist recruitment in the other nerves. These antidromic currents in themselves would not be enough to cause recruitment in the other nerves. However if the nerve is partially stimulated, *i.e.* some of its afferents are spiked, these antidromic currents may assist in bringing the nerve over the 'trigger threshold'.

These results are suggestive of nerves assisting each other in recruitment. Nerves may be communicating with each other through the lateral giant. For example, if one nerve is stimulated sufficiently for all its afferents to be recruited, and the other nerves are partially stimulated but not to the trigger threshold necessary for recruitment, the fully stimulated nerve would cause antidromic currents to spread into the unstimulated nerves which may bring them up to the trigger threshold.

7.3 The afferent networks response to phasic input

In section 7.2 we tested the afferent network's response to synchronous input. We now address the question of how the afferent network responds to asynchronously fired afferents. The questions we consider are:

- Does the network's performance decline as the input becomes more asynchronous?
- Does the afferent network assist in coincidence detection?

To address these questions, an afferent network was set up with midpoint parameters ($g_{\text{aff}} = 0.1\mu$S, LG $g_{\max} = 0.2\mu$S, $\frac{g_{\max}}{g_{\min}} = 100$, 6 afferent-afferent connections, see section 7.2.) The performance of the network was measured

Table 7.4: The values of the input parameters used to test the afferent network's response to phasic input.

Parameter	Values tested	Number of points
Stimulated afferents	0, 1, 2, ... , 50	51
Spread of firing times (ms)	0, 0.025, 0.05, 0.1, 0.2, 0.4, 0.8	7

in response to input that could be characterised by two parameters: the number of afferents stimulated and the synchronisation of the input. The specific afferents that were stimulated were picked at random. The timing of afferents fired was spread about a central point. The amount by which the afferent timings were spread out was varied over different simulation runs. This was done by distributing the firing times about a central point and varying the standard deviation of the spread. Table 7.4 shows the values used to test these parameters. Each input parameter combination was tested 100 times. For the results the mean and standard deviation were calculated.

Figure 7.18 shows the how the gain curves change as the afferents are stimulated with different degrees of synchronisation. The input can be made 0.05 ms asynchronous and not have any statistically significant effect on recruitment, as compared with perfectly synchronised input. There is a slight decline when one makes the input 0.1 milliseconds asynchronous. There is no statistically significant difference in recruitment until the asynchrony is made 0.1 milliseconds or greater. (A Student t-test was used to compare the recruitment for 0, 0.025 and 0.05 ms. The greatest statistical distance any two points were apart was $t = 1.48$.) When the asynchrony is increased to 0.4 milliseconds the peak recruitment gain is less than half the synchronised maximum recruitment gain, and the peak is reached later. When the input is desynchronised by 0.8 milliseconds, very little recruitment occurs.

Figure 7.19 shows how the PSP in the dendrites declines as the stimulus becomes desynchronised. These PSPs are part of the α-component of the LG depolarisation. The PSPs in the LG decline as the input is desynchronised.

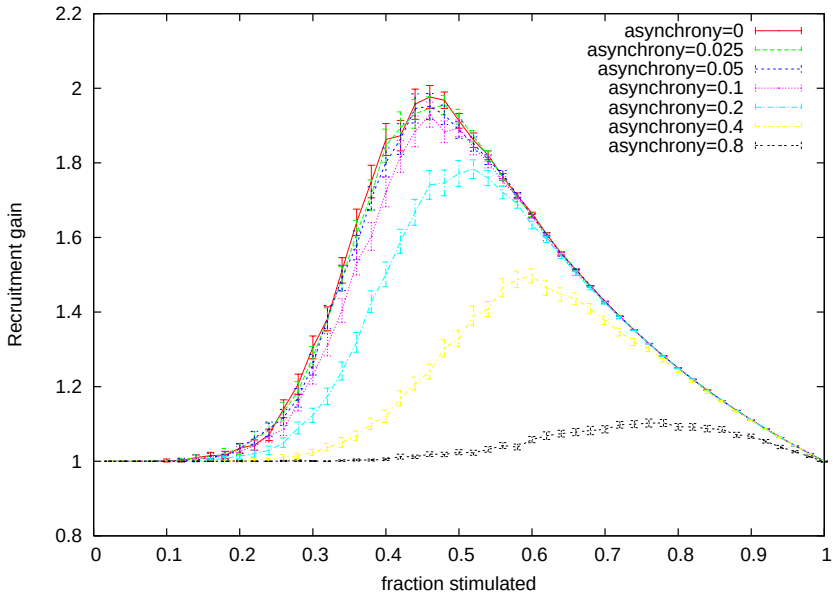

Figure 7.18: The gain curves of the afferent network for input of different synchronisation. Each point is the mean of 100 simulation runs. The error bars show the 95% confidence interval.

The decline is bell shaped. The decline is steepest between 0.05 milliseconds and 0.4 milliseconds. When the stimulus is spread out with a standard deviation of 0.8 milliseconds, the PSP in the dendrites declines by as much as 47% of the synchronised response.

These simulations show that the model afferent network responds better to synchronised stimuli than desynchronised stimuli. Recruitment is higher and the α-PSPs in the dendrites are higher when the afferents are fired simultaneously. The input needs to be spread out with a standard deviation of about 0.4 milliseconds to observe a major declines in recruitment and PSPs. Whilst the afferent network favours phasic input, the decline is not sharp enough to explain the coincidence detection observed experimentally. In experiments (Edwards et al., 1998), input separated by as little as 0.1 milliseconds was sufficient to cause a major decline in the recorded PSPs. Our model of the afferent network does not explain that phenomenon.

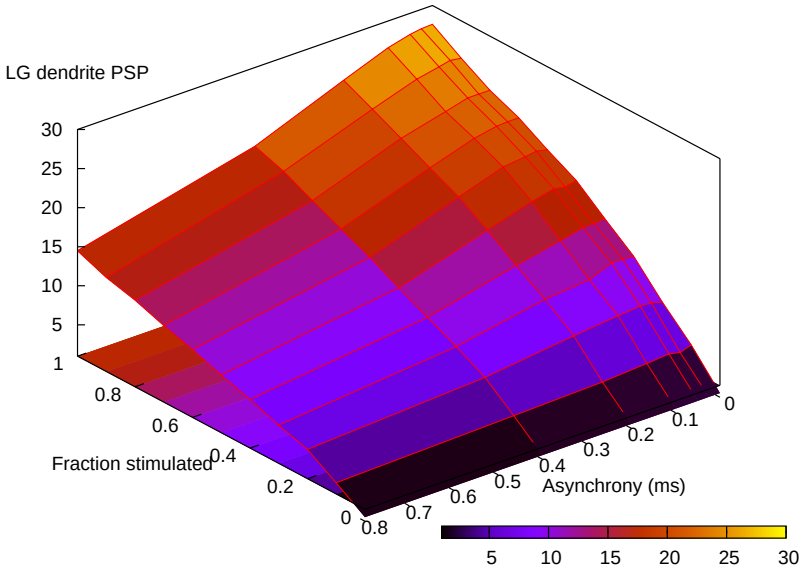

Figure 7.19: The PSP in the LG dendrites as a function of the fraction of afferents stimulated and input synchronisation. Each point is the mean of 100 simulation runs.

7.4 Connection patterns and stimuli patterns

The response of the afferent network is not just a function of the number of afferents stimulated but also which specific afferents are stimulated. How the network responds to a specific set of stimulated afferents is determined by the organisation of the afferent connections.

To demonstrate this, imagine a group of fifty afferents arranged in a hexagonal grid. Consider two alternative connection organisations. In the first, the connections are highly ordered and regular, afferents form one connection to each of their adjacent neighbours in the hexagonal grid. In the second connection organisation, the connections are randomised. That is, we have the same number of total connections but an afferent is just as likely to be connected to an afferent on the other side of the grid as it is to connect to its neighbours.

Now we excite each network with two different stimulus patterns. For each stimulus we excite 24 afferents. For the first stimulus pattern, the excited afferents are picked at random. For the second pattern, we choose the excited afferents to be clustered on one side of the network.

What is the response of the two networks to these two different stimulus patterns? Figure 7.20 shows the results. The regularly organised network responds very well to the random input, all afferents are recruited. However it responds very poorly to the clustered input, only two afferents are recruited. The randomly organised network responds reasonably well to both input patterns. The random input recruits 22 afferents. The clustered input recruits 17.

This rather contrived example demonstrates that the connection regime of the network can have a big effect on how it responds to different input stimuli.

7.4.1 Network connections

The connections between the afferents facilitate spiking afferents to recruit previously non-excited afferents. The recruitment of additional afferents has the effect of amplifying the input given to the lateral giant and the interneurons. The pattern by which the afferents are connected together determine the nature of the amplification and the input patterns which will cause this amplification effect to be realised.

One task for the model is to examine the effect that the connection pattern between neurons exacts on the behaviour of the network. The exact nature of the connection pattern between the afferents is still the subject of research and debate. The findings of (Antonsen and Edwards, 2003) and (Herberholz et al., 2002) indicate that the connection pattern is nonuniform. Afferents

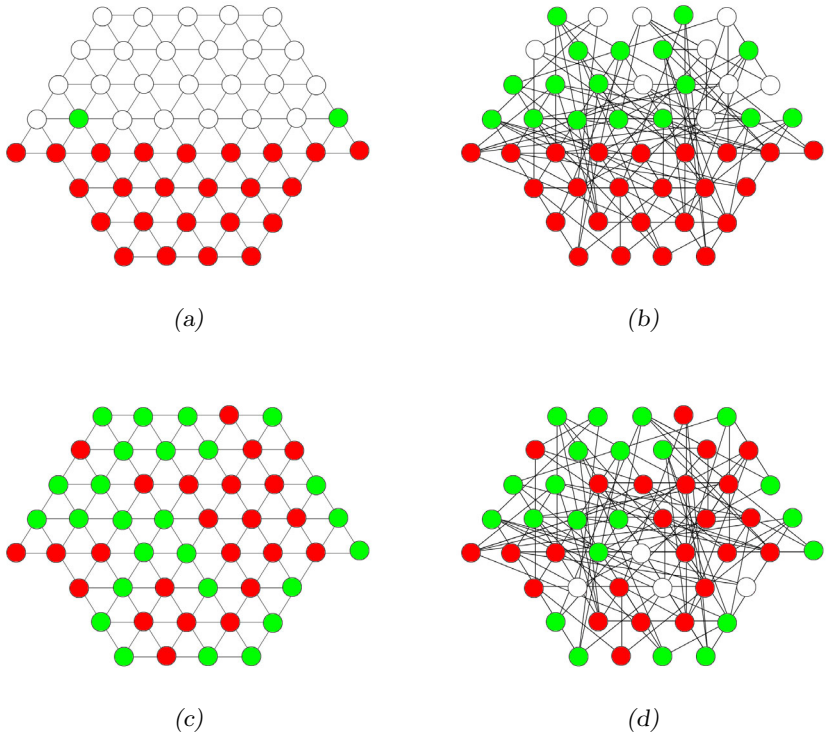

Figure 7.20: The response of two different networks with differently organised connections to two different stimulus patterns. Red circles represent the initially stimulated afferents. Green circles represent afferents that are recruited. (a) A regularly connected network's response to random to clustered stimuli, (b) a randomly organised network's response to clustered stimuli, (c) a regularly organised network's response to random stimuli, and (d) a randomly organised network's response to random stimuli. (Parameter settings used: $g_{\text{aff}} = 0.2\mu S$, *RJ* $g_{\max} = 0.2\mu S$, $\frac{g_{\max}}{g_{\min}} = 100$

are connected to between 2 and 14 other afferents. Aside from the number of connections, there is also the question of whether there is any grouping of afferents. There is some evidence that afferents are more likely to connect to another afferent close by than one far away. But is there a more distinct grouping order than this? One hypothesis is that the afferents are grouped into 'bundles', and each afferent within the bundle is connected to all the other afferents in the bundle. However it has been shown that, in some cases at least, if afferent A is connected with afferent B, and afferent B is connected with afferent C, then afferent A does not necessary make a direct connection with afferent C (Antonsen and Edwards, 2003). This shows that if there is bundling it is not perfect, but does not exclude the possibility of imperfect bundling.

We have already demonstrated in section 7.4 that extreme differences in the stimulus pattern can give extremely different results for different connection patterns. It is possible to design the network connections to be specifically biased for certain combinations of input. Without direct evidence of such a deliberate organisation let us examine the response of the afferent network to stimulus patterns of randomly selected afferents. We have seen that differently selected stimuli can have extremely different responses. It would be useful to measure the variance of the response of the network to randomly distributed stimulus patterns. Three alternative network connection regimes are considered.

The first connection pattern simulated is a uniform one, where afferents each make a standard number of connections to other afferents, *i.e.* afferents each have the same number of connections. The afferents are placed in a hexagonal grid, and attempt to connect to a fixed number of neighbouring afferents. Afferents placed on the edge of the grid may not be able to make the full number of connections, but all afferents not on the edge will

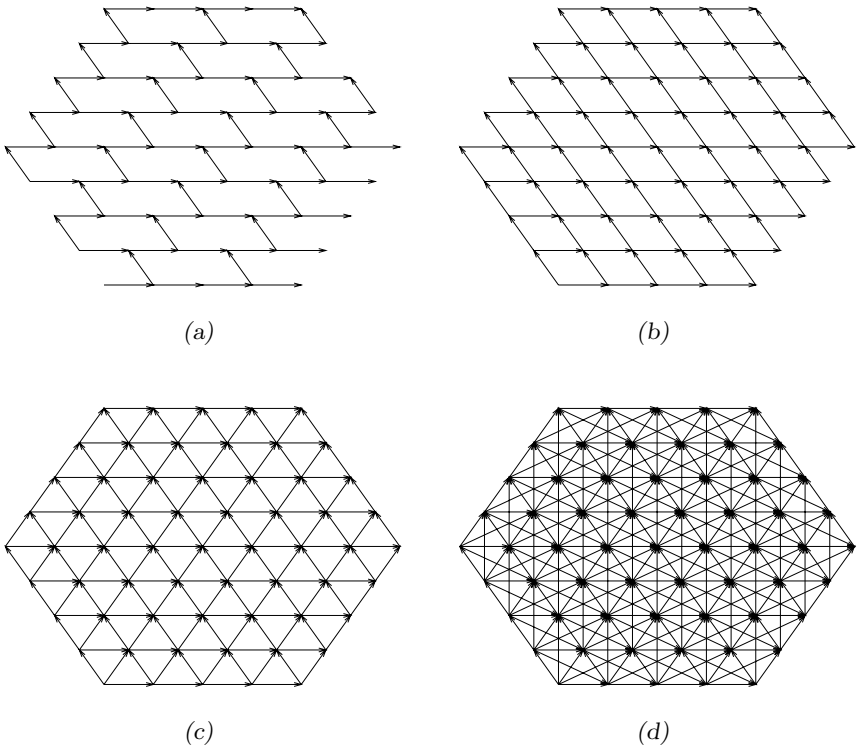

Figure 7.21: Some uniform connection patterns used for the afferent network. For each pattern, each afferent attempts to make a fixed number of connections to adjacent afferents. Shown are the patterns for (a) 3, (b) 4, (c) 6 and (d) 12 connections.

make the same number of connections. Variations were made on how many connections each afferent makes. Different patterns were set up for between 3 and 12 connections. Figure 7.21 shows the connection patterns for some of these variations.

An alternative configuration of the connections of the network where the connections are chosen probabilistically, with connections more likely between afferents that are closer. Such a connection scheme is biologically plausible. We choose an exponential function to represent the probability

of ith and jth afferents forming a connection:

$$p(i,j) = p_0 \exp(\frac{r_{ij}}{L}) \tag{7.11}$$

Where r_{ij} is the distance between afferents i and j, p_0 is the maximum probability and L is the length constant. Figure 7.22 shows the connections between afferents for a network formed with this method.

The final connection regime considered is where connections are randomised. Each afferent makes c connections, however the target of each connection is chosen from the entire population of afferents with each afferent having an equal probability as being selected as the target. The network shown in Figures 7.20b and 7.20d is an example of a network organised to with this regime. Biologically, such a connection scheme is unlikely as there is likely to be biases which make some connections more likely than others. We consider this organisation because it is the extreme of randomisation. The biological network is likely to be randomised to some extent, although we do not know to what degree. It is useful to know the behaviour at the extreme.

For each connection organisation we determine the variance of the recruitment gain at various initial stimulus fractions. The same total number of connections were used for all three networks. This is the number of connections in a uniformly connected network with 6 connections per afferent. For each configuration and at each stimulus fraction 100 simulations were run. For each simulation the initially stimulated afferents were picked at random. We then calculate the standard deviation for each configuration.

Figure 7.22: A probabilistic afferent network, where the probability of two afferents forming a connection is inversely proportional to the distance between them.

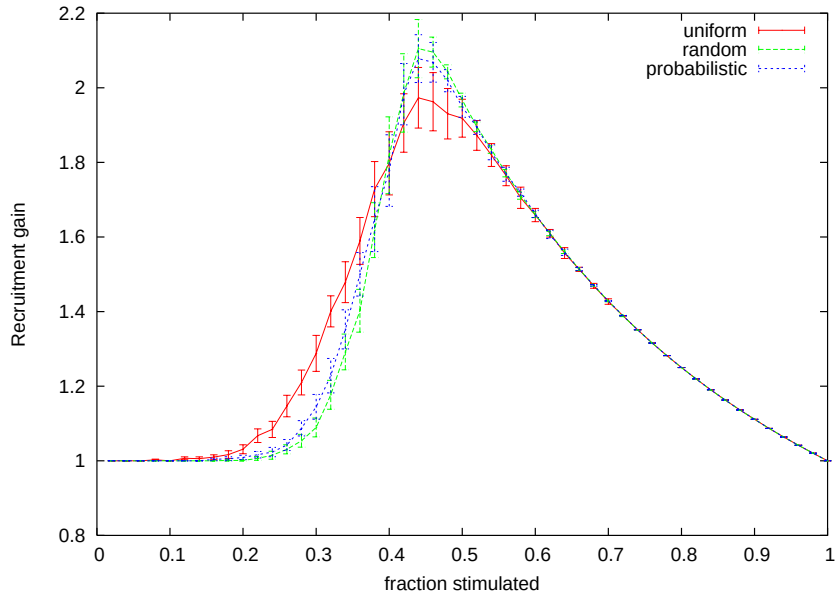

Figure 7.23: The recruitment gain curves for different connection organisations. The standard deviation of the gain is plotted as error bars. The probabilistically and randomly connected networks have a similar response. The uniformly connected network has a less steep transition to maximum recruitment and a larger standard deviation for most fractions stimulated. (Parameters used: $g_{\text{aff}} = 0.1\mu S$, LG rectifying junction $g_{\max} = 0.2\mu S$, $\frac{g_{\max}}{g_{\min}} = 100$

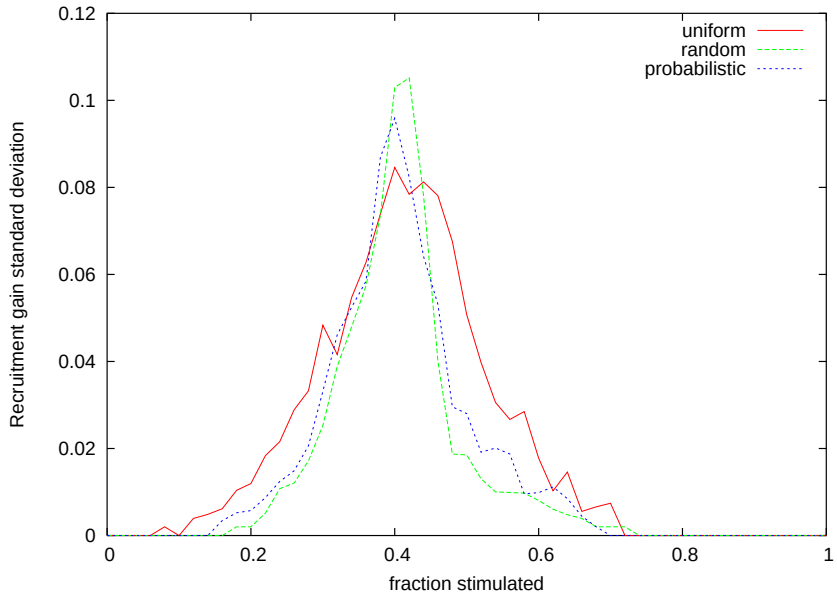

Figure 7.24: The standard deviation of the recruitment gain as a function of the fraction of afferents stimulated for different connection organisations. Variation is the highest when recruitment gain is at a maximum. The uniformly connected network has higher standard deviations for most fractions stimulated.

Figure 7.23 compares the recruitment gain curves for the different connection organisations. The probabilistic and random organisations are of a similar shape. The uniformly connected network makes a less steep transition to maximum recruitment. Figure 7.24 compares the standard deviations of the recruitment gains for the differently connected networks. For most fractions stimulated the uniform network has the highest variance. However the random and probabilistic networks have the highest peak standard deviation. The peak variance for all networks is near to where they approach their respective maximum gains. At this point the standard deviation is of a similar magnitude for all three networks.

Overall the standard deviation is not very great. In absolute terms the standard deviation never gets much above 0.1. This maximum is reached as recruitment approaches its peak. At this point the standard deviation is equal to about 10% of the recruited afferents. This indicates that random variation is not expected to have a great effect on the recruitment gain at the critical threshold where recruitment takes hold. That is, recruitment is reliable at the critical threshold.

Figure 7.25 plots the standard deviation of recruitment as a fraction of the total recruited afferents. The standard deviation in relative terms is higher at sub-threshold stimulation levels. This shows that sub-threshold recruitment is less reliable. Given that during this region the number of afferents recruited is low and thus does not add much to the gain of the circuit, this is not a very important result.

As can been seen in figure 7.24 the variance appears to be greatest for the uniform network, at least for most fractions stimulated. This reinforces the anecdotal result of section 7.4. Uniformly connected networks are more vulnerable to random variation than randomly connected networks.

The performances of the probabilistically and randomly connected networks

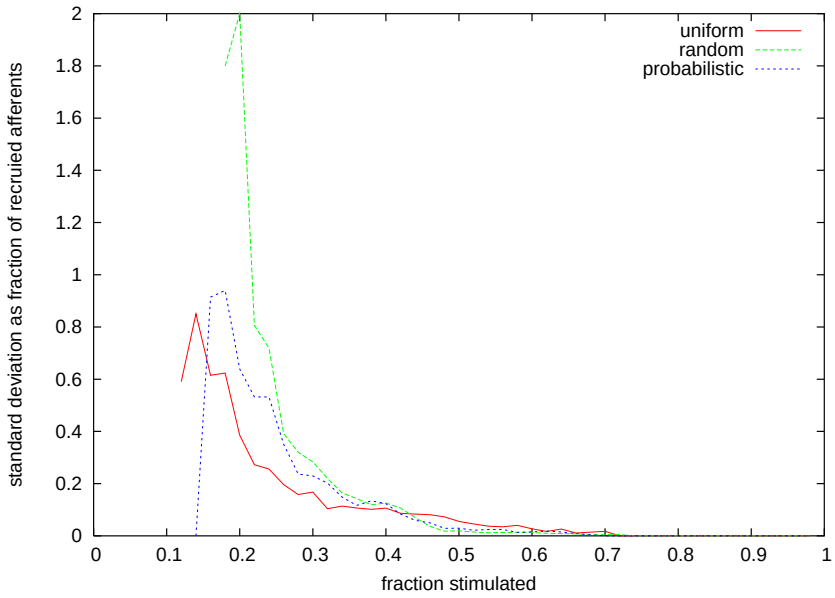

Figure 7.25: The standard deviation of recruited afferents as a fraction of the afferents recruited. The standard deviation is high during sub-threshold recruitment but low at and post threshold.

are very similar. This result complements the findings of (Watts and Strogatz, 1998) which show that adding a small amount of randomness to the linking of a network greatly increases the network's connectivity.

Adding an element of randomisation steepens the recruitment transition. This is a biologically desirable property as it sharpens the threshold of when recruitment starts and reduces the risk of an accidental sub-threshold escape.

7.5 Role of the interneurons

The beta component of the excitation to the LG (see chapter 3) comes from the interneurons. If the LG is to spike (and thus the escape response is initiated) it is the beta component that takes it over threshold (except for in very small animals.) The primary afferents excite the interneurons as well as the lateral giant directly. Thus the interneurons play a vital role in the

circuit.

We present examining of the nature of role of the interneurons as a topic for further work.

The question to examine is how might the β response caused by the afferents interact with the α caused by the direct afferent-LG connections. We would like our model LG to respond to input with a nice α-β shaped PSP, as is observed in the biological system (see figure 3.12.) It is critical in order for LG to spike for the α and β inputs to sum effectively. The β input needs to arrive onto the LG arrives before the excitation from the α input has subsided. The β input's path is much longer than the α input as it needs to pass through a chemical synapse and the interneuron. What mechanisms are necessary to keep the time delay caused by the longer path to a minimum presents itself as an open question.

We propose to model interneuron in 'cartoonish' detail. This is to reduce the complexity of the model. There are different kinds of interneurons and they (probably) each integrate their input in different ways. This obviously introduces a lot of extra complexity, which is ignored in our model. In any case there is insufficient detail available to accurately model all the different kinds of interneurons. In our model all interneurons are identical. The morphology of the interneurons modelled at a crude level, to be approximately similar to the sizes of interneuron A and interneuron C. The membrane properties used are similar to the afferents. There are chemical synapses between the afferents and the interneurons. The model has rectifying electrical synapses between the interneurons and the LG.

The interneurons add an extra layer of processing to the lateral giant circuit. The interneurons receive their input from the afferents and pass on the processed signal on to the lateral giant. In these simulations the interneurons use a very simple method to integrate the afferent input. It should be noted

that the processing done by the interneurons can be made very complex.

7.6 Discussion

Until recently it was assumed that the decision to escape or not to escape was made by the lateral giants and the interneurons. The afferents were assumed to be messengers conveying the stimuli. They were not thought to have any part in the decision making process.

The lateral giant circuit makes a critical (literally life or death) decision as to whether or not to escape. The circuit needs to make the decision very quickly and also correctly. The simulations in this chapter show that the lateral excitatory network is capable of being discriminate in selecting specific combinations of stimuli.

The simulations in this chapter show that the lateral excitatory network may play a dual role in the operation of the LG circuit. One function is to act as an amplifier with a nonlinear gain curve. This role is identified in previous work (Herberholz et al., 2002). Stimulated afferents recruit non-stimulated afferents thus increasing the amount of input that the lateral giant and the interneurons receive. The direct connections onto the lateral giant enhances the α-response. However in simulations the magnitude of the α-PSP in LG is not increased by very much by the recruited afferents. This is due to the recruited spikes being slightly out of phase with the spikes of directly stimulated afferents combined with the LG's selectivity for phasic input. The recruited spikes do however sustain the α-response which may cause it to sum more effectively with the β. Furthermore, the recruitment means that the interneurons receive more input. This is likely to cause an increased number of interneurons to be spiked. Firing extra interneurons increases the β-PSP in LG. In larger animals, the β is the largest component of the PSPs

in LG and the LG fires it fires off the β. The extra interneurons fired as a consequence of recruitment in the the lateral excitatory network may bring the lateral giant to firing threshold.

The afferent network does not just recruit extra afferents but the nature of this recruitment is nonlinear. In simulations we have seen that when one stimulates a fraction of the afferents up to a threshold there is very little recruitment. However when the threshold is exceeded recruitment makes a very rapid transition from minimal recruitment to maximum recruitment. There is a 'hair trigger' for turning recruitment on or off. This property of the afferent network agrees with the experimental observations (Herberholz et al., 2002). A steep transition from minimal to maximum recruitment is a robust property of the model afferent network. This property was present for a large fraction, indeed the majority, of the simulated networks within the parameter space explored. The architecture of having a group of afferents connected to another appears well suited to providing a 'hair trigger' amplification mechanism. This provides a desirable property in the escape circuit for the animal: when afferent activity passes the threshold of a predator attack, an escape response should be initiated. The afferent network sharpens the threshold and reduces the risk of sub-threshold activity accidentally triggering an escape response.

There are many other possible mechanisms for achieving amplification of the input. For example, the LG could be given a lower threshold or the synapses exciting the lateral giant could be made stronger. The afferent network appears to be a more complicated way of achieving amplification compared to these other mechanisms. However, these alternatives would give a more linear amplification the input. They would not give the sharp threshold 'all-or-nothing' amplification that the afferent network provides.

Our simulations predict that the lateral excitatory network causes the es-

cape circuit to respond better to some specific combinations of input. The response depends not only the number of afferents that are initially stimulated but also on the specific afferents that are stimulated. Which specific combinations of afferents firing give more or less recruitment depends on how the network is connected.

The escape system needs to detect when a predator is attacking it. The crayfish is receiving continuous sensory input. If needs to discern a predator attack from other sensory stimuli. The classical way this was assumed to happen is simply that a predator attack sets off more afferents firing simultaneously than any other sensory stimuli. Perhaps this is all the escape circuit detects. The afferent network has the capability to do more than this. It is capable of being wired to respond better to specific combinations of input. For example, perhaps an approaching predator sets off different combinations of firing afferents than a sudden surge in the current. And when one considers that the afferent network output still needs to be processed through the interneurons before reaching the lateral giant, then we have a system with potentially very rich pattern discrimination capabilities. The lateral giant escape system may be doing more than just counting the number of simultaneously firing afferents.

The architecture of the LG escape circuit resembles a two layered recurrent artificial neural network. The primary afferents form the first layer and the interneurons are the second layer. Two layered recurrent neural networks have been demonstrated to be very powerful computational tools (D. and J., 1986a; D. and J., 1986b; Wasserman, 1989). Neural networks can be trained so that the network will respond positively to certain inputs and ignore others. Artificial neural networks are very powerful at pattern recognition. The biological escape circuit is much more complex than an artificial neural network. We have shown that depending on how the network is wired it will

respond differently to different inputs. Whether the wiring is coincidental or deliberate is unknown. We hypothesise that the lateral excitatory network is deliberately wired to respond better to certain stimuli. This hypothesis is difficult to directly test experimentally. There is however some indirect experimental evidence which is suggestive of this. For example, the lateral giant response is sensitive to direction (Wine and Krasne, 1982), the afferent network may be the mechanism that causes this. Assuming that the bias of network is deliberate, how such wiring occurs is unknown. It may be 'hard-wired', or it may be learned. When the crayfish is young the threshold for escape is much lower and it escapes very frequently. The crayfish may learn which triggers are false alarms and wire the afferent network to ignore such stimuli.

The real afferent network is much more complex than what we have modelled. In our simple model we have assumed that all afferents and all afferent connections are identical. This is of course false. There is likely to be much diversity between the afferents. This adds an extra level to the discrimination capabilities of the network. The afferent connections would effectively be weighted. Even with our simple model we have demonstrated that the afferent network is capable of discriminating specific input. With the added complexity of the real system, we have the necessary neural machinery for some complex decision making.

In our model we have made a lot of assumptions and simplifications. We do not have a realistic model. Our model explains some of the biologically observed phenomena of the lateral excitatory network and it makes some predictions. However, the effects of the complexities we have ignored is unknown.

The simulations in this chapter only only a starting point for a complete and well defined computational model of the lateral excitatory network. Some

of the questions that present themselves for further work include:

- The relationships between the parameters explored and the response of the network have been qualitatively described. An extension would be to model these relationships with mathematical equations. The final goal would be to have a model equation which describes recruitment as a function of the model parameters.

- In simulations we have shown that the afferent connection pattern plays an important role in the response of the circuit. It would be useful to simulate the biologically plausible afferent connection patterns and characterise their response to different inputs.

- The afferents do not exist exclusively for the lateral giant. The afferents convey sensory information to many other neural systems in the crayfish. It is not possible for the crayfish to distinguish whether an afferent was recruited or whether it was innervated by a sensory stimulus. How the recruited afferents affect other systems is an open question. Investigating the effect recruitment has on other systems is a topic for further research.

CHAPTER 8

Conclusion

Animal behaviour is robust and adaptive. Properties that have so far been largely elusive to artificial systems. This thesis concerns how you might incorporate robustness into artificial intelligence. We argue that studying the mechanisms of animal behaviour is instructive for building intelligent artificial systems.

An example of robust behaviour is the crayfish escape response. Almost all animals possess some form of startle response. This is compelling evidence that reflex behaviours are an integral part of animal adaptation. For robots one can imagine that reflex behaviours are also important. For example, a mine clearing robot needs to be focused on implementing its general strategy of clearing the mine field but needs to be able to respond quickly when unexpectedly faced with a mine. A robot soccer player needs to execute the team's strategy of setting up a goal but also needs to rapidly evade when tackled by an opposing player.

In our artificial world simulation, we have shown that adding escape be-

haviours to a system performing a simple task in a hostile environment not only increases the survivability of the agent but also impacts the strategy of how the agent undertakes its main activity. An advanced escape system allows the agent to ignore the threats of the environment while it is undertaking its primary task. If the escape system is weak, then the agent must incorporate a strategy for avoiding danger into its approach to the primary task. An advanced escape system decouples the problems of an agent achieving its main goals from the problem of dealing with the danger in the environment. Adding reflex behaviours to artificial systems, such as robots, will be an important component of a robust architecture.

The crayfish escape circuit is an interesting case study of decision making in animals. The circuit serves as a concrete example of an animal behaviour where the controlling neural circuitry is mapped out and well understood. One expects that some of the mechanisms that drive decision making in the escape circuit generalise to other instances of decision making.

In order to fully understand how a biological system functions one needs a computational model. The understanding neuroscientists have of neural circuits is often qualitative. Some neuroscience models are largely descriptive. To really understand how biological neural networks perform their cognition, quantitative models are necessary. The field of computational neuroscience is concerned with making quantitative models of neuroscience descriptions. In the past computational neuroscience has mostly been of interest to neuroscientists. We argue that it is also of interest to computer science. If one were to construct a computational model of the crayfish escape circuit that replicates the circuit's functions, this would demonstrate an understanding of the underlying decision making.

Biological systems are of immense complexity. They are much more complex than computer programs. A useful model would produce the interesting

characteristics of biological systems with the simplest possible design. The problem is it is difficult to distinguish *a priori* what parts of the complexity in biological systems can be ignored from the parts that are essential to their adaptive properties. If one models superficially one can easily leave out the essential features of the system.

The goal of this thesis was to make a complete model of the crayfish lateral giant escape circuit to demonstrate that its robust qualities can be reproduced on a computer and to identify the specific properties that give rise to robustness.

We have succeeded in creating a computational model of a subset of the sensory part of the lateral giant escape circuit, comprising of the lateral giant and a network of primary afferents. Our model is faithful to the known biological mechanisms of the circuit. The neurons have been modelled as R-C electrical circuits with electrical properties that correspond to biological measurements.

In this thesis we have demonstrated that we can reproduce some of the crayfish escape circuit's interesting properties in computer models. Some of the key components in the robust decision making of the lateral giant escape circuit have been identified.

Coincidence detection is a key function of the lateral giant escape circuit. It is used to detect abrupt massive stimuli which are the signature of a predator attack. We have shown that coincidence detection can be effectively simulated using rectifying junctions between the afferents and the lateral giant. This work complements the previous studies of (Edwards et al., 1999). The rectifying junctions are have nonlinear electrical conductivity which is fundamental to coincidence detection. The rectifying junctions are in a state of high conductivity in response to simultaneous inputs, causing them to sum effectively. In contrast, asynchronous inputs are ineffective because the

early input partially excites the lateral giant causing reduced conductivity across the rectifying junctions. This renders late inputs to be largely ineffectual. The coincidence detection properties of the rectifying junctions are dependent on their parameter settings. For effective coincidence detection the rectifying junctions need to have fast closing latency, a high conductance at rest and a low conductance under reverse bias. Given other properties the rectifying junctions perform a different function.

The direct interconnections between the afferents have also been shown to play a fundamental role in the escape circuit's operation. Stimulated afferents recruit non-stimulated afferents thereby amplifying the sensory input (Herberholz et al., 2002). Our computer simulations show that networking the primary afferents sharpens the escape threshold and lessens the risk of sub-threshold activity accidentally triggering an unwanted escape response. Furthermore we predict that the afferent network adds the ability of the escape circuit to not only detect temporal patterns of activation but also spatial patterns.

The work undertaken in this thesis uncovered a number of issues. Our initial approach to studying the crayfish escape circuit was through a high level model. This approach however was too far removed from the biology and had no interesting underlying complexities. This directed us to taking a more detailed consideration of the underlying biology.

The biological details of the circuit do give rise to many interesting complexities. The complexity is indeed overwhelming. Biological systems are very detailed. Furthermore, the deeper one looks the more details that are uncovered. At each level of detail there are complicated behaviours that cannot be completely ignored at a higher level. For example, one may model a neuron as a single compartment, but this ignores that neurons are complex structures made up of dendrites and axons. One may include the

dendritic structure and use uniform membrane properties, but this ignores that the membrane contains many different kinds of ion channels which are non-uniformly distributed. It is a challenge to decide what details should be ignored and which details should be included. Furthermore, the details of the biological mechanisms are only partially discovered. The modeller must make assumptions and fill in gaps. It is possible that a model mimic a biological behaviour in a way that is unfaithful to reality. Arguably however, if the model satisfactorily replicates the behaviour this is ultimately unimportant.

Human beings like to isolate the properties of the individual components in the design. The problem of taking this approach to understanding biology is that nothing exists in isolation. There are so many interconnections and everything influences everything else. As a simple example from the lateral giant escape system, creating an isolated model of a primary afferent ignores that the antidromic feedback it receives from the lateral giant. Studying one component in isolation ignores that it behaves differently when it is part of the system.

What is clear is that the primitives used in biological systems are fundamentally different to those in computer systems. In electronic systems the fundamental building block is the transistor, or the NAND gate at a slightly higher level. In making models of biological systems it is unclear what the basic building block should be. It has often assumed to be the neuron but, as demonstrated by the work in this thesis, an individual neuron is an immensely complicated cell. Furthermore the building blocks in biology are novel compared to the components that make up computer systems. Consider the rectifying junctions which have been shown to play a pivotal role in the function of the crayfish escape circuit. These are novel devices which perform complex nonlinear calculations. Establishing what the necessary

basic primitives are for creating biologically accurate models is an important area that needs to be resolved.

In looking at a biological system in too much detail there is a possibility in getting trapped into looking at irrelevant properties of the system. There may be more than one way of achieving the same outcome. The biology works to achieve a satisfactory solution out of what it has. Its solution may not be the most optimal. Biology might have taken a different solution if other components were available to it. As shown by (Beer et al., 1999) it is possible to achieve similar outcomes with vastly different connection organisations between neurons.

It is interesting to note that the escape circuit is also used in social interaction in the establishment of dominance hierarchies (Issa et al., 1999; Herberholz et al., 2001). This illustrates how biology re-uses the components it has, which is instructive in itself. In artificial systems we do not know how to effectively overload the same behaviour with different purposes.

There exists a wide gulf between our biological simulations and the high level simulation described in chapter 2. The ultimate goal is to obtain an understanding of what exactly gives biological systems their robust properties. With such knowledge we could create high level artificial systems with biological like robustness.

The fundamental question of what biological systems possess that artificial ones do not remains unanswered. The models presented in this thesis are in the direction of the kind of work that we believe is necessary to uncover an understanding of how biological systems function.

The crayfish escape circuit serves as a useful domain for simulations whose aim is to unlock the secrets of how nervous systems function.

APPENDIX A

Constants used in simulated environment simulation

Constant	Value	Equations used
A	0.001	(2.1)
B	0.03	(2.1)
A_P	1	(2.2)
L	40	(2.2)
A_M	1	(2.3)
τ	15	(2.3)
T_P	0.1	Figure 2.1b
V_{pred}	12	Section 2.2
R_{pred}	120	Section 2.2
V_{prey}	10	Section 2.2
R_{prey}	50	Section 2.2

APPENDIX B

Constants used in compartmental model simultations

B.1 Abstracted morphology LG coincidence detection constants

Constant	Value
Temperature: T	20°C

B.1.1 LG constants

Constant	Symbol	Value
Axial resistivity	r_a	60 Ωcm
Membrane resistivity	r_m	10000 Ωcm^2
Membrane capicitance	c_m	1 $\mu\text{F.cm}^{-2}$
Hodgkin-Huxley channel constants		
Leak conductance	g_L	0.001 $S\text{cm}^{-2}$
Leak reversal potential	E_L	-80 mV
Maximum potassium conductance	$\bar{g}_\text{K}$	0.036 $S\text{cm}^{-2}$
Sodium reversal potential	E_Na	45 mV
Maximum sodium conductance	$\bar{g}_\text{Na}$	0.120 $S\text{cm}^{-2}$
Potassium reversal potential	E_K	-62 mV

B.1.2 Afferent constants

Constant	Symbol	Value
Axial resistivity	r_a	60 Ωcm
Membrane resistivity	r_m	3000 Ωcm^2
Membrane capicitance	c_m	1 $\mu\text{F.cm}^{-2}$
Hodgkin-Huxley channel constants		
Leak conductance	g_L	0.001 $S\text{cm}^{-2}$
Leak reversal potential	E_L	-70 mV
Maximum potassium conductance	$\bar{g}_\text{K}$	0.036 $S\text{cm}^{-2}$
Sodium reversal potential	E_Na	45 mV
Maximum sodium conductance	$\bar{g}_\text{Na}$	0.120 $S\text{cm}^{-2}$
Potassium reversal potential	E_K	-62 mV

B.2 Accurate morphology LG coincidence detection constants

Constant	Value
Temperature: T	20°C

B.2.1 LG and afferent constants

Constant	Symbol	Value
Axial resistivity	r_a	60 Ωcm
Membrane resistivity	r_m	6000 Ωcm^2
Membrane capicitance	c_m	1 μF.cm^{-2}
Hodgkin-Huxley channel constants		
Leak conductance	g_L	0.0001 Scm^{-2}
Leak reversal potential	E_L	-75 mV
Maximum potassium conductance	$\bar{g}_{\mathrm{K}}$	0.036 Scm^{-2}
Sodium reversal potential	E_{Na}	45 mV
Maximum sodium conductance	$\bar{g}_{\mathrm{Na}}$	0.120 Scm^{-2}
Potassium reversal potential	E_{K}	-70 mV

B.3 Afferent network simultation constants

Constant	Value
Temperature: T	20°C

B.3.1 LG and afferent constants

Constant	Symbol	Value
Axial resistivity	r_a	60 Ωcm
Membrane resistivity	r_m	6000 Ωcm^2
Membrane capicitance	c_m	1 μF.cm^{-2}
Hodgkin-Huxley channel constants		
Leak conductance	g_L	0.0001 Scm^{-2}
Leak reversal potential	E_L	-75 mV
Maximum potassium conductance	$\bar{g}_{\mathrm{K}}$	0.036 Scm^{-2}
Sodium reversal potential	E_{Na}	45 mV
Maximum sodium conductance	$\bar{g}_{\mathrm{Na}}$	0.120 Scm^{-2}
Potassium reversal potential	E_{K}	-70 mV

B.3.2 Rectifying junction constants

The following table shows the default constants used to govern the rectifying junctions between the afferents and the LG (as per equation 3.1.)

Constant	Value
$g_{\max}$	0.33 μS
$g_{\min}$	0.005 μS
V_0	0 mV
A	0.15 mV^{-1}
τ_{open}	0.5 ms
τ_{close}	0.5 ms

APPENDIX C

Parts of the simulation source code

C.1 NEURON rectifying junction objects

The source for post-synaptic rectifying junction object.

```
TITLE Rectifying junction

COMMENT
    Implementation of a rectifying gap junction as a point process.
    The steady state conductance of the synapse is governed by this equation:

                                          G_max - G_min
    g(V_pre, V_post) = G_min + --------------------------------
                               1 + exp(-A(V_pre - V_post - V0))

    where: V_pre and V_post are the pre and post synaptic potentials,
    G_min and G_max represent the minimum and maximum conductances
```

```
    of the synapse and A and V0 are constants.

    The instantaneous conductance of the synapse moves exponentially
    towards this steady state value with the time constants tau_open
    and tau_close.

    rj2 should be used with an gapmir point process.  rj2 is used at the
    post synaptic section, and gapmir is used at the pre-synaptic section.
ENDCOMMENT

UNITS {
    (mV) = (millivolt)
    (nA) = (nanoamp)
    (umho) = (micromho)
}

NEURON {
    POINT_PROCESS rj2
    NONSPECIFIC_CURRENT i
    RANGE i, g, gmin, gmax, A, V0, tau_open, tau_close, g_steady, tau
    POINTER vpre
}

PARAMETER {
    gmin  = 0.002 (umho)
    gmax  = 0.2 (umho)
    A = 0.15 (/mV)
    V0 = 0 (mV)
    tau_open = 0.75 (ms)
```

```
    tau_close = 0.75 (ms)
    v (mV)
    vpre (mV)
}

ASSIGNED {
    i (nA)
    g_steady (umho)
    g_diff (umho)
    tau (ms)
}

STATE {
    g (umho)
}

BREAKPOINT {
    SOLVE states METHOD cnexp
    i  = -(g * (vpre - v))
}

INITIAL {
    calcvars()
    g = g_steady
}

DERIVATIVE states {
    calcvars()
    g' = (g_steady - g) / tau
```

```
}

PROCEDURE calcvars() {
    g_steady = calcginst(vpre - v)
    g_diff = g_steady - g
    if (g_diff > 0.0) {
        tau = tau_open
    } else {
        tau = tau_close
    }
}

FUNCTION calcginst( vd (mV) ) (umho) {
    calcginst = gmin + (gmax - gmin) / (1 + exp(-A * (vd - V0)))
}
```

The source for pre-synaptic rectifying junction object.

```
TITLE gapmir.mod        reverse end of rectifying junction

COMMENT
 This is the dummy end of a rectifying or gap junction.  This model
 contains only a current variable that is set by the other end of the
 gap junction.  Basically the idea is that the current at the
 one end of the gap junction is the negative of the other so you only
 need to calculate it once.  The business end of the gap junction
 assigns the current to this model.
 See rj.mod
ENDCOMMENT
```

```
NEURON {
    POINT_PROCESS gapmir
    NONSPECIFIC_CURRENT i
    POINTER g, vpre
    RANGE i
}
PARAMETER {
    v (millivolt)
}
ASSIGNED {
    i (nanoamp)
    g (micromho)
    vpre (millivolt)
}
BREAKPOINT {
    i = -g * (vpre - v)
}
```

APPENDIX D

Glossary

α-PSP: the PSP that the lateral giant receives that is caused by the direct primary afferent connections.

abdomen: the tail of the crayfish.

acetylcholine (ACh): An amine that is used as a neurotransmitter at many synapses, including neuromuscular junctions. In the LG escape circuit, acetylcholine is used at the primary afferent to sensory interneuron synapses and at the neuromuscular junctions.

action potential: a sharp rise and fall in the neuron's membrane potential. Action potentials are mediated by the opening and closing of sodium and potassium channels (see section 3.1.2.) Action potentials propagate actively along axons and are an important mechanism for passing signals to other cells. Action potentials are also called 'spikes'.

active channel: channels in the membrane that are of variable conductance, usually voltage-dependent variable conductance.

active membrane: a part of the membrane containing active channels. The active channels cause inputs in a cell to be integrated nonlinearly. Active membrane may amplify depolarisations and generate action potentials.

afferent: an axon projecting inward to the contextual neural structure. For example, the axons from mechanosensory neurons, projecting onto the LG, are afferents.

agonistic behaviour: the behaviour of a crayfish confronting another crayfish, including threat displays, defensive posturing and fighting.

animat: an artificial animal, either simulated by a computer or embodied in a robot, which must survive and adapt in progressively more challenging environments (Cliff et al., 1994).

anterior: the direction towards the rostrum (front) of the animal. Equivalent to rostral. (See figure 3.8.)

antidromic: a direction of synaptic transmission where transmission is from the post-synaptic terminal to the pre-synaptic terminal. Antidromic transmission is against the usual flow of synaptic transmission.

axon: a neurite specialised to conduct action potentials away from the neuron's cell body to the target cells of the neuron.

berry: A female carrying eggs on her abdomen, encased by the swimmerets, is said to be 'in berry.'

β-PSP: the PSP in the lateral giant that is caused by the interneurons.

carapace: the hard exterior shell of the crayfish that encases the head and thorax. (See figure 3.7.)

caudal: the direction towards the rear of the animal (down the spinal cord.) Equivalent to posterior. (See figure 3.8.)

cell: in the context of this thesis, a cell means 'neuron'.

chemical synapse: a connection between neurons where the method of transmission is mediated by the release of neurotransmitters. Neurotransmitters are released at the pre-synaptic terminal and cause a response (for example the opening of channels) at the post-synaptic cell.

CNS: Central nervous system. The CNS encapsulates the brain and spinal cord and all connecting nerve cells.

cheliped: a crayfish's claw. (See figure 3.7.)

channel: a pore in the cell membrane, which allows the passage of ions. Channels are often permeable only to specific kinds of ions.

connexin: the proteins from which a connexon is composed.

connexon: a channel through the membrane of two neurons at a electrical synapses, allowing ions to pass from one cell into another. A connexon is composed of two hemi-channels.

contralateral: in anatomy, being on the opposite side (of the mid-line.)

crayfish: a fresh water dwelling decapod closely related to the lobster.

cytosol: the fluid portion of a cell. Cytosol is a continuous aqueous solution that is conductive to electricity.

cytoplasm: everything that is contained within the confines of the cell membrane, excluding the nucleus. It consists of a continuous aqueous solution (cytosol) and the organelles. Most of the cell's chemical activity takes place in the cytoplasm.

depolarise: a rise in the membrane potential from its resting potential.

dendrite: a neurite specialised to receive inputs from other neurons (Bear et al., 1996).

detritus: decomposing plant material.

distal: the direction along a neurite away from the cell body.

disynaptic: a neural pathway that progresses through two synapses.

efferent: an axon projecting outward (from the contextual neural structure.)

electrical synapse: a direct electrical connection between two cells. Electrical synapses may be ohmic or rectifying.

electrogenic membrane: a membrane that actively contributes to an electric current that it is conducting.

electrophysiology: the study of how living organisms function with regard to electric phenomena.

equilibrium potential:

EPSP: Excitatory post-synaptic potential.

exoskeleton: a hard encasement on the surface of an animal such as an insect or crustecean.

fire: the action of a cell generating an action potential (or spike.)

GABA: gamma-aminobutyric acid. An inhibitory neurotransmitter used in the nervous systems of many animals, including crayfish.

ganglion: a cluster of neurons.

gap junction: an electrical synapse. Gap junctions may be ohmic or rectifying.

giant: a very large neuronal cell.

GMS: giant motor synapse. A giant synapse in the crayfish lateral giant (LG) escape circuit that connects the LG to motor giant (MoG). This is an electrical synapse has been shown to be rectifying.

hemi-channel: a channel in an electrical synapse in one of the two cells being connected.

hemi-ganglion: half of a ganglion. A ganglion is usually divided into symmetrical halves: one on each side of the body.

homologue: when referring to neurons, describes the neuron that performs the same function but is located on the contralateral side of the animal.

hyperpolarise: to cause a fall in the membrane potential.

integrating segment: the spike initiation zone (see spike initiation zone.)

interneuron: a neuron that is neither a primary afferent or a motorneuron. If one thinks of the nervous system as receiving input from the primary afferents and interacting with the environment through motorneurons, the interneurons are all the cells in between responsible for processing and storing the information. Most neurons in the nervous system are interneurons.

ipsilateral: in anatomy, being on the same side (of the mid-line.)

IPSP: Inhibitory post-synaptic potential.

lateral: a direction meaning away from the mid-line. (See figure 3.8.)

lateral giants: A group of giant connected neurons located in the crayfish abdomen. The lateral giants are command neurons. When the lateral giants are spiked an escape tail-flip propelling the crayfish forwards is initiated. (Refer to section 3.3.1.)

LG: lateral giants. See lateral giants.

mechanosensory neuron: A neuron attached to a sensory hair, carrying touch signals. In the crayfish tailfan, mechanosensory hairs detect water movement and touch.

medial: a direction meaning towards the mid-line. (See figure 3.8.)

membrane: the barrier that encloses a cell.

membrane time constant:

membrane potential: the voltage difference between inside of the cell and outside of the cell.

mid-line: an imaginary plane, running from the front to the rear of the crayfish, that dissects the crayfish into two mirroring halves.

MG: Medial giant. A command neuron in the crayfish escape circuit. The MG responds to stimulus to the front of the animal and effects tail-flip propelling the crayfish backwards if the MG is fired. (Refer to section 3.3.1.)

morphology: the anatomical shape of a neuron.

monosynaptic: a neural pathway between two neurons that are directly connected by a synapse.

motorneuron: a neuron that is connected to and drives a muscle cell.

neurite: any thin tube that extends from the cell body. A neurite can be either a dendrite or an axon.

neurotransmitter: a chemical that used in synaptic transmission between cells. At chemical synapses, it is released at the pre-synaptic terminal and acts on receptor sites at the post-synaptic terminal.

nerve: In the context of the crayfish tailfan, a nerve is a group of primary afferent axons originating from a section of the tailfan.

non-giant circuit: the neural circuits in crayfish that are responsible for generating tail-flips that are not LG or MG tail-flips. Non-giant tail-flips

nucleus: a roughly spherical organelle within a cell containing the chromosomes, which contain the cell's DNA.

octopamine: a neurotransmitter.

ohmic gap junction: an electrical synapse that is non-rectifying.

organelle: a structure within the cell which is enclosed within its own membrane. Organelles perform a specific functions within the cell.

orthodromic: a direction of synaptic transmission where transmission is from the pre-synaptic terminal to the post-synaptic terminal. (The usual direction of synaptic transmission.)

passive membrane: a membrane that passively conducts an electric current. In classical models of neurons, dendrites have passive membranes.

pereiopods: the four pairs of thoracic limbs that serve as crayfish's walking legs. (See figure 3.7.)

phasic: in phase. Typically used in the context of meaning highly synchronised inputs.

pleopods: small appendages on the crayfish's abdomen that assist in swimming. Also called swimmernets. (See figure 3.7.)

polysynaptic: a neural pathway that progresses through multiple synapses.

posterior: the direction towards the rear of the animal. (See figure 3.8.)

post-synaptic: refers to the neuron that receives synaptic input.

post-synaptic terminal: the part of the cell membrane in the post-synaptic neuron where synaptic transmission takes place.

pre-synaptic: refers to the neuron that transmits a signal to a target neuron (*i.e.* the post-synaptic cell.)

pre-synaptic terminal: the part of the cell membrane in the pre-synaptic neuron where synaptic transmission takes place.

primary afferent: a sensory cell.

proxal: the direction along a neurite towards the cell body.

PSP: post-synaptic potential. The change in depolarisation at the post synaptic cell. Usually refers to an EPSP.

receptor: a specialised protein that detects chemical signals, such as neurotransmitters, and initiates a cellular response (Bear et al., 1996).

rectifying gap junction: an electrical synapse that has variable conductance that is a function of the voltage difference between the pre-synaptic cell and the post-synaptic cell. There are rectifying gap junctions on the connections from the primary afferents and the interneurons onto the lateral giant.

recurrent inhibition: inhibition applied to the lateral giant escape circuit after an escape tail-flip has been invoked to prevent repeated reactivation of the lateral giant circuit.

resting potential: the potential a neuron's membrane converges to if it receives no stimulation. A typical resting potential of a neuron is around 65-70 millivolts.

rostral: the direction towards the rostrum (front) of the crayfish. Equivalent to anterior. (See figure 3.8.)

rostrum: a sharp spine extending from the front of the crayfish's carapace.

serotonin (5-HT): an amine neurotransmitter that is often used in modulatory effects.

soma: the cell body of the neuron, which contains the nucleus (refer to figure 3.1).

spike: an action potential (see action potential).

spike initiation zone: the part of the

spiking neuron: a neuron which is able to generate action potentials (also know as spikes.)

spiking threshold: the minimum size of stimulus that needs to be applied to a neuron to cause it to spike.

swimmernets: small appendages on the crayfish's abdomen that assist in swimming. Also called pleopods. (See figure 3.7.)

synapse: A connection between neurons, by which an excitation is conveyed from one to the other. Synapses may be chemical or electrical.

telson: the tailfan at the end of the crayfish abdomen. (See figure 3.7.)

thorax: the middle region of the body of an arthropod between the head and the abdomen.

tonic: continuous or long term.

tonic inhibition: distally applied inhibition to the lateral giant, mediated by a GABA channel from rostral ganglia, that has the effect of raising the lateral giant threshold during feeding and restraint.

Bibliography

Agmon-Snir, H., Carr, C. E., and Rinzel, J. (1998). The role of dendrites in auditory coincidence detection. *Nature*, 393(6682):207–208.

Antonsen, B. L. and Edwards, D. H. (2003). Differential dye coupling reveals lateral giant escape circuit in crayfish. *Journal of Comparative Neurology*, 466(1):1–13.

Atwood, H. L. and Pomeranz, B. (1977). Dendritic bottlenecks of crustecean motoneurons. *Journal of Neurocytology*, 6:251–268.

Beall, S. P., Langley, D. J., and Edwards, D. H. (1991). Inhibition of escape tailflip in crayfish during backward walking and the defense posture. *Journal of Experimental Biology*, 152:577–582.

Bear, M. F., Connors, B. W., and Paradiso, M. A. (1996). *Neuroscience: exploring the brain*. Williams & Wilkins, Baltimore, USA.

Beer, R. D. (1990). Intelligence as adaptive behaviour: An experiment in computational neuroethology. In *Perspectives in Artificial Intelligence*, volume 6. Academic Press.

Beer, R. D. and Chiel, H. J. (1993). Simulations of cockroach locomotion and escape. In *Biological Neural Networks in Invertebrate Neuroethology and Robotics*, pages 267–285. Academic Press, San Diago.

Beer, R. D., Chiel, H. J., and Gallagher, J. C. (1999). Evolution and analysis of model cpgs for walking: Ii. general principles and individual variability. *Journal of Computational Neuroscience*, 7(2):119–147.

Beer, R. D., Chiel, H. J., Quinn, R. D., Espenschied, K., and Larsson, P. (1992). A distributed neural net architecture for hexapod robot locomotion. *Neural Computation*, 4:356–365.

Beer, R. D., Chiel, H. J., and Sterling, L. (1989). Heterogeneous neural networks for adaptive behavior in dynamic environments. In Touretzky, D. S., editor, *Advances in Neural Information Processing Systems 1, [NIPS Conference, Denver, Colorado, USA, 1988]*, pages 577–585. Morgan Kaufmann.

Bennett, M. V. L. (1984). Escapism: Some startling revelations. In *Neural Mechanisms of Startle Behavior*, pages 353–363. New York: Plenum Press. edited by R. C. Eaton.

Bennett, M. V. L., Contreras, J. E., Bukauskas, F. F., and Sez, J. C. (2003). New roles for astrocytes: Gap junction hemichannels have something to communicate. *Trends in Neurosciences*, 26(11):610–617.

Berthoz, A., Floreano, D., J, J. M., Roitblat, H., and Wilson, S., editors (2000). *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*. MIT Press.

Bower, J. M. and Beeman, D. (1998). *The Book of GENESIS: Exploring Realistic Neural Models with the GEneral NEural SImulation System.* Springer-Verlag, New York.

Brent, R. P. (1973). *Algorithms for minimization without derivatives*. Prentice-Hall, Englewood Cliffs, New Jersey.

Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23.

Bukauskas, F. F., Bukauskiene, A., Verselis, V. K., and Bennett, M. V. L. (2002). Coupling asymmetry of heterotypic connexin 45/ connexin 43-egfp gap junctions: Properties of fast and slow gating mechanisms. *Proceedings of the National Academy of Science*, 99(10):7113–7118.

Bush, P. C. and Sejnowski, T. J. (1993). Reduced compartmental models of neocortical pyramidal cells. *Journal of Neuroscience Methods*, 46(2):159–166.

Carnevale, N. T. and Hines, M. L. (2003). The NEURON book. Yale University, New Haven, CT.

Carr, C. E. and Boudreau, B. E. (1993). Organization of the nucleus magnocellularis and the nucleus laminaris in the barn owl: encoding and measuring interaural time differences. *Journal of Comparative Neurology*, 333(3):337–355.

Cattaert, D., Libersat, F., and Manira, A. E. (2001). Presynaptic inhibition and antidromic spikes in primary afferents of the crayfish: A computational and experimental analysis. *Journal of Neuroscience*, 21(3):1007–1021.

Chan, P., Lee, R., and Kramer, D. (1998). *The Java Class Libraries*. Addison-Wesley, Reading, Massachusetts.

Cliff, D., Husbands, P., Meyer, J.-A., and Wilson, S. W., editors (1994). *From Animals to Animats 3: Proceedings of the Third International*

Conference on Simulation of Adaptive Behavior, Cambridge, Massachusetts.

Cliff, D. and Miller, G. F. (1996). Co-evolution of pursuit and evasion II: Simulation methods and results. In *Animals to Animats IV: Proceedings of Fourth International Conference on Simulation of Adaptive Behaviour*, pages 506–515. MIT Press.

D., R. and J., M. (1986a). *Parallel Distributed Processing - Explorations in the Microstructures of Cognition - Foundations*, volume 1. MIT Press.

D., R. and J., M. (1986b). *Parallel Distributed Processing - Explorations in the Microstructures of Cognition - Psychological and Biological Models*, volume 2. MIT Press.

Dorigo, M. and Colombetti, M. (1994). Robot shaping: developing autonomous agents through learning. *Artifical Intelligence*, 71:321–370.

Eaton, R. C., editor (1984). *Neural mechanisms of startle behaviour*. Plenum Press, New York.

Ebihara, L. (2003). Physiology and biophysics of hemi-gap-junctional channels expressed in xenopus oocytes. *Acta Physiologica Scandinavica*, 179(1):5–8.

Edgar, G. (2000). *Australian marine life: the plants and animals of temperate waters*. Reed New Holland, Sydney.

Edwards, D. H. (1990a). Mechanisms of depolarizing inhibitino at the crayfish giant motor synapse. I. Electrophysiology. *Journal of Neurophysiology*, 64(2):532–540.

Edwards, D. H. (1990b). Mechanisms of depolarizing inhibition at the crayfish giant motor synapse. II. Quantative reconstruction. *Journal of Neurophysiology*, 64(2):541–550.

Edwards, D. H. (1991). Mutual inhibition among neural command systems as a possible mechanism for behavioral choice in crayfish. *Journal of Neuroscience*, 11(5):1210–1223.

Edwards, D. H., Fricke, R. A., Barnett, L. D., Yeh, S.-R., and Leise, E. M. (1994a). The onset of response habituation and the growth of the lateral giant neuron of crayfish. *Journal of Neurophysiology*, 72:890–898.

Edwards, D. H., Heitler, W. J., and Krasne, F. B. (1999). Fifty years of a command neuron: the neurobiology of escape behavior in the crayfish. *Trends in Neuroscience*, 22(4):153–161.

Edwards, D. H., Heitler, W. J., Leise, E. M., and Fricke, R. A. (1991). Postsynaptic modulation of rectifying synaptic inputs to the lg escape command neuron in crayfish. *Journal of Neuroscience*, 11(7):2117–2129.

Edwards, D. H. and Mulloney, B. (1987). Synaptic integration in excitatory and inhibitory motoneurons. *Journal of Neurophysiology*, 57:1425–1445.

Edwards, D. H., Yeh, S.-R., Barnett, L. D., and Nagappan, P. R. (1994b). Changes in synaptic integration during the growth of the lateral giant neuron of crayfish. *Journal of Neurophysiology*, 72(2):899–908.

Edwards, D. H., Yeh, S.-R., Barnett, L. D., and Nagappan, P. R. (1994c). Changes in synaptic integration during the growth of the lateral giant neuron of crayfish. *Journal of Neurophysiology*, 72:899–908.

Edwards, D. H., Yeh, S.-R., and Krasne, F. B. (1998). Neuronal coincidence detection by volage-sensitive electrical synapses. *Proceedings of National Academy of Science USA*, 95:7145–7150.

Elliot, T. and Shadbolt, N. R. (2003a). Developing a robot visual system using a biologically inspired model of neuronal development. *Robotics and Autonomous Systems*, 45(2):111–130.

Elliot, T. and Shadbolt, N. R. (2003b). Developmental robotics: manifesto and application. *Philosophical Transactions of the Royal Society of London A*, 361:2187–2206.

Espenschied, K. S., Quinn, R. D., Chiel, H. J., and Beer, R. D. (1993). Leg coordination mechanisms in stick insect applied to hexapod robot locomotion. *Adaptive Behavior*, 1(4):455–468.

Espenschied, K. S., Quinn, R. D., Chiel, H. J., and Beer, R. D. (1996). Biologically-based distributed control and local reflexes improve rough terrain locomotion in a hexapod robot. *Robotics and Autonomous Systems*, 18:59–64.

Ficici, S. G. and Pollack, J. B. (1998). Coevolving communinative behavior in a linear pursuer-evasion game. In *Animals to Animats V: Proceedings of Fifth International Conference on Simulation of Adaptive Behaviour*, pages 557–561. MIT Press.

Floreano, D., Hallam, B., Hallam, J., Hayes, G., and Meyer, J., editors (2002). *Animals to Animats 6: Proceedings of Fifth International Conference on Simulation of Adaptive Behaviour*. MIT Press.

Fraser, P. J. (1982). Views on the nervous control of complex behaviour. In *The Biology of Crustacea*, volume 4, pages 293–319. Academic Press.

Fricke, R. A. (1986). Structure-function considerations in the developmental expression of crayfish behavioral plasticity. *IEEE Transactions on Systems, Man and Cybernetics*, 1:513–518.

Fricke, R. A., Block, G. D., and Kennedy, D. (1982). Inhibition of mechanosensory neurons in the crayfish. II. Inhibition associated with proprioceptive feedback from locomotion. *Journal of Comparative Physiology*, 149:251–262.

Furshpan, E. J. and Potter, D. D. (1959). Transmission at the giant motor synapses of the crayfish. *Journal of Physiology*, 145:289–325.

Giaume, C., Kado, R. T., and Korn, H. (1987). Voltage-clamp analysis of a crayfish rectifying synapse. *Journal of Physiology*, 386:91–112.

Glantz, R. M. and Viancour, T. (1983). Integrative properties of crayfish medial giant neuron: steady-state model. *Journal of Neurophysiology*, 50(5):1122–1142.

Glanzman, D. L. and Krasne, F. B. (1983). Serotonin and octopamine have opposite modulatory effects on the crayfish's lateral giant escape reaction. *Journal of Neuroscience*, 3(11):2263–2269.

Goddard, J. S. (1988). Food and feeding. In Holdich, D. M. and Lowery, R. S., editors, *Ecology in Freshwater Crayfish: Biology, Management and Exploitation*. Croom Helm Australia.

Grau-Serrat, V., Carr, C. E., and Simon, J. Z. (2003). Modeling coincidence detection in nucleus laminaris. *Biology Cybernetic*, 89:388–396.

Grefenstette, J. J., Ramsey, C. L., and Schultz, A. C. (1990). Learning sequential decision rules using simulation models and competition. *Machine Learning*, 5(4):355–391.

Heitler, W. J., Fraser, K., and Edwards, D. H. (1991). Different types of rectification at electrical synapses made by a single crayfish nerone investigated experimentally and by computer simulation. *Journal of Comparative Physiology A*, 169:707–718.

Heitler, W. J. and H.Edwards, D. (1998). Effect of temperature on a voltage-sensitive electrical synapse in crayfish. *Journal of Experimental Biology*, 201:503–513.

Herberholz, J., Antonsen, B. L., and Edwards, D. H. (2002). A lateral excitatory network in the escape circuit of crayfish. *Journal of Neuroscience*, 22(20):9078–9085.

Herberholz, J., Issa, F. A., and Edwards, D. H. (2001). Patterns of neural circuit activation and behavior during dominance hierarchy formation in freely behaving crayfish. *Journal of Neuroscience*, 21(8):2759–2767.

Herberholz, J., Sen, M. M., and Edwards, D. H. (2003). Parallel changes in agonistic and non-agonistic behaviors during dominance hierarchy formation in crayfish. *Journal of Comparative Physiology. A, Sensory, neural, and behavioral physiology*, 189(4):321–325.

Herberholz, J., Sen, M. M., and Edwards, D. H. (2004). Escape behaviour and escape circuit activation in juvenile crayfish during prey-predator interactions. To appear in *Journal of Experimental Biology*. Accepted March 2004.

Hill, A. A. V., Edwards, D. H., and Murphey, R. K. (1994). The effect of neuronal growth on synaptic integration. *Journal of Computational Neuroscience*, 1:239–254.

Hines, M. (1984). Efficient computation of branched nerve equations. *International Journal of Bio-Medical Computation*, 15:69–76.

Hines, M. (1993). Neuron–a program for simulation of nerve equations. In Eeckman, F., editor, *Neural Systems: Analysis and Modeling*, pages 127–136. Kluwer Academic Publishers, Norwell Massachusetts.

Hines, M. (1994). The neuron simulation program. In Skrzypek, J., editor, *Neural Network Simulation Environments*, pages 147–163. Kluwer Academic Publishers, Norwell Massachusetts.

Hines, M. (1998). The neurosimulator neuron. In Koch, C. and Segev, I., editors, *Methods in Neuronal Modeling*, pages 129–136. MIT Press.

Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation nerve. *Journal of Physiology*, 117:500–544.

Hogger, J. B. (1988). Population, biology and behaviour. In Holdich, D. M. and Lowery, R. S., editors, *Ecology in Freshwater Crayfish: Biology, Management and Exploitation*. Croom Helm Australia.

Holdich, D. M. and Reeves, I. D. (1988). Functional morphology and anatomy. In Holdich, D. M. and Lowery, R. S., editors, *Ecology in Freshwater Crayfish: Biology, Management and Exploitation*. Croom Helm Australia.

Howard, A. (1999). *Probabilistic Navigation: Coping with Uncertainty in Robot Navigation Tasks*. PhD thesis, Department of Computer Science and Software Engineering, University of Melbourne.

Hoy, R. R. (1993). Acoustic startle: an adaptive behaviour act in flying insects. In *Biological neural networks in invertebrate neuroethology and robotics*, pages 139–158. Academic Press, San Diago.

Huber, R. and Kravitz, E. A. (1995). A quantitative analysis of agonistic behavior in juvenile american lobsters (homarus americanus l.). *Brain Behaviour and Evolution*, 46(2):72–83.

Huxley, T. H., F. R. S. (1896). *The Crayfish: An introduction to the study of Zoology*. Kegan Paul, Trench, Trubner & Co., London.

Issa, F. A., Adamson, D. J., and Edwards, D. H. (1999). Dominance hierarchy formation in juvenile crayfish procambarus clarkii. *Journal of Experimental Biology*, 202(24):3497–3506.

Jackson, D. J. and MacMillan, D. L. (2000). Tailflick escape behavior in larval and juvenile lobsters (homarus americanus) and crayfish (cherax destructor). *Biological Bulleltin*, 198(3):307–318.

Joy, B., Steele, G., Gosling, J., and Bracha, G. (2000). *Java Language Specification*. Addison-Wesley, 2nd edition. Available at `http://java.sun.com/docs/books/jls/`.

Katz, B. (1966). *Nerve, Muscle and Synapse*. McGraw-Hill, New York.

Kennedy, D., McVittie, J., Calabrese, R., Fricke, R. A., Craelius, W., and Chiapella, P. (1980). Inhibition of mechanosensory interneurons in the crayfish. I. Presynaptic inhibition from giant fibers. *Journal of Neurophysiology*, 43(6):1495–1509.

Kirk, M. D. (1985). Presynaptic inhibition in the crayfish cns: pathways and synaptic mechanisms. *Journal of Neurophysiology*, 54(5):1305–1325.

Kirk, M. D. and Wine, J. J. (1984). Identified interneurons produce both primary afferent depolarization and presynaptic inhibition. *Science*, 225:854–856.

Koza, J. (1991). Evolution and co-evolution of computer programs to control independently-acting agents. In *Animals to Animats: Proceedings of First International Conference on Simulation of Adaptive Behaviour (SAB92)*, pages 366–375. MIT Press.

Krasne, F. B. and Glanzman, D. L. (1986). Sensitisation of the crayfish lateral giant escape reaction. *Journal of Neuroscience*, 6(4):1013–1020.

Krasne, F. B. and Lee, S. C. (1988). Response-dedicated trigger neurons as control points for behavioural actions: selective inhibition of lateral giant command neurons during feeding in crayfish. *Journal of Neuroscience*, 8(10):3703–3712.

Krasne, F. B., Shamsian, A., and Kulkarni, R. (1997). Altered excitability of the crayfish lg escape reflex during agnostic encounters. *Journal of Neuroscience*, 17(2):709–716.

Krasne, F. B. and Teshiba, T. M. (1995). Habituation of an invertebrate escape reflex due to modulation by higher centers rather than local events. In *Proceedings of the National Academy of Science, USA*, volume 92, pages 3362–3366.

Krasne, F. B. and Wine, J. J. (1975). Extrinsic modulation of crayfish escape behaviour. *Journal of Experimental Biology*, 63:433–450.

Krasne, F. B. and Wine, J. J. (1984). The production of crayfish tailflip escape response. In *Neural Mechanisms of Startle Behaviour*, pages 179–211. New York: Plenum Press. edited by R. C. Eaton.

Kravitz, E. A. (1988). Hormonal control of behavior: amines and the biasing of behavioral output in lobsters. *Science*, 241(4874):1775–1781.

Kupfermann, I. and Weiss, K. R. (1978). The command neuron concept. *Behavioural Brain Science*, 1:3–39. cited in (Krasne and Wine, 1978).

L. Davis, J. and Lorente de Nó, R. (1947). Contribution to the mathematical theory of the electrotonous. *Studies from the Rockefeller Institue for Medical Research*, 131:442–496.

Lang, F., Govind, C. K., Costello, W. J., and Greene, S. I. (1977). Developmental neuroethology: changes in escape and defensive behavior during growth of the lobster. *Science*, 197:682–685.

Laurent, G. (1993). Integration by spiking and nonspiking local neurons in the locust central nervous system. importance of cellular and synaptic properties for network function. In *Biological neural networks in inver-*

tebrate neuroethology and robotics, pages 69–85. Academic Press, San Diago.

Lee, S. C. and Krasne, F. B. (1993). Ultrastructure of the circuit providing input to the crayfish lateral giant neurons. *Journal of Comparative Neurology*, 327(2):271–288.

Livingstone, M. S., Harris-Warrick, R. M., and Kravitz, E. A. (1980). Serotonin and octopamine produce opposite postures in lobsters. *Science*, 208:76–79.

Lockery, S. R. (1993). Voyages through weight space: Network models of an escape reflex in the leech, biological neural networks in invertebrate neuroethology and robotics. In *Biological neural networks in invertebrate neuroethology and robotics*, pages 251–266. Academic Press, San Diago.

Lowery, R. S. (1988). Growth, moulting and reproduction. In Holdich, D. M. and Lowery, R. S., editors, *Ecology in Freshwater Crayfish: Biology, Management and Exploitation*. Croom Helm Australia.

MacMillan, D. L. and Patullo, B. W. (2001). Insights for robotic design from studies of the control of abdominal position in crayfish. *Biological Bulletin*, 200(2):201–205.

Maddison, D. R. (2001). The tree of life. Web database hosted by the College of Agriculture, University of Arizona. Available at `http://phylogeny.arizona.edu/tree/`.

Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J., and Wilson, S. W., editors (1996). *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, Cambridge, Massachusetts.

Merrick, J. R. (1993). *Freshwater Crayfishes of New South Wales*. Linnean Society of New South Wales.

Meyer, J.-A., Roitblat, H. L., and Wilson, S. W., editors (1992). *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, Cambridge, Massachusetts. MIT Press.

Meyer, J.-A. and Wilson, S. W., editors (1990). *From Animals To Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, Cambridge, Massachusetts. MIT Press.

Microbrightfield Inc. (2003). Neurolucida.

Miller, G. F. and Cliff, D. (1994). Protean behaviour in dynamic games: co-evolution of pursuit-evasion tactics. *From Animals to Animats 3*, pages 411–420.

Miller, M. W., Lee, S. C., and Krasne, F. B. (1987). Cooperativity-dependent long-lasting potentiation in the crayfish lateral giant escape reaction circuit. *Journal of Neuroscience*, 7(4):1081–1092.

Minsky, M. and Papert, S. (1969). *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, Massachusetts, USA.

Morse, T. M., Feree, T. C., and Lockery, S. R. (1998). Robust spatial navigation in a robot inspired by chemotaxis in caenorhabditis elegans. *Adaptive Behavior*, 6(3/4):393–410.

Nagayama, T., Aonuma, H., and Newland, P. L. (1997). Convergent chemical and electrical synaptic inputs from proprioceptive afferents onto an identified intersegmental interneuron in the crayfish. *Journal of Neurophysiology*, 77(5):2826–2830.

Nolfi, S. (1997). Using emergent modularity to develop control systems for mobile robots. *Adaptive Behavior*, 5(3/4):343–363.

Olszewski, P. (1980). *A Salute to the humble yabby*. London Angus & Robertson.

Peña, J. L., Viete, S., Funabiki, K., Saberi, K., and Konishi, M. (2001). Cochlear and neural delays for coincidence detection in owls. *Journal of Neuroscience*, 21(23):9455–9459.

Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S. W., editors (1998). *From Animals to Animats 5: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior*. MIT Pres.

Pfeifer, R. and Scheier, C. (1999). *Understanding Intelligence*. MIT Press.

Plum, A., Hallas, G., Magin, T., Dombrowski, F., Hagendorff, A., Schumacher, B., Wolpert, C., Kim, J.-S., Lamers, W. H., M. Evert, P. M., Traub, O., and Willecke, K. (2000). Unique and shared functions of different connexins in mice. *Current Biology*, 10(18):1083–1091.

Prescott, T. J., Redgrave, P., and Gurney, K. (1999). Layered control architectures in robots and vertebrates. *Adaptive Behaviour*, 7(1):99–127.

Rall, W. (1957). Membrane time constant of motoneurones. *Science*, 126:454.

Rall, W. (1989). Cable theory for dendritic neurons. In Koch, C. and Segev, I., editors, *Methods in Neural Modelling: from synapses to networks*, pages 9–62. MIT Press.

Rall, W. and Agmon-Snir, H. (1998). Cable theory for dendritic neurons. In Koch, C. and Segev, I., editors, *Methods in Neural Modelling: from ions to networks*, pages 27–91. MIT Press.

Rall, W. and Rinzel, J. (1973). Branch input resistance and steady attenuation for input to to one branch of a dendritic neuron model. *Biophysical Journal*, 13:648–688.

Ritzmann, R. E. (1993). The neural organisation of cockroach escape and its role in context-dependant orientation. In *Biological neural networks in invertebrate neuroethology and robotics*, pages 113–137. Academic Press, San Diago.

Roberts, A. (1968). Recurrent inhibition in the giant-fibre system of the crayfish and its effect on the excitability of the escape response. *Journal of Experimental Biology*, 48(3):545–567.

Rosenblatt, F. (1962). *Principles of Neurodynamics*. Spartan Books.

Russell, S. J. and Norvig, P. (2003). *Artifical Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, New Jersey, USA, 2nd edition.

Schrameck, J. E. (1970). Crayfish swimming: alternating motor output and giant fiber activity. *Science*, 169(4304):698–700.

Segev, I. and Burke, R. E. (1999). Compartmental models of complex neurons. In Koch, C. and Segev, I., editors, *Methods in Neural Modelling: from ions to networks*, pages 93–136. MIT Press.

Simon, J. Z., Carr, C. E., and Shamma, S. A. (1999). A dendritic model of coincidence detection in the avian brainstem. *Neurocomputing*, 26–27:263–269.

Teshiba, T., Shamsian, A., Yashar, B., Yeh, S.-R., Edwards, D. H., and Krasne, F. B. (2001). Dual and opposing modulatory effects of serotonin on crayfish lateral giant escape command neurons. *Journal of Neuroscience*, 21(12):4523–4529.

Versteeg, S., Antonsen, B. L., Agran, J., Heberholz, J., and Edwards, D. H. (2003). Simulation of the lateral excitatory network in crayfish based on anatomical and physiological data. Program no. 270.8. In *Neuroscience 2003 Abstract Viewer/Itinerary Planner*. Washington, DC: Society for Neuroscience.

Versteeg, S. C. (2001). Getting the job done in a hostile environment. In *Proceedings of the 14th Australian Joint Conference on Artificial Intelligence (LNAI 2256)*. Springer.

Vescovi, P. (1995). Integration and processing of sensory information in a neural network. Master's thesis, University of Melbourne.

Vu, E. T., Berkowitz, A., and Krasne, F. B. (1997). Postexcitatory inhibition of the crayfish lateral giant neuron: A mechanism for sensory temporal filtering. *Journal of Neuroscience*, 17(22):8867–8879.

Vu, E. T. and Krasne, F. B. (1993a). Crayfish tonic inhibition: Prolonged modulation of behavioural excitability by classical gabaergic inhibition. *Journal of Neuroscience*, 13(10):4394–4402.

Vu, E. T. and Krasne, F. B. (1993b). The mechanism of tonic inhibition of crayfish escape behaviour: distal inhibition and its functional significance. *Journal of Neuroscience*, 13(10):4379–4393.

Vu, E. T., Lee, S. C., and Krasne, F. B. (1993). The mechanism of tonic inhibition of crayfish escape behaviour: distal inhibition and its functional significance. *Journal of Neuroscience*, 13(10):4379–4393.

Wahde, M. and Nordahl, M. G. (1998). Evolution of protean behavior in pursuit-evasion contests. In *Animals to Animats V: Proceedings of Fifth International Conference on Simulation of Adaptive Behaviour*, pages 557–561. MIT Press.

Wasserman, P. D. (1989). *Neural Computing: Theory and Practice*. Van Nostrand Reinhold, New York.

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393:440–442.

White, T. W. (2003). Nonredundant gap junction functions. *News in Physiological Sciences*, 18(3):95–99.

White, T. W. and Bruzzone, R. (2000). Gap junctions: Fates worse than death? *Current Biology*, 10(18):R685–R688.

Wiersma, C. A. G. (1947). Giant nerve fibre system of the crayfish. a contribution to comparative physiology of synapse. *Journal of Neurophysiology*, 10:23–38.

Wiersma, C. A. G. (1961). Comparitive neurophysiology, its aims and its present status. *Fetschr. Kh. Koshtoyants*, pages 82–88. In Russian. Cited in (Fraser, 1982).

Wine, J. J. (1984). The structural basis of an innate behavioural pattern. *Journal of Experimental Biology*, 112:283–319.

Wine, J. J. and Krasne, F. B. (1982). The cellular organisation of crayfish escape behaviour. In *The Biology of Crustacea*, volume 4, pages 241–292. Academic Press.

Withnall, F. (2000). Biology of yabbies (cherax destructor). Aquaculture Notes, Notes Series No AS0004, Department of Natural Resources and Environment, Melbourne, Australia. ISSN 1441-1199. Available at `http://www.nre.vic.gov.au/web/root/domino/infseries/infsheet.nsf/1ac21787989b9a504a2566500008ceca/9acd9b441cac230d4a25665000839c5d?OpenDocume%nt`.

Wooldridge, M. (1999). Intelligent agents. In Weiss, G., editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pages 27–78. The MIT Press, Cambridge, MA, USA.

Wooldridge, M. and Jennings, N. R. (1995). Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2):115–152.

Yeh, S.-R., Fricke, R. A., and Edwards, D. H. (1996). The effect of social experience on serotonergic modulation of the escape circuit of crayfish. *Science*, 271(5247):366–369.

Yeh, S.-R., Jeng-Wei, Y., Li-Yun, T., and Edwards, D. H. (2002). Ltp at a rectifying electrical synapse in crayfish. Program no. 347.1. In *Neuroscience 2002 Abstract Viewer/Itinerary Planner*. Washington, DC: Society for Neuroscience.

Yeh, S.-R., Musolf, B. E., and Edwards, D. H. (1997). Neuronal adaptations to changes in the social dominance status of crayfish. *Journal of Neuroscience*, 17(2):697–708.

Zucker, R. S. (1972a). Crayfish escape behaviour and central synapses. I. Neural circuit exciting lateral giant fiber. *Journal of Neurophysiology*, 35:599–620.

Zucker, R. S. (1972b). Crayfish escape behaviour and central synapses. II. Physiological mechanisms underlying behavioural habituation. *Journal of Neurophysiology*, 35:621–637.