



Stochastic Processing for Enhancement of Artificial Insect Vision

by

Gregory P. Harmer

B.Sc. (Applied Maths & Computer Science), Adelaide University, Australia, 1996
B.E. (First Class Honours), Adelaide University, Australia, 1997


Thesis submitted for the degree of

Doctor of Philosophy

in

Department of Electrical and Electronic Engineering,
Faculty of Engineering, Computer and Mathematical Sciences
Adelaide University, Australia

November, 2001



© 2001
Gregory P. Harmer
All Rights Reserved



Contents

Heading	Page
Contents	iii
Abstract	ix
Statement of Originality	xi
Acknowledgments	xiii
Publications	xv
List of Figures	xvii
List of Tables	xxiii
Chapter 1. Introduction and Motivation	1
1.1 Thesis Layout	2
1.2 Introduction	3
1.3 Background	4
1.3.1 Motion Detection	6
1.4 Motivation	7
1.4.1 Millimetre Wave Antenna Array	9
1.5 Stochastic Processing	12
1.5.1 Brownian Ratchets	12
1.5.2 Stochastic Resonance	13
1.6 Current State of Knowledge	14
1.7 Original Contributions	15
Chapter 2. Brownian Ratchets	17
2.1 Brownian Motion	18

2.2	Rectifying Brownian Motion	19
2.2.1	Laws of Thermodynamics	20
2.2.2	Ratchet and Pawl Device	21
2.2.3	Maxwell's Demon	24
2.2.4	Utilising Thermal Fluctuations	27
2.3	Brownian Ratchets	28
2.3.1	Flashing Ratchets	30
2.3.2	A Very Brief Review	31
2.3.3	Examples of Ratchets	31
2.4	Summary	35
Chapter 3. Parrondo's Games		37
3.1	Introduction	38
3.2	Construction and Characteristics	39
3.2.1	Construction of the Games	39
3.2.2	Playing the Games	41
3.2.3	Fairness	41
3.2.4	Distributions and Behaviour	45
3.2.5	Mixing Strategies	47
3.2.6	Explanation of the Games	49
3.2.7	Observations	52
3.3	Analysis of the Games	54
3.3.1	Discrete-Time Markov-Chains (DTMCs)	54
3.3.2	Modelling the Games as DTMCs	57
3.3.3	Constraints of Parrondo's Games	63
3.3.4	Rate of Winning	67
3.4	Entropy	69
3.4.1	Entropy and Probability Space	71
3.4.2	Entropy Paradox	73
3.5	History Dependent Games	74
3.5.1	Construction	74

3.5.2	Results and Distributions	75
3.5.3	Analysis of the Games	76
3.5.4	Probability Space	81
3.6	Other Phenomena	83
3.6.1	Co-operative Parrondo’s Games	83
3.6.2	Fractal Properties of Parrondo’s Games	83
3.6.3	Parrondo’s Games are Ubiquitous	86
3.6.4	Related Works	88
3.7	Summary	88
 Chapter 4. Stochastic Resonance		 91
4.1	Introduction	92
4.1.1	The Basics	93
4.1.2	Examples of SR	98
4.2	Systems Demonstrating SR	109
4.2.1	Threshold Devices	110
4.2.2	Neuron Models	111
4.2.3	Bistable Dynamics	115
4.2.4	Other Nonlinear Systems	120
4.3	Noise	120
4.3.1	White – Gaussian	120
4.3.2	Coloured – Ornstein-Uhlenbeck	121
4.3.3	Generation of Random Numbers	122
4.3.4	Simulating Stochastic Systems	124
4.4	Quantifying SR	124
4.4.1	Signal-to-Noise Ratio	125
4.4.2	Residence Time Distribution	127
4.4.3	Correlation Coefficients	128
4.4.4	Average Mutual Information	132
4.4.5	Channel Capacity	139
4.5	Noise Related Phenomena	144
4.5.1	Noise Linearisation	144
4.5.2	Phase Space	145
4.6	Summary	147

Chapter 5. Motion Detection	149
5.1 Introduction	150
5.1.1 Local and Widefield Configuration	151
5.1.2 Classification of Detection Schemes	152
5.1.3 Criteria for a Bidirectional Detector	153
5.2 Biological Processing Schemes	154
5.2.1 Reichardt's Delay and Compare Detector	154
5.2.2 Barlow and Levick's Inhibitory Mechanism	160
5.2.3 Directionally Selective Local Inhibitory Motion Detector	165
5.2.4 Horridge's Template Model	167
5.2.5 Gradient Schemes	172
5.3 Rudimentary Noise Analysis	174
5.4 Summary	174
Chapter 6. Stochastic Motion Processing	177
6.1 Schemes Investigated	178
6.2 Architecture of Detectors	178
6.2.1 Template Model Configurations	179
6.3 Effectiveness of Additive Noise	181
6.3.1 Individual Widefield Characteristics	182
6.3.2 Parallel Network of Detectors	187
6.3.3 Detector Linearity	187
6.3.4 Suboptimal Parameter Settings	188
6.3.5 Distributed Threshold Settings	192
6.3.6 Oversampling	193
6.4 Limitations of SR in Motion Detection	194
6.5 Summary	195
Chapter 7. Conclusions and Directions	197
7.1 Summary	198
7.2 Future Directions	200
7.3 Closing Comments	202

Appendix A. Parrondo's Games Analysis	203
A.1 Constraints for the Generalised modulo M Game	203
A.2 Calculating the Equilibrium Distribution	207
A.3 Rate of Winning	209
A.4 Quasi-Birth-and-Death Analysis	210
Appendix B. SR Calculations	213
B.1 Random Number Generators	213
B.2 Numerically Integrating SDEs	215
B.3 AMI Expressions for Gaussian Signals and Noise	217
B.3.1 Preliminary Calculations	217
B.3.2 Simplifying the average mutual information	220
B.4 Shannon's Channel Capacity	221
Appendix C. Lowpass & Highpass Filters	223
Bibliography	229
Glossary	247
Index	249
Resume	253

Abstract

Traditionally, noise has been an enemy of engineers – its unpredictability limiting the performance of many types of systems. It places an upper limit on the data transfer rate of communications systems, plays havoc with sensors and is difficult to avoid. The traditional approach to improve performance is to minimise the effects of noise and build redundancies into a system. However, within the last 10-20 years, new fields have emerged that use noise to enhance a system's performance. Two such fields that we will examine are stochastic resonance and Brownian ratchets.

This thesis explores these noise enhancing phenomena with the intent of applying them to motion detection schemes. We consider the classical Reichardt detector, an inhibitory scheme based on a shunting inhibition neurons, and the Horridge template model. All the schemes have a biological background; this is to take advantage of the robust and simple algorithms produced by many years of evolution.

The first half of this thesis looks at Brownian ratchets – devices that use noise to generate directed motion of Brownian particles. The concept of Brownian ratchets dates back to the discovery of Brownian motion, where a number of schemes were proposed to harness the random fluctuations. One of the proposed rectification devices, referred to as the “ratchet and pawl machine,” showed that when an external energy source is available, thermal noise is able to generate directed motion. It was this device that eventually inspired the development of Brownian ratchets.

Many types of Brownian ratchets have since been constructed and the theory developed. Based on this theory, we expound a set of mathematical games that mimic the behaviour of the Brownian ratchet. Called “Parrondo's games,” they are two gambling games that have the counter-intuitive property of losing when played individually, but when played in a random order produce a winning expectation. The games are explained heuristically and via a mathematical analysis, which show this result is indeed the case. Though they do not have any direct link to utilising noise for motion detection, the concept was an instructive example of using noise or randomness to enhance the performance of a system and led us to look more closely at stochastic resonance.

The second half of this thesis deals with stochastic resonance and its application to motion detection. Stochastic resonance (SR) is the phenomenon where the performance

of a nonlinear system is optimised with the addition of noise. The noise improves the coherence between the input and output signals. We consider a variety of nonlinear systems and explain the different measures, including the channel capacity of SR. Different types of SR are also described, in particular suprathreshold SR, which has potential for the application of motion detection. However, whatever the measure or type of SR used, a system operating optimally cannot benefit from the addition of noise.

The motion detection schemes used are simple, biologically inspired algorithms that specifically detect edges. By using only local information, the detector can be implemented as a smart sensor. We present a quantitative noise analysis of the three schemes, and employ the framework used in SR. The results show that for an optimally configured scheme, noise does not help. However, by utilising the network structure of suprathreshold SR, the performance is improved. By constraining the threshold in one of the schemes, the performance of the detector is improved by adding noise.

In many real systems, there are trade-offs, and unavoidable internal noise can cause suboptimal operation of a system. It is in these cases where stochastic resonance becomes beneficial. We know that biological sensory neurons are very noisy and utilise SR, which they have achieved after millions of years of evolution. Thus, not all noise is bad.

Statement of Originality

This work contains no material that has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent to this copy of the thesis, when deposited in the University Library, being available for loan and photocopying.

Signed _____

Date

5/11/01

Acknowledgments

Many people have made undertaking this Ph.D. an enjoyable and memorable experience. This includes fellow students, colleagues and friends.

Firstly, I would like to extend my thanks to my supervisor, Dr Derek Abbott, who has supported my work throughout my Ph.D. He has been particularly helpful in bringing to light ideas from a range of diverse fields that I may not have otherwise been aware of. He has also made a concerted effort to encourage the publishing of my work in journals, and the presentation of my work at conferences.

I would also like to thank Assoc. Prof. Peter Taylor (Adelaide University, Australia), Prof. Juan Parrondo (Universidad Complutense, Spain) and Assoc. Prof. Charles Pearce (Adelaide University, Australia). They have offered much assistance in formalising the theory of Parrondo's games and continuously dealing with related questions. Thanks also to Prof. Laszlo Kish (Uppsala University, now Texas A & M University, USA) for bringing to light the channel capacity metric of stochastic resonance, and for our continuous and useful communication since.

I would also like to thank Dr Nigel Stocks (Warwick University, UK), Dr Al Nuttall (NUWC, Rhode Island), Dr Ross Dawe (DSTO, Australia), Prof. Charles Doering (University Michigan, USA), Prof. Ron Pyke (University Washington, USA), Prof. Riccardo Mannella (Università di Pisa, Italy) and Assoc. Prof. Bruce Davis (Adelaide University, Australia) for useful communication during my Ph.D. Thanks to Andrew Allison for the use of his first-rate personal library and communications regarding the state-space representation of Parrondo's games.

During my time at the EEE department, I have had the pleasure of progressively moving through three different laboratories, thus there are a number of fellow students whom I have enjoyed time with. Typical activities of tea & coffee breaks (though I do not drink either!), Friday afternoon drinks and the occasional hit of golf have left many good memories. Special thanks to Brad Ferguson, Sam Mickan, Mark McDonnell and Jade Duncan for proof reading this thesis and offering many comments and corrections.

Acknowledgments are due to Leonard Hall and Dr Ross Dawe for making electronic copies of their figures available.

Acknowledgments

Finally, most thanks go to my long time partner, now wife, Jade Harmer (née Duncan). Jade's devotion and encouragement have allowed me to finish this thesis within a finite time. Also many thanks to my parents for their support and for teaching me that though persistence and hard work, success will prevail.

– Greg Harmer

“So little done, so much to do!”

– John Rhodes

Publications

- GOODFELLOW-D. C, HARMER-G. P AND ABBOTT-D (1998). Millimeter-wave collision avoidance sensors: future directions, *Proceedings of SPIE*, Vol. 3525, pp. 352–362.
- HARMER-G. P, ABBOTT-D AND TAYLOR-P. G (2000a). The paradox of Parrondo's games, *Proceedings of the Royal Society A*, **456**(1994), pp. 247–260.
- HARMER-G. P, ABBOTT-D, TAYLOR-P. G AND PARRONDO-J. M. R (2000b). Parrondo's paradoxical games and the discrete Brownian ratchet, in D. Abbott and L. B. Kish (eds.), *Second International Conference on Unsolved Problems of Noise and Fluctuations*, Vol. 511, American Institute of Physics, Adelaide, Australia, pp. 189–200.
- HARMER-G. P, ABBOTT-D, TAYLOR-P. G AND PARRONDO-J. M. R (2001). Parrondo's games and Brownian ratchets, *Chaos*, **11**(3), pp. 705–714.
- HARMER-G. P, ABBOTT-D, TAYLOR-P. G, PEARCE-C. E. M AND PARRONDO-J. M. R (2000c). Information entropy and Parrondo's discrete-time ratchet, in D. S. Broomhead, E. Luchinskaya, P. McClintock and T. Mullin (eds.), *Stochastic and Chaotic Dynamics in the Lakes*, Vol. 502, American Institute of Physics, Melville, NY, USA, pp. 544–549.
- HARMER-G. P AND ABBOTT-D (1999a). Parrondo's paradox, *Statistical Science*, **14**(2), pp. 206–213.
- HARMER-G. P AND ABBOTT-D (1999b). Parrondo's paradox: losing strategies cooperate to win, *Nature*, **402**, p. 864.
- HARMER-G. P AND ABBOTT-D (1999c). Simulation of circuits demonstrating stochastic resonance, *Proceedings of SPIE*, Vol. 3893, pp. 195–203. Also appeared in Harmer and Abbott (2000a).
- HARMER-G. P AND ABBOTT-D (2000a). Simulation of circuits demonstrating stochastic resonance, *Microelectronics Journal*, **31**(7), pp. 553–559.
- HARMER-G. P AND ABBOTT-D (2000b). Smart sensor motion detection schemes in a noisy environment, *Proceedings of SPIE*, Vol. 4236, pp. 25–35.

Publications

- HARMER-G. P AND ABBOTT-D (2001). Motion detection and stochastic resonance in noisy environments, *Microelectronics Journal*, 32(12), pp. 959–967.
- HARMER-G. P AND ABBOTT-D (2002). Computer simulations of Parrondo’s capital and history dependent games, in A. Nowack (ed.), *Annals of the International Society on Dynamic Games*, Birkhäuser. Submitted.
- HARMER-G. P, DAVIS-B. R AND ABBOTT-D (2002). A review of stochastic resonance: Circuits and measurement, *IEEE Transactions on Instrumentation and Measurement*. Accepted.
- KISH-L. B, HARMER-G. P AND ABBOTT-D (2001a). How fast can a neuron transfer information: Bandwidth is the real issue, *16th International Conference on Noise and 1/f Fluctuations*, World Scientific Publishing, Gainesville, Florida, USA. In press.
- KISH-L. B, HARMER-G. P AND ABBOTT-D (2001b). Information transfer rate of neurons: stochastic resonance of Shannon’s information channel capacity, *Fluctuation and Noise Letters*, 1(1), pp. L13–L19.
- PARRONDO-J. M. R, HARMER-G. P AND ABBOTT-D (2000). New paradoxical games based on Brownian ratchets, *Physical Review Letters*, 85(24), pp. 5226–5229.

List of Figures

Figure		Page
1.1	Layout of thesis	2
1.2	Range and time-to-collision diagrams	5
1.3	A section of the electromagnetic spectrum	7
1.4	Passive millimetre wave (PMMW) imaging at Shafter airfield, California	9
1.5	Average atmospheric absorption of millimetre waves	10
1.6	The original Rotman lens used for 15 and 37 GHz	11
1.7	The current version of the beamforming lens	12
<hr/>		
2.1	Trajectory of a Brownian particle	18
2.2	The ratchet and pawl machine	22
2.3	Maxwell's demon creating a temperature differential	25
2.4	The trapdoor demon creating a pressure differential	26
2.5	Ratchet with retractable teeth	28
2.6	Linearised ratchet with Brownian particles	29
2.7	The flashing Brownian ratchet	30
2.8	Molecular motors using ATP hydrolysis	33
2.9	Flashing ratchet system to separate different sized particles	34
<hr/>		
3.1	Construction of Parrondo's games	40
3.2	Progress when playing Parrondo's games deterministically and stochastically	42
3.3	Transient properties of capital in game B	42
3.4	Probability density functions of the games	46
3.5	Standard deviations of the games	47
3.6	Using different deterministic mixing strategies between the games	48

List of Figures

3.7	Varying the mixing parameter of the randomised games	49
3.8	Sawtooth shape representing the ratchet potential	50
3.9	Ratchet potential formed by game B	51
3.10	The breaking of the individual PDFs by the randomised game	52
3.11	Discrete-time Markov chain representing game A	58
3.12	Discrete-time Markov chain representing game B	58
3.13	Discrete-time Markov chain corresponding to to the modulo game B	59
3.14	Probability space of game B	66
3.15	Gain for all the fair probabilities of game B	67
3.16	Probability space of all the games showing the paradoxical region	68
3.17	The relation between the fairness of the games and the entropy	71
3.18	Entropy and probability space for game A	72
3.19	Entropies for game B	73
3.20	Entropies for the randomised game	74
3.21	Construction of Parrondo's history dependent games	75
3.22	Progress when playing Parrondo's history-dependent games	76
3.23	Expected returns from deterministic sequences of the history-dependent games	77
3.24	Probability density function of the history dependent games	77
3.25	Discrete-time chain for game B'	78
3.26	Discrete-time Markov chain for game B' using previous histories	79
3.27	Probability space of the history-dependent games	82
3.28	State space and transformation of Parrondo's games	85
3.29	Fractals in Parrondo's games	86
3.30	The probability space of Parrondo's games in p_1 and p_2	87
<hr/>		
4.1	Stochastic resonator	94
4.2	Example input signals and response to a STR	95
4.3	Power spectral density and the signal-to-noise ratio from Figure 4.2	96
4.4	Improvement due to SR for different signal to threshold distances	97

4.5	Crayfish mechanoreceptors and the SNR in the presence of noise	99
4.6	Visual perception of SR in image dithering	101
4.7	Dithering in a 1 bit ADC with uniform and Gaussian noise	102
4.8	Additive dither block diagram	103
4.9	Probability density distributions and ROC curves	105
4.10	ROC curves showing enhanced signal detection	107
4.11	Sea trials of ISHTAR showing improved signal detection	108
4.12	Level crossing circuit and characteristic	111
4.13	Examples of thresholding schemes	112
4.14	Typical structure of sending and receiving neurons	113
4.15	Response of the integrate-and-fire neuron model	114
4.16	Response of the FitzHugh-Nagumo neuron model	115
4.17	The quartic double well potential	117
4.18	Schmitt trigger circuit and transfer characteristic	118
4.19	S-type circuit diagram and transfer characteristics	119
4.20	Time series and power spectral densities of coloured noise	123
4.21	The stochastic resonator in the frequency domain	125
4.22	Residence time distributions for the double well potential	128
4.23	Interspike interval distribution for the FitzHugh-Nagumo neuron model	129
4.24	Correlation coefficients for the FHN neuronal model	131
4.25	Distribution of ensembles forming the correlation coefficients	132
4.26	Correlations of parallel neuron ensembles and output signals	133
4.27	A summing network of N devices	134
4.28	Average mutual information for a network of threshold devices	136
4.29	Setting the optimal thresholds for a Gaussian distributed signal	138
4.30	Transmitted information for suboptimal settings of thresholds	139
4.31	The two components of the channel capacity	142
4.32	Channel capacity and signal-to-noise ratio	143
4.33	Noise induced linearisation for a sawtooth and sinusoidal wave	145
4.34	Phase space of a particle in the DWP	146

List of Figures

4.35	Phase space of the FHN neuronal model	147
<hr/>		
5.1	Local and widefield detectors	151
5.2	Classification hierarchy of motion detection schemes	152
5.3	Preferred and null directions of a correlation detector	155
5.4	The Reichardt motion detector (RMD) consisting of two EMDs back-to-back	156
5.5	Response of the Reichardt detector to a step input	157
5.6	The bandpass filter and response to step edges	159
5.7	Two methods for directionally sensitive motion detectors	160
5.8	Feedforward and feedback local inhibitory motion detectors (LIMD)	161
5.9	Response of the LIMDs to a step input	162
5.10	Steady state response of the feedforward LIMD	163
5.11	Adaptive properties of the feedforward LIMD	164
5.12	Adaptive bandpass filter architecture and response	165
5.13	Asymmetrical LIMD and responses to a step edge	166
5.14	The DSLIMD in the widefield configuration	168
5.15	Response of the DSLIMD to all possible edges	169
5.16	Horridge template model in a widefield configuration	169
5.17	Look-up table of the templates	170
5.18	Response of the template model to moving objects	171
5.19	Spacing of templates with the threshold placement	171
5.20	Gradient operations applied to a test image	173
5.21	Gradient operations for the template model	174
5.22	Simple noise comparison between several motion detection schemes	175
<hr/>		
6.1	Network of widefield detectors	178
6.2	Network configurations for the template model	179
6.3	Simple noise analysis of template configurations from Figure 6.2	180

6.4	Test patterns for widefield detectors	182
6.5	Correlation coefficient for Reichardt detector	183
6.6	Correlation coefficient for DSLIMD	184
6.7	Correlation coefficient for Template Model	185
6.8	Correlations for a network of Reichardt and DSLIMDs	187
6.9	Linearity of the detectors against contrast	188
6.10	SR in the template model	189
6.11	SSR in the template model	190
6.12	SR and SSR in the template model	191
6.13	Correlations for the template model under changing conditions	192
6.14	Distributed threshold settings in the template model	193
6.15	Oversampling in the template detector	194
—————		
A.1	Portion of a DTMC	204
—————		
B.1	A decreasing symmetric density function	215
B.2	Relationships between variables for the threshold network	220
—————		
C.1	Simple RC circuits for lowpass and highpass filters	223
C.2	Magnitude responses of lowpass, bandpass and highpass filters	227

List of Tables

Table	Page
3.1 The relationship between quantities in Parrondo's games and the Brownian ratchet	54
5.1 Requirements for a directionally selective motion detector	153

Introduction and Motivation

THE aim of this research is to investigate the signal processing of noisy images and to devise a reliable and robust scheme to detect motion in these images. The few solutions that are available are generally provided by the computer science community and tend to opt for full image processing, which requires considerable hardware and cost.

One driving force behind this thesis is the move from acquiring images in the visible spectrum to the millimetre wave region. Millimetre waves offer substantial advantages in outdoor applications and the antenna arrays are small and can be fabricated inexpensively. However, the challenge is that passive millimetre wave detection is generally noisy.

In this thesis, we bring together two seemingly unrelated fields of signal processing: the constructive role of noise (via Brownian ratchets and stochastic resonance), and motion detection. To minimise the computational needs of motion detection, the schemes investigated are inspired by biological visual systems — in particular, insect vision.

1.1 Thesis Layout

Figure 1.1 is a diagrammatic view of the thesis layout showing the dependencies between the chapters. The grouped or overlapped boxes imply the material is similar in nature.

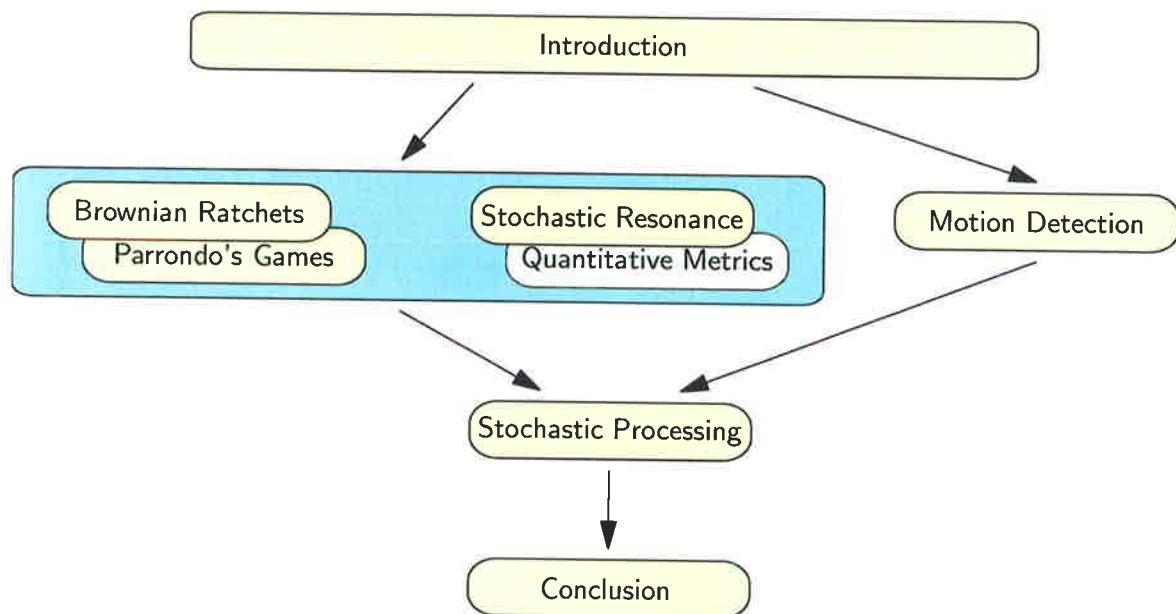


Figure 1.1. Layout of thesis.

In Chapter 1, (this chapter), the introduction, motivation and background are given, along with the current state of knowledge in the field and the key contributions provided in this thesis.

In Chapter 2 we introduce noise in terms of Brownian motion. Strategies on how one could harness work from these random movements of particles are discussed, which lead to some interesting apparent paradoxes involving the second law of thermodynamics. One of these ideas, the ratchet and pawl machine serves as the basis for forming a transport mechanism called the Brownian ratchet, which is the main point of the chapter.

Chapter 3 closely follows Chapter 2 as a new variation of a flashing Brownian ratchet called Parrondo's games are described. The games take a mathematical form that differs from the flashing ratchet as they are discrete in both time and space. Even experienced mathematicians find the counter-intuitive result of Parrondo's games paradoxical in the first instance. The paradoxical results are due to the games' constructive use of noise. Several variations and properties of the games are discussed. Nevertheless,

since the games are based on the Brownian ratchet, the apparent paradox can be explained mathematically – a simple derivation is given in Chapter 3, while a more strict version is given in Appendix A.

The constructive use of noise via Brownian ratchets is discussed in Chapter 2, whilst in Chapter 4 stochastic resonance is described, which also constructively utilises noise. Chapter 4 deals with different stochastic resonant systems and the metrics used to quantify their performance. Several different types of stochastic resonance are discussed, in particular suprathreshold stochastic resonance, which becomes useful in Chapter 6.

Chapter 5 describes several motion detection schemes that have been inspired by biological models based on insect vision and motion sensitive neurons. The advantage being there is no traditional full image processing, only edge detection that can be performed locally and implemented as a smart sensor.

Chapter 6 amalgamates the motion detection schemes of Chapter 5 and the stochastic resonance phenomena of Chapter 4 to investigate the processing of noisy images. It is shown that while the addition of noise generally does not improve performance, in some circumstances it can assist. This is when the system is suboptimal, as shown in Chapter 4, and using the network approach used for suprathreshold stochastic resonance.

The final conclusions and directions are summarised in Chapter 7.

1.2 Introduction

Detection of movement is one of the most basic tasks performed by the visual system. This elementary task is vital for the survival of animals; avoiding obstacles, chasing prey, escaping from predators, and simply navigating within the environment. One of the first tasks performed, even by the simplest animals is to process receptor outputs by selective movement detection units. The signals are then processed in a variety of ways, including directly driving reflex muscles, to more difficult tasks of target tracking and object recognition in complex visual systems.

Many schemes have been developed to artificially reproduce motion detection. A common approach is to model schemes based on biological sources such as insect vision – insects contain far fewer neurons than humans and possess a relatively simple neurophysiology, yet demonstrate remarkable abilities in motion detection. A conceptual

1.3 Background

understanding of motion detection is important as it provides an insight into the dynamics of biological motion detectors, which are reliable and fault tolerant.

The one unavoidable step in artificial motion detection schemes is converting the signal into a workable form. The signal can be from whatever part of the spectrum, and processing can be in whatever technology is appropriate to provide a signal in useable form. When the signal is from light (i.e. the visible part of the spectrum) it is converted to an electrical signal via a photodetector, which is usually realised in CCD or CMOS technology. This technology is relatively straight forward and produces signals with high signal-to-noise ratios. However, utilising other parts of the spectrum can lead to signals that have very low signal-to-noise ratios.

The detection of motion in a fluctuating environment is a central problem in vision research. The aim of this work is to investigate biologically inspired detection schemes in the presence of noise, particularly utilising the concepts of stochastic resonance.

1.3 Background

In the early 1990s an artificial insect vision project dubbed the “bugeye project” was initiated at Adelaide University to develop a *smart sensor* capable of motion detection. A smart sensor simply refers to a device that has detectors and processing circuitry integrated on a single chip. The project was proposed to implement elementary discrimination mechanisms to measure the direction of motion.

The approach taken was to model the smart sensor on the early visual stage of insects, in particular the template model that was devised by Professor Adrian Horridge [Horridge and Sobey, 1991; Horridge, 1990]. Insects contain only a fraction of the neurons of humans, yet their visual manoeuvrability skills are unsurpassed. To quote Horridge’s paper on insect vision [Horridge, 1992],

“The lesson for robot vision is obvious: considering how underdeveloped is our technology for visual processing, almost anything we can learn from insect vision is likely to be of use.”

This detection scheme and others are discussed in more detail in Chapter 5, but we will follow the progress of the template model as it sets the scene for the work undertaken in this thesis.

The driving application of the sensor was a blind spot detector to be placed on the rear vision mirrors of cars, where it could also perform secondary tasks such as white line tracking (lane tracking for smart cruise control), and braking distance warning (calculating the time-to-impact). A smart sensor for this application has several advantages over conventional CCD camera systems [Abbott et al., 1994]. A CCD system typically processes all the pixels of an image which leads to vast amounts of data that need to be processed. This means it is bulky and has limited success for mobile and real time applications. For smart sensor systems, most of the processing is performed in parallel, on chip and a cheap microcontroller is sufficient for making decisions on the sensor output. This makes it compact and cheap. Lastly, since the photoreceptor has a logarithmic response, no focussing is necessary. Hence, no iris is required and a gradient index (GRIN) lens can be simply glued onto the chip.

Several versions of the bug-eye chip have been implemented [Moini et al., 1997], each focussing on specific issues. Without going into detail, the bug-eye chip successfully provided a directionally selective output that indicated the position of moving edges.

The separation between the photoreceptors corresponds to one degree of visual angle once the light has been focussed through the cylindrical GRIN lens. Using this bit of information, the output can be used to calculate quantities such as range, velocity and time to impact.




Figure 1.2. Range and time-to-collision diagrams. The detector is moving forward with velocity v relative to the object, which is at a distance of R from the detector. The left diagram shows an object that is offset to the velocity of the detector, and the diagram on the right shows an object that is looming towards the detector.

1.3 Background

For a relative velocity v between the detector and an object at θ degrees to the direction of motion, as shown in Figure 1.2, the range is given by

$$R = \frac{v \sin \theta}{\dot{\theta}} \quad \text{or} \quad \dot{\theta} = \frac{v \sin \theta}{R},$$

where $\dot{\theta}$ is the angular velocity of the object. Since each pixel represents a fixed angle, the quantities θ and $\dot{\theta}$ are easily determined. This relation is also a known result for determining the *blur zone*, the maximum angular velocity at which features can be made out from an aircraft [Whiteside and Samuel, 1970].

The range for a looming object, which has expanding edges is given by

$$R = \frac{v \sin 2\theta}{2\dot{\theta}},$$

where θ is half the angle subtended by the looming object as shown in Figure 1.2. When calculating the time-to-impact of a looming object, the velocities of the object and detector are irrelevant and the time-to-impact is

$$T_{\text{impact}} = \frac{\sin 2\theta}{2\dot{\theta}},$$

which simplifies to $\theta/\dot{\theta}$ for small θ . The fact that we do not require any velocity may not be intuitive at first, as illustrated in the fictional story *The Black Cloud* [Hoyle, 1957, p. 19] – though the mathematical reasoning is correct [Lee, 1976; Lee, 1980]. The character Dave Weichart shows by setting $\theta = d/R$ [using the small angle approximation], differentiating with respect to time to give

$$\frac{d\theta}{dt} = -\frac{d}{R^2} \frac{dR}{dt},$$

then using the relations $v = -dR/dt$ and $v = R/T$ gives $d\theta/dt = d/RT$, which simplifies to $T = \theta/\dot{\theta}$.

1.3.1 Motion Detection

In order to use the range and time-to-impact formulas, a suitable motion detection algorithm is required. We are required to detect the edges in order to find the angular displacement θ , and hence the angular velocity $\dot{\theta}$ of an object.

Many of the biological schemes use intensity based delay and compare algorithms. The easiest and perhaps the simplest is Reichardt's correlation scheme. This simply correlates two spatially separated channels, one of which is delayed. An alternative method,

in principle, that uses an inhibitory mechanism, is Barlow and Levick's scheme. The scheme used for the bugeye project is the template model, which has a mixed biological and engineering heritage, but nevertheless uses spatiotemporal relations. This scheme is advantageous as it is easy to implement and has some desirable properties that allow the templates to be tracked, to allow easy estimation of θ .

A common property of the aforementioned schemes is that they only utilise local processing, thus can be implemented as a smart sensor. These schemes are discussed in detail in Chapter 5.

1.4 Motivation

The bugeye chips were successful in providing a motion detection smart sensor, but their applicability as a blind spot detector is limited. What happens when it is raining, dark or a thick fog has settled in? The answer is they do not work. Consequently an alternative *all-weather* detection sensor is required, though the algorithms that implement the collision avoidance can remain essentially unchanged. This highlights the shift from technology centred solutions to technology independent solutions [Abbott et al., 1995].

One solution is to use the millimetre wave band instead of the visible range of the spectrum. The frequencies, wavelength and bands of the electromagnetic spectrum are shown in Figure 1.3.




Figure 1.3. A section of the electromagnetic spectrum. The spectrum of interest is the millimetre wavelength band between 30 and 100 GHz. Note that the boundaries overlap and the ranges shown are guides only.

1.4 Motivation

The proposed sensor utilises *radiometry*, which is the science of using passive detection techniques to detect background radiation. Unlike radar, which transmits a signal and then receives the backscattered radiation to measure reflection from objects, a radiometer merely receives naturally occurring radiation. This radiation is the black body radiation that is emitted by all objects.

By appropriately selecting the radiometer parameters, such as wavelength, polarization and viewing angle, relations between the magnitude of energy received and specific parameters of the scene are found. The main parameter of interest is the *brightness temperature* T_B , which represents the intensity of the radiation emitted by the scene under observation. The brightness temperature may vary from zero Kelvin (for a non-emitting medium) to a maximum value equal to the physical temperature T_0 of the scene (for a perfect blackbody emitter).

The question then remains, why use millimetre waves and not some other wavelength? The answer is simple; in the frequency range 30-100 GHz, the waves are able to propagate over aerosol obscured paths. This includes aerosols such as dust, smoke, rain and fog. Figure 1.4 shows a comparison of the Shafter airfield in California using visible and millimetre wavelengths. In the visible spectrum on a foggy day no detail can be made out, but using passive millimetre wave (PMMW) imaging the fog becomes almost transparent revealing most of the detail seen in the visible spectrum. This makes it suitable for take-off and landing systems for an all-weather aircraft, television imaging through fog, rain and smoke, enhanced vision for fire fighters in smoky conditions and for use in land, air and sea navigational aids [Lawrence, 2001].

Though millimetre waves are able to propagate through rain and fog, the main source of their attenuation is water vapour (H_2O) and oxygen (O_2) in the atmosphere. Thus, there is a limit to how much fog or rain can be tolerated, or alternatively the distance that can be viewed with millimetre waves. The attenuation as a function of frequency is shown in Figure 1.5 – the absorption peaks caused by H_2O and O_2 are indicated. The windows at 15 and 37 GHz were used as initial testing frequencies with the aim to move to 94 GHz for the final system. By increasing the operating frequency, the physical size of the detector proportionally decreases.




Figure 1.4. Passive millimetre wave (PMMW) imaging at Shafter airfield, California. The top two pictures are taken in the visible spectrum, the top picture on a clear day and the middle on a foggy day with 50 m visibility. The bottom is taken using PMMWs at 94 GHz on the same foggy day as the middle picture. Reproduced from Lawrence (2001).

1.4.1 Millimetre Wave Antenna Array

While millimetre wave (mm-wave) imaging is not new, it remains bulky and expensive in both computation and cost. The proposed mm-wave sensors use the fact that only low resolution images are required for motion detection, as opposed to full image processing. This, and the ability to monolithically fabricate the passive sensors, make the system smaller, simpler, cheaper and easier to fabricate [Abbott and Parfitt, 1997].

Preliminary designs by Abbott and Parfitt (1997) have been tested at microwave frequencies, which show that millimetre wave sensitivity is suitable for a test system at 37 GHz. The lens design is shown in Figure 1.6. The folded dipole antenna (top) provides phase corrected inputs to the Rotman beamformer lens (centre). The structures just below the dipoles are impedance matching “balun” networks. The outputs of the Rotman lens are then fed via impedance matching networks to Schottky barrier diodes (bottom), to provide a direct conversion receiver. The dc level of the outputs provide a measure of the radiometric intensity of the beams. The idea of the flared “batman

1.4 Motivation

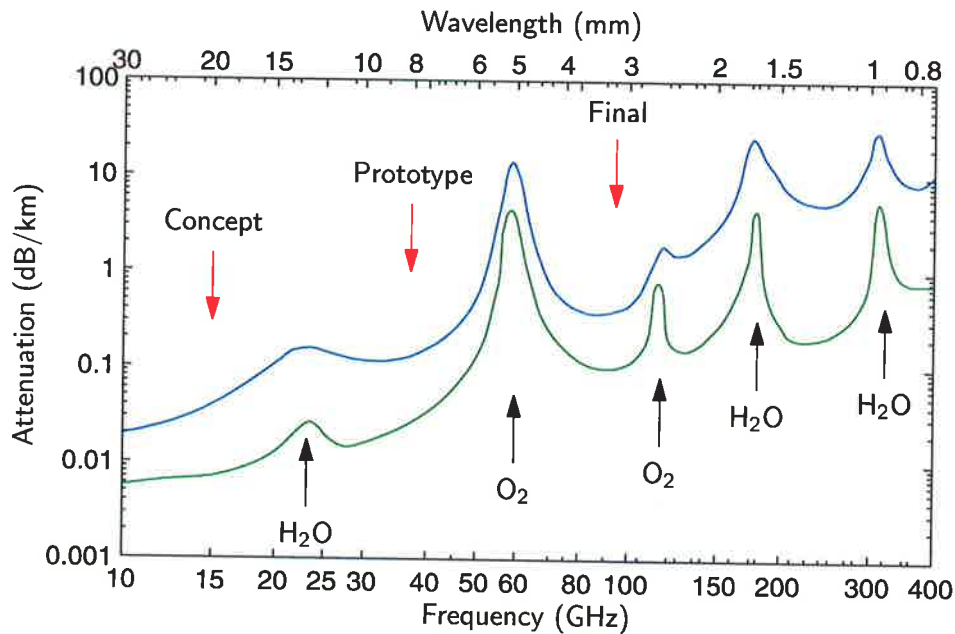


Figure 1.5. Average atmospheric absorption of millimetre waves. The peaks are caused mainly by absorption of H_2O and O_2 as indicated by the upward pointing arrows. The pointed arrows are the frequency windows that have desirable propagation properties. The top line is at sea level with $T = 20^\circ\text{C}$ and $\text{H}_2\text{O} = 7.5\text{ g/m}^2$ and the bottom line is at 4 km above sea level with $T = 20^\circ\text{C}$ and $\text{H}_2\text{O} = 1\text{ g/m}^2$. Redrawn from Currie and Brown (1987).

wings” each side of the lens is to absorb excess radiation in unwanted directions by preventing reflections inside the lens. This lowers the sidelobe levels in the antenna gain pattern.

The dipole antennas absorb radiation in the plane of the lens which allows them to be stacked on top of each other to produce 2-D coverage, which is a main advantage of this design.

However, several problems were identified with the Rotman lens in Figure 1.6. Firstly, the bent phase delayed input lines lead to the electrical lengths being difficult to predict accurately, and the corners act as points of reflections within the line. Secondly, the complexity of designing the absorption wings proves problematic [Hall, 2001].

An alternative design shown in Figure 1.7 eliminates these problems. The main differences between this lens and the Rotman lens are the absence of the batman wings and the use of patch antennas. The patch antennas radiate perpendicularly to the area of the patch (in the z direction in Figure 1.7), and are able to be moved closer or further away from the lens (in the x direction) without disturbing the performance in the yz




Figure 1.6. The original Rotman lens used for 15 and 37 GHz. The dipole antennas at the top absorb radiation in the plane of the page where it is beamformed to the output ports at the bottom. The “batman wings” attenuate excess sidelobe radiation. Reproduced from Abbott and Parfitt (1997).

plane. This allows the antenna array to conform to a surface, which at higher frequencies is only a small area. In place of the batman wings are dummy ports that do the same job of decreasing sidelobe levels, but more predictably.

There is no strong relation between the number of patch antennas and output ports. The number of patches controls the beamwidth of the central beam, the more patches the smaller the beamwidth. The number of output ports determines the angular resolution, the more ports the higher the resolution.

The expectation is to use monolithic fabrication which should make an eleven port array achievable with a total width of 41 mm at 37 GHz and a total width of 16 mm at 94 GHz.

Although millimetre waves appear to offer the solution to the requirement for an all-weather sensor, there is one major problem. The source of millimetre waves is from blackbody radiation and hence the process is inherently noisy. It is for this reason that we need to consider a signal processing approach that is robust against noise, or even use it to our advantage. A description of two such schemes follows.

1.5 Stochastic Processing




Figure 1.7. The current version of the beamforming lens. Compared with the lens shown in Figure 1.6, the patch antennas absorb radiation perpendicular to the plane of the page and the “batman wings” are replaced by dummy ports. Reproduced from Hall (2001).

1.5 Stochastic Processing

In traditional signal processing much of the emphasis in improving a communication system is to minimise the amount of noise in the system.

However, there are some techniques where noise can be added and used constructively by a system to improve its performance. They can be classified under the two broad categories of Brownian ratchets and stochastic resonance. These are explained in detail in Chapters 2 and 4 so only a brief introduction of the two processes follows.

1.5.1 Brownian Ratchets

To introduce Brownian ratchets, consider the following problem. Ever wondered why when you shake a bag of mixed nuts of varying sizes the big ones always rise to the top even though they may be more dense? The answer is in the shaking [Rosato et al., 1987]. When the nuts are shaken, small air gaps form between all the nuts which allow the smaller ones to wedge in. The net effect is the larger ones move to the top while the smaller ones move to the bottom. This is an example of a ratcheting effect, the shaking – which is essentially noise – allows certain nuts to move in one direction and the others in the opposite direction. Ratchets rely on asymmetry for operation and in this example, one asymmetry is due to the fact that gravity only acts on one direction. This experiment will not work in free space.

Another mechanism is that shaking particles produces convection currents in the nuts, the large nuts travel easily along the rising currents, but the downward currents trap the large nuts against the container walls [Kestenbaum, 1997]. Both viewpoints illustrate a ratchet mechanism.

If the dimensions of the problem are reduced to microscopic size, small particles or molecules can be used in place of the nuts where shaking is provided by the thermal fluctuations of the particles, i.e. Brownian motion. Hence the term *Brownian ratchets*. Of course, an appropriate field needs to be used that exerts a force on the particles.

The key ingredient for Brownian ratchets, or any other type of ratchet to function is asymmetry in the system. Also, if (mechanical) work is to be done by the system then energy must be supplied. Most types of familiar ratchets involve asymmetry in the spatial dimension, but other types of ratchets also exist. An example is given in Chapter 2 (on page 32) where asymmetry exists between the phases of two events.

1.5.2 Stochastic Resonance

To explain stochastic resonance (SR) simply, consider a nonlinear “black box” system where the output infers some information about the input signal. In the absence of additive noise to the input no information is present at the output. However, adding noise to the input signal can produce some information at the output. Upon adding further noise the output starts to produce less information about the input signal. The characteristic of SR is that there is a nonzero value of noise intensity that transmits the maximum amount of information to the output about the input signal, that is, the coherence between the input and output is maximal.

Simple linear systems will not produce this type of characteristic. They are well defined by linear response theory with noise at the input proportionally transmitted to the output. For SR to occur, the system under investigation must be nonlinear. Initially the belief was that the system had to be either bistable or contain a threshold of some description, though recent results demonstrate that not all of these features are necessary [Bezrukov, 1998]. As interesting as this probing is to find the simplest set of underlying principles, it is not the goal of this work. Suffice to say SR can be generally viewed as a class of interactions between disorder (noise) and asymmetry (a nonlinear system) [Abbott, 2001].

1.6 Current State of Knowledge

Another issue regarding SR is how to suitably quantify it. Which metrics should be used? Following the evolution of SR from narrow to broadband input signals, the metrics have matured from the simple signal-to-noise ratio to information theoretic measures. The subject of metrics is discussed in more detail in Sec. 4.4.

1.6 Current State of Knowledge

Much of the literature that deals with the processing of noisy images originates from the computer vision community where full image processing is usually undertaken [Ullman, 1981; Poggio et al., 1985]. Thus there is little literature published that considers the motion detection schemes discussed here and the influence of noise. A large portion of the literature on biological visual systems concentrates on the physiological (neuroscience) aspects, which are concerned with how well the models mimic their real biological counter part and how accurate they are. The following paragraphs describe work that has been previously undertaken with respect to the processing of noisy images.

Work by Koivunen (1992), describes a noise insensitive motion detection algorithm for television signals. The drawback is that it uses a 5-point processing window that may require several passes, hence it is suspected that this scheme is reasonably computationally intensive.

Keirstead and Huberman (1986) present a method to accurately detect moving spots in one dimension in the presence of noise. It uses a digital detection scheme with arrays of identical processors that are connected to their nearest neighbours, the number depending on the velocity to be detected. Each processing unit is incremented if motion is detected in the appropriate direction, then the outputs are sampled and compared to a threshold value after a fixed number of time steps.

The algorithm was generalised by Ceccatto and Schreckenberg (1989) to accommodate arbitrary motion with velocities in the continuous range in any direction. However, it still remains a digital system that detects the presence of moving spots, which is outside the scope of this thesis.

Sezgin et al. (1995) describe a visual target tracking algorithm using motion-energy and template location based approaches. A morphological filter is used to remove the registration noise. However, due to necessary filtering, high speed vision hardware is

required. They show the motion-energy method is the fastest, but poor against noise, whereas the correlation coefficient method gives the best results in noisy environments.

Recently, Harrison and Koch (2000) have implemented a robust VLSI Reichardt motion sensor. It is a typical Reichardt motion sensor with integrated photodetectors. The robustness against noise is investigated by moving fixed pattern gratings at a constant velocity for spatial and temporal noise.

The noise analysis undertaken in Harrison and Koch (2000) plots the response to the input signal-to-noise ratio while the contrast is adjusted to keep the dynamic range of the input signal constant. Though they still get direction discrimination at high noise levels, it is difficult to quantitatively compare the results at different noise intensities due to the large dependence of the response on the signal contrast.

1.7 Original Contributions

During the course of work undertaken in this thesis a broad range of topics were investigated as is the nature of noise, and contributions were made in several areas. The key contributions are pointed out below.

- When exploring Brownian ratchets for a possible adaptation, a variation referred to as Parrondo's games were explored [Harmer and Abbott, 1999a; Harmer and Abbott, 1999b]. This is a discrete time and space version of the flashing Brownian ratchet. Much of the initial groundwork and problem formulation that was established during this work has resulted in dedicated sessions at conferences and follow-up papers by others [Pearce, 2000b; Moraal, 2000; Costa et al., 2002; Fuh and Yeh, 2001; Meyer and Blumer, 2001].
- Exploring the other avenue for constructive use of noise, stochastic resonance, progress was made in quantifying measures. Historically, the metrics used progressed from signal-to-noise ratios, input-output cross-correlation coefficients and information content. In some instances these are not appropriate and we describe stochastic resonance in terms of the channel capacity for the first time [Kish et al., 2001].
- A quantitative noise analysis was performed on three motion detection schemes, namely the Reichardt detector, the inhibitory (shunting neuron) detector and the Horridge template model using stochastic resonance: [Harmer and Abbott, 2001]

1.7 Original Contributions

- This gave comparative results between the long time established Reichardt detector and the relatively new inhibitory detector.
- Exploration of the template model in a generalised network configuration, to apply the concept of suprathreshold stochastic resonance, was carried out.

Chapter 2

Brownian Ratchets

WE start by describing Brownian motion, the thermal fluctuations of microscopic particles. In the early days many ideas and thought experiments were proposed to enable the rectification of these thermal fluctuations. The flawed implication being that this would provide a source of energy in a thermal equilibrium environment — a perpetual motion machine.

In 1885, before the formalisation of Brownian motion, Maxwell hypothesised a thought experiment termed Maxwell's demon to illustrate the limitations of the second law of thermodynamics. At the time, the flaws in the aforementioned scenarios puzzled scientists. Then, in 1912, the Polish physicist Smoluchowski expounded a ratchet and pawl machine on a molecular scale that could harness the thermal fluctuations of gas molecules in nonequilibrium conditions, and explained it in thermodynamic terms.

From the second law, rectification in equilibrium can never occur and energy must be supplied to achieve any useful work. Moreover, this process can never occur without loss — consider a fridge or combustion engine for example. Similarly the Brownian ratchet consumes energy from a source in return for directed motion of Brownian particles, at an efficiency less than unity.

2.1 Brownian Motion

If you suspend a small particle in solution, viewing it at the appropriate time scale will reveal that it moves erratically and unpredictably — this is Brownian motion. Although this random motion was first reported by the Dutch physician Jan Ingenhousz in 1785 by placing powdered charcoal on the surface of alcohol [Klafter et al., 1996], it was the Scottish botanist Robert Brown who first carried out extensive investigations leading up to 1828 [Abbott, 2001]. He observed the completely irregular motion of pollen and spores suspended in water. When it was first observed, it was not clear how such a motion could be possible. This was prior to knowledge of the existence of atoms.

It was initially believed by Brown that the ‘molecules’ of plants retained their vitality long after the plant’s death [Lavenda, 1985]. Brown ruled out this explanation by studying a drop of water that had been in a piece of quartz millions of years ago. He correctly attributed the movements of the particles to a physical phenomenon rather than a biological one.

It is now well known that the random motion of particles is due to constant bombardment of smaller ‘invisible’ particles, in this case, water molecules. If the resolution of the viewing device is increased, the random path appears even more jagged (see Figure 2.1), and the smaller previously invisible particles are observed to undergo Brownian motion themselves.

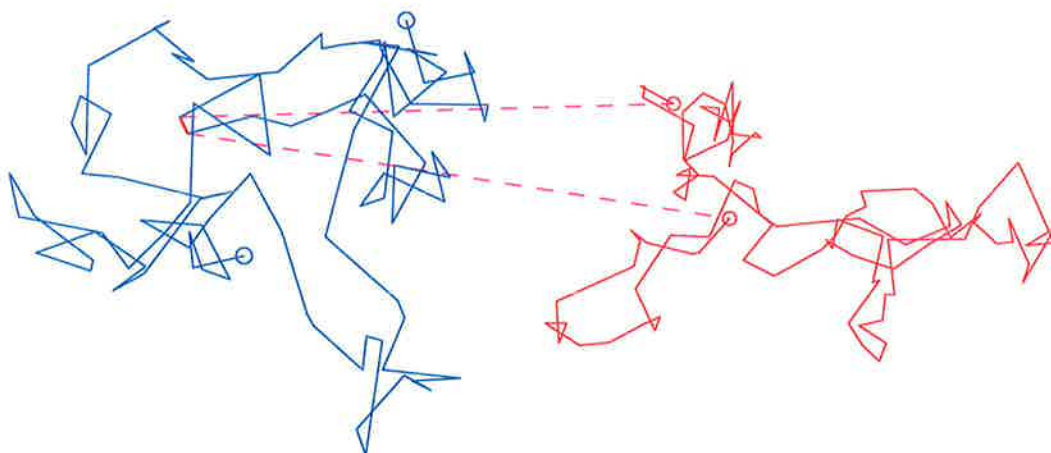


Figure 2.1. Trajectory of a Brownian particle. The path of a particle undergoing Brownian motion. The zoomed trajectory on the right is one of the straight line segments, which exhibits further Brownian motion on a smaller time scale.

It was not until around 1905 that the first ground breaking description of Brownian motion was given by the famous Albert Einstein [Einstein, 1905], although he was unaware of Brown's work at the time. Independently Smoluchowski arrived at an equivalent description [von Smoluchowski, 1906]. Einstein determined a formula for the mean square displacement of a Brownian particle in terms of Avogadro's number. This led to Jean Perrin winning the 1926 Nobel Prize for determining this number [Klafter et al., 1996]. It is seldom known though that Louis Batchelier independently derived several mathematical properties of Brownian motion in his 1900 Ph.D. thesis on stock market fluctuations [Shlesinger et al., 1999].

Two years later the French physicist Paul Langevin independently derived the same result for the mean squared displacement of a Brownian particle. In Langevin's 1908 paper [Lemons and Gythiel, 1997, English translation], which is considered simpler than Einstein's, he divided the forces that a Brownian particle experiences into the micro and macroscopic world. A macroscopic particle experiences a viscous drag force over a large time scale, and is unaffected by molecular bombardments. A smaller microscopic particle experiences the rapid fluctuating forces and is unaffected by any drag force. The total force is equal to the mass of the particle multiplied by the acceleration caused by the two components. This is Langevin's " $F = ma$ " of stochastic physics that is now referred to as the Langevin equation [Lemons and Gythiel, 1997].

The fluctuations forming Brownian motion depend on the number and intensity of collisions between the particles in the suspension fluid, which in turn depends on temperature. Hence, the fluctuations are often called *thermal noise*. As this thermal noise is always present (in any material above 0 K) it is tempting to try and harness it. There have been numerous efforts to rectify this energy, which will be briefly discussed in the following section.

2.2 Rectifying Brownian Motion

Since the emergence of Brownian motion, or thermal noise, many attempts have been made to harness random fluctuations. If one could harness this energy when a system is in equilibrium, one would have energy for free! Unfortunately, this can never be the case as it violates the second law of thermodynamics.

However, this fact has not hindered the many questions, devices, circuits and thought experiments to rectify thermal fluctuations [Brillouin, 1950; Marek, 1959; McFee, 1971;

2.2 Rectifying Brownian Motion

Sokolov, 1998; Abbott et al., 1996]. In fact, people still propose ideas on how to defeat the second law. Most of the flaws in the reasoning are due to incorrect interpretation and understanding, or the problem falling between the boundary of two or more disciplines, and leading to an incorrect analysis. Nevertheless, many of the thought experiments have provided an excellent theoretical launching pad to investigate the implications of the second law [Astumian and Moss, 1998].

2.2.1 Laws of Thermodynamics

The thermodynamic laws consist of purely macroscopic statements that make no reference to the microscopic properties of a system. The laws consist of a number of statements and are given below [Reif, 1985, p. 122].

The first statement, referred to as the Zeroth law was introduced after the other laws, but it is thought to be more fundamental. It simply relates thermodynamic systems together.

Zeroth law: If two systems are in thermal equilibrium with a third system, they must be in thermal equilibrium with each other.

The next statement is essentially the conservation of energy for thermodynamics.

First law: An equilibrium macrostate of a system can be characterised by a quantity \bar{E} (called *internal energy*), which has the property that for an isolated system, $\bar{E} = \text{constant}$.

If the system is allowed to interact and thus goes from one macrostate to another, the resulting change in \bar{E} can be written in the form $\bar{E} = -W + Q$, where W is the macroscopic work done by the system, and Q is the heat absorbed by the system.

The next statement, which is the most relevant in the context of this chapter introduces the entropy, S .

Second law: An equilibrium macrostate of a system can be characterised by a quantity S (called *entropy*), which has the properties that:

- (i) In any process in which a thermally *isolated* system goes from one macrostate to another, the entropy tends to increase, i.e. $\Delta S \geq 0$.

- (ii) If the system is not isolated and undergoes a quasi-static infinitesimal process in which it absorbs heat dQ , then $dS = dQ/T$, where T is a quantity characteristic of the macrostate of the system, the *absolute temperature*.

Finally, the last statement.

Third law: The entropy S of a system has the limiting property that as $T \rightarrow 0_+$, then $S \rightarrow S_0$ where S_0 is a constant independent of all parameters of a particular system.

The three quantities (\bar{E} , S and T) have been introduced for each macroscopic state of the system. We will now look at the implications of these laws in regard to the rectification of Brownian motion.

Implications of the Laws

The first law requires that energy is to come from some well defined source, that is, not from a reservoir. If the energy were to come from a reservoir, its entropy would decrease and the entropy of the universe would decrease in the cyclic process, hence violate the second law [Leff and Rex, 1990]. The word *cyclic* is important as it is possible to convert heat completely into work in a noncyclic process – the isothermal expansion of gas for example [Tipler, 1991, p. 567]. This also implies that we can never gain 100% efficiency from a heat engine.

The net result is that if the entropy decreases in one part of a system, then this must be compensated by an increase in entropy elsewhere in the system. The second law then guarantees that this change will be non-negative. The entropy of the universe is then increased by the same amount as the change in the system – the entropy of the universe always increases [Feynman et al., 1963].

Thus, according to the second law, at equilibrium, the effect of thermal noise is symmetric, even in an anisotropic medium. That is, whatever the construction of the medium, thermal noise exerts no bias.

2.2.2 Ratchet and Pawl Device

The ratchet and pawl machine was originally analysed by Marion von Smoluchowski, where he referred to it as *Zahnrad mit einer Sperrlinke* in German [von Smoluchowski,

2.2 Rectifying Brownian Motion

1912]. The device is considered to be of molecular scale and is shown in Figure 2.2. The aim is to harness the thermal Brownian fluctuations of gas molecules, by a process of rectification.




Figure 2.2. The ratchet and pawl machine. The two thermal baths at T_1 and T_2 are thermodynamically isolated, but are connected mechanically by an axle that has a pulley at its midpoint. The vane in the right box can only rotate in one direction due to the ratchet and pawl in the left box.

The device works in the following manner. Let the temperature of the thermal baths contained in the boxes be equal so $T_1 = T_2 = T$, i.e. they are in equilibrium. The energy of the thermal baths, which is directly related to the temperature, is equal for each bath. Due to the bombardments of gas molecules on the vane, it jiggles back and forth and since the ratchet in the other box is linked mechanically, the wheel will only turn in one direction. Thus the wheel will turn slowly and may even be able to lift some weight. This creates a paradox, the ratchet and pawl will apparently work in perpetual motion at equilibrium, when $T_1 = T_2$.

This is clearly a violation of the second law of thermodynamics. As discussed earlier, thermal noise cannot do any work in equilibrium conditions. The structural forces alone cannot bias Brownian motion as has been suggested with the ratchet and pawl device.

The reason for the existence of the paradox is that it combines two aspects of physics, macro and microscopic analysis. Typically, the casual observer views the vane as a microscopic device (i.e. interacting with molecules), but the ratchet and pawl as a macroscopic device (i.e. not taking into account molecular interactions). One must consider *all* of the components on a microscopic scale. The answer to this paradox was provided by Smoluchowski in 1912 [von Smoluchowski, 1912], and later in *The Feynman Lectures on Physics* [Feynman et al., 1963]. The explanation is rather complex, and it is argued whether Feynman's analysis is completely accurate with regard to the Carnot efficiency – though this detail does not affect us [Parrondo and Español, 1996]. A brief explanation follows, a more detailed transcript can be found in Feynman et al. (1963).

To consider everything in microscopic detail, more attention must be directed to the spring loaded pawl – the only moving part. A detail that has not been mentioned is that the spring cannot be perfectly elastic. If it were, after falling from the top of a tooth on the ratchet wheel, it would bounce and continue to do so forever. Thus, there must be some type of dampening mechanism that stops the bouncing. This mechanism transfers the energy possessed by the pawl into heat.

The other consideration is that the pawl and wheel are at temperature T_2 and also experience Brownian motion. This motion causes the pawl to lift itself up accidentally and at the same time a tooth may pass underneath it due to the jiggling from the vane. If this occurs, the force of the pawl on the ratchet tooth actually causes the wheel to move backwards until the pawl is at the pit of the tooth. The hotter it gets, the more often this occurs.

This is why the device does not do any net work in perpetual motion. Bombardments on the vane sometimes cause the pawl to lift up and go over a tooth, but also bombardments in the opposite direction occur at a time when the pawl is up. The net result is that there is no motion, only jiggling, but on average, nothing. It is shown that the number of times there is enough energy to turn the wheel forward when the pawl is down is equal to the number of times the wheel turns freely backwards and the pawl is up [Feynman et al., 1963].

An experimental verification of Feynman's theoretical analysis has been presented by Kelly and colleagues [Kelly et al., 1998] using a synthesized organo-molecular ratchet with triptycene as the vane. It is linked to a four ring helicene as the ratchet and pawl [Astumian and Moss, 1998], showing the frequencies of clockwise and anticlockwise rotations to be exactly equal.

2.2 Rectifying Brownian Motion

Now, consider the device in nonequilibrium. By creating a thermal gradient between the two baths, typically we set T_2 less than T_1 . Since the wheel is cold, the number of accidental fluctuations of the pawl will be small, while the hotter vane will have energy often enough to rotate and cause the pawl to be lifted. Thus, the device will turn in one direction and lift the weight as intended. When the ratchet and pawl is in nonequilibrium it is functioning as an engine.

When the ratchet and pawl engine is *running*, every time a molecule of sufficient energy hits the vane it causes it to rotate and the energy is used to lift the pawl over the tooth and lift the weight. The net result is that the molecule is moving slower and the vane has stopped rotating, and since $\bar{E} \propto T_2 \propto v_{\text{rms}}$ it is cooler. Conversely, in the right box the ratchet is rotated by one tooth which the pawl fell from. The dampening mechanism then attenuates the bouncing of the pawl by producing heat. Consequently the temperature of T_1 increases. Given that these two processes are continuously occurring, the system will eventually reach equilibrium where there is no net motion. That is, heat is transferred mechanically from one box to the other. Therefore, we have to supply the engine with energy in order to maintain the temperature differential. By supplying this energy and keeping the system in nonequilibrium, the wheel will continue to lift weight [Abbott et al., 2000]. Moreover, reversing the temperature differential causes the ratchet to work in the opposite direction [Feynman et al., 1963; Astumian and Moss, 1998].

2.2.3 Maxwell's Demon

We present an overview of Maxwell's demon. An in depth analysis is beyond the necessary scope here and can be found in Leff and Rex (1990), which contains a comprehensive commentary and reprints of significant articles from the time of the demon's conception to the current day.

The demon originated in 1871 from the Scottish physicist James Clerk Maxwell in his book, *Theory of Heat* [Maxwell, 1885].

“...if we conceive a being whose faculties are so sharpened that he can follow every molecule in its course, such a being, whose attributes are still as essentially finite as our own, would be able to do what is at present impossible to us. For we have seen that the molecules in a vessel full of air at uniform temperature are moving with velocities by no means uniform, though the mean velocity of any great

number of them, arbitrarily selected, is almost exactly uniform. Now let us suppose that such a vessel is divided into two portions, A and B, by a division in which there is a small hole, and that a being, who can see the individual molecules, opens and closes this hole, so as to allow only the swifter molecules to pass from A to B, and only the slower ones to pass from B to A. He will thus, without expenditure of work, raise the temperature of B and lower that of A, in contradiction to the second law of thermodynamics."

The *being* became known as Maxwell's demon, due to the implications it has on the second law and is accordingly defined by *Webster's Third International Dictionary*, "[after J. C. Maxwell, its hypothecator]: a hypothetical being of intelligence but molecular in size imagined to illustrate limitations of the second law of thermodynamics." This demon attempts to create a temperature difference between the compartments, shown in Figure 2.3.

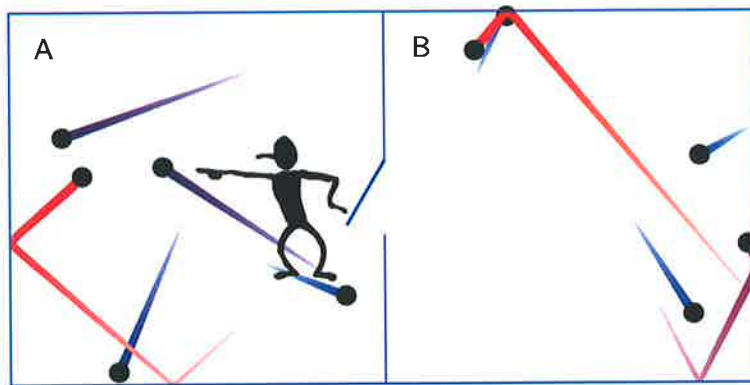


Figure 2.3. Maxwell's demon creating a temperature differential. The demon observes individual molecules, allowing fast ones to pass from A to B while allowing slow ones to pass from B to A, thus, creating a temperature difference from thermal equilibrium.

The fact that this construction uses an "intelligent being" created added confusion in the analysis of the devices in the early days. So much that in 1914 Smoluchowski wrote [Bennett, 1987],

"As far as we know today, there is no automatic permanently effectively perpetual-motion machine, in spite of the molecular fluctuations, but such a device might perhaps, function regularly if it were appropriately operated by intelligent beings."

To simplify the problem we can consider a similar proposition, the trapdoor demon. This is another form of Maxwell's demon that creates a pressure difference, but without the use of intelligent beings.

2.2 Rectifying Brownian Motion

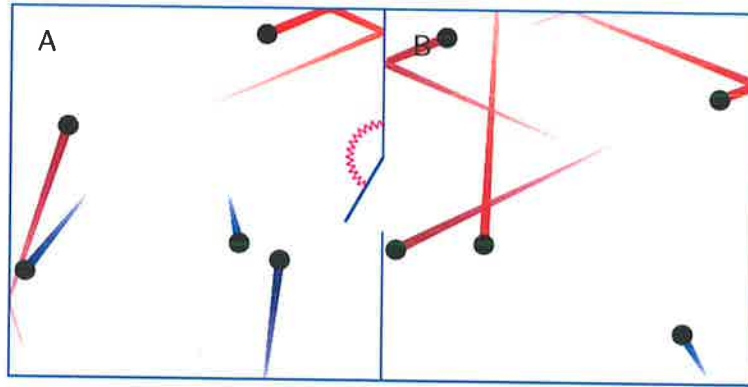


Figure 2.4. The trapdoor demon creating a pressure differential. The demon is in the form of a one way spring loaded trapdoor. As the molecules in B hit the trapdoor, it deflects open and allows the molecules to move into A. But when the opposite occurs, the door does not budge and the molecules remain in A, hence creating a pressure difference.

The trapdoor demon is shown in Figure 2.4. By allowing molecules to move from B to A but preventing them from passing A to B via the use of a one way valve, the particles will be concentrated in A and a partial vacuum created in B. Even if such a device could be built, why would it not work? Smoluchowski pointed out in 1912 that the trapdoor is repeatedly bombarded with gas molecules and accordingly (the spring) heats up [Bennett, 1987]. That is, the door now moves with Brownian motion. This motion is comparable to the molecules since the door is very small and the spring is relatively weak. This is similar to the pawl heating up and fluctuating with Brownian motion in the aforementioned ratchet and pawl device. As the trapdoor fluctuates open it does not act as a one way valve and molecules can move freely back to B. Worse still, it is possible that the door will knock slow molecules back into B upon closing. Thus, the ability of the door to function as a one way valve is exactly zero, and the pressure on average in both components is equal.

Maxwell's original demon can be explained in a similar fashion. The problem is that the demon itself gets warmer due to the act of taking measurements, and after a while it is shaking from Brownian motion so much that it can not determine if the molecules are moving fast or slow, backwards or forwards, hence it does not function [Feynman et al., 1963]. This is not a rigorous explanation which essentially treats the intelligent being as a piece of machinery.

An interesting question to determine is how much energy the demons would consume in their operation? To answer such a question is not trivial, and a great number of papers have been written on the topic. A complementary problem that arises, especially in a thermodynamic context, is what change in entropy does such a device cause?

It was identified by Szilard (1929) that there are three central issues relating to Maxwell's demon: measurement, information and memory [Leff and Rex, 1990]. He further identified the *fundamental amount* of entropy generated by these events as $k_B \ln 2$. These two observations established the underpinnings of information theory and its connections with physics. However, it is unclear from Szilard (1929) whether the cost is from the measuring, remembering or forgetting [Leff and Rex, 1990]. The act of taking the measurement can be done with zero cost [Bennett, 1987]. This was also demonstrated by Szilard by use of the now known Szilard heat engine [Szilard, 1929]. In this engine the expanding gas (an individual molecule) does mechanical work, which is balanced by the energy required to reset one bit of information from two possible values [Magnasco, 1996; Bennett, 1987]. Thus, there is no decrease in entropy, and no perpetual motion or energy. The previously stored measurement needs to be erased to return the system to a standard state – this is where the cost is incurred [Bennett, 1982; Landauer, 1961].

The implications of this work to modern day information theory were quite revolutionary, and one should consult Leff and Rex (1990) for further commentary and references. There are possibly many other types of demons that *prima facie* appear to work, though when correctly analysed do not. Further questions and criticisms are raised in Parrondo (2001) and [Ishioka and fuchikami, 2001], though this is beyond our introductory treatment here.

2.2.4 Utilising Thermal Fluctuations

The previous examples have shown in accordance with the second law, that energy cannot be generated from the random fluctuations of Brownian motion for free. If we do supply energy in some form, it is possible for the device to transfer energy to get mechanical work done.

2.3 Brownian Ratchets

In the next section we consider a device that stems from the concepts of the ratchet and pawl machine. It is essentially a linearised ratchet that contains Brownian particles instead of a pawl, and by supplying energy to the system it is possible to output mechanical work.

2.3 Brownian Ratchets

We have shown, in accordance with the second law, that work cannot be extracted from thermal fluctuations at equilibrium. But, when a system is out of equilibrium, that is, when some external energy is supplied, useful work can be extracted. In the ratchet and pawl system, by keeping a temperature gradient between the two thermal baths ($T_1 > T_2$ say) some weight can be lifted by the wheel, i.e. mechanical work. In fact, this device forms the basis for the idea of Brownian ratchets.

Imagine the ratchet and pawl system, but the teeth are able to move in and out [Bier, 1997], see Figure 2.5, the vane is not drawn but does exist. The energy required to lift the pawl over the teeth is greater than the kinetic energy of the gas molecules in the vanes thermal bath. When the teeth are out (left diagram), the wheel does not move. But when the teeth are in (right diagram) the wheel is free to move with Brownian motion, assuming there is no friction between the wheel and pawl, or it is negligible compared with the molecules kinetic energy. The teeth later emerge and the wheel moves so that the pawl is pushed to the pit. The average net rotation is half a tooth length in the direction shown by the arrow on the left ratchet wheel in Figure 2.5.

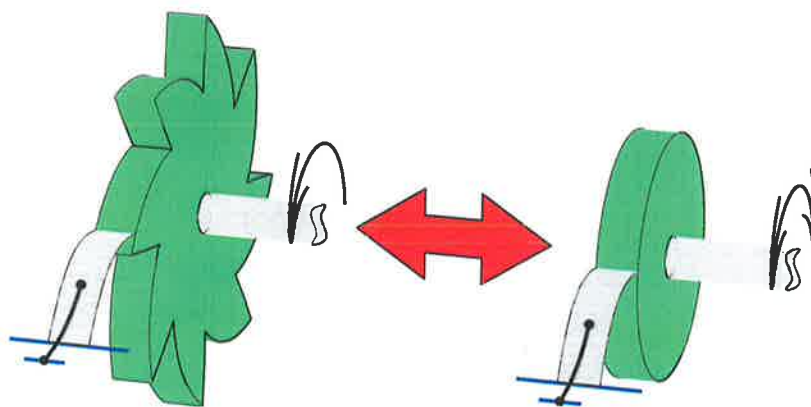


Figure 2.5. Ratchet with retractable teeth. The ratchet teeth move in and out, represented by switching between the two above diagrams. The vane to the right is not shown.

A linearised version of the ratchet wheel can easily be created by suspending the pawl mechanism from a fixed point and allowing a sawtooth ratchet to move underneath it. Alternating between the sawtooth and a flat surface produces movement relative to the pawl. Conversely, if the sawteeth are fixed there will be directed movement of the pawl mechanism. For a Brownian ratchet, the pawl is simply replaced by Brownian particles as shown in Figure 2.6.




Figure 2.6. Linearised ratchet with Brownian particles. The Brownian particles take the place of the pawl, which are able to move along the ratchet.

By alternating the ratchet on and off, directed motion of the particles is achieved – the details are given in the following section.

Referring to the sawtooth shape and flat surface as potentials hints at the flexibility available in the system. This means that the particles and potential can be realised in any form as long as they interact with each other. So, instead of having gravity as the force that makes the particles fall into the pits of the teeth by virtue of their mass, it is possible and common to use charged particles in an electrostatic ratchet potential. The particles interact with the ratchet formed by the field by virtue of their charge. Examples of this are given in Sec. 2.3.3.

So far we have only alluded to one mechanism that allows Brownian ratchets to function, namely, the *flashing ratchet*. Another common mechanism, referred to as the *rocked ratchets* switches between $U = U_{\text{saw}} \pm x|F_{\text{max}}|$ where U_{saw} is the sawtooth potential and F_{max} is the average macroscopic slope of the potential [Astumian, 1997; Bier, 1997; Bartussek et al., 1994]. A similar mechanism is the *fluctuating barrier ratchet* that switches between two in-phase sawtooth potentials of different magnitudes [Astumian and Bier, 1994].

2.3.1 Flashing Ratchets

The flashing Brownian ratchet works in the following manner. Consider a system where there exists two one-dimensional potentials, U_{saw} and U_{flat} , as shown in Figure 2.7. The asymmetry of the triangular potential U_{on} is determined by α , where $0 \leq \alpha \leq 1$. Having $\alpha = 1/2$ creates a symmetric potential, for other values the potential is asymmetrical like U_{saw} in Figure 2.7 where $\alpha < 1/2$. The Brownian particles present in the potential diffuse to positions of least energy. In equilibrium, if the potential height is larger than the thermal noise, the particles are localised in a potential minima and their concentration is peaked (Figure 2.7a). Switching the potential

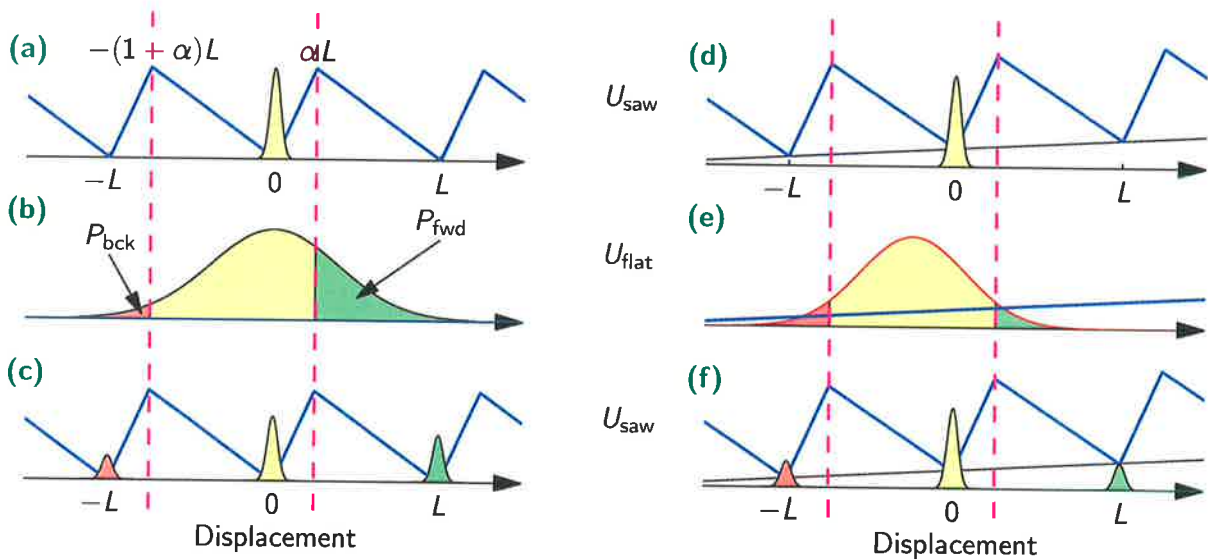


Figure 2.7. The flashing Brownian ratchet. The diagrams on the left, (a)-(c) shows when there is no macroscopic gradient present and the net movement of particles is in the forward direction (defined by arrow). On the right, (d)-(f) introduces an opposing macroscopic gradient to the net flow of particles.

to U_{flat} allows the particles to diffuse freely so that the distribution spreads, centred around the minima (Figure 2.7b). When U_{saw} is switched on again (Figure 2.7c) there is a probability P_{fwd} that some particles are to the right of αL and move forward to the minima located at L . Similarly, there is a probability P_{bck} , and that some particles are to the left of $-(1 - \alpha)L$ and move to the left minima located at $-L$. Since $\alpha < 1/2$, then $P_{\text{fwd}} > P_{\text{bck}}$, and the net motion of the particles is to the right. We can define the probability current as $J = P_{\text{fwd}} - P_{\text{bck}}$ for a particle diffusing forward one step in the potential.

Even if a small macroscopic force is presented against the net movement of the particles, motion against the force is still observed [Hänggi and Bartussek, 1996]. When the macroscopic force and U_{flat} are applied, the particles still diffuse, but also drift backwards. This reduces the effectiveness when U_{saw} is switched on. Figure 2.7d-f shows when the motion to the right due to the flashing ratchet is cancelled by the motion to the left due to the force, then $P_{\text{fwd}} = P_{\text{bck}}$ and $J = 0$.

2.3.2 A Very Brief Review

While the ratchet and pawl machine and Brownian motion were known long ago, the appearance of Brownian ratchets in literature is only quite recent. Adjari and Prost (1992) and Magnasco (1993) formulated the Brownian ratchet mechanism as a possible separation method that is free of a macroscopic driving field.

A flood of papers then proceeded describing different applications of Brownian ratchets. Two basic themes emerged; separating particles and a possible explanation of the motion of proteins and enzymes in biology, both are shown as examples in the next section.

In biology, the flashing Brownian ratchet possibly explains how proteins and enzymes move along muscle fibres, referred to as molecular motors. Svoboda et al. (1993) used silica beads carrying single molecules of the protein kinesin on microtubules to investigate the motion. Their results were consistent with a model formed by Astumian (1997), though other models also explain the ratchet mechanisms [Peskin and Oster, 1995]. Many other papers on molecular motors have been published [Magnasco, 1994; Bier, 1997; Kitamura et al., 1999; Astumian, 2001].

An optical realisation of the ratchet uses optical tweezers to modulate the potential which accommodates a single Brownian particle [Faucheux et al., 1995; Travis, 1995]. By using only a single particle, a tiny polystyrene sphere, the hydrodynamic interactions that occur between particles are eliminated.

More recently, quantum ratchets have been investigated [Reimann et al., 1997; Roncaglia and Tsironis, 1998; Linke et al., 1999].

2.3.3 Examples of Ratchets

General types of macroscopic ratchets occur in many varied areas.

2.3 Brownian Ratchets

Mixed Nuts. This example simply consists of a container of mixed nuts of varying sizes. Upon shaking the container the larger nuts tend to rise to the top even though they may be denser than the smaller ones [Rosato et al., 1987; Kestenbaum, 1997]. This is part of the field of granular physics [Jaeger et al., 1996; Makse et al., 1997; Hayakawa and Hong, 1997].

The reason that the larger nuts rise to the top is that the smaller nuts fill the gaps that are formed when the larger nuts become air-born due to the shaking. This continual wedging action of the smaller nuts is the ratchet mechanism, and is reliant on gravity for it to work. Another viewpoint is that when convection currents are formed, the big nuts get jammed by the sides at the top of the container. [Kestenbaum, 1997]. Either viewpoint nicely illustrates a ratchet.

Self Winding Watches. Some styles of 'windless' wrist watches have a ratcheting mechanism that winds the watch during everyday movements of the hand. This mechanism has also been suggested as an energy source for future 'wearable computers'.

Phase Ratchet. This example employs a type of phase ratchet [Abbott, 2001] based on an old probability brainteaser [Perelman, 1967; Mosteller, 1965, p. 7]. A train station is located between two towns *A* and *B* with trains leaving for each town every hour. Each day the same person walks to the station arriving at a random time and catches the first train that arrives. After a number of days the person finds they have gone to town *A* 20% of the time and town *B* 80% of the time. How?

If one were to see the timetable the answer would be immediately obvious – the two trains do not come at half hour intervals, but are asymmetrically staggered. For example, let the train to town *A* arrive ten minutes past the hour while the train to *B* arrives 22 minutes past the hour. This only leaves a 12 minute window to get to town *A*, which is a probability of 20%. Consequently there is a 48 minute window (or 80% probability) of travelling to town *B*.

Although these ratchets are not based on the flashing ratchet, they all utilise noise or randomness together with an asymmetry in order to operate. The next two examples, based on the flashing ratchet, illustrate biological and artificial realisations.

Molecular Motors

A molecular motor is the term used to describe how proteins move along a biopolymer. The biopolymer can be thought to consist of an array of dipoles, as shown in Figure 2.8. The proteins have a positive charge of $q = +2$ and become neutralised through the hydrolysis reaction. The reaction consists of two stages, the protein binds with ATP (adenosine triphosphate) and then at some time later releases ADP (adenosine diphosphate) and P_i (inorganic phosphate) to regain its charge. The reaction can be summarised by $\text{ATP} \rightleftharpoons \text{ADP} \times P_i$, though the conditions are very unidirectional [Magnasco, 1994]. The nonequilibrium is provided by the biological power plant that removes used ADPs and P_i s and replaces them with ATPs. This continual renewal process along with the hydrolysis reaction causes the molecule to continually switch between $q = 0$ and 2. Though the ratchet potential is always present, its effects are only experienced by the molecule when $q = +2$.




Figure 2.8. Molecular motors using ATP hydrolysis. A sawtooth potential is created by the linear array of cells in the biopolymer. A protein molecule undergoes ATP hydrolysis to change its charge, which has the effect of switching the potential on and off.

Thus, with the aid of Figure 2.8 one can see how the molecular motion functions as a flashing ratchet. Of course, the level of the internal noise must be small enough that the ratchet potential is not insignificant in comparison, but large enough for sufficient diffusion when the molecule's charge is zero.

Note that this is a very simplified version of the mechanics [Bier, 1997]. Biologically, the motor protein is not a pure point charge nor is the biopolymer made up of simple dipoles. They consist of a distribution of charges and dipoles respectively. ATP hydrolysis also involves many subsequent chemical reactions.

Separating Particles

By using the flashing ratchet technique it is possible to separate small and large particles, with a system set up as shown in Figure 2.9 [Astumian, 1997]. A mixture of charged particles in solution are placed at the centre of the inclined plate. The electrodes etched onto the plate provide the electrostatic sawtooth potential, which can be modulated on and off by a square wave voltage generator. The external macroscopic force is provided by gravity due to the plate inclination.




Figure 2.9. Flashing ratchet system to separate different sized particles. By applying a potential difference to the electrodes etched onto the surface a ratchet potential is created. This allows the heavier particles to move down while the smaller particles move up.

By modulating the potential at the correct frequency the small particles can be made to move uphill, while the larger particles move downhill. The smaller particles feel a small effect due to gravity and have a large amount of Brownian motion, thus have an efficient ratchet mechanism. They only move downhill a little when the potential is off, and move uphill more when the potential is switched on. On the other hand the larger particles feel a greater effect due to gravity and have only a small ratchet effect. That is, the smaller particles have more microscopic characteristics whereas the large particles behave more macroscopically.

It is also possible to separate particles without the use of a macroscopic force by using a three state signal to modulate the potentials [Bier and Astumian, 1996; Bier, 1997]. Another useful factor is that the switching frequency influences the net velocity of particles, which is also dependent on physical properties such as size and mass.

The application of separating charged particles according to particular characteristics using Brownian ratchets has sparked possible commercial interest and consequently, patents [Ajdari et al., 1997; Bader et al., 1999; Soane and Soane, 1992]. The method is very suited to sizing and separating DNA fragments, generating DNA fragment length polymorphism patterns, and sequencing DNA [Duke and Austin, 1998; Ertas, 1998].

2.4 Summary

In this chapter, we have described the origins of Brownian motion and several ideas of how to rectify this motion to generate a continuous supply of energy. As highlighted by the second law of thermodynamics, this can never occur at equilibrium and the best we can achieve is to construct some type of engine utilising the Brownian particles.

Engines have been constructed using the flashing Brownian ratchet, which was inspired by an earlier ratchet and pawl machine. We described the mechanism of the flashing ratchet and two of the main areas of interest: a separation mechanism of microscopic particles, and a transport mechanism of proteins in biological systems.

The question now is, can we extend the mechanics of the Brownian ratchet to other systems? The answer is yes, and is described in the next chapter. It is realised that the two potentials used to create the ratchet (U_{flat} and U_{saw}) can be mapped to games of chance to mimic mathematically the mechanics of the Brownian ratchet. The next chapter sets much of the groundwork for this model and considers further adaptations.

Parrondo's Games

INSPIRED by the flashing Brownian ratchet, Parrondo's games present an apparently paradoxical situation. The games can be modelled as coin tossing events. Game *A* uses a single biased coin while game *B* uses two biased coins and has a state dependent rule based on the player's current capital. Playing each of the games individually causes the player to lose. However, a winning expectation is produced when randomly mixing games *A* and *B*. This phenomenon is investigated and mathematically analysed to give explanations on how such a process is possible. The games are expanded to become dependent on other properties rather than the capital of the player. Fractal properties and the ubiquity of Parrondo's games are also mentioned.

3.1 Introduction

The study of probability dates back to the seventeenth century. It arises from games of chance, originating from the ancient game of throwing bones — the forerunners of dice. Strongly associated with probability is gambling, from dice to actuarial tables and risk-benefit analysis, gambling has always been at the forefront of expanding probability theory [Shlesinger, 1996]. This dates back to correspondence between Pascal and Fermat in 1654 when a problem was posed to Pascal by a French gambler.

One of the simplest “games of chance” is the tossing of a biased coin to decide between possible outcomes, usually referred to as ‘heads’ or ‘tails’. Information theorists have long studied these mechanisms [Roche, 1992], which trace back to the work of von Neumann (1951). Using rational biased coins it is possible to generate any n -sided dice or simulate different biased coins, if n is dyadic only one coin is necessary [Feldman et al., 1993; Itoh, 1996; Gargamo and Vaccaro, 1999; Hoeffding and Simons, 1970]. Furthermore, we can make a fair roulette from biased coins [Dijkstra, 1990], and Durrett et al. (1991) discuss making money from fair games.

With Parrondo’s games, which are also constructed with simple biased coins, we go one step further to generate a positively biased outcome from a combination of two negatively biased processes. This type of behaviour is not impossible and exists in other fields. In control theory, the combination of two unstable systems can cause them to become stable [Allison and Abbott, 2001]. In granular flow, drift can occur in a counter intuitive direction, as mentioned in Chapter 2. Lastly, Pinsky and Scheutzow (1992) show that switching between two transient diffusion processes in random media can form a positive recurrent process, which can be thought as a continuous-time version of Parrondo’s games.

The inspiration of Parrondo’s games came from the Spanish physicist Juan M. R. Parrondo. He used the games as a pedagogical illustration of the Brownian ratchet. In one of his fields of interest, Brownian motors, he was able to see a link between mathematical games of chance and the Brownian ratchet. Using the equations for detailed balance, analogous to chemical reaction rates [Onsager, 1931a], enabled Parrondo to devise some working probabilities. However, no material was published, and after communications with Parrondo results were first published by Harmer and Abbott (1999a). The games were named after their creator as “Parrondo’s games”, and the results referred to as “Parrondo’s paradox.”

Some scientists have criticised the use of the term “paradox” — however we use it in the sense of an apparent paradox and this is comparable to existing terminology, such as in “Simpson’s paradox” [Klay and Foulis, 1990; Neufeld, 1995], “Braess paradox” [Kameda et al., 2000; Korilis et al., 1999] and “the renewal paradox” [Pearce, 2000b].

3.2 Construction and Characteristics

The games about to be described are remarkably simple to construct. They only involve either a decision, or the result of a random event. In all cases, the random event can be implemented with biased coins [Harmer and Abbott, 1999b].

The following subsections explain the games and the paradoxical result that occurs when the games are played in certain sequences. This is followed by a description of the characteristics of the games to enable comparisons with Brownian ratchets, described in Chapter 2.

3.2.1 Construction of the Games

The games can be formed using elementary probability rules, that is, we win with a probability p or lose with probability $1 - p$. Such games or processes are well known and can be likened to going on a biased random walk or tossing a biased coin. Since each step of the games consists of a win-loss decision they can be formed by using a set of biased coins.

These games are not typically associated with game theory in the von Neumann sense [Neumann, 1944], the difference being we do not have a choice of what happens or a strategy to play – we are merely observers of a game of chance. This is, however, game theory in the Blackwell sense [Blackwell and Girshick, 1954].

Game A is straight forward and consists of a single decision branch. The probability of winning and losing is p and $1 - p$ respectively.

Game B is a little more complex, as a decision is first required to decide which coin to toss. It is described by the following statement. If the present capital is a multiple of an integer M , then the chance of winning is p_1 , if the capital is not a multiple of M , then the chance of winning is p_2 . The losing probabilities are respectively $1 - p_1$ and $1 - p_2$. Thus, game B uses two coins, the decision of which one to use depends on the value

3.2 Construction and Characteristics

of the current capital. For future reference let the coins be B_1 and B_2 have winning probabilities p_1 and p_2 respectively.

The two games are represented diagrammatically in Figure 3.1, using branching elements to represent decision conditions or win/loss probabilities. The notation (x, y) at the top of the branch gives the probability (i.e. biasing of the coin) or condition for taking the left and right branch respectively.




Figure 3.1. Construction of Parrondo's games. The games are formed using only simple branching elements. Game A consists of winning (W) or losing (L) with respective probabilities p and $1 - p$. Game B involves a decision followed by a coin toss. Both games can be implemented using a total of three biased coins.

If we require control of the three probabilities p , p_1 and p_2 via a single variable, a biasing parameter ϵ can be used to represent a subset of the parameter space with the transformation

$$\begin{aligned}
 p &= 1/2 - \epsilon, \\
 p_1 &= 1/10 - \epsilon \quad \text{and} \\
 p_2 &= 3/4 - \epsilon.
 \end{aligned}
 \tag{3.1}$$

This parameterisation along with $M = 3$ gives Parrondo's original numbers for the games [Harmer and Abbott, 1999a]. For simplicity, most of the simulations and analysis of the games in this chapter use $M = 3$, although it is not difficult in most cases to generalise for larger values of M . Where appropriate, details about trends for larger values of M are given.

3.2.2 Playing the Games

The games are played between two parties where one wins at the expense of the other. To eliminate confusion we will consider playing against a common opponent, the *house* say. It is also assumed that we can have negative amounts of capital or even play with zero capital, in practice this can be avoided by providing enough initial capital to offset the desired maximum amount of loss to play up to.

The games are traditional gambling games where a bet of one unit is made and we either win a unit (plus our original bet) or lose the unit bet, the net result being either an increase or decrease in our capital by one unit for each game that is played.

It can be deduced by a detailed balance (explained later in Sec. 3.3.2) and simulations, that both game *A* and game *B* lose when ϵ is greater than zero. Consider the scenario if we start switching between the two losing games, play two games of *A*, two games of *B*, two of *A*, and so on for example. The result, which is quite counter intuitive, is that we start winning! That is, we can play the two losing games *A* and *B* in such a way as to produce a winning expectation. Furthermore, deciding which game to play next by tossing a fair coin also yields a winning expectation. Figure 3.2 shows the average progress when playing games *A* and *B* individually, switching deterministically and stochastically between them. The switching sequence affects the gain as shown by the different finishing capitals in Figure 3.2.

This apparent paradoxical result immediately raises questions. Firstly, the reason for the word *apparent* is that this is not an actual paradox as we shall see later, it is completely explained mathematically. However, the word paradox does imply a seemingly contradictory situation that may nonetheless be true, and is the case here. We will discuss a few issues regarding the paradox.

3.2.3 Fairness

So far, we have been careful not to mention a game as being fair, only winning or losing. The reason for this is that the behaviour of game *B* differs from game *A* as we are likely to win or lose a small amount depending on the starting capital. If the starting capital is a multiple of M then we will lose a little, or conversely gain a little when the starting capital is not a multiple of M . The deviations from different values of starting capital after 30 games are shown in Figure 3.3.

3.2 Construction and Characteristics

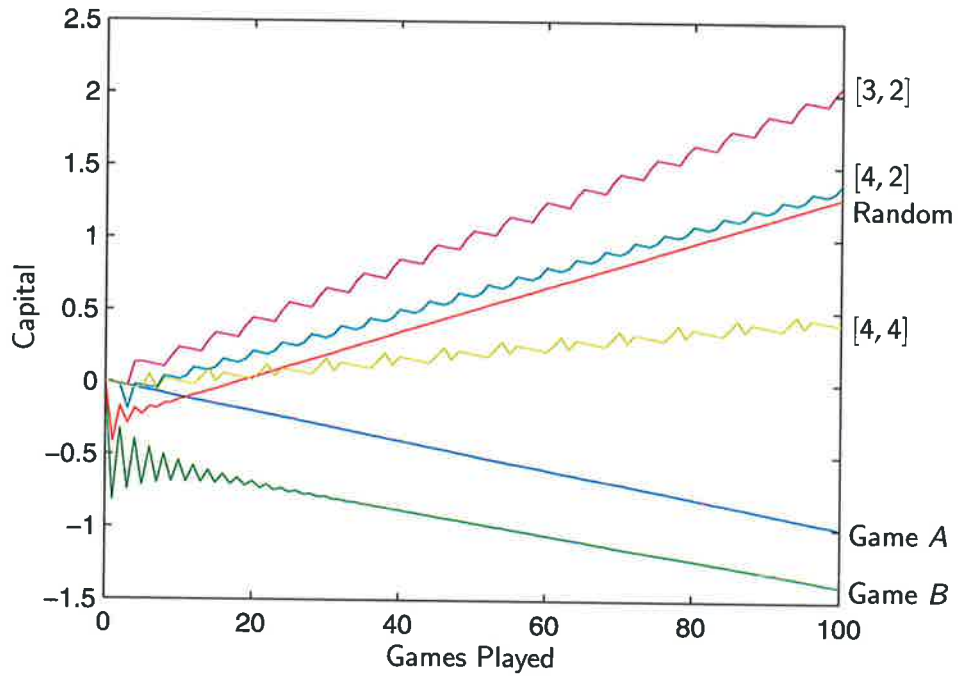


Figure 3.2. Progress when playing Parrondo's games deterministically and stochastically.

The simulation was performed by playing game A a times, game B b times and so on until 100 games were played, averaged over one million trials. The values of a and b are shown by the vectors $[a, b]$.

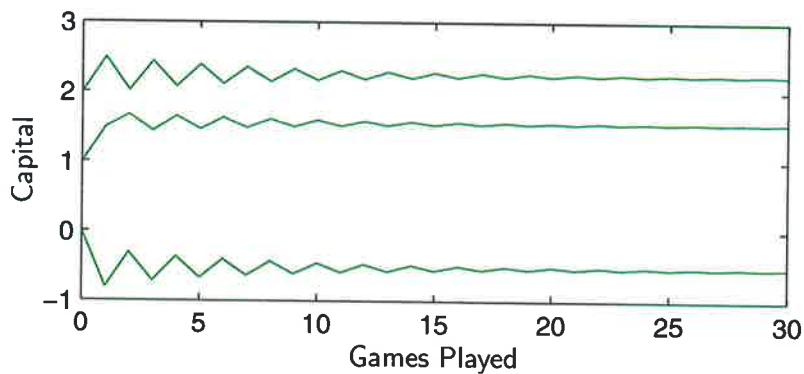


Figure 3.3. Transient properties of capital in game B. Depending on the starting capital in game B, a small amount is either lost or gained. This transient effect quickly dies away after playing a small number of games.

A brief discussion of fairness follows, and a more detailed mathematical formulation relating to Parrondo's games is given by Costa et al. (2002) and Key et al. (2001). Consider a gambler repeatedly playing a game. After the n th game the gambler has capital $X(n)$, or X_n for short. Classically, a fair game satisfies [Doob, 1953, p. 299]

$$E[X_{n+1}|X_0, X_1, \dots, X_n] = X_n \quad (3.2)$$

for $n \in \mathbb{Z}_+$. That is, the game is a *martingale*, where the expected value of capital after playing a game is the same as the present value.

The difficulty with game B , as shown by Figure 3.3, is that when X_0 is a multiple of M , $E[X_1|X_0] < X_0$ and correspondingly when X_0 is not a multiple of M , $E[X_1|X_0] > X_0$. This makes it troublesome to classify game B as either winning, losing or fair [Costa et al., 2002]. Suffice to say it is argued in Costa et al. (2002) that fairness can be defined in terms of drift rates, which were assumed in previous literature [Harmer and Abbott, 1999a; Harmer et al., 2000a]. Thus, if the capital tends to drift toward infinity it classifies as winning ($\epsilon < 0$), or if it drifts toward negative infinity it is classified as losing ($\epsilon > 0$). If there is no drift, then the game is fair ($\epsilon = 0$).

Therefore, using the above criterion, both games A and B are fair when ϵ is set to zero in equations (3.1). This is true of game A because the probabilities of moving up and down in capital are equal for all n . It is also true of game B even though the value of starting capital influences the probability of going up and down for small values of n . As n increases, it is clear from Figure 3.3 there is no gain (i.e. drift) in capital.

Although there is some concern over whether game B is technically fair, it is not too important in the context of the paradoxical nature of the games as they definitely lose when $\epsilon > 0$. This is satisfactory since the only prerequisite we have for the paradox is that games A and B lose when $\epsilon > 0$.

Game B Appears to be Winning?

When investigating game B *prima facie*, it can be mistakenly interpreted as a winning game, thus invalidating the paradoxical result. This is due to taking the wrong path of analysis by considering the games statistically. For example, when $M = 3$ this approach assumes that the chance $X_n \bmod 3$, which is either 0, 1 or 2, occurs with equal probability – a third each. Hence we conclude that coin B_1 is used a third of the time and coin B_2 used the remaining two thirds of the time. Then from the probabilities

3.2 Construction and Characteristics

(3.1) with $\epsilon = 0$ the winning probability is

$$p_{\text{win}} = \frac{1}{3} \cdot \frac{1}{10} + \frac{2}{3} \cdot \frac{3}{4} = \frac{16}{30}, \quad (3.3)$$

which is greater than a half. This implies that the game B is winning, which is incorrect – it is actually fair.

The correct analysis involves employing discrete-time Markov chains, described later in Sec. 3.3.1, which reveals the probability in each state is not a third, but $5/13$, $2/13$ and $6/13$ respectively. Using the correct probabilities for the coins yields the winning probability as

$$p_{\text{win}} = \frac{5}{13} \cdot \frac{1}{10} + \frac{2}{13} \cdot \frac{3}{4} + \frac{6}{13} \cdot \frac{3}{4} = \frac{1}{2}, \quad (3.4)$$

which correctly dictates the game is fair.

Trivial Paradoxical Games

There have been a number of claims made that Parrondo's games are not paradoxical and can be easily replicated. Such claims that artificially create the games usually have two points in common.

- (i) They are constructed with $M = 2$. This is equivalent to deciding what coin to use in game B based on the capital being odd or even. Thus it is possible to create a version of game B that does not use one of its branches when playing alternatively, but used when played individually, which causes the game to lose.
- (ii) They have payoffs that are not simply ± 1 . This makes it even easier to construct games when $M = 2$.

Due to the construction of such games, they usually work with only a few different switching combinations. A good test is to play the games randomly or reverse the order. Typically after a short number of tests a sequence is found that breaks the games.

Using only a single unit payoff it is *impossible* to generate Parrondo's games with $M = 2$. Physically, this creates a symmetric ratchet potential and can never support directed motion. Mathematically, it is possible to show that the constraints for Parrondo's paradox to exist cannot be satisfied, and this is shown in Sec. 3.6.3.

Cashing in at the Casino?

An obvious application of the games would be to head to the closest casino and get rich. Is this possible? The short answer is no.

When the rules of the games are first read, it needs to be realised that the games are not mutually exclusive. This may not be apparent at first inspection. The reason is that the games are linked through the capital. Playing one game changes the capital, which may, or may not affect the probability that is used for the next game. This value of capital can be thought of as memory. The problem is that all the games at the casino are mutually exclusive and definitely do not have any memory — playing one game does not affect other or subsequent games.

An alternative approach is to find three games that model the three coins needed for the probabilities of (3.1). Finding a game that has slightly less than half a chance of winning is easy, as is finding a game that only wins about 1/10 of the time. However, there are no casino games to my knowledge that win 3/4 of the time as required by one of the coins in game *B*.

3.2.4 Distributions and Behaviour

This section looks at the distributions of capital after a number of games have been played. This is to check their behaviour does not rapidly diverge or have some other undesirable effect. The distributions also reinforce the ideas of fairness considered in Sec. 3.2.3.

Several probability density functions (PDFs) have been plotted in Figure 3.4. Since we must win or lose at each game, the PDF will consist of only odd or even values depending on how many games have been played and where we started. To compensate for this misleading effect a centred average has been applied denoted by the over-hat and given by

$$\hat{p}(x, n) = \frac{p(x, n-1) + 2p(x, n) + p(x, n+1)}{4}, \quad (3.5)$$

where x represents the capital and n the number of games played.

The PDFs in Figure 3.4 show games *A*, *B* and randomised with $\epsilon = -0.1, 0$ and 0.1 after 100 games have been played. It is clear that the drifts of the distributions are dependent on the biasing parameter ϵ . When $\epsilon > 0$ the drift is to the left (the game is losing), when $\epsilon < 0$ the drift is to the right (the game is winning), and when $\epsilon = 0$ the

3.2 Construction and Characteristics

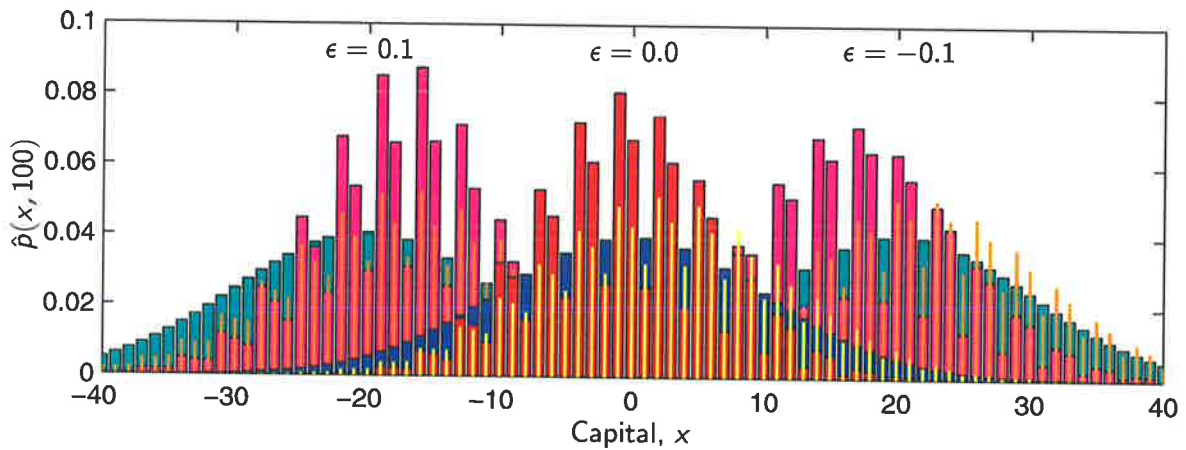


Figure 3.4. Probability density functions of the games. The PDFs of the games using the probabilities in (3.1) with $\epsilon = -0.1, 0$ and 0.1 . Games *A* and *B* are the bars with the smooth and jagged envelopes respectively, the thick lines are the randomised games.

PDF is driftless (the game is fair). As expected, the longer the games are played, the flatter the PDFs become as the standard deviation increases.

As game *A* is well known, we can use its characteristics as a benchmark to judge the other two games. It can be shown that game *A* follows a normal distribution [Harmer and Abbott, 1999a] with the following parameters,

$$\mathcal{N}(n(p - q), 4npq), \quad (3.6)$$

where $q = 1 - p$ is the losing probability. Using the probabilities of (3.1), gives a mean of $-2n\epsilon$, in agreement with the PDFs in Figure 3.4.

From (3.6), the standard deviation for game *A* is $2\sqrt{npq}$, which is proportional to \sqrt{n} . Using sample paths from Figure 3.2 the standard deviations are plotted in Figure 3.5 to show the proportionalities. The linear proportionality with \sqrt{n} for all the games is to be expected since the PDFs resemble normal distributions. The striking result is the standard distribution of game *B* and the randomised game are smaller (i.e. the distributions are tighter) than that of game *A* even though the distributions appear more jagged. Thus, one may conclude that game *B* and the randomised game are as well behaved, if not more, than game *A*, which is considered reasonably well behaved. As will be discussed later in Sec. 3.2.6, game *A* serves to break the “pattern” in the PDF of game *B*, which explains the proportionalities (i.e. slopes) of the standard deviation.




Figure 3.5. Standard deviations of the games. Comparing the standard deviations of the games using 100 000 sample paths from Figure 3.2. The standard deviations of the games are all proportional to \sqrt{n} .

3.2.5 Mixing Strategies

This section explores different strategies for mixing games A and B . Firstly, we consider simple deterministic combinations of games A and B , $\{AABAAB \dots\}$ for example. There are more complex deterministic switching strategies that can reach higher rates of return, such as $\{ABBABABBAB \dots\}$ for example, they will not be considered. Secondly, different stochastic mixes of games A and B are considered.

Figure 3.6 shows the results when mixing deterministically according to $[a, b]$. This notation refers to playing game A a times, game B b times and so on. Thus, for example $[2, 3]$ refers to the sequence $\{AABBBAABBB \dots\}$.

Clearly from Figure 3.6 the larger the switching period the smaller the returns. This is due to playing large stretches of the same game, which as we have seen is bad for the capital. Thus, quickly switching between the games produces the best result – it is the mixing of the games that produces the gain. This is evident by the game $[1, 2]$ having the highest rate of return, followed by other fast switching games.

When $a = 0$ or $b = 0$ we expect zero capital, which is true for $b = 0$ (only playing game A), but not for $a = 0$ (only playing game B). Recall that when only playing game B a small amount is lost or gained. Since we start with zero capital, which is a multiple of M we lose a little, hence the line $a = 0$ is slightly below zero.

3.2 Construction and Characteristics

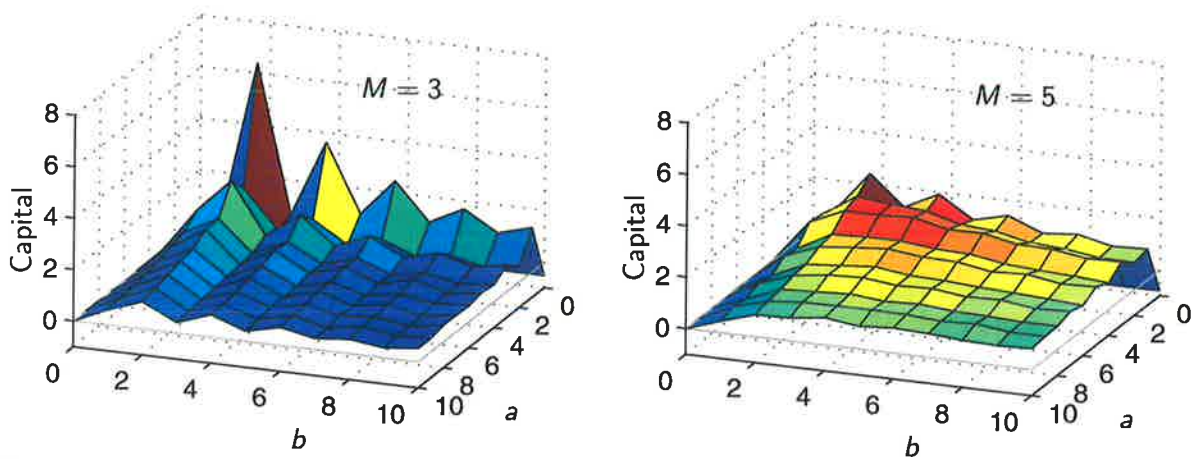


Figure 3.6. Using different deterministic mixing strategies between the games. The value of capital after the 100th game was played using the sequence determined by $[a, b]$. The probabilities of (3.1) with $\epsilon = 0$ were used for $M = 3$. For $M = 5$: $p = 0.5$, $p_1 = 0.1$ and $p_2 = 0.634$.

When mixing the games stochastically we have chosen which game to play next with fixed probability, namely a half. Let us introduce a *mixing parameter* γ with $0 \leq \gamma \leq 1$, which is defined as the probability of playing game A. The probability of playing game B is $1 - \gamma$. Using this parameter we can vary the proportions that each game is played; if $\gamma = 0$ then only game B is played and if $\gamma = 1$ then only game A is played. From Figure 3.2 both games A and B lose individually ($\gamma = 1$ and 0), but there is a mix that wins, thus there must be a optimal value of γ that gives the best rate of return.

Figure 3.7 shows the rate of return (i.e. capital after 100 games) against the mixing parameter γ . Two methods are used to determine the capital. Firstly, the expected capital when directly playing the games labelled 'Direct', and secondly using the slopes from Figure 3.2 (after transients have died away) to extrapolate to 100 games labelled 'Slope'. Methods for finding the slopes are dealt with in Sec. 3.3.4. The difference being that the former method takes into account the initial transient behaviour and is therefore affected by how many games are played, while the latter ignores all the transient effects and is independent of the number of games played. Thus, the difference between the optimal values vanishes as the number of games played approaches infinity. The effect of the biasing parameter is to shift the plots vertically. Note that although the shift is almost linear with ϵ , there are some higher order terms that vary the optimal γ by a small amount.

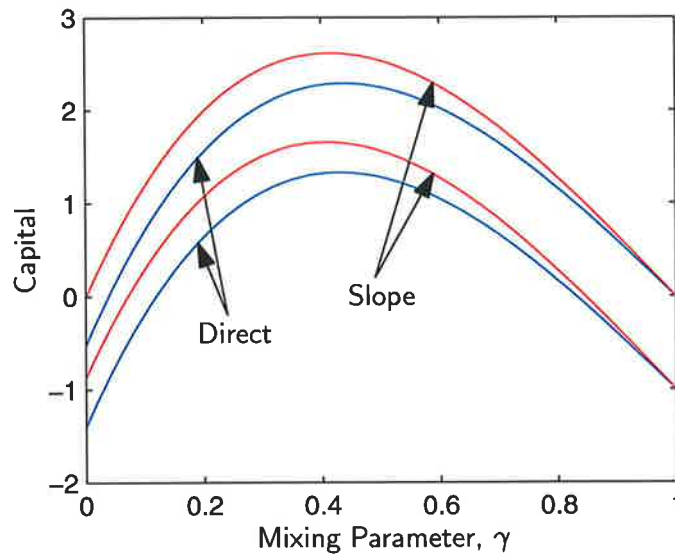


Figure 3.7. Varying the mixing parameter of the randomised games. There is an optimal value of γ that gives the highest capital after playing 100 games. The blue curves indicate directly playing the games and the red curves using the analytically derived slopes. The probabilities of (3.1) were used with $\epsilon = 0$ (top two curves) and $\epsilon = 0.005$ (bottom two curves).

3.2.6 Explanation of the Games

The games were originally formed as an illustration of the Brownian ratchet, so it would be expected that the games could be explained in terms of the Brownian ratchet. This section explains the games intuitively, without the use of mathematics, simply using the characteristics discussed in the previous section.

Formation of the Games

The aim is to convert the physical system of the Brownian ratchet to a mathematical system. One method is to use simple gambling games. The quantities of the Brownian ratchet must be mapped to analogous mathematical variables. The capital can be used in lieu of particles in the ratchet and the probabilities in lieu of energy potentials. Initially this will be adequate to transform the system.

The easy job is to convert the flat potential, which allows diffusion equally in both directions. This is simply a fair game using a single probability, p say.

The sawtooth potential appears tricky at first, but on closer inspection is not too difficult. The sawtooth shape shown in Figure 3.8 can be considered piece-wise linear,

3.2 Construction and Characteristics

consisting of a steep positive slope (segment 1) and a gentle negative slope (segment 2). This suggests that in order to replicate this mathematically two games are required, one for each of the two segments, remembering that this represents a potential and particles will fall downwards. Segment 1 requires a game with a strong losing probability ($1 - p_1$ say) and segment 2 requires a game with a weak winning probability (p_2 say). In the Brownian ratchet the particles know what segment to follow due to their spatial location, there is spatial dependency. This attribute can be mapped to the games by use of an *if* statement. If the capital is within a certain range play the first game (game 1), otherwise play the second game (game 2). This needs to be periodic, which invites the use of a modulo operator. The capital is an integral amount so the modulo can be formed over a range of capital, say M . Then within a certain range play game 1, or if in the remaining part of the range play game 2.




Figure 3.8. Sawtooth shape representing the ratchet potential. A piece-wise linear construction consisting of two segments, which can be replicated using probabilities.

To summarise, if the capital is a multiple of M we play game 1, otherwise we play game 2. This forms the basis of the games. All that is required is to find values for the probabilities to make it work. A quick way is to use a detailed balance, which is explained later in Sec. 3.3.2.

Ratchet Potential of the Games

It should be of no surprise that there exists some type of ratchet potential for game B . This is evident in Figure 3.4 where the sawtooth shape in the PDFs are modulated by the Gaussian function. The Gaussian shape can be removed by superimposing distributions with different starting values. The resulting distribution for games A and B are shown in Figure 3.9. To emphasise the ratchet shape, $M = 7$ was used with probabilities to make the games fair, and the central average $\hat{p}(x, n)$ from (3.5) was employed.

From the distributions in Figure 3.9 it is possible to deduce the shape of the ratchet. Consider an unknown electrostatic potential – electrodes etched onto a substrate and

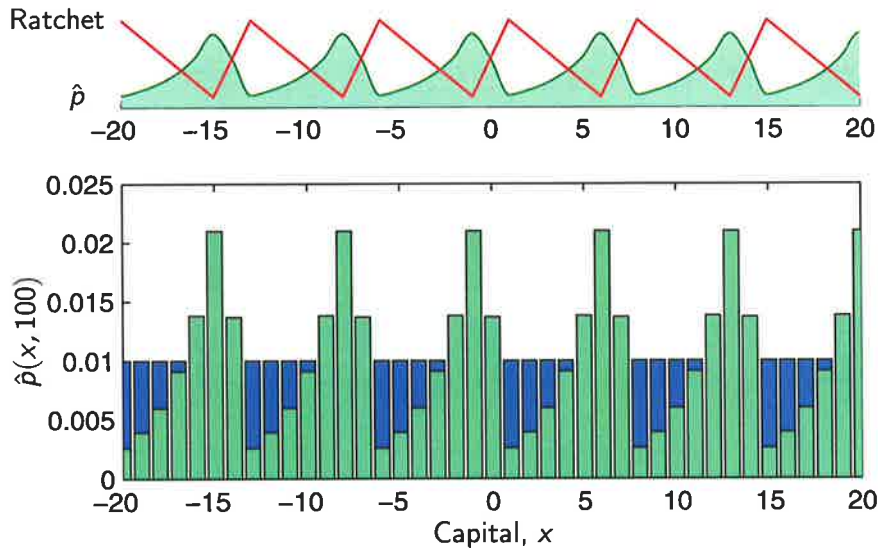


Figure 3.9. Ratchet potential formed by game B. The bottom plot is the superposition of PDFs for games A and B, the flat PDF is for game A and the other game B. The top plot shows the distribution of capital in the ratchet formed by game B. Parameters are $p = 0.5$, $p_1 = 0.075$, $p_2 = 0.6032$ and $M = 7$.

painted over for example. By placing a solution of charged colloidal particles on the surface, the forces cause the particles diffuse to a point of maximum attraction. The density of particles indicates directly the strength of the potential. Similarly in this system, the distribution of capital indicates directly the ratchet shape of the games, shown in the top part of Figure 3.9.

Distribution Localisation

An alternative explanation of the two systems can be given in terms of the localisation of particles or capital at system 'ceilings' [Harmer et al., 2001]. Consider game B, the capital tends to localise between the $Mn - 1$ and Mn states for an integer n that represents the subsystems. This is due to the chosen probabilities of p_1 and p_2 . At $Mn - 1$, there is a high probability (p_2) the capital will increase to Mn and at that state there is an even higher probability ($1 - p_1$) the capital will be pushed back down to $Mn - 1$. Due to the oscillatory behaviour between these two states and the tendency for capital to drift upwards when in states $(M(n - 1) + 1)$ to $(Mn - 1)$, there is a localisation of capital at the Mn ceilings. In the same way, the particles in the ratchet teeth are localised to the pits, a small movement is met with an opposing long gentle or short steep edge. Thus, in each subsystem an extra 'kick' is required to move the capital or particles backwards or forwards one period.

3.2 Construction and Characteristics

Adding game A to the playing sequence improves the chance of moving up to the next subsystem because most of the capital is localised at these ceilings. Switching to an approximately fair game allows almost half of the capital at these ceilings to move up to the next subsystem (i.e. from Mn to $M(n + 1)$), while the other half moves down a few steps only to be brought back to the Mn ceiling when game B is played again. This is exactly what happens when the ratchet teeth are made to disappear in the Brownian ratchet – about half of the particles can easily move over the steep edge into the next pit while the remaining particles fall back into the same pit via the gentle edge when the ratchet teeth appear again.

3.2.7 Observations

Breaking the Equilibrium Distribution

We deduce from the previous discussions that in order for game B to function properly, it is dependent on the ratchet shape. We will consider the games in terms of coins. For game B we have a *bad* coin B_1 , and a *good* coin B_2 . Even though the state-dependency allows the good coin B_1 to be played more often (i.e. when not in state Mn), the bad coin B_2 is sufficiently bad to cause the game as a whole to be losing [Harmer and Abbott, 1999b].

To make this into a winning game, the distribution needs to be perturbed to alter the proportions of how many times each coin is played, that is, we want more of B_1 and

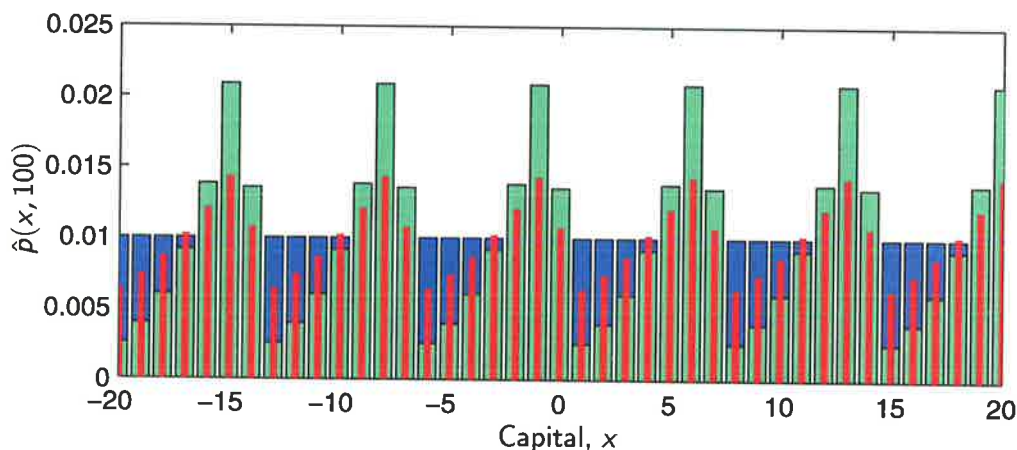


Figure 3.10. The breaking of the individual PDFs by the randomised game. The randomised game is shown by the thick lines. Adding game A to B causes degradation of the ratchet shape, which increases the probability of winning.

less of B_2 . This is accomplished by adding in game A , which creates the new PDF as shown by the thick lines in Figure 3.10. This breaks the distributions formed by game B alone, and allows coin B_1 to be used more often, thus creating an overall winning game [Key et al., 2001]. This new game needs to win sufficiently to offset the slight hindrance caused by the introduction of game A .

Analogous Quantities

Observing that Parrondo's games and the Brownian ratchet use the same transport mechanism, we can form analogies between the variables for each of the systems. A summary of the analogies is shown in Table 3.1.

The source of the potential in Brownian ratchets can be provided by a variety of means, but typically an electrostatic potential is used. For the games, the potential source is generated by the rules, they define the shape of the potential. The capital can be thought of as the transport quantity analogous to the particles in the Brownian ratchet.

Further depth between these two analogies is found by considering the dynamics of the two systems. The Brownian ratchet is continuous in time and space. The particles can exist at any real displacement along the potential, which is flashed on and off at any real time. This is in contrast to Parrondo's games that are discrete in both the analogous time and space. The capital of the games must be an integral number of units and a game is an indivisible operation. This is highlighted by the mode of analysis – the Brownian ratchet is analysed via continuous variables in the Fokker-Planck equation, whereas Parrondo's ratchet is via discrete-time Markov chain analysis.

When we consider the ratchet and pawl machine, directed motion is only achieved when energy is added to the system, like in a heat engine. Similarly for a flashing Brownian ratchet, energy is taken up by switching between two states to produce directed motion of Brownian particles. From the simulations and mathematical analysis of Parrondo's games, the two losing games can yield a winning expectation, without any apparent cost. This creates a paradox, "money for free." Where is the energy coming from in Parrondo's games? Of course, the money itself is conserved since the winnings of the player are at the expense of the losing opponent.

In stock market models, switching energy could be thought of as the buying and selling transaction cost. For the flashing ratchet, the work done by the system is less than the external energy used to flash the potentials (i.e. efficiency is less than unity). This is costly in the physical system as dictated by the laws of thermodynamics.

3.3 Analysis of the Games

Table 3.1. The relationship between quantities in Parrondo's games and the Brownian ratchet.

Quantity	Brownian Ratchet	Parrondo's Paradox
Source of Potential	Electrostatic, Gravity	Rules of games
Switching	U_{on} and U_{off} applied	Games A and B played
Switching Durations	for τ_{on} and τ_{off}	a and b
Duration	Time	Number of games played
Biasing	Macroscopic field gradient	Parameter ϵ
Transport Quantity	Brownian particles	Capital
Measurable Output	Displacement x	Capital amount X_n
External Energy	Switching U_{on} and U_{off}	None
Potential Shape	Depends on α	Probabilities p_1 , p_2 and M
Mode of Analysis	Fokker-Planck equation	Discrete-time Markov chains

However, it is perhaps more accurate to say that the analogy between Parrondo's discrete-time ratchet and the conventional physical flashing ratchet simply break down at this point. An expanded discussion on this issue is given in Harmer et al. (2001).

3.3 Analysis of the Games

As has been hinted throughout the previous sections, the mode of analysis for the games is via discrete-time Markov chains. Each value of capital is represented by a state, and the transition probabilities between the states are determined by the rules of the games. The next subsection gives a brief introduction to discrete-time Markov chains. Note though, that the analysis presented here is not a rigorous one and uses basic Markov chain theory. More in depth mathematical approach can be found in Pyke (2001) among others [Pearce, 2000b; Pearce, 2000a].

3.3.1 Discrete-Time Markov-Chains (DTMCs)

The following section is an introduction to DTMCs. We will cover what they are, their dynamics, equilibrium distributions and some state classifications. Most of the following information is sourced from Yates and Goodman (1999), which provides an excellent introductory level theory to stochastic processes. More in depth information can

be found in Norris (1997). The reader familiar in DTMC theory can skip to Sec. 3.3.2 on page 57.

Definition and Properties

We consider a discrete time, discrete value random sequence

$$\{X_n = n = 0, 1, 2, \dots\},$$

which is not an independent and identically distributed (iid) sequence. A *Markov chain* has X_{n+1} dependent on X_n but not any earlier values X_0, \dots, X_{n-1} . The value X_n is referred to as the *state* at time n , and the sample space of X_n is the *state space*. Thus, at any point in time we are considered to be in one of the states in the chain. With respect to the games, at every time increment the chain only changes by one state and thus can be referred to as a *birth-and-death* process or a *skip-free* chain.

The movement between states is controlled by the *transition probability*, p_{ij} . This is the probability of moving from the i th state (the current state) to the j th state. Thus, from the total law of probability, that is there is a probability of 1 that something will happen, we have

$$\sum_{j=0}^{\infty} p_{ij} = 1, \quad (3.7)$$

where $0 \leq p_{ij} \leq 1$. This means that the sum of all the transitions out of a state must be unity.

The transitions between all of the states in a DTMC can be conveniently written with use of the *state transition matrix*

$$\mathbb{P} = \left. \begin{array}{c} \overbrace{\left[\begin{array}{ccc} p_{00} & p_{10} & \cdots \\ p_{01} & p_{11} & \\ \vdots & & \ddots \end{array} \right]}^{\text{from } i} \\ \left. \vphantom{\left[\begin{array}{ccc} p_{00} & p_{10} & \cdots \\ p_{01} & p_{11} & \\ \vdots & & \ddots \end{array} \right]} \right\} \text{to } j, \end{array} \right\} \quad (3.8)$$

which contains all the one step transition probabilities, p_{ij} . From the relations in (3.7) all the columns in (3.8) must sum to unity. Many mathematical texts define the transition matrix as the transpose of \mathbb{P} , however this notation is commonly used in physics and engineering. The matrix \mathbb{P} can be represented graphically, examples of which are shown in later sections.

3.3 Analysis of the Games

DTMC Dynamics

The *dynamics* of a DTMC refers to the variation of the state during a short time interval from a given initial state. Predicting a future state becomes difficult due to the many paths that can be taken to get to that state. However, using \mathbb{P} , which contains the transition probabilities between all of the states, future states can be determined from simple matrix algebra.

First, we define the *state probabilities* at time n by

$$\{\pi_j(n) | j = 0, 1, 2, \dots\} \quad \text{where} \quad \pi_j(n) = P[X_n = j].$$

That is, $\pi_j(n)$ is the probability that we are in the j th state at time n . We can then further define the *state probability vector* or *distribution* as

$$\boldsymbol{\pi} = [\pi_0(n), \pi_1(n), \dots, \pi_K(n)]^T.$$

With a little thought, we see this must satisfy

$$\sum_{j=0}^K \pi_j = 1, \tag{3.9}$$

for a DTMC that contains K states. This is referred to as the normalisation condition.

For a Markov chain with transition matrix \mathbb{P} , the *n -step transition matrix* is given by \mathbb{P}^n . Thus, the state probability vector for n -steps in the future can be determined by

$$\boldsymbol{\pi}(n) = \mathbb{P}^n \boldsymbol{\pi}(0), \tag{3.10}$$

where $\boldsymbol{\pi}(0)$ is the initial state vector.

Equilibrium Properties

An important property of DTMCs is the evolution of $\boldsymbol{\pi}(n)$ as n becomes large, that is as $n \rightarrow \infty$. The *equilibrium distribution* is defined as

$$\boldsymbol{\pi} = \lim_{n \rightarrow \infty} \boldsymbol{\pi}(n) = \lim_{n \rightarrow \infty} P[X_n = j]. \tag{3.11}$$

This is also commonly referred to as the *stationary distribution* or the *limiting state probabilities*. $\boldsymbol{\pi}$ is not guaranteed to exist, and if it does, it may or may not depend on the initial condition, $\boldsymbol{\pi}(0)$.

To calculate π , consider the transition of an arbitrary time n where we have $\pi(n+1) = \mathbb{P}\pi(n)$. Taking the limit when $n \rightarrow \infty$ and according to (3.11) where $\pi(n) \rightarrow \pi$, we have

$$(\mathbb{I} - \mathbb{P})\pi = 0, \quad (3.12)$$

where \mathbb{I} is the identity matrix of the appropriate size. This is now a typical, well known eigenvalue problem. However, since π is associated with a Markov chain, it has an eigenvalue $\lambda_1 = 1$ with the remaining eigenvalues $|\lambda_i| < 1$. A method for finding π is given in Appendix A.2.

State Classification

The structure of Markov chains can be described by classifying the states within the chain. We are interested in the conditions of *recurrence* and *transience*. A recurrent state is one where we are certain to eventually return to, whereas for a transient state we may never return to. Defining V_{ij} as the event that we eventually visit state j given we started in state i , then state j is either,

- (i) *transient* if $P[V_{jj}] < 1$. That is, we may never return to state j , or
- (ii) *positive recurrent* if $P[V_{jj}] = 1$ (we will return to state j) and occurs within a finite number of transitions, or
- (iii) *null recurrent* if $P[V_{jj}] = 1$, but takes an infinite number of transitions to do so.

From these conditions, a finite Markov chain always has a set of recurrent states.

These classifications are useful when we wish to determine the *hitting times* [Norris, 1997, p. 12], which is the expected time to reach a particular state. For example, what is the probability that we ever reach state β given that we started in state α ?

3.3.2 Modelling the Games as DTMCs

The DTMC representing game A is shown in Figure 3.11, where the states represent the value of capital. Since the range of capital can extend to $\pm\infty$ it is referred to as *doubly infinite*. We move up a state with probability p and down with $1 - p$.

3.3 Analysis of the Games




Figure 3.11. Discrete-time Markov chain representing game A.

This has the transition matrix \mathbb{P}_A , which has zeros everywhere except for the two diagonals that are offset from the leading diagonal by ± 1 .

$$\mathbb{P}_A = \begin{bmatrix} 0 & 1-p & & (p) \\ p & 0 & 1-p & \\ & \ddots & \ddots & \ddots \\ & & p & 0 & 1-p \\ (1-p) & & & p & 0 \end{bmatrix}. \quad (3.13)$$

The DTMC in Figure 3.11 extends to $\pm\infty$, hence so do the dimensions of \mathbb{P}_A . However, in practice the size of \mathbb{P}_A can be restricted to the range of interest. For example, if we are only playing N games in a row then a $2N + 1$ square matrix will suffice as the capital will not exceed the bounds of the matrix. Hence, the bracketed terms that cater for the boundary conditions are not required in practice, and are shown for completeness.

The DTMC shown in Figure 3.12 represents game B. This consists of n periodic subsystems of length M , which is taken as three. This representation clearly shows the state dependency that is exhibited by game B – the probabilities leaving the M th state are p_1 and $1 - p_1$, while the probabilities leaving all the other states are p_2 and $1 - p_2$.




Figure 3.12. Discrete-time Markov chain representing game B.

This has the transition matrix \mathbb{P}_B of

$$\mathbb{P}_B = \begin{bmatrix} 0 & 1 - p_2 & & & & & (p_2) \\ p_1 & 0 & \dots & & & & \\ & p_2 & \dots & 1 - p_2 & & & \\ & & \dots & 0 & 1 - p_1 & & \\ & & & p_2 & 0 & 1 - p_2 & \\ & & & & p_1 & 0 & \dots \\ (1 - p_1) & & & & & & \dots \end{bmatrix}, \quad (3.14)$$

where the losing and winning probabilities in every M th column are $1 - p_1$ and p_1 respectively. The bracketed terms are included for completeness.

By extracting the periodic subsystem from the DTMC representation in Figure 3.12, the dynamics of the games can be more easily studied. The subsystem is defined by

$$Y_n \equiv X_n \text{ mod } M. \quad (3.15)$$

Though this representation does not reveal the absolute value of capital, meaningful trends can be easily calculated. The DTMC defined by Y_n has the states $\{0, \dots, M - 1\}$, and is cyclic. That is, if we win at the highest state $M - 1$ we go back to state 0 and vice versa from state 0 to $M - 1$. The corresponding DTMC to Y_n is shown in Figure 3.13.

Y_n can now also be used for game A by letting $p_1 = p_2 = p$. Having the games represented by non-infinite DTMCs allows useful calculations to be performed, as shown




Figure 3.13. Discrete-time Markov chain corresponding to the modulo game B . The cyclic DTMC defined by (3.15) represents one of the subsystems of the DTMC representing game B in Figure 3.12.

3.3 Analysis of the Games

in subsequent sections. For example, from Onsager's Nobel prize winning work on reversible chemical reactions, we can use a detailed balance to quickly determine which direction has the greatest drift: clockwise or counter-clockwise [Onsager, 1931a; Onsager, 1931b; Van den Broeck et al., 1999]. The detailed balance simply entails finding the product of probabilities in one direction, and comparing this to the product of the probabilities in the opposite direction. For the DTMC in Figure 3.13 clockwise is the winning direction, thus from the detailed balance game B is winning if

$$p_1 p_2^{M-1} > (1 - p_1)(1 - p_2)^{M-1}. \quad (3.16)$$

Although mathematicians may feel hesitant about this "back of the envelope" type of calculation, it does give the same result as a more formal analysis shown in Appendix A.1 and is more physically intuitive.

The transition matrix representing the modulo game B in Figure 3.13 is

$$\mathbb{P}_B = \begin{bmatrix} 0 & 1 - p_2 & & & p_2 \\ p_1 & 0 & \ddots & & \\ & p_2 & \ddots & 1 - p_2 & \\ & & \ddots & 0 & 1 - p_2 \\ 1 - p_1 & & & p_2 & 0 \end{bmatrix}. \quad (3.17)$$

This is restricted to $M \times M$ in size and the previously bracketed terms now become important as they enforce the modulo rule.

The Randomised Game

Dealing with the randomised game is not as difficult as it first appears. Recall that the mixing parameter γ (from Sec. 3.2.5) gives the relative probability of playing game A , where we have assumed the value of a half. Therefore, when the capital is a multiple of M the probability of winning is

$$q_1 = \gamma p + (1 - \gamma)p_1. \quad (3.18a)$$

This is the chance of playing game A , multiplied by the chance of winning, added to the chance of playing game B , multiplied by the chance of winning it. Alternatively, when the capital is not a multiple of M , the probability of winning is

$$q_2 = \gamma p + (1 - \gamma)p_2. \quad (3.18b)$$

The respective losing probabilities are $1 - q_1$ and $1 - q_2$. Using these probabilities we can treat the randomised game exactly the same as game B , except replacing the p_i s with q_i s. This does not affect the DTMC analysis as a combination of two DTMCs simply form another DTMC, which accordingly also obeys Markov chain theory.

Playing the Games Analytically

Having determined the transition matrices representing the games and using the distribution vector giving the capital in each of the states, playing the games n times gives

$$\boldsymbol{\pi}_n = \mathbb{P}^n \boldsymbol{\pi}_0, \quad (3.19)$$

where bpi_n is short for $\boldsymbol{\pi}(n)$, not to be confused with π_j .

If we wish to determine the distribution of capital we need to use the transition matrices that represent the doubly infinite chains. As mentioned earlier the size can be reduced to a $2N + 1$ square matrix. Starting with zero capital gives $\boldsymbol{\pi}_0 = [\dots, 0, 1, 0, \dots]^T$. To play either game A or B , \mathbb{P}_A or \mathbb{P}_B is accordingly substituted into (3.19). To play a mix of the games, the appropriate transition matrix to be substituted is dependent on n . Thus, we could have

$$\boldsymbol{\pi}_n^{[a,b]} = \mathbb{P}_X^n \boldsymbol{\pi}_0,$$

using the notation $[a, b]$ described earlier and where

$$\mathbb{P}_X = \begin{cases} \mathbb{P}_A & \text{if } ((n - 1) \bmod (a + b)) < a \\ \mathbb{P}_B & \text{otherwise,} \end{cases}$$

where $n = 1, 2, \dots$

For the random mix of games, we have

$$\mathbb{P}_R = \gamma \mathbb{P}_A + (1 - \gamma) \mathbb{P}_B,$$

which can be substituted in (3.19) to easily replicate the randomised games when averaged over an infinite number of trials. That is, the expected result according to the central limit theorem.

The deterministically mixed games can be implemented using a single transition matrix by grouping the periodic sequence. For example, $\mathbb{P}_{2,2} = \mathbb{P}_A^2 \mathbb{P}_B^2$ represents the equivalent transition matrix of playing $\{AABB\}$. Applying $\mathbb{P}_{2,2}$ is then equivalent to playing four consecutive games. Due to the multiple paths the capital can take within

3.3 Analysis of the Games

those four games, the algebra becomes tedious – a symbolic programming language is most advantageous to generate an expression for $\mathbb{P}_{2,2}$.

Using the technique of (3.19) we can determine statistical properties of the games, namely the mean μ , and standard deviation σ . We define the vector

$$\mathbf{x} = [-N, \dots, -1, 0, 1, \dots, N], \quad (3.20)$$

which contains all the values of capital possible (i.e. states) when playing N games. The mean is then given by

$$\mu_n = \sum_{i=-N}^N x_i \pi_i(n),$$

noting that $\sum_i \pi_i(n)$ is unity for all n . In vector notation we have

$$\mu_n = \mathbf{x} \boldsymbol{\pi}_n, \quad (3.21)$$

where the matrix multiplication provides the summing. Similarly, the standard deviation is given by

$$\sigma_n = \sqrt{(\mathbf{x} - \mu_n) \circ (\mathbf{x} - \mu_n) \boldsymbol{\pi}_n},$$

where the Hadamard (\circ) is an element-wise multiplication operator. These results agree with the sample paths shown in Figure 3.2 and standard deviations in Figure 3.5.

Equilibrium Distribution

The equilibrium (or stationary) distribution of a DTMC occurs when the distribution of capital in the states does not change from one game to the next. That is, the distribution is invariant under the action of \mathbb{P} . From Sec. 3.3.1, it is shown we need to solve (3.12). Although there are many ways to solve this typical eigenvalue problem, a systematic method states that the stationary distribution is proportional to the diagonal cofactors of $\mathbb{I} - \mathbb{P}$ [Pyke, 2001]. That is

$$\boldsymbol{\pi} = \frac{1}{D} \text{diag}(\text{cofac}(\mathbb{I} - \mathbb{P})), \quad (3.22)$$

where D is the normalisation constant. The function ‘diag’ returns the main diagonal of a matrix and ‘cofac’ gives the cofactors of a matrix, an example is given in Appendix A.2.

Alternatively one can use the global balance equations with \mathbb{P} [Pearce, 2000a]. By either method, the stationary distribution for $M = 3$ is

$$\pi^B = \frac{1}{D} \begin{bmatrix} 1 - p_2 + p_2^2 \\ 1 - p_2 + p_1 p_2 \\ 1 - p_1 + p_1 p_2 \end{bmatrix}, \quad (3.23)$$

where $D = 3 - p_1 - 2p_2 + 2p_1 p_2 + p_2^2$. If we let $p_1 = p_2 = p$ to represent game A , then the stationary distribution simplifies to

$$\pi^A = \frac{1}{3} [1, 1, 1]^T, \quad (3.24)$$

as expected for a three state chain with identical transition probabilities.

Using the probabilities of (3.1) with $\epsilon = 0$, the stationary distribution for game B is found to be

$$\pi^B = \frac{1}{13} [5, 2, 6]^T, \quad (3.25)$$

which are the quantities used in (3.4). For the randomised game with $\gamma = 1/2$, we have

$$\pi^R = \frac{1}{709} \begin{bmatrix} 245 \\ 180 \\ 284 \end{bmatrix} = \begin{bmatrix} 0.346 \\ 0.256 \\ 0.401 \end{bmatrix}.$$

3.3.3 Constraints of Parrondo's Games

It would be desirable, given a set of parameters, to find constraints to determine if Parrondo's paradox would be exhibited. An intuitive approach is finding the probability of winning using the stationary distribution, given by

$$p_{\text{win}} = \sum_{j=0}^{M-1} \pi_j p_j, \quad (3.26)$$

where p_j is the winning probability in state π_j . The games are winning, losing or fair when p_{win} is greater, less or equal to a half, which implies that $\langle X_n \rangle$ is a decreasing, increasing or constant with respect to n respectively.

From (3.26) we require $p < 1/2$ for game A , or alternatively

$$\frac{1-p}{p} > 1. \quad (3.27)$$

3.3 Analysis of the Games

For game B , the winning probability of (3.26) becomes

$$\begin{aligned} p_{\text{win}}^B &= \pi_0 p_1 + \pi_1 p_2 + \cdots + \pi_{M-1} p_2 \\ &= \pi_0 p_1 + (1 - \pi_0) p_2. \end{aligned} \quad (3.28)$$

Using $M = 3$ to simplify the algebra and using the stationary distribution (3.23) with $p_{\text{win}}^B < 1/2$ yields

$$\frac{(1 - p_1)(1 - p_2)^2}{p_1 p_2^2} > 1. \quad (3.29)$$

This is the condition that needs to be satisfied for game B to be losing.

For the randomised game we use the expression for game B except replacing the p_i s with q_i s from (3.18) and setting $p_{\text{win}}^R > 1/2$ we get

$$\frac{(1 - q_1)(1 - q_2)^2}{q_1 q_2^2} < 1. \quad (3.30)$$

This is the condition for the randomised game to win. Therefore, in order for Parrondo's paradox to be exhibited we require probabilities and parameters to satisfy (3.27), (3.29) (i.e. to make game A and B lose) and (3.30) (i.e. make the randomised game win).

This type of analysis becomes tedious as M becomes larger due to the necessity of finding the equilibrium distribution. A more formal analysis considers the conditions of recurrence of the corresponding DTMC for the games and is given in Appendix A.1.

However, also notice that the numerator is the product of all the losing probabilities over one period and the denominator is the product of all the winning probabilities. This is the reverse drift over the forward drift, as also found using Onsager's detailed balance in Sec. 3.3.2. This approach allows us to generalise the results by replacing the exponent of 2 with $M - 1$ in (3.29) and (3.30). Thus, for Parrondo's paradox to be exhibited in the general modulo M game the following inequalities need to be satisfied,

$$\frac{1 - p}{p} > 1, \quad (3.31a)$$

$$\frac{(1 - p_1)(1 - p_2)^{M-1}}{p_1 p_2^{M-1}} > 1 \quad \text{and} \quad (3.31b)$$

$$\frac{(1 - q_1)(1 - q_2)^{M-1}}{q_1 q_2^{M-1}} < 1. \quad (3.31c)$$

Range of Biasing Parameter

In Sec. 3.2.3, the biasing parameter was shown to control whether a game was winning or losing. In terms of Parrondo's paradox, we have shown that randomising the games improves the performance. However, if ϵ is too large then all the games lose, albeit the randomised game does not lose by as much. Conversely if ϵ is too small (negative), then all the games win. Thus, ϵ needs to be chosen such that it biases games A and B to lose, but the improvement gained by mixing is greater than the offset made by ϵ .

By substituting the probabilities of (3.1) into the equations (3.31) we deduce a range of ϵ for which Parrondo's paradox exist. The equations (3.31) are respectively,

$$\epsilon > 0, \quad (3.32a)$$

$$\epsilon(80\epsilon^2 - 8\epsilon + 49) > 0 \quad \text{and} \quad (3.32b)$$

$$320\epsilon^3 - 16\epsilon^2 + 229\epsilon - 3 < 0. \quad (3.32c)$$

For the quadratic part of (3.32b), $\Delta = b^2 - 4ac < 0$, so the roots are imaginary, meaning that $80\epsilon^2 - 8\epsilon + 49 > 0$ for all ϵ , which leaves $\epsilon > 0$. For (3.32c) we can numerically find the roots or use Cardan's method for cubic polynomials to deduce that there is one real and two imaginary roots. Either way the real root is, $\epsilon_{\max} \approx 0.0131$, which gives the possible range of the biasing parameter as

$$0 < \epsilon < 0.0131. \quad (3.33)$$

To approach the upper limit of this range ϵ_{\max} , n needs to be large to offset the initial transient behaviour.

Probability Space

The probability space refers to the reachable space in \mathbb{R}^3 by the point formed by (p, p_1, p_2) . Using the constraints of (3.31), the regions of probability space where Parrondo's paradox exists are defined.

Game A only depends on p , the winning and losing regions are separated by the plane $p = 1/2$.

Game B depends on p_1 and p_2 for a given M and the curves constituting fair games can be derived from the equality (3.31b). This is shown for various values of M in Figure 3.14. The straight line shows the region of space that is used with the parameterisation

3.3 Analysis of the Games

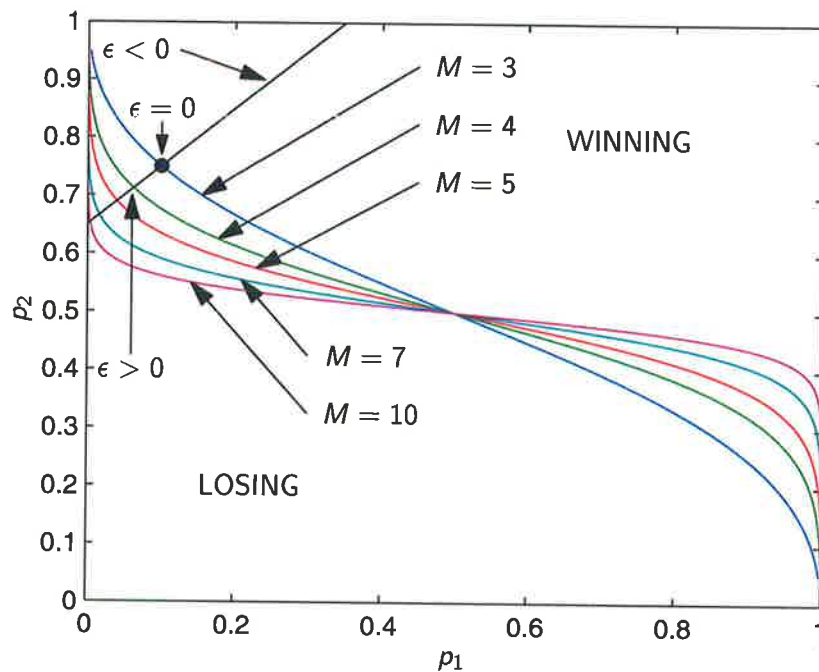


Figure 3.14. Probability space of game B. The curves for various values of M divide the winning and losing regions of probability space. The straight line is the parameterisation given in (3.1).

given in (3.1). From the trends related to ϵ discussed in Sec. 3.2.3, we can deduce that the region above the curves is winning, and the region below losing.

The only fair version of game B considered so far is the one represented by the large dot in Figure 3.14, but there are a continuous range of probabilities (p_1, p_2) . The capital after 100 games is plotted in Figure 3.15 while transversing along the $M = 3$ curve in Figure 3.14, the rate of return clearly is not linear. At the extreme values of p_1 we lose or win by the greatest amount. This is due to the increasing asymmetry of the teeth in the ratchet potential, the magnitude of the slope increases at the limits of the probability space.

For game B in Figure 3.15 fair probabilities were used for p_1 and p_2 , but there is a slight deviation from a zero return. This is due to the initial transient effects discussed earlier. Also, when $p_1 > 0.5$ the randomised game produces a losing result, which is due to the probabilities not satisfying the constraints of (3.31).

In Figure 3.16 the surface Π_R separating the winning and losing volumes for the randomised game has been plotted, as well as the surfaces Π_A and Π_B for games A and B respectively. The arrows indicate the side of the surface where a point is required to be to satisfy (3.31). The point needs to be below Π_A and to the right of Π_B for games

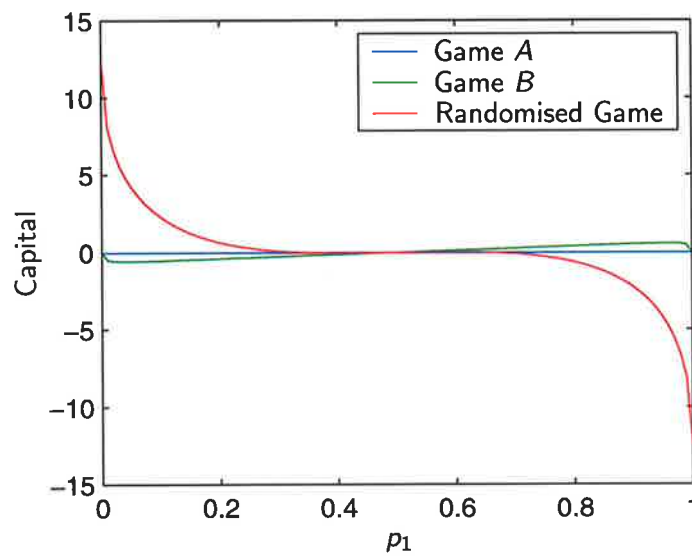


Figure 3.15. Gain for all the fair probabilities of game *B*. The capital after playing 100 games using the probabilities $p = 1/2$, p_1 and p_2 from the $M = 3$ curve in Figure 3.14.

A and *B* to lose, but above Π_R for the randomised games to win. A small volume that satisfies all these constraints exists in the front left side of the plot – any point in this region will give rise to Parrondo's paradox being exhibited.

There is also a region at the opposite side of the volume that gives rise to the opposite of Parrondo's paradox. Two winning games can combine to form a losing game. This is also observed in Figure 3.15 when $p_1 > 0.5$. In a practical sense this is like changing the perspective observation of the games, whether from the player's or the bank's point of view. The region in probability space that exists for Parrondo's paradox is not large, only a mere 0.032% of the whole probability space.

3.3.4 Rate of Winning

From the distributions of the games it is possible to find the rate of winning as a function of the number of games played, $r(n)$. The method given here simply uses stationary distributions, a more formal approach is shown in Appendix A.3.

Using the stationary distributions we can find the rate of winning by subtracting the probability of losing from the probability of winning. Thus, we have

$$\frac{d\langle X_n \rangle}{dn} \equiv r = \sum_{i=0}^{M-1} 2\pi_i p_i - 1. \quad (3.34)$$

3.3 Analysis of the Games




Figure 3.16. Probability space of all the games showing the paradoxical region. The surfaces Π_A , Π_B and Π_R represent the boundaries between winning and losing games. The small volume also bounded by the $p_1 = 0$ plane at the front left of the plot is where Parrondo's paradoxical games exist. Parameters are $M = 3$ and $\gamma = 1/2$.

For game A, $\pi_i = 1/M$ and $p_i = p$ for all i and the expression (3.34) reduces to

$$r_A = 2p - 1, \quad (3.35)$$

which is the expected result.

For game B the rate of winning is

$$r_B = 2p_2 - 1 + 2\pi_0(p_1 - p_2), \quad (3.36)$$

which is valid for all M , though one still needs to find the stationary distribution to obtain π_0 . For the probabilities in (3.1) with the stationary distribution for $M = 3$ given in (3.23), the slope of game B is

$$\begin{aligned} r_B &= \frac{3(p_1 p_2^2 - (1 - p_1)(1 - p_2)^2)}{3 - 2p_2 + p_2^2 + 2p_1 p_2 - p_1} \\ &= \frac{-6\epsilon(49 - 8\epsilon + 80\epsilon^2)}{169 - 16\epsilon + 240\epsilon^2} \\ &= -1.740\epsilon + 0.119\epsilon^2 + O(\epsilon^3), \end{aligned} \quad (3.37)$$

using a Taylor series expansion with centre $\epsilon_0 = 0$. This is negative for small $\epsilon > 0$, which indicates game B is losing. As suspected earlier in Figure 3.7, this contains higher order nonlinearities with respect to ϵ .

The same calculation holds for the randomised game by appropriately changing the probabilities to give

$$\begin{aligned} r_R &= \frac{6(3 - 229\epsilon + 16\epsilon^2 - 320\epsilon^3)}{709 - 32\epsilon + 960\epsilon^3} \\ &= 0.0254 - 1.937\epsilon + 0.0136\epsilon^2 + O(\epsilon^3), \end{aligned} \quad (3.38)$$

which is positive for small ϵ , but becomes negative for larger ϵ .

When comparing simulations with this analysis the initial transient game needs to be ignored by determining the slope from the 100th to 600th game for example.

3.4 Entropy

In this section the analysis of Parrondo's games is linked to information theory. We investigate how this relates to the entropy rates of the games.

Now let us consider the sequences of wins and losses of the games in terms of information theory [Harmer et al., 2000c]. Denote by X_j the random variable which represents the outcome at time point j when playing any of the games A , B or R , where the latter denotes the randomised game. If the game wins at the j th time point, $X_j = 1$, if it loses $X_j = 0$.

For game A , the sequence X_1, X_2, \dots is an ergodic stationary sequence. In the case of games B and R , the sequence X_1, X_2, \dots is not stationary but it can be regarded so after an initial period. Hence, by the Shannon-MacMillan-Breiman Theorem [Durrett, 1991, p. 314],

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log p(X_0, \dots, X_{n-1}) = H \quad (3.39)$$

almost surely, where

$$H \equiv \lim_{n \rightarrow \infty} E[-\log p(X_n | X_{n-1}, \dots, X_0)] \quad (3.40)$$

is the *entropy* of the sequence and

$$p(x_0, \dots, x_{n-1}) \equiv P(X_0 = x_0, \dots, X_{n-1} = x_{n-1})$$

is the probability measure of the sequence. It is convenient to take logarithms with respect to base 2 to give entropy in bits. For the case of game A , the right hand side of equation (3.40) reduces to

$$H^A = H(p) \equiv -p \log p - (1 - p) \log(1 - p) \quad (3.41)$$

3.4 Entropy

for $0 < p < 1$. For the limiting case when $p = 0$, we note that $\lim_{p \rightarrow 0} \log p = 0$, thus let $\log 0 = 0$. Intuitively when $p = 0$ the event should never occur, thus it should not contribute to the total entropy, hence $0 \log 0 = 0$ [Nielsen and Chuang, 2000, p. 500]

The random variables X_j for games B and R are correlated and so the expressions for their entropies are more complicated. To find the entropy of game B , we need to calculate the equilibrium distribution π^B so that the capital is congruent to i modulo M for $i = 1, \dots, M$. The right hand side of equation (3.40) then becomes,

$$H^B = \pi_0^B H(p_1) + (1 - \pi_0^B) H(p_2). \quad (3.42)$$

Similarly for the randomised game, the entropy is, in the obvious notation,

$$H^R = \pi_0^R H(q_1) + (1 - \pi_0^R) H(q_2). \quad (3.43)$$

For both games B and R , we can calculate the entropy rates considering the random variables X_j to be independent. This results in

$$H_i^B = H(\pi_0^B p_1 + (1 - \pi_0^B) p_2) \quad (3.44)$$

and

$$H_i^R = H(\pi_0^R q_1 + (1 - \pi_0^R) q_2) \quad (3.45)$$

for game B and the randomised game respectively. The arguments of H on the right hand side of (3.44) and (3.45) are simply the stationary probabilities of producing a 1 at a given time point of the games. Since $H(p)$ is concave, it is easy to see that

$$H_i^B \geq H^B \quad (3.46)$$

and

$$H_i^R \geq H^R \quad (3.47)$$

with equality only occurring when $p = p_1 = p_2 = 1/2$.

The entropy is a measure of how ordered a sequence is, the higher the entropy, the less order there is. Thus H^A is maximised when $p = 1/2$ and the sequence is completely random. Given this, we may expect the entropy of a game to be related to the fairness of the game. A fair game, one that has a net gain of zero, is likely to have an equal number of wins and losses, hence 0s and 1s in the chain. It might be reasonable to expect that this is the most unordered state and has a maximal value for its entropy rate.

From the Shannon-MacMillan-Breiman Theorem, it follows that by making n large enough, we can get an estimate of the entropy of a stochastic process which generates a sequence from a sample path via the left hand side of equation (3.39). We did this for the sample paths generated by the simulations reported above, using expressions (3.41), (3.44) and (3.45). The average values of these estimates are plotted along with the average gain for the various games in Figure 3.17. Note that the entropy rates found in Figure 3.17 are calculated, *not* taking correlations into account. The effects of taking correlation into account will be dealt with in the following section.




Figure 3.17. The relation between the fairness of the games and the entropy. The bottom three lines are the final capital after the 100th game and the top three curves are the entropy rates. Simulations were performed using (3.1) with $\epsilon = 0.05$.

3.4.1 Entropy and Probability Space

In this section we explore the relations between the entropy rates generated from Parrondo's games and how they relate to each other and the probability space discussed in Sec. 3.3.3. As previously noted above, the successive values of X_j in games B and R are correlated, as can be seen from a simple illustration. If $p_1 = 0.9$ and $p_2 = 0.1$, then the capital will oscillate between two adjacent states most of the time. That is, given that we are in one state, we can predict with high probability what state we will be in next, and which digit will be generated. Hence there is high correlation, even though there are approximately the same numbers of 0s and 1s.

3.4 Entropy




Figure 3.18. Entropy and probability space for game A. The maximal entropy indicates where the game is fair.

The entropy for game A can be calculated directly from (3.41), as shown in Figure 3.18. The maximum entropy occurs at $p = 1/2$, which divides the winning and losing games.

The uncorrelated entropy for game B , H_I^B is shown in Figure 3.19a, while Figure 3.19b shows the correlated entropy H^B . Comparing the plots in Figure 3.19 we can verify (3.46). The fairness of the game is actually reflected by H_I^B , when $H_I^B = 1$. The fair game is shown by the curve along the ridge in Figure 3.19a, which is the same as the $M = 3$ curve in Figure 3.14.

Even though $H_I^B = 1$ occurs for *all* fair games, we can have $H^B < 1$ for a fair game [Pearce, 2000a; Pearce, 2000b]. From Figure 3.19b, we can conclude that the usual entropy of a game allowing for correlations, is not related in any simple way to fairness.

The entropies for the randomised game are set naturally in four dimensions, which make them difficult to visualise. However the same traits can be carried from games A and B to game R . By fixing $p = 1/2$, the uncorrelated (H_I^R) and correlated (H^R) entropies are plotted in Figure 3.20a and 3.20b respectively. As before $H_I^R = 1$ occurs for all the fair games and we can verify (3.47) as shown in Figure 3.20a when $p = 1/2$.

Comparing the entropy rates from games B and R (considered as sources), when $p = 1/2$, $H^R \geq H^B$ as shown in Figure 3.20b. With $p = 1/2$, H^R has maximal entropy at $p_1 = p_2 = 1/2$, but perturbations of p away from $1/2$ causes the surface H^R to sink and the peak of H^B , which is invariant to p , raises above the surface H^R . Hence,

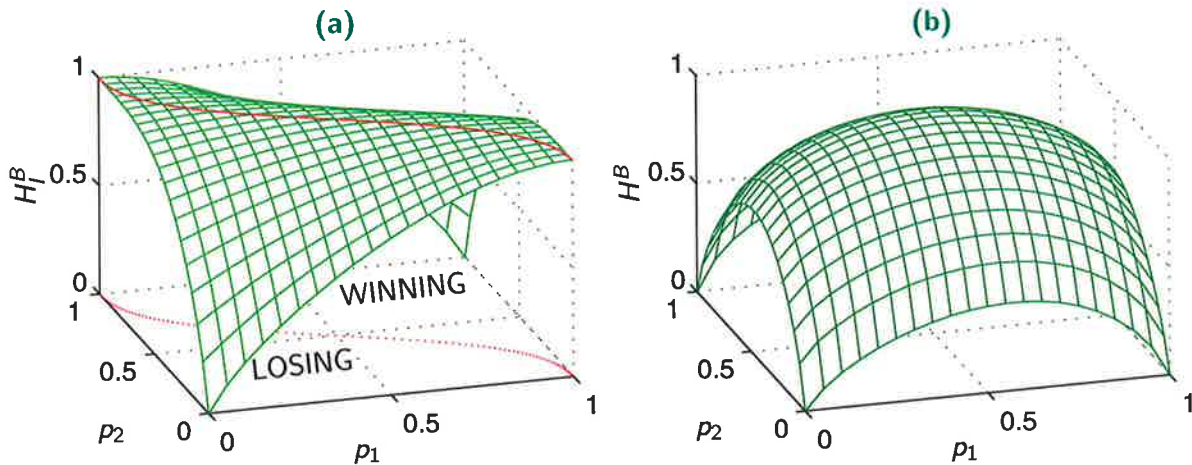


Figure 3.19. Entropies for game B . (a) The uncorrelated entropy $H_I^B = 1$ indicates where the game is fair along the ridge at $H_I^B = 1$. (b) The entropy taking into account correlations, which give no information about fairness.

$H^R < H^B$ for values of p_1 and p_2 in the neighbourhood of $1/2$ when $p \neq 1/2$. This makes sense if we consider game R to be game B plus another source. If this other source is completely random ($p = 1/2$), then we are adding disorder to game B and cannot decrease the entropy.

If we consider entropies taking into account the last λ bits generated, when $\lambda = 0$, we have the uncorrelated entropy H_I and when $\lambda = \infty$ we have the correlated entropy H . Then as λ approaches infinity, the entropy asymptotically approaches H from above. This was noticed by Shannon for passages of English text [Shannon, 1951].

3.4.2 Entropy Paradox

The uncorrelated entropy rates are closely related to the probability space, in particular fair games have an uncorrelated entropy of 1 bit. In fact, it is easy to see that measuring the uncorrelated entropy rate is just another way of counting the proportion of zeros and ones, and that the uncorrelated entropy is maximised when zeros and ones occur in equal proportion. The entropies allowing for correlations do not have a simple relationship to the fairness of the games. Instead of simply counting, they take into account the order of the sequence, i.e. its “pattern.”

One way to think of this is a new paradox in terms of uncorrelated entropy, the fair games A and B separately create sequences with maximum uncorrelated entropy. However, mixing A and B creates a sequence with a smaller uncorrelated entropy. This

3.5 History Dependent Games

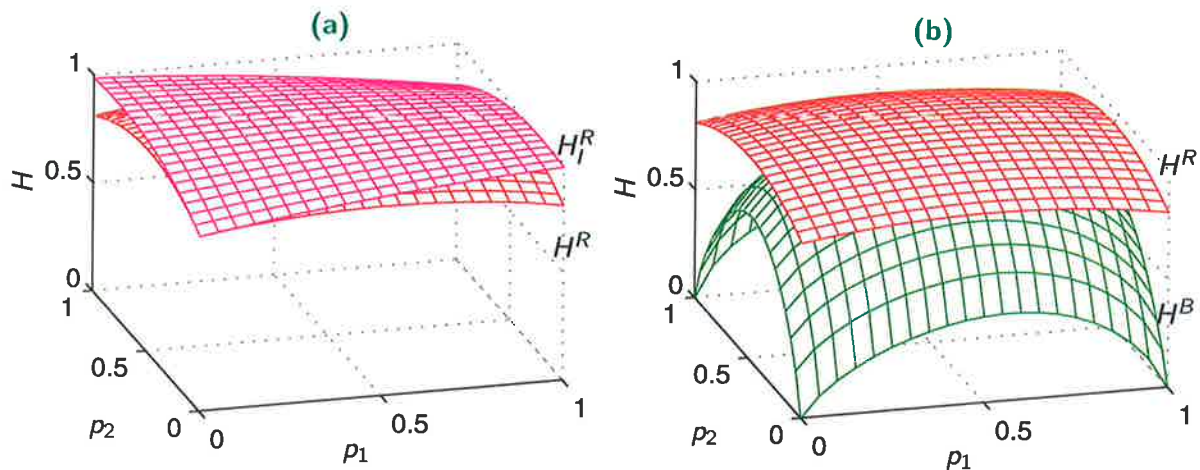


Figure 3.20. Entropies for the randomised game. (a) The uncorrelated (H_I^R) and correlated (H^R) entropies. (b) Comparing the entropies of game B and the randomised game when $p = 1/2$.

paradox does, however, have a very easy solution. There is no reason to think that mixing games with maximal uncorrelated entropies should produce another game with maximal uncorrelated entropy in the presence of correlations.

3.5 History Dependent Games

It has now been shown that two losing capital dependent games can win, but are there any other types of games that have this characteristic? Although state dependent games are applicable in some areas, it may be desirable to have a version of the games independent of capital. It turns out that we can answer the aforementioned question in the affirmative, in the form of history dependent games. Such games were also devised by Parrondo in Parrondo et al. (2000), although other implementations are possible [Costa et al., 2002].

3.5.1 Construction

The probability of winning and losing each game depends on the result of the previous two games, as shown diagrammatically in Figure 3.21.

Game A is identical to the state dependent games, hence the same name. Again, game B' , the counter part to game B , is a little more complex and the probabilities depend on the two previous results. The subscript $t - 1$ refers to the previous game and $t - 2$

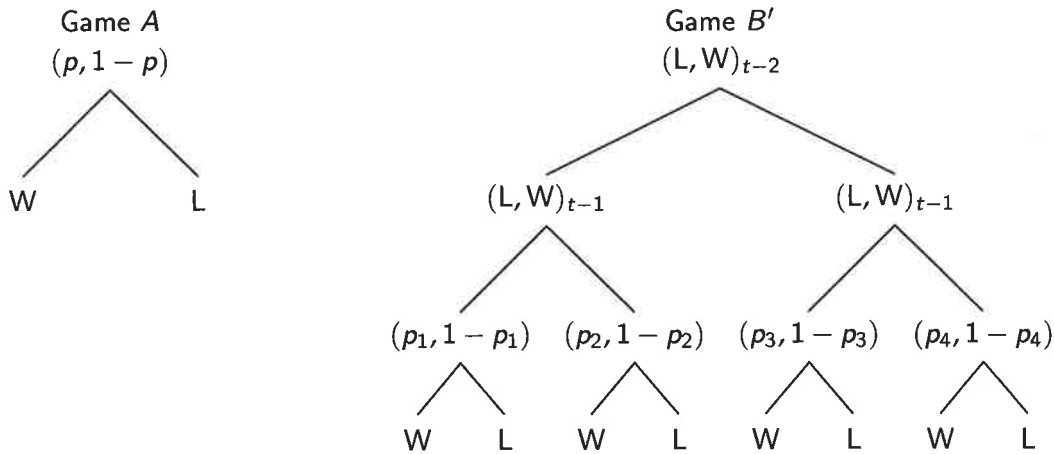


Figure 3.21. Construction of Parrondo's history dependent games. The decisions are based on previous results, either winning (W) or losing (L). The games can be implemented as biased coins.

to the game prior to that. If we had previously lost then won, the next game would be played with a winning probability of p_2 according to Figure 3.21.

These probabilities can be parameterised using the following transformation,

$$\begin{aligned}
 p &= 1/2 - \epsilon, \\
 p_1 &= 9/10 - \epsilon, \\
 p_2 &= p_3 = 1/4 - \epsilon \quad \text{and} \\
 p_4 &= 7/10 - \epsilon.
 \end{aligned}
 \tag{3.48}$$

This parameterisation gives Parrondo's original numbers for the history dependent games [Parrondo et al., 2000], which behave very similarly to the parameterisation of the capital dependent games in (3.1). That is, the games can be considered fair when $\epsilon = 0$, losing when $\epsilon > 0$ and winning when $\epsilon < 0$. The method of analysis for the games will follow closely to that of the original capital dependent games.

3.5.2 Results and Distributions

The same counter-intuitive result occurs when playing games A and B', that is, when playing the games individually they lose, but switching between them creates a winning expectation. The switching can be either stochastic or deterministic as shown by the selected games plotted in Figure 3.22. Similarly, there are some initial starting transients; the magnitude and shape depend on the initial conditions used, i.e. LL, LW,

3.5 History Dependent Games




Figure 3.22. Progress when playing Parrondo's history-dependent games. Simulations were performed using the probabilities in (3.48) with $\epsilon = 0.003$. A total of one million sample paths were averaged using each of the four initial conditions a quarter of the time.

WL or WW. The sequences shown in Figure 3.22 are averaged from each of the four possible starting conditions, thus eliminating much of the transient behaviour.

Figure 3.23 shows the outcome after the 100th game was played using different deterministic sequences. The general shape follows closely to that of the equivalent plot in Figure 3.6 for the capital dependent games. The faster the rate of switching, the better the rate of return.

The PDFs for the history dependent games are plotted in Figure 3.24. The capital is divided into four groups representing each of the possible histories the games could have taken. The PDFs plotted have $\epsilon = 0$, so do not show any observable drift. In fact, the PDFs do not appear very distinctive when compared to the capital dependent PDFs. They are in general fairly smooth and symmetric with only slight differences between them.

3.5.3 Analysis of the Games

To analyse the history dependent games we need to manipulate their representation into a workable DTMC. Upon doing so the same method can be employed as was used for the capital dependent games. An alternative approach using a quasi-birth-and-death process to represent the games is given in Appendix A.4.




Figure 3.23. Expected returns from deterministic sequences of the history-dependent games. The value of capital after playing the 100th game using the sequence denoted by $[a, b]$. Simulations used the probabilities of (3.48) with $\epsilon = 0$.




Figure 3.24. Probability density function of the history dependent games. The PDFs of the games using the probabilities of (3.48) with $\epsilon = 0$.

3.5 History Dependent Games

The discrete-time chain for game B' is shown in Figure 3.25. The states are periodic, that is, each column-wise slice is identical. The rows represent the history of the previous two games, which are $\{LL, LW, WL, WW\}$. It is the role of the rows to keep track of the history for 'each capital' so that it knows what probability (i.e. coin) to use for the next game. The amount of capital we possess is indicated by the column index, which simply needs to be summed. Therefore, for every game that is played we move back or forth one column and move to the appropriate row depending on the outcome of the game.

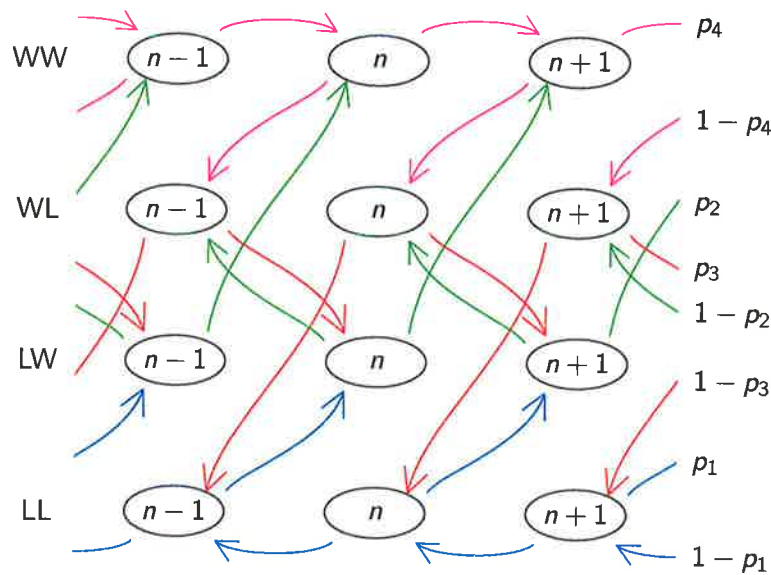


Figure 3.25. Discrete-time chain for game B' . The rows encode the previous results (LL, LW, WL and WW) and the column indices hold the current value of capital X_n .

However, the chain shown in Figure 3.25 is not a standard Markovian process since the next state depends not only on the previous states (i.e. the columns), but the rows as well. This is remedied by taking the periodic components of the chain in Figure 3.25. This only records the past histories and not the value of capital. We can define this using the previous two states, giving the vector notation

$$Y'(n) = [X(n-1) - X(n-2), X(n) - X(n-1)]. \quad (3.49)$$

This gives four states as $[-1, -1]$, $[-1, +1]$, $[+1, -1]$ and $[+1, +1]$, where $+1$ represents a win and -1 a loss. This representation forms a discrete-time Markov chain as shown in Figure 3.26.

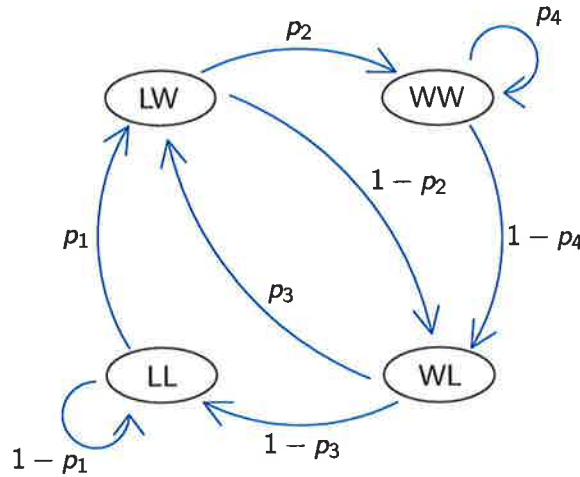


Figure 3.26. Discrete-time Markov chain for game B' using previous histories. The DTMC is formed by $Y(n)$ defined in (3.49).

The corresponding transition matrix for the chain in Figure 3.26 is

$$\mathbb{P}_{B'} = \begin{bmatrix} 1-p_1 & 0 & 1-p_3 & 0 \\ p_1 & 0 & p_3 & 0 \\ 0 & 1-p_2 & 0 & 1-p_4 \\ 0 & p_2 & 0 & p_4 \end{bmatrix}, \quad (3.50)$$

with the rows and columns representing the four states LL, LW, WL and WW, labelling from the top left corner. This matrix is always 4×4 since we only ever record the results of the previous two games.

When randomly mixing the games, the probabilities are given by

$$q_i = \gamma p + (1-\gamma)p_i \quad (3.51)$$

for $i = 1, \dots, 4$ and γ is the mixing parameter.

Equilibrium Distribution

Having represented the history dependent games as a DTMC and formed the corresponding transition matrix, the standard DTMC analysis can be performed.

The equilibrium distribution is found using (3.22), which gives

$$\pi^{B'} = \frac{1}{D'} \begin{bmatrix} (1-p_3)(1-p_4) \\ (1-p_4)p_1 \\ (1-p_4)p_1 \\ p_1 p_2 \end{bmatrix}, \quad (3.52)$$

3.5 History Dependent Games

where the normalisation constant $D' = p_1 p_2 + (1 + 2p_1 - p_3)(1 - p_4)$. Using Parrondo's original probabilities of (3.48) with $\epsilon = 0$ gives $\pi^{B'} = \frac{1}{22}[5, 6, 6, 5]$.

If one naively assumed that each state was occupied for a quarter of the time, the probability of winning appears to be

$$p_{\text{win}} = \frac{1}{4} \left(\frac{9}{10} + \frac{1}{4} + \frac{1}{4} + \frac{7}{10} \right) = \frac{21}{40} > \frac{1}{2},$$

which is winning. Like the similar analysis in the capital dependent games, this is also incorrect. The correct probability of winning using $\pi^{B'}$ is

$$p_{\text{win}} = \frac{5}{22} \cdot \frac{9}{10} + 2 \left(\frac{3}{22} \cdot \frac{1}{4} \right) + \frac{5}{22} \cdot \frac{7}{10} = \frac{1}{2},$$

which is exactly fair.

Paradoxical Games Test

Since the construction of the games only involves using the past two results, the calculations remain tractable. Thus, we can simply use the probability of winning to find constraints for the paradox to exist.

Using

$$p_{\text{win}} = \sum_{i=1}^4 \pi_i p_i \quad (3.53)$$

with the stationary probabilities of game B' in (3.52) we get

$$p_{\text{win}} = \frac{p_1(1 + p_2 - p_4)}{p_1 p_2 + (1 - p_4)(1 + 2p_1 - p_3)}. \quad (3.54)$$

Subjecting this to the constraint $p_{\text{win}} > 1/2$ for a winning game or $p_{\text{win}} < 1/2$ for a losing game, we have the following conditions,

$$\frac{1-p}{p} > 1, \quad (3.55a)$$

$$\frac{(1-p_3)(1-p_4)}{p_1 p_2} > 1 \quad \text{and} \quad (3.55b)$$

$$\frac{(1-q_3)(1-q_4)}{q_1 q_2} < 1 \quad (3.55c)$$

for game A and B' to lose and the randomised game to win.

Range of Biasing Parameter

Using the probabilities in (3.48), the equations (3.55) become

$$\begin{aligned} \epsilon &> 0, \\ \epsilon &> 0 \quad \text{and} \\ \epsilon - 1/168 &< 0. \end{aligned}$$

The range of biasing parameters possible for the paradox to exist are

$$0 < \epsilon < 1/168. \quad (3.56)$$

3.5.4 Probability Space

Similarly to the capital dependent games there exists a probability space defined by the inequalities (3.55). Unfortunately these relations depend on five variables, which makes it difficult to visualise. To combat this we fix $p_2 = p_3$ and $p = 1/2$ to leave a visualisable probability space, plotted in Figure 3.27. The surfaces $\Pi_{B'}$ and $\Pi_{R'}$ represent the boundaries between winning and losing for game B' and the randomised game respectively. The four bounded volumes that are created by the surfaces are denoted Q_i for $i = 1, 2, 3, 4$.

In Figure 3.27, points below the surfaces are losing, and above are winning. Hence, the volumes denoted by Q_1 and Q_3 represent regions where the games are paradoxical – two losing games combine to form a winning game. The effect of making game A losing (rather than fair as shown $p = 1/2$) is to cause $\Pi_{R'}$ to lift, while $\Pi_{B'}$ is unaffected. This reduces the volumes of Q_1 and Q_3 , as expected. Note, there is almost always a region in the top right corner of the probability space where the paradox appears to always exist. This is an exception as it is an extreme case where the probabilities are very near 0 or 1.

Conversely, in the other regions, Q_2 and Q_4 , the opposite to the paradox occurs – two winning games combine to form a losing game. Adjusting game A to win slightly causes $\Pi_{R'}$ to fall, and the volumes of Q_2 and Q_4 decrease as expected.

Comparing Figure 3.27 with the probability space of the capital dependent games (Figure 3.16 on page 68), a striking difference is its lack of symmetry. Although there are regions that have opposite properties, they are not geometrically related to each other.

3.5 History Dependent Games




Figure 3.27. Probability space of the history-dependent games. $\Pi_{B'}$ and $\Pi_{R'}$ are the surfaces separating winning and losing regions for game B' and the randomised game respectively. The paradox occurs in the volumes marked Q_1 and Q_3 and the reverse effect occurs in Q_2 and Q_4 . Parameters are $\gamma = 1/2$ and $p = 1/2$.

However, there does exist some symmetry along the boundaries of $\Pi_{B'}$ and $\Pi_{R'}$. The dashed lines at the bottom of the plot are the projected lines of interception between $\Pi_{B'}$ and $\Pi_{R'}$, although this symmetry breaks down when $p \neq 1/2$.

The other noticeable departure between the two probability spaces is the volume occupied by the paradoxical regions. Using a numerical approach, approximately 0.032% of the probability space of the original games is occupied by the paradoxical region. This compares to a much larger 1.76% probability space for the history dependent games. The volumes are also dependent on other parameters (γ and M for example) and conditions (setting $p_2 = p_3$), but they do indicate about a 50 times difference. Caution must be exercised when using the comparisons as the entities are not that similar, i.e. \mathbb{R}^3 and \mathbb{R}^5 . However, a trade-off may be in the rate of return, which is larger for the original games (approximately 1.25 in Figure 3.2 on page 42) than the history dependent games (approximately 0.8 in Figure 3.22). The increased volume of the paradoxical regions are a trade-off against the smaller rate of return. This is also evident in the deterministic outcomes. See Figure 3.6 and Figure 3.23 for example.

3.6 Other Phenomena

The following sections briefly cover other selected phenomena inspired by Parrondo's games.

3.6.1 Co-operative Parrondo's Games

We have described games that are state dependent and history dependent. Inspired by this work, cooperative Parrondo games have been formed by Toral (2001). These games rely on the state of a player's neighbours. We refer to this state as whether a player has either lost or won the previous game.

A group of N players are arranged in a circle so each player has two neighbours. Game A remains unchanged, thus does not have any dependencies. Game B depends on the state of the neighbours to the left and right of a player. This gives the possible states as $\{LL, LW, WL, WW\}$, where each pair is the previous state of the left and right neighbour respectively. The games are classified by

$$C_n = \sum_i C_n(i),$$

where $C_n(i)$ is the capital of the i th player after the n th game.

The results shown in Toral (2001) demonstrate that these games indeed exhibit the properties of Parrondo's paradox. Toral also considers the variance to determine if the average winning is due to a few players becoming excessively wealthy at the expense of all the other players. This is not the case as the variance grows with n as expected. Moreover, the trends in the variance shown by the cooperative games closely resemble the standard deviation of the capital dependent games in Figure 3.5. Namely, the variance of the mixed games lies between that of the individual games.

3.6.2 Fractal Properties of Parrondo's Games

Recent work by Allison et al. (2002) has shown that Parrondo's games exhibit fractal patterns in their state space.

Parrondo's games with $M = 3$ are considered so the state space is restricted to \mathbb{R}^3 . The state space is defined by the equilibrium distribution π . This is visualised in Cartesian

3.6 Other Phenomena

space by letting π_0 , π_1 and π_2 represent the three orthogonal axes, c.f. x , y and z . From the total law of probability we have

$$\pi_0 + \pi_1 + \pi_2 = 1, \quad (3.57)$$

which is the equation of a plane. The portion of the plane that satisfies the other constraint of $0 \leq \pi_i \leq 1$ is shown in Figure 3.28. This shows all the possible state vectors that are allowed by the games.

Two sample paths are shown on the state space at the top of Figure 3.28. Since we have used the starting distribution of $\pi = [1, 0, 0]$ there are only two initial paths that can be taken as opposed to the usual eight. These are both shown in Figure 3.28. Even after the small sequence of 40 games that were used, both trajectories reduce to similar paths. The average of these similar paths reduces to a point – the stationary distribution.

Since the state space shown in Figure 3.28 is planar, it is possible to easily transform it to more appropriate axes. This is desirable to improve the visualisation of the state space.

The transforms are achieved using elementary rotations and translations. The steps are listed below and shown in Figure 3.28.

- (i) The aim is to place the state space in the xy plane, that is, to have no z dependency. Thus, we need to place the centroid of the state space, which is the *foot of the normal* in this case, at the origin of the axes. The foot is a distance of $1/\sqrt{3}$ from the origin, so we require a translation of $-1/3$ in each direction.
- (ii) Next, it is rotated about the z -axis so the base of the state space (as shown in Figure 3.28) is perpendicular to the y -axis. To maintain a reference to direction, it is rotated so the 'top' of the state space is in the positive y direction. Simple geometry reveals a $-3\pi/4$ rotation is required.
- (iii) The last step flattens the state space into the xy plane via a rotation about the x -axis. Making a rotation of $-(\arctan \sqrt{2})$ leaves the state space oriented upwards (with respect to the y axis), with its centroid at the origin as desired.




Figure 3.28. State space and transformation of Parrondo's games. The shaded area in the top plot is the valid state space \mathcal{S} for the games. Two sample paths of the randomised game are shown by the lines on \mathcal{S} . The three bottom plots are the required transformations to put \mathcal{S} in the xy plane. From left to right they are a translation, a rotation and another rotation.

The steps are summarised by the following transformation,

$$\begin{aligned}\tilde{\pi} &= R_x R_z T \pi = R \pi \\ &= \frac{1}{6} \begin{bmatrix} -3\sqrt{2} & 3\sqrt{2} & 0 & 0 \\ -\sqrt{6} & -\sqrt{6} & 2\sqrt{6} & 4\sqrt{6}/3 \\ 2\sqrt{3} & 2\sqrt{3} & 2\sqrt{3} & 2/\sqrt{3} \\ 0 & 0 & 0 & 6 \end{bmatrix} \pi, \end{aligned} \quad (3.58)$$

using the points $\pi = (\pi_0, \pi_1, \pi_2, 1)$, which takes into account the translation. R_x and R_z are rotation matrices about the x and z axes respectively and T is the translation matrix. $\tilde{\pi}$ is the transformed state space, which has the property that $\tilde{\pi}_2 = 0$ for all π .

3.6 Other Phenomena

Alternatively, the top left 3×3 submatrix can be used, and the translation carried out separately.

Using the transformation of (3.58), the fractals for Parrondo's games can be viewed, as shown in Figure 3.29. The zoomed sections demonstrate the self-similarity property of the fractal. As we magnify the fractal, the resolution diminishes. This is due to the limited number of games used (i.e. 10 000), and can be improved by playing more games to increase the number of sample points.




Figure 3.29. Fractals in Parrondo's games. The left plot shows the transformed triangular state space with points from the randomised game using (3.1) with $\epsilon = 0.005$. The smaller plots to the right show zoomed sections to demonstrate the self-similarity of the fractal.

It is interesting to note that changing the biasing parameter ϵ does not vary the fractal greatly; only slightly translating and rotating it. Also for much larger values of ϵ , so the games are not paradoxical, rotations in the zoomed sections are clearly noticeable. However, changing the probability values of p_1 and p_2 produces completely new fractal patterns.

3.6.3 Parrondo's Games are Ubiquitous

A number of authors have used the fact that the randomised games are a linear combination of the other two games to explain the paradox [Costa et al., 2002; Moraal, 2000; Pyke, 2001]. In a generalised form the games are described by

$$G(\rho_1, \rho_2, M).$$

This describes game A by setting $\rho_1 = \rho_2 = p$, game B by setting $\rho_1 = p_1$ and $\rho_2 = p_2$ and the randomised game with $\rho_1 = \gamma p + (1 - \gamma)p_1$ and $\rho_2 = \gamma p + (1 - \gamma)p_2$. The parameter γ varies the relative strengths of game A and B . Using this notation we can plot the parameter space as shown in Figure 3.30.




Figure 3.30. The probability space of Parrondo's games in p_1 and p_2 . The manifold shown is with $M = 3$ and the line segment joining P_1 and P_2 are randomised games formed by varying γ . Since the manifold is convex Parrondo's paradox exists. The manifold when $M = 2$ and the possibilities for game A are also shown.

Game A only exists on the straight line indicated, whereas game B exists everywhere in the plot area. The curved line represents the manifold that separates the winning and losing games for B when $M = 3$.

If we select game A and B as the points P_1 and P_2 , the randomised game can exist anywhere along the straight connecting line by varying γ . It is not hard to see that any randomised game in the highlighted section of the line will produce Parrondo's paradox. Furthermore, we can choose any two games of B and as long as the connecting line crosses the manifold twice, Parrondo's paradox will be exhibited. Thus, we can clearly see how two winning games can form a losing game.

The reason why this occurs is because the randomised game is formed by a *convex combination* of two other games. It is also pointed out by Costa et al. (2002) that any

3.7 Summary

game where the manifold is not a hyperplane, which are all but the simplest games, should be expected to exhibit Parrondo's paradox.

This is the reason that games with $M = 2$, previously discussed in Sec. 3.2.3, cannot exhibit the paradox. The manifold for $M = 2$ is shown by the dashed line in Figure 3.30, which is a hyperplane in \mathbb{R}^2 .

3.6.4 Related Works

After the publications of Parrondo's game in literature, papers appeared that related the games to other disciplines. Several brief descriptions are given below.

Van den Broeck et al. (1999) uses the games to give an example of a discrete "pulsating ratchet". They describe playing the games against Maxwell's demon to highlight their apparent paradoxical nature.

Several variations of the games have also been investigated. Arena et al. (2001) plays the games using chaotic switching sequences and mixing capital and history dependent games together.

The games can be related to physical systems such as lattice gas automata [Meyer and Blumer, 2001], spin models [Moraal, 2000], random walks and diffusions [Pyke, 2001], as quantum games [Ng and Abbott, 2002; Meyer and Blumer, 2001], and other general articles [Astumian, 2001; Klarreich, 2001]. The signal-to-noise ratio has also been used to describe the counter intuitive phenomenon [Fuh and Yeh, 2001].

3.7 Summary

We have given a detailed explanation of Parrondo's games and the circumstances that surround their apparent paradoxical nature. The mechanism behind the games was first described heuristically in terms of the Brownian ratchet. In the same way the ratchet directs the motion of random particles, Parrondo's games use the randomised switching of the games to direct flow of capital. This analogy between the games and the ratchet allows us to compare the respective role of the variables in each of the systems. The question is, what information can the games provide about discrete time and space ratchets that cannot be extrapolated from the standard continuous ratchet? Moreover, can the games infer anything back to the continuous ratchet?

The ratchet explanation is based on the underlying physical structure. However, as one may expect, the apparent paradoxical nature of the games can be fully explained mathematically with the use of discrete-time Markov chains. This requires no knowledge of the physical origins of the games. The mathematical analysis allows us to generalise the games and determine trends that can characterise them.

The original capital dependent games were then modified to be history dependent. That is, the rules depend on the results of previous games. This type of rule maybe more appropriate for gambling games as the results are public information whereas the capital of a player is not.

Other aspects of the games were also raised, as well as some physical systems where the games can be applied.

The study of Parrondo's games did not appear to have a direct application to managing noise for motion detection sensors. However, as with any toy model, what is important is the application of the principles illustrated by the model, rather than the model *per se*. One key principle is that noise or randomness can become useful when in the presence of a nonlinearity or an asymmetry. Furthermore, a random mixture of game *A* with game *B* could be thought of as a noisy process that breaks up the state dependence in game *B* causing it to favour the good coin. Too little of game *A* in the mixture or too much game *A* is detrimental. However, there is an optimum amount of game *A* that is mixed with *B* to maximise winnings. This points us towards the study of stochastic resonance. As we shall see in the following chapter, (i) the interaction of noise with a nonlinearity and (ii) optimised output for nonzero noise levels, are precisely the ingredients of stochastic resonance. It is these principles that we will carry through in later chapters to attempt to manage the noise in motion detectors based on insect vision.

Stochastic Resonance

NOISE in dynamical systems is usually considered a nuisance. However, in certain nonlinear systems including electronic circuits and biological sensory systems, the presence of noise can enhance the detection of weak signals. The phenomenon is termed stochastic resonance and is exhibited in many types of systems from a large number of fields.

We first introduce stochastic resonance (SR) by use of a simple “black box” system. Several detailed examples are given to gain an appreciation of SR and the wide range of fields it covers.

The required components of a stochastic resonant system are described, these include the system dynamics and noise. The quantifying measures are also described in detail.

4.1 Introduction

Engineers have normally sought to minimise the effects of noise in electronic circuits and communication systems. Nowadays however, it is acknowledged that noise or random motion is beneficial in breaking up the quantisation pattern in a video signal [Netravali and Haskell, 1988], in the dithering of analog to digital converters (ADCs) [Gammaitoni, 1995], in the area of Brownian ratchets [Bier, 1997] and in the physics of granular mixtures [Rosato et al., 1987; Jaeger et al., 1996; Makse et al., 1997]. Also it is known that when training a neural network, adding noise to the training data set can improve *network generalisation*, i.e. the neural network's ability to fit real data outside of the initial training set [Bishop, 1996, p. 346]. Noise even plays a role in game theory [Harmer et al., 2000a] and number theory [Chaitin, 1999; Planat, 2001]. There are many other fields where noise can be of benefit – an extensive review is given by Gammaitoni et al. (1998).

Stochastic resonance was originally proposed by Benzi (1980) as an explanation of the behaviour of the Earth's ice ages which exhibit a 100 000 year periodicity. This was closely followed by a number of papers [Benzi et al., 1982; Nicolis, 1982; Benzi et al., 1981] describing a general dynamical mechanism whereby small periodic perturbations could be greatly amplified by large environmental fluctuations.

Experimentally, SR was first demonstrated by Fauve and Heslot (1983) with a noise driven electronic circuit based on the Schmitt trigger. This work was also the first to characterise the phenomenon in terms of the signal-to-noise ratio. It took more than five years before the interest of physicists was sparked by a demonstration of SR in a bistable ring-laser experiment [McNamara et al., 1988]. Stochastic resonance has been reported in a variety of physical systems, and the classical theory is well established [Gammaitoni et al., 1989a; Fox and Lu, 1993; Jung and Hänggi, 1991; McNamara and Wiesenfeld, 1989].

Today SR has crossed disciplinary boundaries: its role in sensory biology is being explored in experiments on single crayfish neurons [Moss and Wiesenfeld, 1995; Wiesenfeld and Moss, 1995], and in perceptive brain functions by experiments on people's ability to resolve ambiguous figures [Simonotto et al., 1997; Gammaitoni, 1995]. SR has been shown to enhance the sensitivity of mechano-electrical transduction in the ears of frogs [Jaramillo and Wiesenfeld, 1998] and humans [Wiesenfeld and Jaramillo, 1998; Moss et al., 1996; Jaramillo and Wiesenfeld, 1998]. These new efforts, together

with attempts to exploit SR for technological advantage are the main trends in current research.

As the system and signals have evolved with SR, the measuring quantities have accordingly changed. SR was originally applied to systems with periodic signals where use of the signal-to-noise ratio is appropriate. Collins and colleagues extended SR with the use of aperiodic input signals, where an input-output cross-correlation measure is utilised. An information theoretic approach was taken by Stocks to measure the average mutual information between the input and output of the system. The metrics of SR are further discussed in Sec. 4.4.

We start this chapter by explaining the basics for those that are not familiar with SR. This is achieved via the use of an example system and showing the effects of SR in terms of the signal-to-noise ratio. This is followed by several applications of SR.

The remainder of the chapter consists of three main parts; a description of common nonlinear systems that exhibit SR, some issues relating to noise, and how to measure SR. These parts are quite detailed, and are organised so that the systems described in Sec. 4.2 are used to show the different measures given towards the end of the chapter.

4.1.1 The Basics

This section is intended to give a simplified “black box” description of SR. We describe the components required to observe stochastic resonance, how to measure it, and give a simplistic analogy.

A nonlinear system is essential for SR to be exhibited. If we consider a linear system, the output is governed by linear response theory (LRT). This states that the signal-to-noise ratio (SNR) at the output is proportional to the SNR at the input. Hence any increase in noise intensity at the input will correspondingly increase the noise at the output. It is common for the nonlinearities of a system to be provided by some type of threshold or barrier function.

For traditional stochastic resonance the amplitude of the input signal is *subthreshold* with respect to the system. This means in the absence of noise the system response will not change, since the signal is always below the threshold. This is important, as we will see later, because if the signal becomes suprathreshold the addition of noise only serves to degrade the response.

4.1 Introduction

The final ingredient is a noise source. This can be realised in any medium that influences the system. For example, the chemical reactions in the media used for spatiotemporal stochastic resonance depend on light intensity [Kádár et al., 1998; Balázsi et al., 2000]. The most common source of noise is electrical noise, as many experiments use electrical circuits to model physical systems. The strength of the noise can be given in terms of its variance σ^2 , or its standard deviation σ . The noise is explained in more detail in Sec. 4.3.



Figure 4.1. Stochastic resonator. The box represents a system that combines the signal and noise using a nonlinear function. The signals are in the time domain.

The nonlinear system is referred to as a *stochastic resonator* (STR) in this context and is shown in Figure 4.1, with the waveforms from an example system shown in Figure 4.2. The threshold is only intended as a visual aid and portrays a limited amount of information about the system. Figure 4.2a shows the subthreshold signal $x(t)$, the threshold θ , and the summed signal and noise $x(t) + \eta(t)$. Figure 4.2b shows the response $y(t)$ of the STR. When the total excitation of the signal and noise exceed the threshold the STR undergoes a transition and the response changes states. For the STR in this example there is only one threshold. However, when the state of the system changes, the threshold moves away from the current operating point, i.e the nonlinear system exhibits a hysteresis. Hence two thresholds are shown in Figure 4.2a. The number of transitions in the response are influenced by the intensity of noise added to the signal.

In the absence of noise there are no transitions, thus no information about the signal is present in the response. Conversely, when there is a large amount of noise added to the signal, the information that is present is swamped by the large number of transitions caused purely by noise. Therefore, there must exist some nonzero value of noise that gives rise to the response containing a maximal amount of information about the input signal. Obviously for the narrowband signal shown in Figure 4.2, the only information we can transfer is the frequency. At this optimal noise intensity there is a maximum *coherence* (i.e. the ‘cooperation’) between the input and output signal. A method to measure this coherence is required.




Figure 4.2. Example input signals and response to a STR. (a) The clean sine wave is the input signal, which is well below the dashed thresholds. The jagged waveform is the input sine wave plus noise. (b) The response from the STR.

For the simple case of a narrowband signal, the coherence is conveniently quantified by the power spectral density (PSD) of the response $S(f)$ in the frequency domain. At the fundamental frequency f_0 , the amplitude of $S(f)$ increases with noise intensity to a maximum, then decreases with further added noise. It is common to define the coherence in terms of the SNR given in decibels as $10 \log_{10}(S/B)$, where B represents the background noise. This is described in more detail in Sec. 4.4.1. This quantity similarly exhibits a peak although shifted to a slightly different noise intensity.

The SNR is shown in Figure 4.3. The calculated SNR emerges from the top at low noise intensities because the example system is continuous [Luchinsky et al., 1999a], and operating in the small signal limit where it is assumed linear. Thus, for low noise intensities, the SNR is high and the input is proportional to the output. The SNR therefore decreases when a small amount of noise is added. However, once enough noise is added transitions begin to occur. This means we are not operating in the small signal limit and the system is nonlinear. It is from this point that noise benefits the system. This behaviour differs for discrete systems where there is no response in the small signal limit, thus the SNR starts from zero.

The calculated SNR emerges from a high value at low noise intensities because the example system is operating in the small signal limit, where a linear response can be assumed. Thus, for a low σ the SNR is high and the SNR at the input and output are directly proportional. Therefore, the SNR decreases when small amounts of noise are added. However, once more noise is added transitions begin to occur. This means we

4.1 Introduction




Figure 4.3. Power spectral density and the signal-to-noise ratio from Figure 4.2. (a) The PSD of the response shown in Figure 4.2b. The measures of the signal and background noise power are indicated by the quantities S and B . (b) The curve originating from the top is the SNR directly calculated from the PSD using many noise intensities. The smooth curve is the scaled predicted SNR for the STR in this example.

are not in the small signal limit and the system is nonlinear. It is from this point that noise benefits the system.

Strictly speaking, the SNR curve does not display a *bona fide* resonance [Gammaitoni, 1995; Bulsara and Gammaitoni, 1996]. By definition, the resonance of the SNR should be a function of the forcing frequency, here it is a function of the noise intensity. A measure that shows a true stochastic resonance in terms of residence times is discussed later, in Sec. 4.4.2. Nevertheless, the definition of SNR is simple to implement, and generally accepted in engineering practice.

Signal-to-Threshold Distance

We will digress for a moment, with reason, to investigate the effects of different input signal strengths. These highlight the limitations of SR and direct us to the type of systems that SR is applicable to. They also serve as a useful tool for understanding other forms of SR.

Using the example STR, the SNRs for increasing signal strengths are shown in Figure 4.4. The signal strengths are increasing from the bottom to top curves. The purpose of Figure 4.4 is to show that once the signal exceeds the threshold (i.e. *suprathreshold*), noise only serves to degrade the output. We must resort back to traditional techniques to reduce the absolute noise in the system to improve the response.




Figure 4.4. Improvement due to SR for different signal to threshold distances. The bottom curve has the maximum signal-to-threshold distance, which is reduced when moving up through the curves. The top curve has a negative distance, i.e. the signal exceeds the threshold. There is no evidence of SR in this case.

The important factor is the signal-to-threshold distance. This can be reduced by either increasing the signal strength or reducing the threshold. In most engineering applications these options are usually fairly simple – much easier than adding noise of the correct intensity to the system.

This brings us to an important point illustrated by the following example. Consider a system with a subthreshold signal with a certain amount of internal noise – the point marked with a cross in Figure 4.4 say. One option to improve the signal-to-noise ratio is by eliminating the signal-to-threshold distance, or alternatively to use SR by adding extra noise to the system. The two options are shown in Figure 4.4 by the arrows moving away from the cross. Although the SNR improves in both cases, the absolute value of the SNR is different. Reducing the signal-to-threshold distance produces a higher SNR. The result: noise should only be added to a system as a last resort, the threshold should be tuned first, if possible [Harmer and Abbott, 1999c; Abbott, 2001]. Thus, SR is appropriate for systems that are constrained so that the threshold cannot be tuned (e.g. hearing aids [Moss et al., 1996; Bruce et al., 1999c]), or where varying the threshold detrimentally affects another part of the system (e.g. signal detection [Dawe, 2000b]). This point has also been highlighted by others [DeWeese and Bialek, 1995; Bulsara and Zador, 1996; Chapeau-Blondeau, 1997b].

4.1 Introduction

Simplistic Model

A simple model used to illustrate SR is a small marble in a two-egg carton [Bulsara and Gammaitoni, 1996; Moss and Wiesenfeld, 1995].

A gentle rhythmic force rocks the carton back and forth. The rocking motion is gentle enough such that the marble does not jump between the two wells. This corresponds to a weak periodic signal. Let us assume that the motion of the marble is only detectable when it jumps between wells. Under the influence of the weak force the marble just rolls around the bottom of one of the wells and the effect of the weak periodic force remains hidden.

The only way to make the marble jump between wells would be to shake the carton, that is adding noise. Initially this would seem to mask the rocking motion of the marble, but this is not the case. The weak periodic forcing coupled with the noise can occasionally give the marble enough energy to jump between the wells. The theory behind SR relies on the fact that the jumps between the wells are not entirely unpredictable. There is a greater chance of the marble switching wells when the weak periodic force is at a peak. Hence for moderate noise, the switchings possess a degree of coherence with the underlying signal.

If too little noise is added there are not enough transitions between wells to make the signal detectable. Similarly, too much noise dominates the transitions, which reveals no information about the signal. There is an optimal amount of noise that conveys the most information about the signal.

4.1.2 Examples of SR

This section reviews some case studies of demonstrate the concepts associated with stochastic resonance and the diversity of fields it covers. The first example examines the sensing systems of the crayfish [Moss and Wiesenfeld, 1995; Wiesenfeld and Moss, 1995; Douglas et al., 1993], the second relates to visual perception, and the third uses SR to help detect weak signals.

Biological – Crayfish

We consider the crayfish *Procambarus clarkii* and how it may be able to detect predators. Remarkably, the crayfish's sense organs can detect faint information bearing signals in noisy environments.

A system of mechanoreceptor cells, which end in fine hairs on the tail of the crayfish (see Figure 4.5) are specialised to detect weak water motion such as that generated by the waving tail of a nearby predatory fish. The hair cells most likely serve as an early warning system, which may explain why crayfish are one of the oldest and most well adapted animals in existence [Moss and Wiesenfeld, 1995].

When moved, the hair generates a nerve impulse, measurable as a short duration electric spike. The impulse travels to a bundle of nerve cells which handle all incoming information from the hair cells. These nerve cells make basic decisions such as activating the animal's escape reflexes. Evidence indicates that the sensory neurons function much like nonlinear threshold detectors. Below a certain level, very weak motion of the hair does not generate a nerve impulse.

An experiment to mimic the noisy environment in which crayfish normally live using an excised piece of crayfish tail is described in Moss and Wiesenfeld (1995). The procedure essentially involves moving the tail periodically in water with an electromechanical transducer until the nerve impulses are nearly undetectable. Noise is then added to the transducer with increasing intensity. This yields the signal-to-noise ratio shown in Figure 4.5, which displays stochastic resonance.




Figure 4.5. Crayfish mechanoreceptors and the SNR in the presence of noise. The crayfish tail on the left showing the fine hairs that act as the mechanoreceptors. The SNR of the mechanoreceptors when exposed to external noise. The squares are from experimental data, the diamonds from neuron simulations and the solid line from theory. Reproduced from Moss and Wiesenfeld (1995).

4.1 Introduction

Figure 4.5 also shows the results for theoretical predictions and an electronic simulation of a threshold detector for comparison [Wiesenfeld and Moss, 1995; Douglas et al., 1993]. The discrepancies at high and low noise levels are due respectively to neurons exhibiting a refractory period and the unavoidable internal noise in sensory systems.

Visual Perception and Dithering

Stochastic resonance can be observed when dealing with visual perception of ambiguous figures [Simonotto et al., 1997; Riani and Simonotto, 1994]. Perception of figures, such as the Necker cube (a wire frame cube) are characterised by noisy bistable dynamics [Bialek and DeWeese, 1995]. Two different interpretations of the Necker cube are alternatively perceived by the observer with a stochastic time course. The bistable perceptual process involved can be modelled by a noisy auto-associative neural network. Using this system, the results indicate that SR can be demonstrated [Riani and Simonotto, 1994; Riani and Simonotto, 1995].

Another example is the human visual system, which derives from threshold theory [Simonotto et al., 1997], and is closely related to the dithering effect [Gammaitoni, 1995]. Consider a system that is capable of transmitting single bits of information, each of which marks a threshold crossing. A visual realisation of this is shown in Figure 4.6, which was generated following the procedure in Simonotto et al. (1997). The original grey scale image of the popular 'Lena' used in signal processing is depressed beneath a threshold and noise added to the grey value in each pixel. The result compared to the threshold value and pixels with intensity above the threshold are made white, the others black. Every pixel contains one bit of information, whether or not the threshold has been crossed. Figure 4.6 shows the result of adding noise of three intensities, increasing from left to right. The optimal noise intensity in the second from the right maximises the information content. Note that the pictures shown in Figure 4.6 only have noise that varies spatially, additional improvement in perceived picture quality can be gained by varying the noise temporally [Simonotto et al., 1997].

The type of stochastic resonance shown in Figure 4.6 is very similar to random dithering, which has been widely used by electrical engineers since the 1950s. Furthermore, it can be described using the same techniques (as SR) of noise induced activated processes [Gammaitoni, 1995].

Random dithering is a well known noise induced phenomenon that is typically used in analog-to-digital converters (ADC). It is a technique where a small amount of noise is




Figure 4.6. Visual perception of SR in image dithering. The original image of Lena on the left is 256 by 256 pixels with 256 grey scales. This image was scaled to 30 shades of grey below the threshold and white noise added to each pixel. Noise intensities increased moving from left to right.

added to the signal before quantisation. Due to the finite storage space of the digitised signal (i.e. word length of ADC) the amplitude becomes quantised to form a stair case like signal. Therefore, the response characteristics of the ADC are nonlinear and the signal becomes distorted. The distortion is referred to as quantisation error, $q_e = y - x$ being the difference between the input x and the output y of an ADC.

Studies of an appropriate dither signal to add have been undertaken during the past 30 years with the following conclusions [Gammaitoni, 1995].

- (i) The addition of a proper dither signal can cause independence and whitening of the quantisation error, resulting in both a reduction of signal distortion and an improvement of the system's dynamic range.
- (ii) The best choice for the dither signal is a random dither, uniformly distributed within an interval of amplitude equal to the quantisation step.

Figure 4.7a-b shows the average output when uniform ($\langle y_u \rangle$) and Gaussian ($\langle y_g \rangle$) distributed noises are added to the system described by

$$y = \begin{cases} 0 & \text{for } x < 1/2, \\ 1 & \text{for } x \geq 1/2. \end{cases} \quad (4.1)$$

This system is essentially a 1 bit ADC. From Figure 4.7a, the average response is approximately linear when the noise intensity matches the quantisation step size. There is always some error present for the Gaussian distributed noise due to the average nonlinear response.

4.1 Introduction

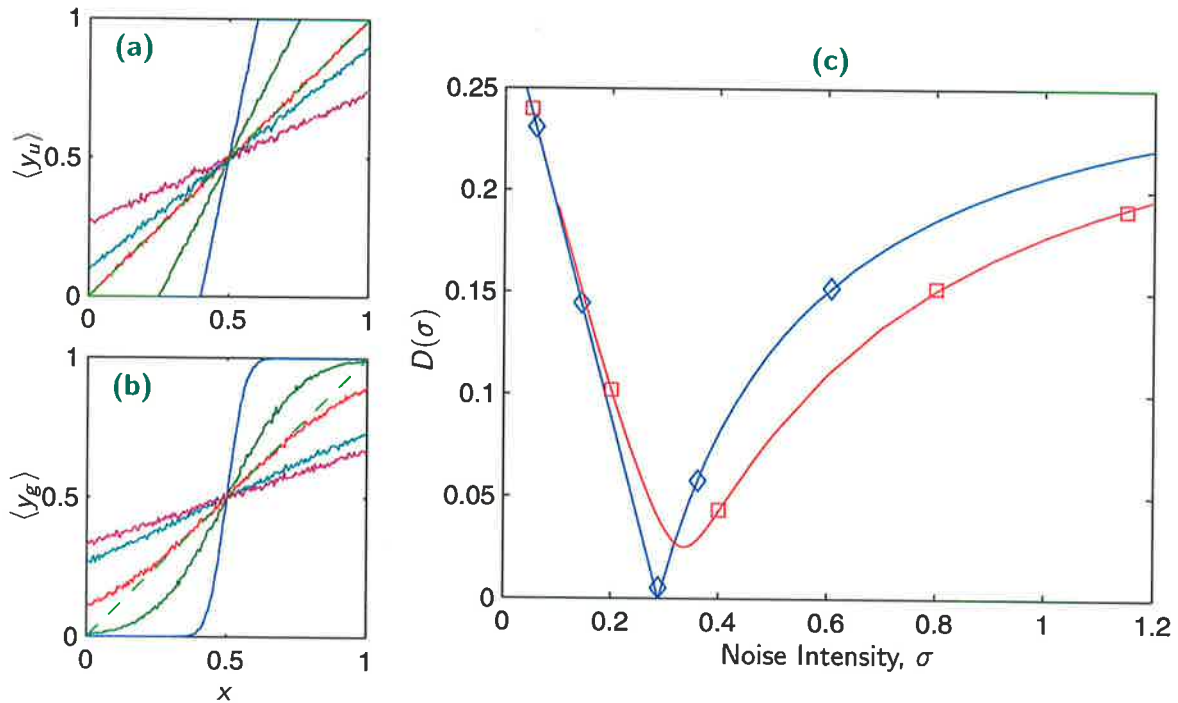


Figure 4.7. Dithering in a 1 bit ADC with uniform and Gaussian noise. (a-b) Average ensemble response to (4.1). In (a) uniform noise is added with limits 0.2, 0.5, 1.0, 1.15 and 2.1 and in (b) Gaussian noise is added with standard deviations of 0.05, 0.2, 0.4, 0.8 and 1.15 both in the clockwise direction. (c) The root mean square of the errors from an ideal response for uniform (\diamond) and Gaussian (\square) noise.

A measure of error between the average response and input can be taken using a root mean square of the differences [Gammaitoni, 1995],

$$D = \sqrt{\int_0^1 (\langle y \rangle - x)^2 dx}. \quad (4.2)$$

This is plotted, as a function of noise intensity σ in Figure 4.7c. Note the standard deviation of uniform noise within amplitude $\pm L$ is $L/\sqrt{3}$, which is approximately 0.3 for $L = 1/2$. This is in agreement of point (ii) stated previously, that when the uniform noise intensity is equal to the quantisation step the best reduction of signal distortion is achieved. In the averaging limit, for perfectly matched noise there is in fact no error. However, for Gaussian noise the error never reduces to zero.

Dithering also serves to increase the dynamic range of the ADC as suggested by point (i). That is, changes in the analog input signal smaller than the quantisation step may have been previously undetected, but with the addition of noise may be detected. The following example is used to explain this. Consider a signal that has an amplitude of 0.3 for the system described by (4.1). The output can either be 0 or 1 and in absence of

noise it will be 0. If we add the optimal amount of noise ($L = 1$), the output will spend on average thirty percent of the time at 1 and seventy percent at 0. In order to achieve this result, the noise PDF must be linear over the entire input range, hence uniform noise must be used. Thus, by averaging at least ten samples the dynamic range can be increased to discriminate a 0.3 amplitude signal. This is referred to as *oversampling* and is used in modern ADC systems [Vanderkooy and Lipshitz, 1984; Hauser, 1991; Stewart and Pfann, 1998]. The higher the internal oversampling frequency, the more accurate the averaged output.

The system described by (4.1) can be extended as in Gammaitoni (1995) to contain two symmetric thresholds at $\pm 1/2$ with a narrowband input signal smaller than the thresholds. The addition of noise then induces transitions at the peaks of the signal. The corresponding output PSD at the signal frequency, when plotted against the noise intensity is of similar shape to that shown in Figure 4.3b.

An alternative arrangement to simply adding noise to the analog input signal is shown in Figure 4.8 [Kikkert and Biddeli, 1997]. The noise is used twice; it is passed through




Figure 4.8. Additive dither block diagram. Noise is firstly generated from a pseudo-random binary generator. It is passed through a DAC and added to the signal, then digitally subtracted from the ADC output to produce the final sampled output. The two negatives at the input summer form the negative sum, which means a difference operation occurs at the output summer.

a DAC and added to the analog signal, and it is also subtracted (digitally) from the digitised signal to give the final sampled output. This is intended for high frequency ADCs where oversampling may be an issue, but higher resolution is also required. The noise source is implemented as a digital pseudo-random noise generator and needs to be passed through a high speed DAC to be added to the input signal. It also needs to be adjusted in magnitude and time (i.e. delayed) before being subtracted from the ADC output. This arrangement achieves the same goals as a standard ADC by linearising the characteristics.

4.1 Introduction

The question remains, what is the difference between SR and dithering? Both are noise added techniques and require the optimisation of certain parameters. The short answer is that dithering is one of the branches of all the phenomena that are labelled as stochastic resonance. Andò and Graziani (2000, p. 70). discuss subtle differences can be made when comparing dithering and SR in general. The basic idea of SR is to determine the level of noise intensity that maximises the interactions between the input and the system to give the optimal coherence between the output and input. On the other hand, dithering is used to modify the input signal to guarantee interactions with the system. This is used to increase the resolution of the system. Finally, the theoretical approaches are different. For SR the characteristic (relaxation) time of the system must be taken into account, whereas they can be ignored for dithering as it uses a static system (assuming the mean signal frequency is smaller than the inverse switching time of the comparator).

Detection of Weak Signals Using Sonar

Recently, stochastic resonance has been attempted in the development of a new type of passive sonar detection [Dawe, 2000b]. The sonar system is designed to detect *stealthy* objects and referred to as 'improved sonar harnessing temporal acoustic resonances' (ISHTAR). This builds on conventional sonar processing, but utilises the injection of artificial signal fluctuations (i.e. noise) into the system. Alternative SR methods have also been developed by others [Asdi and Tewfik, 1995; Nafie and Tewfik, 1998; Inchiosa and Bulsara, 1996].

To detect a weak signal, or target, we need to make a decision – this invariably involves the use of a threshold. The setting of this threshold is crucial. If it is too high only strong targets will be detected, too low and too many false alarms will be raised. Such information is typically plotted on receiver-operating-characteristic (ROC) curves, see Figure 4.9 [Urlick, 1967, p. 308]. The *detection probability* ($p(D)$, the probability that if a signal is present, it will be detected) is plotted against the *probability of a false alarm* ($p(FA)$, the probability that the threshold is crossed due to noise). Hence, as the threshold is varied a curve is traced out.

The ROC curves are completely determined by the probability density functions (PDF) of the signal and noise at the receiver where the threshold is set [Urlick, 1967]. This is shown in Figure 4.9 for a Gaussian distributed signal and noise.




Figure 4.9. Probability density distributions and ROC curves. (a) The area to the right of the threshold of the noise PDF is the probability of detection while the area to the right of the signal plus noise is the probability of false alarm. (b) Creating the ROC from (a) by varying the threshold.

Some details for Gaussian statistics are given to show the improvement of the $p(D)$ for a given $p(FA)$ of the ISTHAR.

For a signal and noise with Gaussian statistics, the PDF is given by

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad -\infty < x < \infty \quad (4.3)$$

where μ and σ^2 are the mean and variance respectively. Dawe (2000b) uses the subscripts s , n and a to refer to the signal, noise and added signal fluctuations, which are all Gaussian, respectively. Then as defined in Urick (1967, p. 309), the detection index is given by

$$d = \frac{(\mu_{sna} - \mu_{na})^2}{\sigma_{sna}^2}, \quad (4.4)$$

where $\mu_{sna} = \mu_s + \mu_n + \mu_a$ and $\sigma_{sna}^2 = \sigma_s^2 + \sigma_n^2 + \sigma_a^2$. If $d = 0$, then $\mu_{sna} = \mu_{na}$ which means the signal mean is vanishingly small. Conversely, a large d means the signal is large compared to the noise. If the detected signal is constant then $\sigma_s^2 = 0$ and $d_n = \mu_s^2/\sigma_n^2$ with no added noise. However, if the signal is fluctuating, we have $d_{sn} = \mu_s^2/(\sigma_s^2 + \sigma_n^2)$, and with the added noise $d_{sna} = \mu_s^2/(\sigma_s^2 + \sigma_n^2 + \sigma_a^2)$. The term *fluctuating signal* simply refers to a signal with added Gaussian noise. For consistency, we will use this terminology as it is used in Dawe's work.

Urick also defines the fluctuation index as

$$k^2 = \frac{\sigma^2}{\sigma_n^2}, \quad (4.5)$$

4.1 Introduction

including the required subscripts as necessary. For example, including signal and noise would give $k^2 = (\sigma_s^2 + \sigma_n^2)/\sigma_n^2$. This quantifies the amount of fluctuations with respect to noise. For $k_{sn} = 1$, $\sigma_s^2 = 0$ which means the signal is not fluctuating at all, and for $k_{sn} = 2$ we get $\sigma_s^2 = 4\sigma_n^2$.

From Figure 4.9a we see that $p(FA) = \int_T^\infty f_n(x) dx$, which we can evaluate using the appropriate μ and σ^2 as

$$p(FA) = \frac{1}{2} \operatorname{erfc} \left(\frac{T - \mu_n}{\sqrt{2}\sigma_n} \right), \quad (4.6)$$

where $\operatorname{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty \exp(-t^2) dt$.

Similarly, we can find the $p(D)$ for a fluctuating signal (i.e. $k_{sn} \neq 1$) in the presence of noise by evaluating $\int_T^\infty f_{sn}(x) dx$, which gives

$$p(D_{sn}) = \frac{1}{2} \operatorname{erfc} \left(\frac{1}{\sqrt{2}k_{sn}} \left(\frac{T}{\sigma_n} - \sqrt{d_n} - \frac{\mu_n}{\sigma_n} \right) \right). \quad (4.7)$$

With added artificial signal fluctuations included, we have

$$p(D_{sna}) = \frac{1}{2} \operatorname{erfc} \left(\frac{1}{\sqrt{2}k_{sna}} \left(\frac{T}{\sigma_n} - \sqrt{d_n} - \frac{\mu_n}{\sigma_n} - \frac{\mu_a}{\sigma_n} \right) \right). \quad (4.8)$$

By plotting (4.7) against (4.6) the ROC curves are generated as shown in Figure 4.10. The solid lines are for a constant signal ($k_{sn} = 1$) and the dashed lines are for added fluctuations ($k_{sn} = 2$). Note that when $\mu_a = 0$, k_{sn} and k_{sna} are interchangeable, seen by comparing (4.7) and (4.8).

From Figure 4.10, it is evident that the addition of noise can increase the $p(D)$ for a fixed $p(FA)$. Take for example when $p(FA) = 0.01\%$ with $d_n = 4$. Without any added signal fluctuations, the probability of detection is about 4% which increases to about 20% by setting $k_{sn} = 2$.

However, as with most systems, an improvement in one area usually compromises another. Here it is the detection of the stronger signals that suffer. Seen from Figure 4.10, for higher values of d_n the addition of noise decreases the $p(D)$ for a large range of $p(FA)$. The solution, as used by Dawe (2000b) is to use the conventional system to effectively detect strong signals in tandem with the new ISHTAR as demonstrated in Figure 4.11. The results of Dawe's algorithm in sea trials show the improvement at the 745 Hz sonobuoy frequency, it is clearly observable in the top plot.

It is shown by Dawe and Galbreath (2000) that improvements in detection probability can be gained for signals and noises of exponential, Rayleigh and chi-squared statistics.




Figure 4.10. ROC curves showing enhanced signal detection. The solid lines are for a constant signal ($k_{sn} = 1$) and the dashed for a fluctuating signal ($k_{sn} = 2$). The crosses show the improvement that can be gained by added fluctuations to the received signal.

Discussion on ISHTAR

Since the publication of Dawe's work there have been a number of inconsistencies and problems raised [Nuttall, 2001]. These are to do with the assumptions of the signal, lack of simulations and problems with mismatched phase during the FFT process.

The initial assumption that is used by Dawe [Dawe, 2000b], is that the received signal is Gaussian. However, the sea trials that are presented use pulses of a pure tone at 745 or 800 Hz. Also, the description surrounding the ROC curves (as shown in Figure 4.10) states that the signal is fluctuating, described as a dc signal in the presence of additive Gaussian noise.

If this is the case, then an optimum likelihood ratio (LR) processor [Helstrom, 1968], also known as a maximum likelihood (ML) detector), says to cross-correlate the received signal against a constant reference and compare the output with a threshold [Nuttall, 2001]. The processor does not require any other processes and it gives the maximum $p(D)$ for a specified $p(FA)$. According to the LR processor, the solid lines in Figure 4.10 cannot be improved. Thus, to get improvement as demonstrated by ISHTAR, the received signal and noise statistics must be different to those described.

It has been stated by Dawe that the ISHTAR processor does not work on simulated data, only on sea trial data. If this were to occur in any application, it would clearly




Figure 4.11. Sea trials of ISHTAR showing improved signal detection. The test signal is a train of CW pulses received at 745 Hz. The pulses last for 0.1 seconds and are separated by 1.0 seconds to represent a fluctuating signal. The bottom plot is a lofargram of time (seconds) versus frequency of a conventional detector, the vertical lines are interference from the generator. The middle plot is an automatic line integration (ALI) histogram that has summed the received power over 120 seconds for each frequency of the conventional detector, i.e. the lofargram. An additional ALI plot which is similar to the conventional ALI uses values that have been individually adjusted using the ISHTAR technique (not shown). The difference at 745 Hz in either the ALI or ISHTAR ALI is not immediately obvious. However, the top plot, which is the difference between the conventional ALI and the ISHTAR ALI (DALI) clearly shows a peak at 745 Hz. Reproduced from Dawe (2000b).

point to some type of flaw in the modelling of the system characteristics. In this case, it is possibly ocean channel effects, otherwise the received improvements on simulated data should be able to be replicated.

Finally we look at the FFT stage and the bin phases in adjacent time segments [Nuttall, 2001; Stocks, 2000a]. First, consider a Gaussian received signal. The complex FFT bins are uncorrelated with each other and the phases are uniformly distributed. With a rectangular window, there is no dependence between adjacent temporal FFT bins, and with a 50% overlapped Hanning window there is only about a 5% correlation. Thus, for a Gaussian process, there is no, or virtually no dependence between the phase in FFT bins of adjacent time segments.

However, for a pure tone as used in the sea trials, there is some carry over of FFT phase from one time segment to the next. This means that the FFT bin phases in adjacent time segments are not the same, and furthermore, the amount of phase shift between time segments is unknown. Therefore, it is unclear how one can average the phases over adjacent time segments.

To explain, consider a pure test tone at a received signal frequency f_0 , which has a period T_0 . Let the window be taken at some time t_n and the next time segment be t_{n+1} . To achieve the same phase in each of the FFT bins requires the difference between t_{n+1} and t_n to be some multiple of T_0 , i.e. kT_0 , where k is an integer. The problem is that f_0 (and consequently T_0) is generally unknown, perhaps due to small Doppler shifts in the signal. Thus, k is unlikely to be an integer and the phases of adjacent FFT bins will differ.

In summary, although Dawe has shown his ISHTAR detector to work in real sea trials, more explanations are required. Unfortunately details are hard to gather as they are classified [Dawe, 2000a]. There does not seem to be much evidence that the signals are Gaussian, but are in fact narrowband.

4.2 Systems Demonstrating SR

There are many systems that have the properties to display SR. The only mandatory specification is the system must be nonlinear. The systems used in this section can be broadly categorised as excitable or bistable.

4.2 Systems Demonstrating SR

For the excitable systems we consider threshold devices and neuron models. We concentrate on the simplified FitzHugh-Nagumo neuron model, which is often simplified crudely with a threshold device.

The double well potential is used as the archetypal bistable system. Discrete and piecewise linear bistable characteristics are formed with simple operational amplifier circuits with feedback.

4.2.1 Threshold Devices

One of the most fundamental nonlinear systems is the threshold device, of which there are several varieties. These include those that generate a fixed width pulse when the threshold is crossed [Jung, 1995a; Jung, 1995b; Gingl et al., 1995; Kiss, 1996; Loerincz et al., 1996], and those that stay excited whenever the signal is greater than the threshold [Chapeau-Blondeau, 1996; Chapeau-Blondeau, 1997b; Chapeau-Blondeau and Godivier, 1997; Godivier and Chapeau-Blondeau, 1997].

Level Crossing Circuit (LCC)

The LCC can be described by the Heaviside function

$$y(t) = \begin{cases} 0 & \text{if } x(t) \leq \theta \\ 1 & \text{otherwise,} \end{cases} \quad (4.9)$$

for an input $x(t)$, output $y(t)$ and threshold θ . If the input exceeds the threshold, for however long, the output remains in the high state. Due to the lack of any feedback elements the device is memoryless, it forms a *static* system. Compare this to the Schmitt trigger (described latter in Sec. 4.2.3) that has memory and forms a dynamic system.

The transfer function (4.9) can be implemented using a single operational amplifier (op amp) as shown in Figure 4.12a; the transfer characteristic is shown in Figure 4.12b. The output voltage of the op amp is $v_o = A(v_i - V_{th})$ where A is the amplification. Due to the typically large value of A (10^5 - 10^7) the output essentially switches between the power supply rails, $\pm V_s$. Since the scale of the device switching time is much faster than any signals of interest, it can be considered for all intents and purposes as a two state device. The transfer function for the circuit is plotted in Figure 4.12b, and described by

$$v_o(t) = \begin{cases} -V_s & \text{if } v_i(t) < V_{th} \\ +V_s & \text{if } v_i(t) > V_{th}, \end{cases} \quad (4.10)$$




Figure 4.12. Level crossing circuit and characteristic. (a) The LCC circuit using an operational amplifier. (b) The corresponding voltage input-output transfer function.

which closely resembles the Heaviside function (4.9).

Level Crossing Detector (LCD)

Another implementation of the threshold device, which is commonly used as a simplification of a neuron, emits a fixed width pulse when excited. The asymmetric version produces a short pulse with amplitude A and duration τ_0 whenever the input crosses the positive threshold in the increasing direction. The response is a random time sequence of positive pulses. A symmetric system contains a negative threshold and produces negative peaks when the negative threshold is crossed in the negative direction.

Figure 4.13 shows the response of the LCC and LCD to a periodic signal plus noise. The magnitudes of the pulses have been varied for demonstration purposes.

4.2.2 Neuron Models

The role of biological sensory systems to process noisy signals is remarkable. Due to the intrinsic internal noise of neurons, their output SNRs are in the order of 0 dB [Bialek et al., 1993]. Even though this is many orders of magnitude lower than a standard “hi-fi” system, their sensitivity, dynamic range and adaptability remain unsurpassed by man-made signal processing systems.

Given that nature has had millions of years of evolution to perfect this detection of weak signals in noisy environments, it makes sense to study these sensory neurons in an effort to create an efficient method to deal with noise. The study of SR in sensory




Figure 4.13. Examples of thresholding schemes. The input signal consists of a sinusoidal wave with added noise. Asymmetric systems only trigger on the upper threshold whereas the symmetric systems trigger on both thresholds. The amplitudes of the responses have been varied for demonstration purposes.

neurons is a hot area of importance to both neurophysiology and biologically inspired smart sensors and detection devices [Koch, 1998].

We briefly describe how a neuron functions, and give several neuronal models that can be used for simulations.

The Physical Structure

The typical structure of a biological neuron is shown in Figure 4.14 [Lavitan and Kaczmarek, 1997]. The neuron is connected to thousands of neighbouring neurons via the dendritic tree. The weighted strengths of all the received signals are spatially and temporally integrated. The aggregate input is then passed via the soma to the entrance of the axon – the axon hillock. If the aggregate input is greater than the axon hillock's threshold value the neuron fires and an impulse is sent along the axon. The strength of the impulse is independent of how much the input exceeds the threshold, and it remains constant through the divisions at the end of the axon.

A new signal is produced at the receiving neuron at the synaptic terminal. The strength and polarity of the new signal are flexible and depend on the neurochemical characteristics of the synapse.

The electrical activity of the neuron can be described as silent, beating or bursting, and the shapes of the impulses (or action potentials) can vary largely for different neurons [Lavitan and Kaczmarek, 1997, p. 59]. The height difference from the resting state of




Figure 4.14. Typical structure of sending and receiving neurons. The signals are collected from the dendritic tree (which could be a sensory organ) and other neurons, which are then spatially and temporally integrated. If this exceeds a threshold the neuron fires and an impulse is sent along the axon. Reproduced and modified from The McGraw-Hill Companies (2001, Psychology Image Bank).

the action potential can vary from tens to hundreds of millivolts. The duration of the action potential is in the order of milliseconds.

From a signal processing point of view, neurons act like voltage controlled oscillators. The larger the aggregated input signal is over the threshold, the quicker another firing will occur due to the temporal integration. Thus, we have an AM (amplitude modulation) to FM (frequency modulation) type of device. This means that the amplitude information about the signal is now encoded by the temporal changes in the firing rate.

Integrate-and-Fire Model

We consider the integrate-and-fire neuron model described by

$$\dot{v} = -av + a - \delta + s(t) + \eta(t), \quad (4.11)$$

where v is the voltage across the model neuron, a and δ are constants with δ/a being the barrier height, $s(t)$ the input signal and $\eta(t)$ the noise. The firing threshold is set for $v = 1$. The response to a periodic input with and without noise is shown in Figure 4.15.

4.2 Systems Demonstrating SR




Figure 4.15. Response of the integrate-and-fire neuron model. The lower trace is the signal attenuated by 10 dB and the trace above is the response, $v(t)$. The output is the spike train shown by the vertical lines.

Hodgkin-Huxley (HH) Model

The HH equations, devised by Hodgkin and Huxley (1952), model the electrical activity in squid axons. The dynamics of the HH model are reasonably sophisticated and are governed by a number of nonlinear ODEs. These contain variables that control the membrane voltage and (in)activation channels that have voltage dependent rates. An application of the HH model is given in Collins et al. (1995a). The HH model is employed when an accurate representation of a sensory neuron is required. Alternatively, the following simplified model can be employed.

FitzHugh-Nagumo (FHN) Model

The dynamics of the FHN neuronal model are a simplification of the HH equations for the electrical activity in an excitable membrane [FitzHugh, 1961]. The FHN system is defined by

$$\epsilon \dot{v} = v(v - a)(1 - v) - w + A + s(t) + \eta(t), \quad (4.12a)$$

$$\dot{w} = v - w - b, \quad (4.12b)$$

where $v(t)$ is the fast membrane potential voltage variable, $w(t)$ is a slow recovery variable and $A = 0.04$ is the constant activation (threshold). The other constants $\epsilon = 0.005$, $a = 0.5$ and $b = 0.15$ are set to commonly used values [Longtin, 1993; Chialvo et al., 1997]. One advantage of this model is the ease in which an analog circuit can be constructed using a resistor, capacitor, inductor and a nonlinear device [Keener, 1991].




Figure 4.16. Response of the FitzHugh-Nagumo neuron model. The lower trace is the slow recovery voltage $w(t)$ and the upper trace is the membrane voltage $v(t)$ with parameters $a = 0.5$ and $\delta = 0.01$. For the first 25 seconds the input is a suprathreshold sinusoidal signal, which is then followed by pure noise. The output is the spike train shown by the vertical lines.

The response of the membrane voltage $v(t)$ is shown in Figure 4.16. The first half of the response is due to a supra threshold sinusoidal signal, while the second half is due to noise only. When the sum of the time varying inputs reaches a sufficient amplitude, the neuron goes into the excited state for a small duration called *firing*, it then returns to the resting state. A small refractory period is present where the neuron is unresponsive. The output of the neuron is given by an *action potential* that is generated when the firing neuron exceeds a threshold, set to 0.5 in Figure 4.16. The action potential has the form of a sharp spike given by the Dirac delta function, as shown by the vertical lines in Figure 4.16. Thus the output of the neuron is a train of action potentials, called a *spike train*.

For simulations where discrete time must be used, we can define the delta function as [Chapeau-Blondeau and Godivier, 1997]

$$\hat{\delta}(kh) = \begin{cases} 1/h & \text{if } k = 0 \\ 0 & \text{otherwise,} \end{cases} \quad (4.13)$$

for some integer k .

4.2.3 Bistable Dynamics

Bistable systems can exist in either one of two states. Unlike the excitable systems that return to a resting state, bistable systems remain in a state until forced into the

4.2 Systems Demonstrating SR

other state. The switching between states can involve a hysteresis or relaxation time. There may be a continuous transition between the states like the double well potential, or there may be a discrete transitions as in the Schmitt trigger. We only consider continuous-time systems, though a discrete-time double well potential is possible [Zozor and Amblard, 1999; Zozor and Amblard, 2001].

Double Well Potential (DWP)

An archetypical model of a bistable system that exhibits SR is the double well potential [Fox, 1989b; Gammaitoni et al., 1989a; Gammaitoni et al., 1989b; McNamara and Wiesenfeld, 1989; Jung and Hänggi, 1991; Hilgers et al., 1995; Mantegna and Spagnolo, 1996; Lanzara et al., 1997]. Using the one dimensional equation of motion

$$m\ddot{x}(t) + \gamma\dot{x}(t) = -\frac{dU(x)}{dx} + f(t), \quad (4.14)$$

with the potential represented by $U(x)$. The equation describes the motion of a particle with a mass m in the presence of friction γ in the potential $U(x)$. The location of the particle, x is considered to be the output of the system. The forcing function $f(t)$ usually consists of a periodic component $s(t)$ and a stochastic component $\eta(t)$. When the system is heavily damped [Lanzara et al., 1997] the $m\ddot{x}$ term can be ignored and rescaling with γ yields

$$\dot{x}(t) = -\frac{dU(x)}{dx} + s(t) + \eta(t). \quad (4.15)$$

Different dynamic systems can be formed by substituting the respective potentials. For example, the analog Hopfield neuron is given by

$$U(x) = \alpha x^2 - \beta \ln(\cosh x),$$

where the state point $x(t)$ denotes a cell membrane voltage, and the SQUID loop is given by

$$U(x) = \alpha x^2 - \beta \cos(2\pi x),$$

where $x(t)$ denotes the magnetic field flux in the loop [Bulsara and Gammaitoni, 1996]. A common electrical component that has bistable characteristics is the tunnel diode [Mantegna and Spagnolo, 1994; Mantegna and Spagnolo, 1995; Mantegna and Spagnolo, 1996; Lanzara et al., 1997; Wacker and Schöll, 1995]. This has similar characteristics to the bistable potential where two stationary points exist.

However, for simplicity we consider the bistable system governed by the quartic potential

$$U(x) = -a\frac{x^2}{2} + b\frac{x^4}{4}, \quad (4.16)$$

which is shown in Figure 4.17. For bistability, we require $a > 0$ which gives an unstable state at $x = 0$ and two stable states at $x_s = \pm\sqrt{a/b}$ that are separated by $\Delta U = a^2/4b$.




Figure 4.17. The quartic double well potential. The central blue curve is the potential (4.16) without any external forcing. It has the minima at $\pm x_s$ and barrier height ΔU . With periodic forcing the potential oscillates between the outer two curves.

In the absence of noise the system is stationary and the particle described by (4.14) remains in one of the potential wells indefinitely. If we add a periodic signal, $s(t) = A \sin(\omega_0 t)$ and include it as part of the potential for visualisation purposes, we have

$$U(x, t) = -a\frac{x^2}{2} + b\frac{x^4}{4} - Ax \sin(\omega_0 t). \quad (4.17)$$

The signal causes the potential to rock back and forth in time, as shown in Figure 4.17. For a sufficiently small amplitude A , the potential will move up and down, but the particle will remain in the same well. Further increases to the signal by adding noise for example, cause the particle to jump between the wells, which gives rise to the SR phenomenon. Written in full, the system is

$$\dot{x}(t) = ax - bx^3 - A \sin(\omega_0 t) + \eta(t). \quad (4.18)$$

4.2 Systems Demonstrating SR

Using the adiabatic approximation [McNamara and Wiesenfeld, 1989; Fox, 1989b] the expression for the signal-to-noise ratio is given as

$$SNR \approx \frac{\sqrt{2}aA^2x_s^2}{(\sigma^2)^2} \exp\left(\frac{-2\Delta U}{\sigma^2}\right),$$

which is plotted in Figure 4.3b against the simulations of the example system, which is in fact the bistable system described by (4.18).

Schmitt Trigger

A simple two state system such as the Schmitt trigger can exhibit a bistable nature and thus show SR characteristics. In fact, the Schmitt trigger, which is a comparator with hysteresis, was the first laboratory demonstration of SR with an electronic device [Fauve and Heslot, 1983]. Later papers described the theory of the system in more detail [McNamara and Wiesenfeld, 1989; Malnikov, 1993].

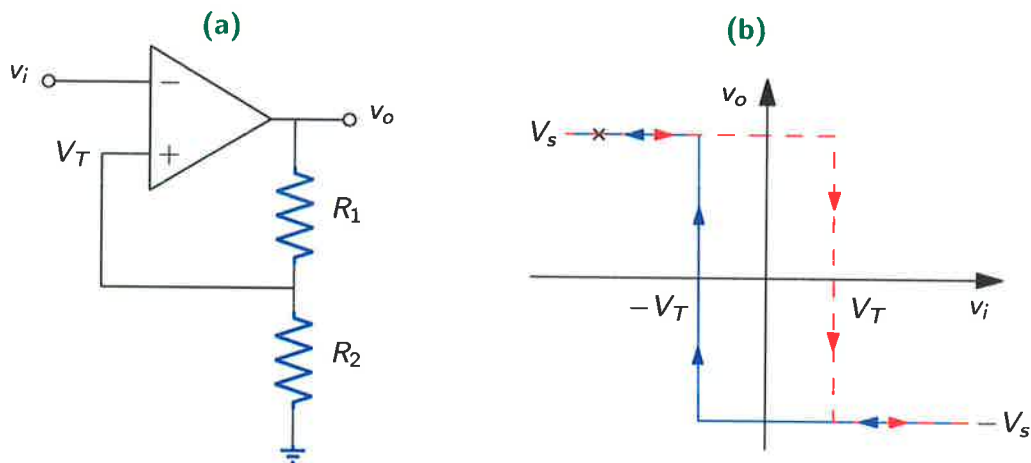


Figure 4.18. Schmitt trigger circuit and transfer characteristic. (a) The circuit using an operational amplifier. (b) The corresponding voltage input-output characteristic that exhibits hysteresis.

A circuit diagram of the Schmitt trigger is shown in Figure 4.18a, where the op amp behaves as described in Sec. 4.2.1. To examine the characteristics, consider v_i to be positive and less than the threshold at $-V_T$ and assume the output is in the high state, marked by a cross in Figure 4.18b. Thus, the output is at the positive supply voltage, i.e. $v_o = V_s$. The threshold voltage is set to $V_T = V_s R_2 / (R_1 + R_2)$ via the voltage divider formed by the feedback loop of R_1 and R_2 . Upon increasing v_i , eventually $v_i > V_T$, and the output switches to the negative supply voltage to give $V_T = -V_s R_2 / (R_1 + R_2)$. The

process repeats when v_i is then decreased. This demonstrates the hysteresis characteristic of the Schmitt trigger, i.e. bistability. The bistable characteristics are shown in Figure 4.18b.

S-Type Characteristic

The Schmitt trigger provides a two state bistable system by utilising the feedback path to one of the terminals. However, by utilising feedback in both terminals of the op amp, the gain can be controlled and a piecewise linear bistable system created. Such a circuit is shown in Figure 4.19a. We refer to this as a S-type circuit due to the shape of the transfer function shown in Figure 4.19b. A similar circuit with the nonlinear components in the other feedback path forms a N-shaped characteristic [Harmer et al., 2002].

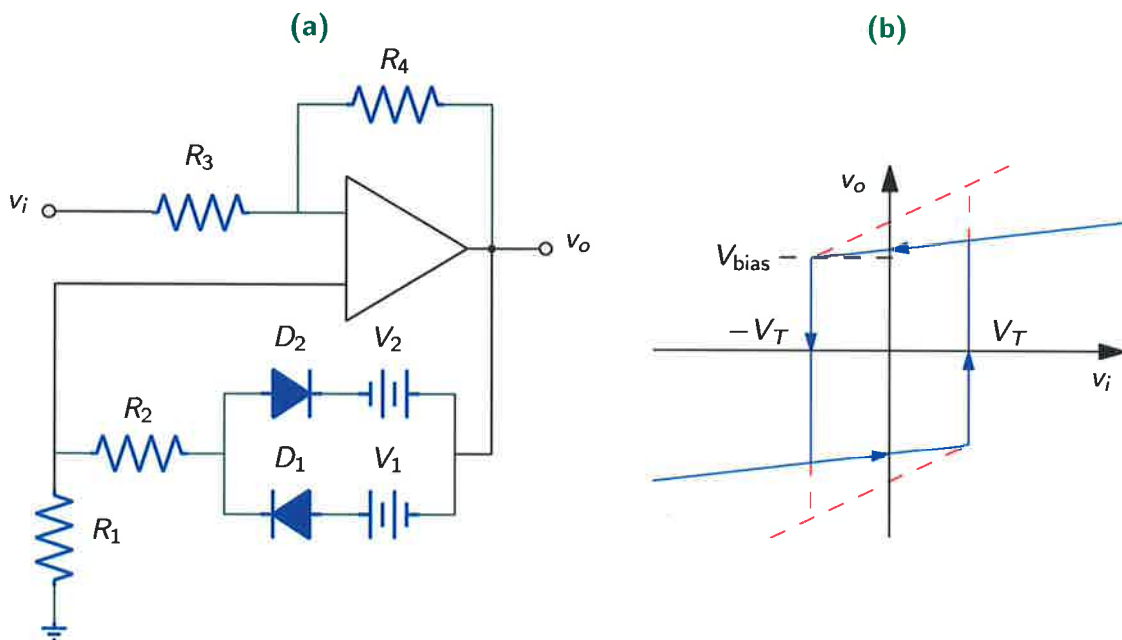


Figure 4.19. S-type circuit diagram and transfer characteristics. (a) The circuit using an operational amplifier and nonlinear components. (b) The corresponding voltage input-output characteristic that exhibits hysteresis and gain.

The main difference compared to the Schmitt trigger is that the output has gain, i.e. the slopes of the characteristics are not zero. This requires components in both feedback paths. The implementation can be achieved using a single voltage source by appropriately rearranging and transforming the components.

4.3 Noise

4.2.4 Other Nonlinear Systems

The coverage of systems in this section is by no means complete. There are multistable systems that have a number of potential barriers and states [Vilar and Rubí, 1997], whereas monostable systems have no potential barrier and hence only have one stable fixed point. An example is the single-well Duffing oscillator [Luchinsky et al., 1999a; Luchinsky et al., 1999b; Stocks et al., 1993]. Coupled systems combine many units of other systems [Inchiosa and Bulsara, 1995; Jung, 1995a; Longtin, 1995]. Chaotic and random systems also show SR, which include Chua's circuit, Lorenz systems and random walks [Leonard, 1992; Mitaim and Kosko, 1998]. Mitaim and Kosko (1998) have shown an example of adaptive SR.

4.3 Noise

In this section we discuss different types of noise that are commonly used with the systems introduced in the previous section. Algorithms are then outlined describing how to reliably generate certain noise distributions and simulate stochastic systems.

The type of noise inflicted on a system is characterised by the noise source. Examples include thermal noise (or Gaussian), $1/f$ noise, shot noise, dichotomous noise and other forms of coloured noise. The noise is characterised by its probability density function $p(x)$ which is associated with the stochastic process $x(t)$. This is called a stationary distribution due to the absence of any temporal arguments, like $P(x, t)$ which is nonstationary. This distribution holds all the information about the moments of the process. Of most relevance are the mean value $\langle x \rangle$ and the variance $\langle (x - \langle x \rangle)^2 \rangle = \langle x^2 \rangle - \langle x \rangle^2$. The variance is usually referred to as the noise intensity.

4.3.1 White – Gaussian

The simplest model for uncorrelated noise is referred to as Gaussian white noise $\xi(t)$. It is called 'Gaussian' because the distribution follows a normal curve, and 'white' because the PDF is flat for all frequencies, like that of white light. This type of noise is an *ideal* model of thermal noise. It is statistically characterised by

$$\langle \xi(t) \rangle = 0 \quad \text{and} \quad \langle \xi(t) \xi(s) \rangle = 2D\delta(t - s),$$

where $\xi(t)$ and $\xi(s)$ are statistically independent for $s \neq t$. The strength of the noise is characterised by the single parameter D , which is a measure of the intensity.

Since this is an idealisation, it can not exist physically – we only need to consider the power to see why. Finding the power over the entire spectrum yields

$$\text{Power} = \int_0^{\infty} S_x(f) df = \infty,$$

since the PSD, $S_x(f)$ is constant. This result is obviously impossible. Nevertheless, white noise can still serve as a good starting point and is sufficient when the time scale of interest is much larger than that of the noise.

Statistically it is common to write a normal distribution as $\xi \sim \mathcal{N}(\mu, \sigma^2)$, where μ is the mean and σ^2 the variance. This can be transformed to an alternative normal distribution by

$$\alpha + \beta \mathcal{N}(\mu, \sigma^2) = \mathcal{N}(\alpha + \mu\beta, \beta^2 \sigma^2), \quad (4.19)$$

which is useful when generating a distribution from the standard distribution $\mathcal{N}(0, 1)$. Another useful property is that the sum of statistically independent normal distributions yields another normal distribution according to

$$\mathcal{N}(\mu_1, \sigma_1^2) + \mathcal{N}(\mu_2, \sigma_2^2) = \mathcal{N}(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2). \quad (4.20)$$

Even if the two distributions are dependent, another normal distribution is still produced – though the means and variances do not combine as simply [Gillespie, 1996; Yates and Goodman, 1999].

4.3.2 Coloured – Ornstein-Uhlenbeck

To deal with this infinite power problem and to model real noise more accurately, we introduce a correlation time τ_c of the noise. This bandlimits the noise at a cut-off frequency $f_c = 1/(2\pi\tau_c)$. Coloured noise is given by the Ornstein-Uhlenbeck (OU) process that is described by the stochastic differential equation (SDE)

$$\dot{\zeta}(t) = -\frac{1}{\tau_c} \zeta(t) + \frac{\sqrt{2D}}{\tau_c} \xi(t), \quad (4.21)$$

and has the statistical properties

$$\langle \zeta(t) \rangle = 0 \quad \text{and} \quad \langle \zeta(t) \zeta(s) \rangle = \frac{D}{\tau_c} \exp\left(-\frac{|t-s|}{\tau_c}\right),$$

4.3 Noise

with variance D/τ_c . The strength of the noise is now measured in terms of D/τ_c , but we still refer to D as the noise intensity [Sancho et al., 1982; Chialvo et al., 1997]. The noise intensity corresponds to the white noise in the limit $\tau_c \rightarrow 0$ in which $\zeta(t) \rightarrow \xi(t)$. The OU process provides control over both noise intensity D , and correlation time τ_c .

Even though an exact solution can be found by integrating (4.21) as (multiply through by e^{t/τ_c} , rearrange, and directly integrate),

$$\zeta(t) = \zeta(0)e^{-t/\tau_c} + \frac{\sqrt{2D}}{\tau_c} \int_0^t e^{(s-t)/\tau_c} \xi(s) ds,$$

it still contains a stochastic term. An exact updating formula for the OU process itself can be written as

$$\zeta(t+h) = \zeta(t)e^{-\alpha} + \sqrt{\frac{D(1-e^{-2\alpha})}{\tau_c}} \xi(t), \quad (4.22)$$

where $\alpha = h/\tau_c$ [Mannella and Palleschi, 1989; Gillespie, 1996]. This formula gives the exponential decay of $\langle \zeta(t)\zeta(s) \rangle$ in an exact way and has been checked for very large time steps [García-Ojalvo and Sancho, 1999, p. 100].

The PSD for several types of coloured noise are shown in Figure 4.20. The corner frequency f_c , marked by the dashed line, represents the 3 dB point, thereafter the PSD drops off at 20 dB/decade which is to be expected as (4.21) has the form of a single stage lowpass filter.

The middle row of Figure 4.20 shows the effect of varying the integration step size h while keeping the noise intensity D and correlation time τ_c fixed. The time series remains relatively unchanged, which is the desired result, though the PSD appears to be of varying strengths. This is due to the N -point FFT used to calculate the PSD using the same value of N whilst changing h . The top row has varying τ_c and D to maintain a constant PSD below the cut-off frequency. In the bottom row h and D are varied to maintain the same shaped PSDs.

4.3.3 Generation of Random Numbers

The generation of purely random numbers on a deterministic machine is impossible. Instead we have to settle for a sequence of *pseudo-random* numbers. These numbers are generated by algorithms that produce sequences that appear to be random, that is, they can pass a series of statistical tests. Several algorithms have been developed since the early days of compilers, some being better than others [Nahin, 2000, p. 175].

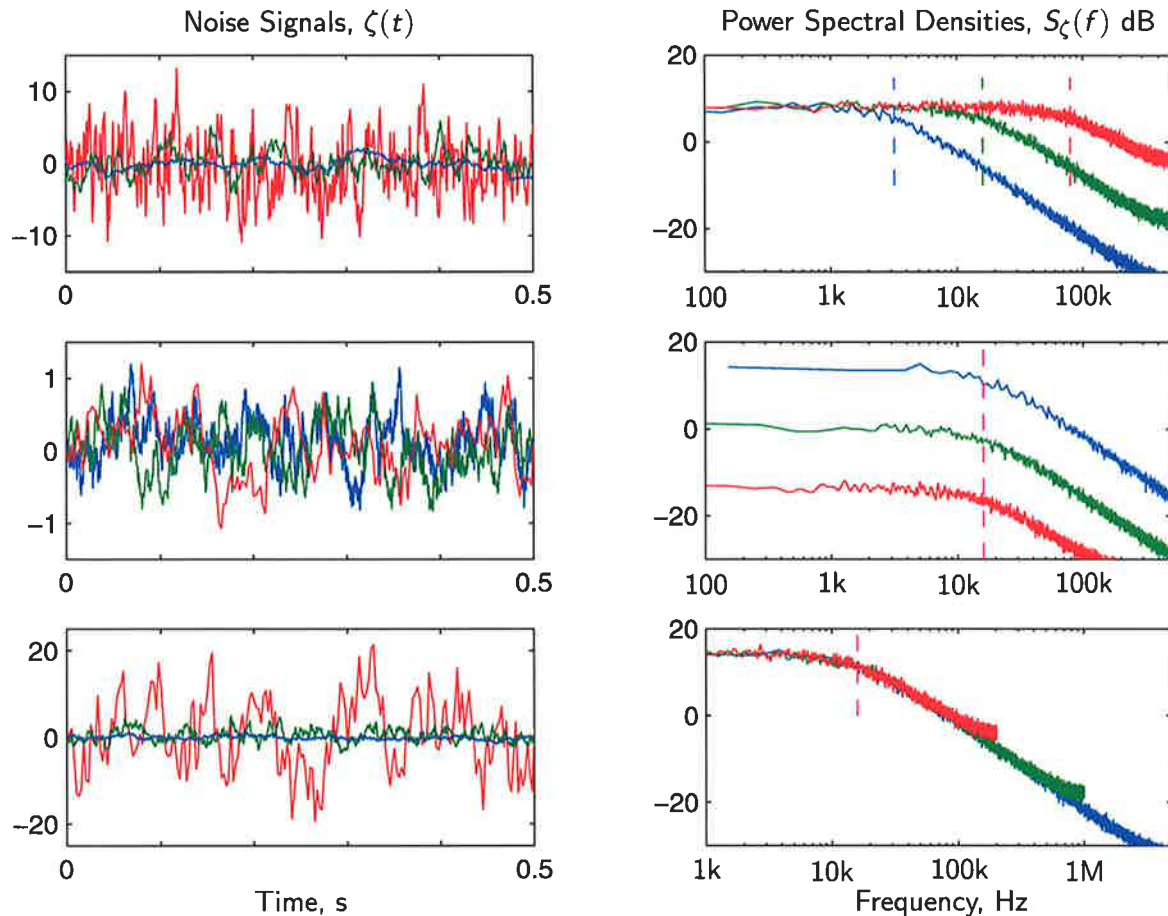


Figure 4.20. Time series and power spectral densities of coloured noise. The PSDs on the right correspond to the time series directly to the left. The dashed lines in the PSDs show the corner frequency, f_c . (Top row) Keeping the integration step size constant at $h = 0.001$ and varying D to maintain constant $S_\zeta(f)$ below f_c while changing τ_c to 0.002, 0.01 and 0.05 from top to bottom. (Middle row) Keeping a constant $D = 15$ with $\tau = 0.01$ while varying the step size h to 0.1, 0.5 and 2.5 ms from top to bottom. (Bottom row) Keeping the correlation time constant ($\tau = 0.01$) while varying h (same as middle row) and D to 15×5^4 , 15×5^2 and 15 from top to bottom.

An important factor is the period, which is how many numbers are produced before the algorithm repeats itself. To illustrate the importance of this, consider length of sequences generated by different versions of MATLAB. MATLAB 4.x has a length of $2^{31} - 2 \approx 10^9$, which may seem large, but at a consumption of one million numbers per second the sequence is exhausted within an hour. However, the random number generator in MATLAB 5.x uses a different algorithm that has a period of 2^{1492} , which is quite large.

4.4 Quantifying SR

Software generators typically only generate uniform deviates between 0 and 1. Any other distribution can be attained via an appropriate mapping function. Much effort goes into developing algorithms that are reliable and fast. Several algorithms are given in Appendix B.1 to generate uniform and normal deviates. Depending on the facilities available, it may be necessary to implement these algorithms, or simply use the *built-in* random number generator.

4.3.4 Simulating Stochastic Systems

When given a system defined by a set of SDEs, such as the FitzHugh-Nagumo neuronal model given by (4.12), an exact solution cannot be found. In these circumstances computer simulations can provide a valuable tool to gain insight into the system. However, numerically integrating SDEs is not straight forward as it is for ODEs. Particular care must be taken with the algorithm used for the integration, especially if supplied as a built-in function of a programming environment. Many papers have been written on how to accurately integrate stochastic systems [Mannella and Palleschi, 1989; Mannella, 1989; Honeycutt, 1992a; Honeycutt, 1992b; Fox, 1989a; Fox, 1991].

The systems of interest here contain deterministic and stochastic parts. Using a fourth order Runge-Kutta algorithm (RK-4) [Kreyszig, 1953, p. 1040], each part can be dealt with separately – in this sense the deterministic system is coupled to the OU process. This allows the integration step size to be chosen to allow the deterministic part to be solved to the required accuracy, with algorithms in Mannella and Palleschi (1989) or Gillespie (1996) providing a noise integration algorithm of $O(h^{3/2})$. The accuracy of the stochastic system can be improved further by utilising the integration methods in Fox (1991). For example, the algorithm used to integrate the FHN system of (4.12) with coloured noise is given in Appendix B.2.

4.4 Quantifying SR

From the time when stochastic resonance was first conceived to the present day, many types of quantities have been proposed and used to measure SR. These range from traditional signal processing concepts to modern information theory. In the following subsections we will describe the signal-to-noise ratio (SNR), residence and interspike

interval times, correlation coefficients, mutual information and finally, channel capacity. Some of the variables used to derive these quantities are shown in Figure 4.21, which is Figure 4.1 redrawn in the frequency domain.



Figure 4.21. The stochastic resonator in the frequency domain. At the input, $B_{n,in}$ is the bandwidth of the input noise and has a PSD of $S_{n,in}$. At the output, $P_{s,out}(f)$ is the power of the periodic component at f and $S_{n,out}(f)$ is the PSD of the output noise at the same frequency.

4.4.1 Signal-to-Noise Ratio

One of the first measures used to quantify SR was the power spectral density from a Lorenz model [Benzi et al., 1981]. By taking the PSD for several noise intensities general trends of the PSD at the signal frequency were found. When SR was first demonstrated experimentally by Fauve and Heslot (1983) the SNR was employed. This takes into account the noise and is defined in terms of PSD by

$$SNR = 10 \log_{10} \left(\frac{S(f_0)}{B(f_0)} \right) \quad [\text{dB}]. \quad (4.23)$$

$S(f)$ is the magnitude of the signal component taken at the fundamental frequency $f = f_0$ and $B(f)$ is the PSD of the background noise taken at nearby frequencies [Inchiosa and Bulsara, 1996]. The quantities $S(f)$ and $B(f)$ are shown in Figure 4.3a on page 96. Both the PSD and SNR exhibit similar profiles in stochastic resonant systems, but with the peaks at slightly different noise intensities. For certain systems, stochastic resonance is observable at higher order harmonics.

More generally, the SNR can be defined for broadband signals by treating it as a function of frequency. It is defined by [Kiss, 1996]

$$SNR_{\text{out}}(f) \equiv \frac{P_{s,\text{out}}(f)}{S_{n,\text{out}}(f)} \quad [\text{Hz}], \quad (4.24)$$

where a major difference is that this is a ratio of power to PSD. The noise power is described by its PSD because it is dependent on bandwidth, whereas a sinusoidal signal is not. The spectrum of the signal cannot be used since it is periodic and would

4.4 Quantifying SR

result with a Dirac delta function. In practice, this means that the height of the spike depends on the frequency resolution of the FFT during measurement, hence affecting the absolute value of the SNR. Conversely, we cannot use the power measure for noise as this would include frequencies far from the signal frequency of interest. The mixed method in (4.24) is valuable because it gives information about the actual SNR at the signal frequency.

Improving the SNR

A common question is if stochastic resonance can improve the SNR at the output compared to the SNR at the input? That is, can we show that the gain $G = \frac{SNR_{out}}{SNR_{in}}$ is greater than 1? Also, can SR improve the SNR compared to that of a linear filter?

In general, for an arbitrary input signal the SNR_{out} cannot be greater than the SNR_{in} [Chapeau-Blondeau, 1997a]. Neiman et al. (1997) considers ensembles of N independent STRs and shows that $G < 1$. When N approaches infinity then G approaches unity, as it should for a linear system.

However, by placing some constraints on the system, G can be made greater than unity. By considering a sufficiently “spiky” signal it has been shown that G can be made greater than unity, for the LCD [Loerincz et al., 1996] and the LCC [Chapeau-Blondeau, 1997a]. These systems are shown in Figure 4.13.

To achieve this improvement in SNR, the system must meet several requirements. These effectively circumvent the small-signal limit and thus manage to get $G > 1$ [Chapeau-Blondeau, 1997a; Loerincz et al., 1996]. In the adiabatic case of the LCD, Kiss (1996) predicts $G \leq 0.85$, but this is invalid for the case of a nonlinear response.

De-chun et al. (1992) use a bistable system to get SNR_{out} larger than the highest SNR obtained by a linear device. They describe the mechanism that increases the SNR via a linear device by the filtering out of certain parts of the noise, whereas for SR the mechanism is the transfer of energy from the noise to the signal under certain conditions. This energy transfer accounts for the large gain in SNR_{out} compared with that of a linear system.

The SNR can also be increased by cascading two or more STRs together. Though there is a limit to the number of STRs that can be cascaded. De-chun et al. (1992) claim the efficiency is reduced by the finite correlation times of the STR outputs, and Asdi and Tewfik (1995) observe that there is an increase in the background noise level.

4.4.2 Residence Time Distribution

It was realised by Benzi et al. (1981) and others [Gammaitoni, 1995; Bulsara and Gammaitoni, 1996] that SR is not strictly a *resonance*, as mentioned earlier. By using residence times, the signature SR peak can be formed as either a function of noise intensity or frequency [Zhou et al., 1990; Longtin et al., 1991]. These are referred to as residence time density functions or distributions.

The *residence time* is defined as the amount of time a system spends in a stable state between consecutive switching events, thus this applies to bistable systems. This is the reciprocal to the other common terminology of the *escape rate*. If we denote t_i as the time of the i th switching event, then

$$T(i) = t_i - t_{i-1} \quad (4.25)$$

is the i th residence time. We can then form the normalised distribution of residence times $N(T)$. This allows characterisation of SR as a synchronisation of the switching mechanism driven by the signal and noise simultaneously [Gammaitoni, 1995].

The bistable system given in (4.18) is considered with the results shown in Figure 4.22. The horizontal scale is the normalised period of the forcing signal and the peaks of the residence times are given by $T_n = (n - 1/2)T_0$.

The residence time metric is sufficient for two state systems, however it is not appropriate for excitable systems such as neuron models. For these systems the *interspike interval* is employed. This is defined as the amount of time between excited states. Thus, we define the interspike interval (ISI) as $T(i)$ given in (4.25) by redefining t_i as the time of the i th excitation.

Figure 4.23 shows the interspike interval distributions, again with the normalised horizontal axis, but the peaks are located at $T_n = nT_0$.

In order to quantify the n th peak of $N(T)$, we sum the values centred around T_n with the following quantity

$$P_n = \int_{T_n - \alpha T_0}^{T_n + \alpha T_0} N(T) dT, \quad (4.26)$$

where $0 \leq \alpha \leq 1/4$ gives the spread that we want to include [Gammaitoni, 1995]. The appropriate T_n is required for the residence or interspike interval times. If we increased α from $1/4$ to a $1/2$ to make regions of P_n join, simulations show that the difference is minimal.

4.4 Quantifying SR

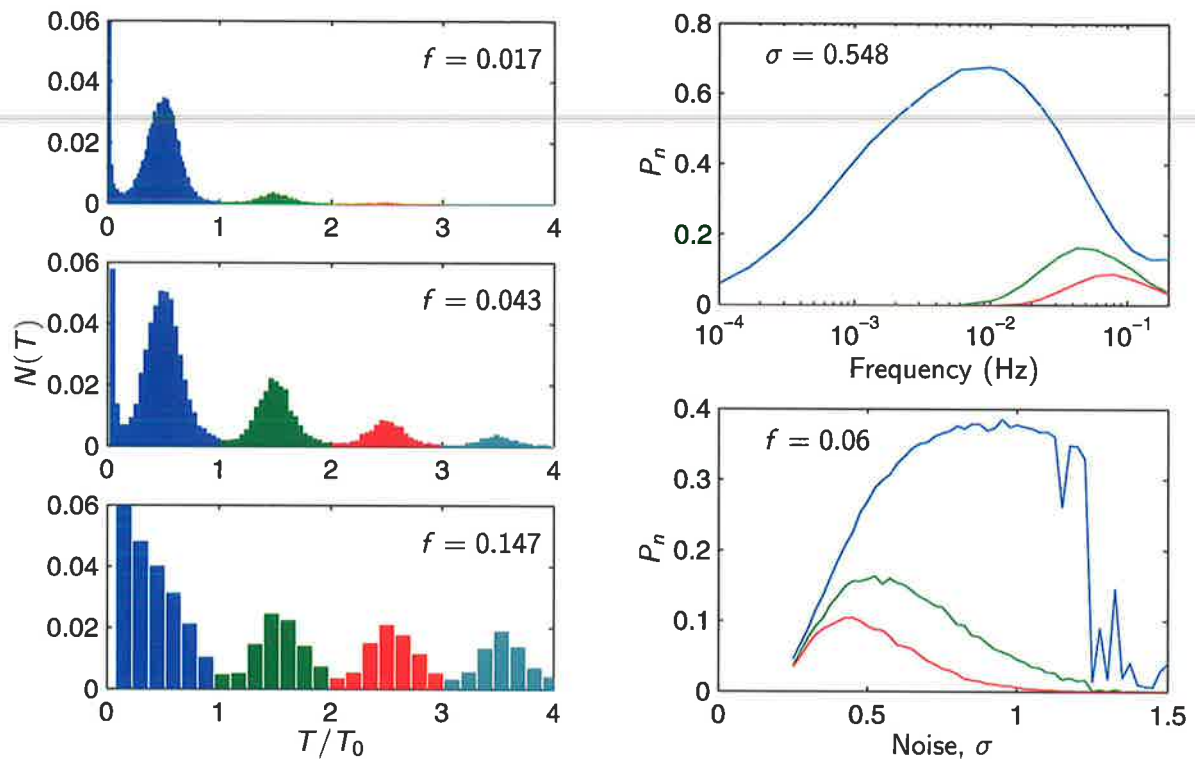


Figure 4.22. Residence time distributions for the double well potential. The plots on the left are residence time histograms against the normalised period. A total of 25 000 jumps were used for each frequency setting. The plots on the right show P_n against frequency (top) and noise intensity (bottom).

For both the plots of P_n against frequency in Figure 4.22 and 4.23, there is a value of frequency that gives the maximal value of P_n . This is true resonant behaviour by the system. The accepted practice of plotting against noise intensity also shows a stochastic resonant peak.

Residence and interspike interval times are mentioned only for completeness, they will not be used any further in this work.

4.4.3 Correlation Coefficients

The previously described metrics are only appropriate when the input signal is periodic. They assess the coherence of the response with the input signal. Therefore, when dealing with aperiodic signals the methods are not suitable, and a new metric is required for SR systems to handle broadband signals.

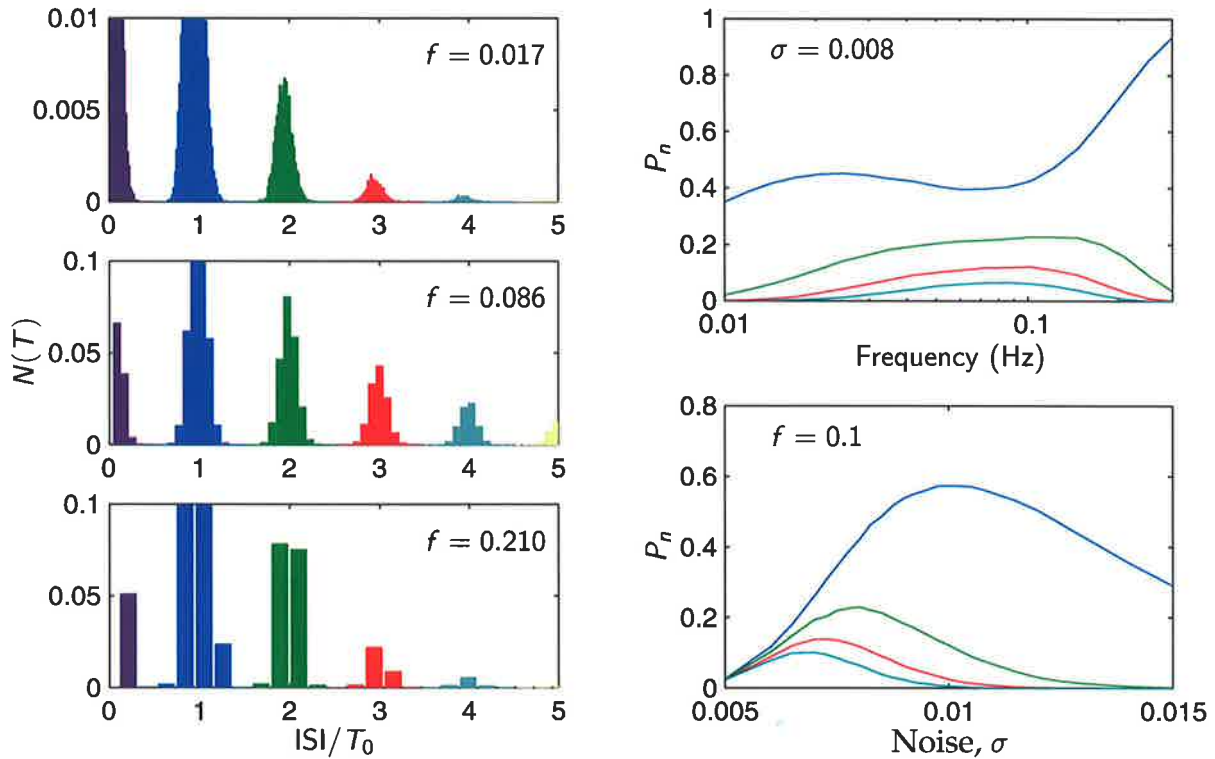


Figure 4.23. Interspike interval distribution for the FitzHugh-Nagumo neuron model. The plots on the left are interspike interval histograms against the normalised period. A total of 50 000 spikes were used for each frequency or noise setting. The plots on the right show P_n against frequency (top) and noise intensity (bottom).

The term aperiodic stochastic resonance (ASR) was coined by Collins et al. (1995a) to describe a SR system with a broadband input signal. The power norm measure was consequently proposed and involves the use of correlations. It is defined as

$$C_0 = \overline{s(t)r(t)}, \quad (4.27)$$

where $s(t)$ is the aperiodic (zero mean) input signal, $r(t)$ is the mean firing rate response of the system, and the overbar refers to an average over time. Collins also gives the normalised power norm as

$$C_1 = \frac{\overline{s(t)r(t)}}{[\overline{(s(t) - \overline{s(t)})^2}]^{1/2} [\overline{(r(t) - \overline{r(t)})^2}]^{1/2}}. \quad (4.28)$$

Maximising C_1 corresponds to matching the shape of $s(t)$ and $r(t)$ while maximising C_0 also takes into account the amplification between $s(t)$ and $r(t)$ [Collins et al., 1995a].

In practical systems there exists an inherent time delay between the time the system receives the input signal, and the time the corresponding response emerges. For ASR

4.4 Quantifying SR

we wish to utilise the measure which gives the peak input-output cross-correlation. Therefore, we ignore any time delay in the system and take the correlation at a time lag of zero [Collins et al., 1995a]. For systems where the time delay is much less than the time scale of the signal it is not as important.

For most types of excitable systems, like the neuron, the output is in the form of a spike train that represents the instantaneous firing rate. To calculate the required mean firing rate $r(t)$, a unit-area symmetric Hanning window is passed over the spike train [Collins et al., 1996a].

This measure may seem unlike the SNR, but the correlations also measure the coherence between the input and response for both periodic and aperiodic signals. It is shown by Chow et al. (1998) that in the infinite time limit, $\overline{[r(t) - \overline{r(t)}]^2} \equiv \hat{r}^2$, which corresponds to the normalised (total) power in the output. Therefore $\hat{r}^2 = \hat{n}^2 + \hat{s}^2$ where \hat{s}^2 is the output power at the signal frequency (and phase) and \hat{n}^2 is the noise power. Thus, we have $SNR = \hat{s}^2/\hat{n}^2$ which leads to $C_1^2 \propto SNR/(1 + SNR)$.

It is also important to note that correlations also take into account phase information [Chow et al., 1998], whereas the SNR does not. This is easy to see by considering a simple example of a sinusoidal signal passing through a filter that changes the phase and has zero insertion loss. The SNR_{out} is therefore very similar to the SNR_{in} . But finding C_1 for the input and output gives a small or even negative coefficient simply due to the phase mismatch.

The correlation metric has been employed in a variety of systems with aperiodic inputs [Collins et al., 1996a; Harmer and Abbott, 1999c; Neiman et al., 1997] and has been applied experimentally in sensory neurons [Collins et al., 1996b]. Aspects such as introducing a refractory period [Chialvo et al., 1997], the effect of coloured noise [Capurro et al., 1998], parallel ensembles of neurons [Collins et al., 1995b; Chialvo et al., 1997] and using chaotic stimuli [Eichwald and Walleczek, 1997] have also been investigated.

Aperiodic Stochastic Resonance

To replicate ASR we consider the FitzHugh-Nagumo neuronal model given by (4.12) with additive coloured noise given by the Ornstein-Uhlenbeck process (4.21). The system is integrated with a step size of $h = 0.001$ s, and coupled to the OU process using the method outlined in Appendix B.2. Figure 4.24 shows the system with and without




Figure 4.24. Correlation coefficients for the FHN neuronal model. The power norm C_0 is plotted in (a) and the normalised power norm C_1 in (b). A total of 512 realisations were taken for each noise value.

refractory periods. The results are averaged over a number of ensembles, denoted by $\langle \cdot \rangle$.

Although these results look promising one must take into account all the ensembles of correlation coefficients that were used [Chialvo et al., 1997]. The individual ensembles are plotted as dots in Figure 4.25. The distribution is very broad for all noise values, even at the optimal noise intensity there are realisations that are anticorrelated, i.e. have a negative correlation coefficient.

Though the distribution is large in Figure 4.25, we can reduce the variance by averaging over several ensembles. This is similar to having the same number of neurons in parallel and averaging the response by

$$r_{\Sigma}(t) = \frac{1}{N} \sum_{i=1}^N r_i(t), \quad (4.29)$$

which is referred to as “SR without tuning” [Collins et al., 1995b; Neiman et al., 1997; Chialvo et al., 1997]. $r_i(t)$ denotes the response of the i th device in the ensemble and $r_{\Sigma}(t)$ the average response *before* the correlation is measured. The C_1 coefficient for various values of N is plotted in Figure 4.26a, and some signals in the time domain are shown in Figure 4.26b.

As the number of neurons in the ensemble increase, C_1 approaches unity, though there is a diminishing rate of return as more neurons are added to the ensemble. Further improvements, especially for larger noise intensities are gained by adding a refractory period after the neuron [Chialvo et al., 1997].




Figure 4.25. Distribution of ensembles forming the correlation coefficients. The points show 256 individual realisations (at each noise intensity), the solid line is the mean (the same as Figure 4.24b) and the dashed lines are one standard deviation each side of the mean. Some of the outer samples have been clipped to show more detail about the mean.

Figure 4.26 shows the signal used for all of the simulations in this section. It was formed by passing a 10 sec unit-area symmetric Hanning window over coloured noise with $\tau_c = 20$ s. Also plotted is the response of a 256 neuron ensemble, and a single noiseless suprathreshold neuron, as done in Chialvo et al. (1997). The relevance of this will become apparent in the next section, but it is interesting to note that the performance of a single suprathreshold neuron is on par with that of a 256 subthreshold neuron ensemble.

Further discussions of parallel STR networks are continued in the next section in tandem with the introduction of suprathreshold stochastic resonance.

4.4.4 Average Mutual Information

Recently, Stocks has written a number of papers that consider a network of stochastic resonators in suprathreshold operation that utilise an information theoretic measure [Stocks, 2000c]. Using this type of network several interesting aspects of SR can be observed.




Figure 4.26. Correlations of parallel neuron ensembles and output signals. (a) Normalised correlation coefficients for a parallel ensemble of N neurons. (b) The original signal (blue), the response from a 256 subthreshold neuron ensemble (green), and the response from a single suprathreshold ($A = 0.125$) neuron (red).

The network under consideration is shown in Figure 4.27. It consists of N STRs in parallel, each having an individual noise source $\eta_i(t)$ and able to have independent parameter settings. The input to each device $x(t)$ is identical, and the outputs are summed to give the total response $y(t)$. It is common to model ensembles of neurons, like the FHN model for example, [Stocks and Mannella, 2000] but in order to simplify the system to allow reasonably straight forward calculations to take place, a simple LCC threshold is used. While removing the dynamic components of the system, the LCC still qualitatively reflects the neuron's behaviour.

Allowing each device to be individually set, the modified Heaviside function from (4.9) is

$$y_i(t) = \begin{cases} 1 & \text{if } x(t) + \eta_i(t) > \theta_i, \\ 0 & \text{otherwise.} \end{cases} \quad (4.30)$$

The output $y(t)$ is then equal to the number of devices that are excited.

Using the input $x(t)$ and output $y(t)$, or some smoothed version thereof, the correlation coefficients described in the previous section can be employed.

However, it is argued that using an information theoretic measure is more appropriate. The cross-correlation measure was originally derived from linear response theory, therefore why use it for a system that is designed from the ground up to be nonlinear [Stocks and Mannella, 2000]? An information theoretic measure makes no assumptions of the type of signal and is thus not biased in any way. We will proceed to define

4.4 Quantifying SR




Figure 4.27. A summing network of N devices. Each device receives an identical input signal but has an individual noise source and parameters.

the information measure used by Stocks [Stocks, 2000c]. Other methods that have been reported are outlined at the end of the section.

We can quantify the amount of information flowing through the network as the average mutual or transmitted information, the *AMI*. This is regarded as a semi-continuous channel, in information theory, and is written as

$$AMI = H(y) - H(y|x) \quad (4.31)$$

$$= - \sum_{n=0}^N P_y(n) \lg P_y(n) - \left(- \int_{-\infty}^{\infty} \sum_{n=0}^{\infty} P_x(x) P(n|x) \lg P(n|x) dx \right). \quad (4.32)$$

The function H is the entropy, or in other words the information content. Therefore, $H(y)$ is the information content of $y(t)$, and $H(y|x)$ is the conditional entropy of the output given the input, that is, the amount of encoded information lost due to noise in the transmission of the signal $x(t)$. The P functions are probability densities; $P_y(n)$ being the probability the response $y(t)$ is numerically equal to n , $P(n|x)$ the conditional probability of the response being in state n given the signal x , $P_x(x)$ and $P_\eta(\eta)$ are the PDFs of the signal and noise respectively. The function \lg is short for the base 2 logarithm, which also implies the units of *AMI* are bits. More accurately, $P_y(n)$ and $P(n|x)$ are probability mass functions since they are functions of discrete variables.

By considering special cases, some of the characteristics of the *AMI* can be deduced. As $\sigma_\eta \rightarrow \infty$ all the information is lost in transmission, giving $H(y|x) = H(y)$ and therefore $AMI = 0$ as expected. Alternatively when $\sigma_\eta = 0$ no transmitted encoded information is lost giving $H(y|x) = 0$, hence $AMI = H(y)$. This is equal to 1 bit when all the thresholds are set to $\langle x \rangle$ [Stocks, 2000b]. This can be derived intuitively; with no

noise the network functions as a one bit ADC as all the devices switch simultaneously giving $y(t)$ to be either 0 or N for half the time each since $\langle x \rangle = 0$. Therefore $P_y(0) = P_y(N) = 1/2$ and $P_y(n) = 0$ for the other values $n = 1, \dots, N - 1$ as these states are unreachable. Then directly from (4.31) we get

$$H(y) = -P_y(0) \lg P_y(0) - P_y(N) \lg P_y(N) = 1. \quad (4.33)$$

One may expect that as σ_η increases from zero, $H(y|x) < H(y)$. If this were true the AMI would continually decrease with no maximum for any nonzero noise intensity. However, $H(y)$ is not constant with respect to σ_η , so a maximum at some nonzero noise intensity may be possible. Thus the AMI is a balance between the gain in information (the $H(y)$ term) and the information lost (the $H(y|x)$ term) during transmission. This makes it possible for the AMI to be maximal for some nonzero value of noise intensity.

The reason that $H(y)$ increases is due to the degrees of freedom that are gained as noise is added to the system. It allows the response to reach more states and consequently the information content of the system increases.

The improvement in AMI with N can be derived from the law of large numbers. The signal response grows with N whilst the error only grows with \sqrt{N} , hence the transmitted information grows with increasing N [Stocks, 2001b].

To numerically calculate the AMI, the calculations shown in Appendix B.3 are used to simplify the AMI in (4.32) to

$$\begin{aligned} AMI = & - \sum_{n=0}^N P_y(n) \lg P'(n) \\ & - \left(-N \int_{-\infty}^{\infty} P_x(x) [P_{1|x} \lg P_{1|x} + P_{0|x} \lg P_{0|x}] dx \right) \end{aligned} \quad (4.34)$$

where $P_y(n) = C_n^N P'(n)$. The above equation only applies when the thresholds are set identically, i.e. $\{\theta_i\} = \theta$. $P_{1|x}$ is the conditional probability of a single device being in state 1 given the signal x and accordingly $P_{0|x}$ for state 0.

Suprathreshold Stochastic Resonance (SSR)

To simulate (4.34) we take both the signal and noise to be Gaussian distributed as

$$P_x(x) = \frac{1}{\sqrt{2\pi\sigma_x^2}} \exp\left(\frac{-x^2}{2\sigma_x^2}\right), \quad (4.35)$$

using η instead of x for the noise distribution.

4.4 Quantifying SR

The only issue remaining is how to set the thresholds for the N devices. The aim is to explore the system in the SSR regime, thus all the thresholds are set to the mean of the input signal, $\{\theta_i\} = \langle x \rangle$. This naive setting of the thresholds, as it turns out, is very suboptimal and SSR is exhibited.

Figure 4.28 shows the results of numerically solving (4.34) and directly simulating the network with all the thresholds set to the mean of the input signal. Clearly, a SSR type of behaviour is observed for $N > 1$. It is required that $N > 1$, otherwise we are looking at a classical system with only a single STR operating in the suprathreshold region, but we have shown in Figure 4.4 (on page 97) that additional noise only degrades performance. With additional STRs in the network, additional noise can improve the transmitted information.




Figure 4.28. Average mutual information for a network of threshold devices. The signal and noise are both derived from Gaussian distributions and the devices have their thresholds set to $\langle x \rangle = 0$. The normalised noise is $\sigma = \sigma_\eta / \sigma_x$. The crosses are from experimental simulations, the solid lines from numerically integrating (4.34), and the diamonds from theory when signal and noise intensities are equal from (4.36).

An important difference of SSR to classical SR is the threshold does not need to be delicately adjusted when the signal strength changes, as long as the noise intensity is scaled accordingly. For a SR system, when the signal exceeds the threshold, noise is bad – the system becomes degraded. However, SSR is observable for any signal and noise strength without tuning as long as the noise and signal scale together. This is

directly noticeable by the x -axis label in Figure 4.28, where the normalised standard deviation is defined as $\sigma = \sigma_\eta/\sigma_x$, the ratio of the standard deviation of noise to signal.

As a special case, consider when the noise and signal strengths are equal, $\sigma = 1$ [Stocks, 2001b]. If the PDFs of the signal and noise are the same the AMI can be analytically determined as

$$AMI\Big|_{\sigma=1} = \frac{1}{2} \lg(N+1) - \frac{1}{2 \ln 2} + O\left(\frac{1}{N}\right), \quad (4.36)$$

which approaches $\frac{1}{2} \lg N$ for large N . This also shows that a noise induced maximum must occur for AMI at a nonzero noise intensity [Stocks, 2001b].

Optimal Placement of Thresholds

The configuration of setting $\{\theta_i\} = \theta$ is far from optimal, hence the large gain seen in Figure 4.28. By extrapolating from the earlier result, using an optimal choice of $\{\theta_i\}$ to match the statistics of the input signal, no gain via SSR should be possible. The problem is to find $\{\theta_i\}$ that maximises the AMI, which is discussed in detail in [Stocks, 2000b]. $H(y)$ is maximised when all of the output states occur with equal probability. With N devices there are $N+1$ output states, so the maximal value is $\lg(N+1)$ and occurs when $\sigma = 1$. In contrast $H(y|x)$ monotonically increases with σ so the maximal AMI must occur for $\sigma \leq 1$ [Stocks, 2001b]. A number of plots of $H(y)$ and $H(y|x)$ are shown in Stocks (2001b) that illustrate this fact.

If all of the states are equally probable, then $P_y(n) = 1/(N+1)$, and in the absence of noise we also have $P_y(n) = \int_{\theta_n}^{\theta_{n+1}} P_x(x) dx$. Equating the two expressions of $P_y(n)$ gives the optimal set $\{\theta_i\}_{\text{opt}}$ by solving

$$\frac{1}{N+1} = \int_{\theta_n}^{\theta_{n+1}} P_x(x) dx, \quad (4.37)$$

for $n = 1, \dots, N$ where θ_0 and θ_{N+1} can be taken as minus and plus infinity respectively.

For signals with a uniform distribution between the limits $\pm A$, the optimal threshold separation is $\Delta = 2A/(N+1)$ [Stocks, 2000c; Stocks, 2001c]. For Gaussian distributions solving (4.37) is not as straight forward, but can be found in terms of the error function. The optimal thresholds are shown in Figure 4.29. We note that the distribution of optimal thresholds is fairly evenly spaced, although not throughout the whole distribution.

4.4 Quantifying SR




Figure 4.29. Setting the optimal thresholds for a Gaussian distributed signal. The PDF of the signal is divided equally between the N thresholds. The k s give the relative spread of the thresholds compared to the optimal configuration when $k = 1$.

Only the relative difference between the thresholds and signal is important, that is we get the same result if either the thresholds are spread out ($k \rightarrow 0$) or as $\sigma_s \rightarrow 0$. The spreading factor k was chosen to mimic the behaviour of Stocks' example with uniform distributions. The plots in Figure 4.30 show the AMI for different spreading factors.

As expected, when $k = 1$ further noise only serves to degrade the AMI – it is the optimal set of thresholds. A noise induced gain is observed in each plot of Figure 4.30 due to two separate mechanisms. When k approaches zero in Figure 4.30a, the signal becomes small enough to fit between two adjacent thresholds, thus making the system subthreshold. This can only occur when N is even. The improvement in AMI is due to traditional SR.

As k increases, the trend of deteriorating AMI continues while the overhang of signal outside the thresholds remains relatively small. Once the signal becomes too large, the thresholds are being used very inefficiently, thus a SSR induced gain is possible. In the limit as k approaches infinity the threshold separation becomes insignificant and the network effectively acts as a single threshold with the AMI approaching the case where $\{\theta_i\} = \langle x \rangle$. These characteristics are consistent with Stocks' uniformly distributed results. In summary, SR extends the dynamic range of the ADC network for subthreshold signals while SSR compliments by assisting in large signal amplification [Stocks, 2001b].




Figure 4.30. Transmitted information for suboptimal settings of thresholds. The curves are from numerically simulating (4.32) with the signal and noise derived from Gaussian distributions. The AMI for $k = 1, 1/2, 1/4, 1/10, 1/18$ and $1/24$ are shown in (a) and $k = 1, 2, 4, 10, 20$ and ∞ in (b) from top to bottom in both.

Other Information Theoretic Techniques

There are also other methods and systems besides the one just investigated that have been analysed in terms of information theory. However, whatever the alternative terminology or notation (e.g. $I(X; Y)$, $I[\{t_i\} \rightarrow s(\tau)]$, T or simply I), the quantity remains the same and is generally defined similarly to (4.31). Systems considered include binary channels [Chapeau-Blondeau, 1997b], ion channels [Goychuk and Hänggi, 2000], neuron spike trains or threshold crossings [DeWeese and Bialek, 1995; Strong et al., 1998], FHN neuron models [Heneghan et al., 1996], integrate and fire neurons [Bulsara and Gammaitoni, 1996], and in a quantum setting [Goychuk and Hänggi, 1999].

Most of the aforementioned references then proceed to give the information rate, as opposed to just the information – the difference being the former is in bits/second. This leads us to the channel capacity, which is discussed in the following section.

4.4.5 Channel Capacity

It has recently been realised by [Goychuk and Hänggi, 1999] and others working in the field (unpublished) that Shannon's information channel capacity is the proper measure to be applied for STRs. The problem is defining the maximal bandwidth of the system. Very recently, this problem was solved [Kish et al., 2001]. The details are as follows.

4.4 Quantifying SR

It had been assumed that the SNR_{out} was a sufficiently good way of characterising the quality of the output signal and that the best coherence between it and the input signal is achieved when the ratio of the SNR at the output versus the input is maximised. That is, the most information about the input signal is transferred through the system to the output, hence we have maximal information transfer. More recently, several new methods of characterisation, which are similar in nature, have been proposed using entropy [DeWeese and Bialek, 1995; Strong et al., 1998; Goychuk and Hänggi, 2000; Stocks, 2000d]. It is common to characterise SR using the entropy difference given by

$$I = H_{output} - H_{lost} \quad [\text{bits}], \quad (4.38)$$

where H_{output} is the entropy of the noisy output signal and H_{lost} is the entropy lost during the transfer of the signal through the STR. This quantity has the same efficiency of output signal characterisation as the SNR_{out} and I has a similar potential for characterisation as $G = SNR_{out}/SNR_{in}$.

According to Shannon, and Nyquist, [Shannon and Weaver, 1949; Shannon, 1949] neither the SNR nor I are sufficient measures of the effectiveness of channel capacity. They only provide information about the entropy of the signal versus the noise, and the degradation of this entropy during transfer. This is directly related to the potential information content of the output, i.e. the resolution of the output. However, it does not say anything about the channel capacity. Simply speaking, these quantities refer to the amount of information, but they do not say anything about how frequently this information is refreshed. This fact is immediately obvious if we look at the dimension of I which is the bit. The proper dimension of the information transfer rate is bits/second. This is obvious from Shannon's formula (and the similar Nyquist formula), which was one of the most important milestones in information theory,

$$C = B_s \lg\left(1 + \frac{P_s}{P_n}\right) \quad [\text{bits/second}]. \quad (4.39)$$

C is the channel capacity, B_s is the maximal bandwidth of the signal, P_s and P_n are the signal and noise power respectively. According to Shannon, (4.39) can be interpreted as follows: half of the logarithmic term is the information entropy and $2B_s$ is the frequency of refreshing this information, since the equation is for continuous signals where Wiener's sampling theorem holds. For the validity of (4.39) in practical cases, any noise outside the frequency bandwidth of the signal is removed by a linear filter. The bandwidth B_s in (4.39) is the key parameter which refers to the rate of refreshing

the information. The logarithmic term refers to the potential amount of information available (resolution) at each refresh time. A low value of the information can be compensated by a high refresh rate, that is by a large bandwidth. The amount of information alone is meaningless for the characterisation of the quality of signal transfer. It is noted in Strong et al. (1998), without using either the Shannon channel capacity or the signal-to-noise ratio, that the information refresh rate is important.

For example, the elements of Morse code can be described by two bits (short beep, long beep, short pause, long pause), so two bits are enough to communicate via this method. The two bits correspond to the base of the logarithmic term in Shannon's formula. The information transfer rate will be determined by the mean frequency of beeps and pauses, which corresponds to the bandwidth B_s in the Shannon formula.

The aim is to estimate the information transfer rate of neurons in the stochastic resonance region using Shannon's formula. In this region, the input signal amplitude is less than the value of the threshold of the neuron. Moreover, the linear response approach will be used, which means that the input signal amplitude is less than the root-mean-square (RMS) noise amplitude. Thus, the signal response remains linear while that of the noise does not. A further assumption needed to ensure a linear response is that the firing rate of the neuron is much lower than the reciprocal of the refractory time.

For the calculations, which are given in Appendix B.4, Kiss' threshold crossing theory [Gingl et al., 1995; Kiss, 1996] is used. This theory describes the SNR and bandwidth of a simple neuron model. From calculations in Appendix B.4, the channel capacity is

$$C = \frac{B_{n,in}}{2\sqrt{3}} \exp\left(\frac{-U_t^2}{2B_{n,in}S_{n,in}}\right) \lg\left(1 + \frac{(2AU_t)^2}{(B_{n,in}S_{n,in})^2}\right), \quad (4.40)$$

where U_t is the excitation threshold the neuron, and A is the RMS amplitude of the input signal. The main difference between this measure and others is that they calculate an information content type of quantity, which is shown to have a maximum with noise intensity. This is simply the signal-to-noise ratio expressed by other measures. When referring to optimised channel capacity, it is assumed that the channel bandwidth does not change.

However, this is not the case as shown by Figure 4.31, where the two terms constituting C are plotted separately. By inspection, the shapes of the curves dictate that there must exist some nonzero value of noise that maximises C . From the interpretation of C given after (4.39), it is a trade-off between how much information can be refreshed and the information entropy (i.e. the content). When $S_{n,in} = 0$, we can have an infinite

4.4 Quantifying SR




Figure 4.31. The two components of the channel capacity. Increasing noise has the trade-off of increasing bandwidth (\diamond) at the expense of the entropy term (\square). This leaves a nonzero noise intensity that maximises C .

amount of information, but unfortunately zero bandwidth. Conversely when $S_{n,in} \rightarrow \infty$, the bandwidth increases but the entropy rapidly diminishes. By selecting the correct amount of noise an optimal compromise can be made.

Alternatively, according to theory [Kiss, 1996] and analog simulations [Gingl et al., 1995], the signal-to-noise ratio at the output is given as

$$SNR_{\text{out}} = \frac{2}{\sqrt{3}} B_{n,\text{in}} \frac{(AU_t)^2}{(B_{n,\text{in}} S_{n,\text{in}})^2} \exp\left(\frac{-U_t^2}{2B_{n,\text{in}} S_{n,\text{in}}}\right). \quad (4.41)$$

Comparing (4.40) and (4.41), it is obvious that both equations display a maximum against the input noise intensity $S_{n,\text{in}}$ given a fixed input noise bandwidth $B_{n,\text{in}}$.

In Figure 4.32, the channel capacity and the output SNR are plotted against noise intensity. For high noise intensities $P_s \ll P_n$ and C approaches zero. The actual value and the shape of C can be modified by linear filtering the output to reduce the bandwidth when the signal is not fully utilising all of the possible bandwidth [Kish et al., 2001]. The stochastic resonance peak for the channel capacity occurs at higher input noise intensities than the SNR_{out} curve.

Discussion

There are several points that need to be further clarified in the previous discussion of channel capacity. They have been omitted until here so as to not interrupt the flow.

One issue relates to the use of the Shannon channel capacity formula (4.39) with the STR. The formula requires that all signals and noise are Gaussian and that the system




Figure 4.32. Channel capacity and signal-to-noise ratio. The channel capacity (\diamond) and output signal to noise ratio (\square) of the neuron model (signal amplitude 0.1 V, threshold 1.0 V, input noise bandwidth 100 kHz).

is linear. Clearly, by design the STR is nonlinear, which implies (4.39) is only valid for weak noise levels. However, we are considering the channel capacity at the output, so it does not matter if the system is nonlinear – by lowpass filtering the nonlinear output of the STR, it is made Gaussian. Setting the cut-off frequency of the LPF to B_s means only the frequency components within the signal bandwidth are retained. The fact that the output is Gaussian is due to the central limit theorem.

Another issue relates specifically to Stocks' SSR networks. The system in question (shown in Figure 4.27), does not use any filtering at the output, thus the response is non-Gaussian and highly nonlinear. Therefore I is not simply related to SNR_{out}/SNR_{in} . This is illustrated by considering the information rate as $R = B_s H'$, where $H' = \lg(1 + SNR)$. Using this formula to calculate H' for the threshold network in the absence of noise gives H' of infinity, which is independent of B_s . This is obviously incorrect, a single comparator cannot transmit an infinite amount of information independent of B_s . The correct answer, in accordance with (4.33), is that $H' = 1$, which is well defined and finite when the noise intensity is zero. Therefore, for the SSR system, H' is not simply $\lg(1 + SNR)$ and B_s is independent of the noise intensity. Thus, it only makes sense to calculate AMI. In reality, the maximal information rate is set by the maximal bandwidth, which is directly related to the switching rate of the comparators.

4.5 Noise Related Phenomena

However, the situation is very different if the signals are subthreshold, in which case the bandwidth becomes a function of noise intensity. Then R and H' are not linearly dependent and the previous discussion on channel capacity becomes relevant. One advantage of AMI though, is it does give some indication about the lost information, however it says nothing about the information rates.

In summary, the appropriate measure is dependent on the system under investigation. The measure utilised must fully characterise the channel capacity. Much of this discussion is based on correspondence between Stocks and Kish resulting from Kish et al. (2001).

4.5 Noise Related Phenomena

4.5.1 Noise Linearisation

When a signal passes through a nonlinear system, the response emerges distorted. The severity of distortion depends on the frequency dispersion of the responses, the system nonlinearity and the shape and size of the signal [Dykman et al., 1994]. However, by adding noise of sufficient intensity, the distortion can often be removed by a process known as (NIL) *noise induced linearisation* [Dykman et al., 1994; Stocks et al., 1996].

Figure 4.33 shows simulations of the bistable system driven by a periodic force. The subthreshold input is shown at the top and consists of sawtooth and sinusoidal waveforms. The output shown is the ensemble average $\langle x(t) \rangle$ with the added noise intensity increasing as we move downwards. For small noise intensities $\langle x \rangle$ is badly distorted. The distortion steadily diminishes as the noise intensity increases – though there is a limit to the amount of noise that can be added.

It is noted in Dykman et al. (1994) and Stocks et al. (1996) that the word linearisation is used in two different senses. Since the sinusoidal signal passes through the system without significant change of shape, it implies linearity in the direct sense, i.e. amplitude proportionally. The sawtooth waveform, which contains higher harmonics, also passes through without significant change, which implies linearity in the “hi-fi” sense. That is, the system becomes nondispersive for certain frequency ranges, which depends on the noise intensity added. The effect of the added noise is to smear the system over a larger region of the phase space so that a variety of different scales and




Figure 4.33. Noise induced linearisation for a sawtooth and sinusoidal wave. The input signal is shown at the top, with ensemble averaged responses $\langle x(t) \rangle$ from the DWP for increasing noise intensity D shown beneath it.

frequencies become accessible to the system [Dykman et al., 1994]. Plots of the phase space are shown next section.

4.5.2 Phase Space

Another method of viewing a system is via its *phase space*. The phase space simply relates the internal or output variables of a system. For the type of systems described in Sec. 4.2, it is common to plot a variable and its derivative. The phase space allows the interactions between the variables to be easily observed. We first consider the DWP as used for the NIL in the previous section, and then the FHN neuronal model with SR.

For the DWP described by (4.18), we plot \dot{x} against x . In physical terms, it is the particle's velocity (in the x -direction) against its position. The phase space for several noise intensities is shown in Figure 4.34, where the forced input signal is a sinusoidal wave with amplitude slightly above the potential barrier between the two wells. The reasonably smooth thicker line is the noiseless trajectory – which is of course identical for each cycle. The trajectory originating from the centre of the plot, more easily observable with low noise, is from the particles initial position. Once the system is started and for low noise intensities, this point is not revisited. As the noise intensity increases the particle starts to explore more of the phase space. With too much noise the particle jumps around the phase space erratically.

4.5 Noise Related Phenomena

It is interesting to note the relationship between the velocity and position of the particle in the absence of noise, noting the particle's trajectory is counter clockwise. Once the particle is at a stable point in one of the wells ($x = \pm 1$), it moves out slowly, noted by the low \dot{x} . But, once it approaches and crosses the potential barrier at $x = 0$ the velocity increases then decreases until the particle is in the other well.




Figure 4.34. Phase space of a particle in the DWP. The phase space is constructed by plotting the particle's velocity \dot{x} against its position x . The thicker, smoother line is the trajectory in the absence of noise. The amplitude of the sinusoidal input signal is slightly above the barrier potential.

The phase space for the FHN model is achieved by plotting \dot{v} against v , as shown in Figure 4.35. This shows the phase space when the neuron is in the SR region by setting the amplitude of the input signal less than the activation level. Initially when the noise is weak, the neuron does not enter the excited state and the phase space is limited to a relatively small area (see top left plot). Once sufficient noise is added, excitations are induced and the phase space cycles around in a much larger loop. This remains the




Figure 4.35. Phase space of the FHN neuronal model. The phase space is constructed by plotting \dot{v} against v . The forcing signal amplitude is subthreshold and the noise intensity increases from left to right.

case even with quite a large amount of noise is added. With too much noise, the cyclic loop becomes indistinguishable.

4.6 Summary

We introduced the basics of SR via the use of a black box system, and a number of examples are given of systems that exhibit SR. These show that noise is only of assistance when the system is constrained in some way. If all the parameters of a system are accessible, the system can usually be made optimal in which case the addition of noise offers no support. We then described the types of noise and systems that are applicable to SR.

With any type of system, in order to study its behaviour, a method to measure its performance is required. Since the conception of SR, a number of measures have been employed. The different measures have arisen in the evolution of the SR field, in some

4.6 Summary

cases certain measures have been superceded, and in other cases some measures have been shown to be inadequate.

The signal-to-noise ratio and residence time distributions are used for classical SR, which uses a subthreshold, periodic input signal. Though the peak observed in the commonly used SNR is not strictly a resonance – by using residence times a *bona fide* resonance against forcing frequency is observed. With the extension of SR using aperiodic signals, correlation coefficients are employed, which can be further generalised using the information content of the output signal. However, all of the aforementioned metrics effectively measure the resolution of the output signal. The useful quantity is the information transfer *rate* of the system, which must take into account the noise dependent bandwidth. This is covered by considering Shannon's channel capacity.

Finally we briefly looked at the related phenomenon of noise induced linearisation, and the resulting phase space generated by a system.

Having given a detailed description of SR, the mechanics of noise enhanced systems are expanded. These mechanisms will be used in successive chapters that deal with motion detection. In the next chapter, several biologically inspired detection schemes are introduced and explained. Only a simple noise analysis is given therein, with a more comprehensive analysis given in the following chapter.

Motion Detection

DETEECTING motion is one of the most fundamental tasks of many biological visual systems. As opposed to full image processing, biological schemes offer a relatively efficient and parallel way to detect motion. In this chapter we consider several schemes specifically targeted at edge detection. These include the Reichardt detector, a detector based on shunting inhibitory neurons, and a template model based on insect vision.

A gradient scheme based on the spatial and temporal derivatives of a signal is used to generate an 'ideal' motion detector response. This is useful for comparing responses from other detectors. We briefly look at responses to noisy signals to gauge the robustness of the systems.

5.1 Introduction

One of the most fundamental abilities of biological visual systems is to sense and perceive visual motion. In this context, there is a distinct difference between sensing and perceiving, words that may otherwise be used interchangeably [Yakovleff and Moini, 1996]. *Sensing* refers to the actual physical process of converting light into an appropriate signal, for example, a camera converts incoming light into levels of luminance. *Perception* on the other hand is the interpretation of the sensed signal. This could be something as simple as thresholding, or as complex as reconstructing a three dimensional environment.

This fundamental ability of biological systems to measure and use visual motion serves many essential functions [Hildreth and Koch, 1987]. In terms of simple systems, like those of insects or small animals, there are certain characteristics that indicate potential disaster, or conversely, opportunity. Sudden movements may indicate the possibility of predators or prey. The frog for example, can effectively detect small dark objects ('bugs') in its visual field, while the house fly can track moving objects when their texture is identical to that of the background [Ullman, 1983]. Work carried out by Srinivasan and colleagues with honeybees has explained their accurate navigation over short distances to be based on the visual flow of information, or *optical flow* [Srinivasan et al., 1996; Srinivasan et al., 1997b]. More generally, motion signals provide input to control eye movements, and these signals form a type of control system that allows objects of interest to be tracked [Hildreth and Koch, 1987].

The examples given above show the effectiveness of biological visual systems. Given the fact that an insects brain only weighs a tenth of a milligram and it contains less than 0.01 % of the neurons of the human brain [Srinivasan et al., 1997a], this achievement is remarkable. In comparison, a computer controlled robot cannot come anywhere near matching the real-time proficiency of even a fly without being tethered to a powerful computer [Franceschini et al., 1992]. Thus, by learning and utilising the mechanics of insect visual systems, a great deal can be gained [Horridge, 1992; Srinivasan et al., 1997a].

An alternative approach to using dedicated hardware to process visual signals is to incorporate processing at the pixel level, using devices referred to as *smart sensors*. That is, it does not just sense, but also performs simple processing. These pixel level computations can take advantage of parallel architectures and use information from the neighbouring resources as required.

Full image processing has a high computational overhead and hence cannot easily be implemented as a smart sensor. However, the goal here is to use simple edge detection algorithms that only use local information – this is well suited to smart sensors. Higher order information can then be derived by post processing the edge detection information. Edge detection can be useful for velocity estimation [Nguyen, 1996]. A problem is that the velocity response is also influenced by the contrast levels, spatial and temporal patterns of the stimulus [Harris et al., 1999; O’Carroll et al., 1997]. Although the contrast levels also effect edge detection responses, a larger response magnitude can be interpreted as a more definite edge.

Next, we set the groundwork for describing and classifying motion detectors, and give the criteria to be satisfied for a detector to be direction selective. We then consider three biological schemes that satisfy the given criteria. They are; the Reichardt detector that uses an excitatory interaction, a scheme based on a shunting neuron that uses a inhibitory mechanism, and a template model that uses spatiotemporal templates. We also describe the gradient scheme for use as a *control*.

5.1.1 Local and Widefield Configuration

We consider a local detector using only information from neighbouring resources. This could be a neighbouring pixel in the image source or other local motion detectors. They can range from simple configurations that only consist of a single interaction to those containing many functional units. The local motion detectors are depicted by the row of central squares in Figure 5.1. The rectangles either side are the image sources and output motion vectors. The widefield configuration is simply a row of local detectors, as shown in Figure 5.1.




Figure 5.1. Local and widefield detectors. The outside rectangles represent the image vectors and motion outputs. The centre squares represent the local motion detectors, which can utilise neighbouring resources.

5.1.2 Classification of Detection Schemes

Motion detection schemes can be divided into two broad categories, artificial and biological, determined by the origins of the scheme. The classification hierarchy is shown in Figure 5.2.

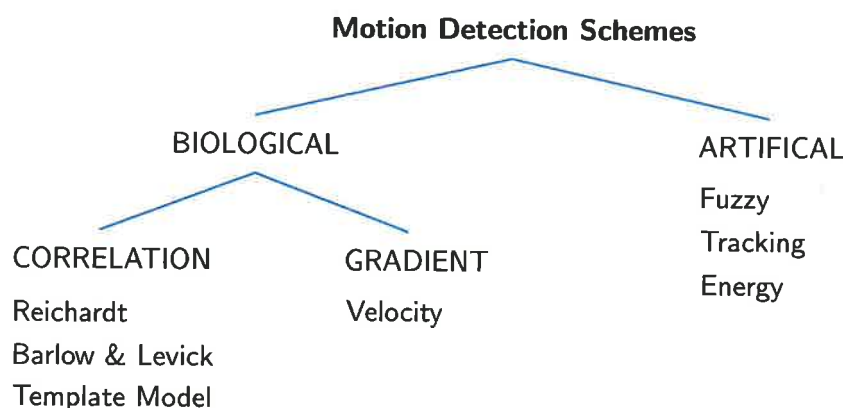


Figure 5.2. Classification hierarchy of motion detection schemes.

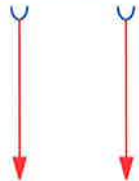
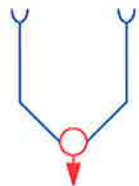
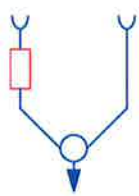
Energy models extract motion information by detecting the oriented spatiotemporal energy. This can be accomplished by using a quadrature pair of spatiotemporally oriented filters [Bouzerdoun and Pinter, 1989]. Alternatively, there is the motion-energy, which in its simplest form is the difference between two successive images. Operations can then be performed to determine the static and dynamic regions [Sezgin et al., 1995].

Tracking methods use some type of *feature* that is tracked from one image to the next [Harrison and Koch, 2000]. Spatial features are typically the edges, which are easy to detect due to the rapid change in intensity over a short distance. The edges can be difficult to track if the feature density is high. Temporal features can also be used, which are typically rapid changes in the image brightness at the pixel level. Whatever features are used though, spurious responses can be generated for weak signals.

The biological models can be further subdivided into gradient and correlation schemes [Borst and Egelhaaf, 1989; Ullman, 1983]. These models use intensity based algorithms, which are derived from the image brightness or luminance. Although some energy models use intensities or gradients, they have been classified as artificial due to their non-biological background.

There are a large class of detectors that employ spatiotemporal correlations, also termed delay and compare. They primarily function by comparing the signals from spatially

Table 5.1. Requirements for a directionally selective motion detector.

	<p>Two inputs. A single receptor (represented by the half circle) can not distinguish an object as moving left or right. Since motion is a vector, two receptors are necessary to be able to detect motion.</p>
	<p>Nonlinear Interaction. Without any nonlinear interaction the time averaged output of the detector is equal to the time averaged input signals. Thus all the temporal information is lost and the detector is not directionally selective.</p>
	<p>Asymmetry. This is required otherwise the direction of the moving object could be changed without affecting the detector output – it is not possible to determine which input was excited first. This is remedied by appropriately modifying one of the inputs.</p>

adjacent channels, one of which is delayed. These are the schemes that are explored in further detail in this chapter.

Principles of Motion Detection

There are certain requirements that need to be satisfied for a motion detector to be directionally selective [Reichardt, 1987; Borst and Egelhaaf, 1989]. It needs at least two inputs that are asymmetrical and interact with some nonlinear function. These requirements are described in more detail in Table 5.1.

5.1.3 Criteria for a Bidirectional Detector

For a system to be defined as directionally selective, certain criteria need to be satisfied [Beare, 1997]. For a local motion detector they are:

- The response must indicate the direction of motion. That is, motion in one direction must lead to a positive response, whereas motion in the counter direction leads to a negative response.

5.2 Biological Processing Schemes

- The response must be independent to the change of contrast between the moving object and the background. A dark object should have the same response as a light object moving along the same trajectory.
- Ideally the response should be robust to noise.
- For a stationary object there should be no response, even when there exists a spatial gradient between two receptors.

For the widefield response, two additional criteria are imposed:

- The position of the response should correspond closely to the position of the edge from the inputs.
- Hence, spatially separated edges should provide spatially separated responses.

However, the actual shape of the response is not that important as long as it follows the above criteria.

5.2 Biological Processing Schemes

In this section the biological schemes shown in Figure 5.2 are described. We first look at the correlation schemes and compare the fundamental differences between them, especially their ability to cope with noisy signals. We then look at a gradient model which will serve as a standard measure.

5.2.1 Reichardt's Delay and Compare Detector

One of the first models of movement computation was proposed by Hassenstein and Reichardt (1956) and is given in more detail by Reichardt (1961). In its simplest form, it consists of only a time delay and a multiplier element, shown in Figure 5.3.

Figure 5.3 shows two elementary motion detectors (EMD) and how the response to a left and right moving pulse is generated. In Figure 5.3a the pulse is first detected by the left receptor where it is delayed for a time interval τ . The pulse is then detected by the right receptor – if the time delay is appropriately tuned, the signals in both input channels will coincide at the multiplication unit, giving rise to a large response at the

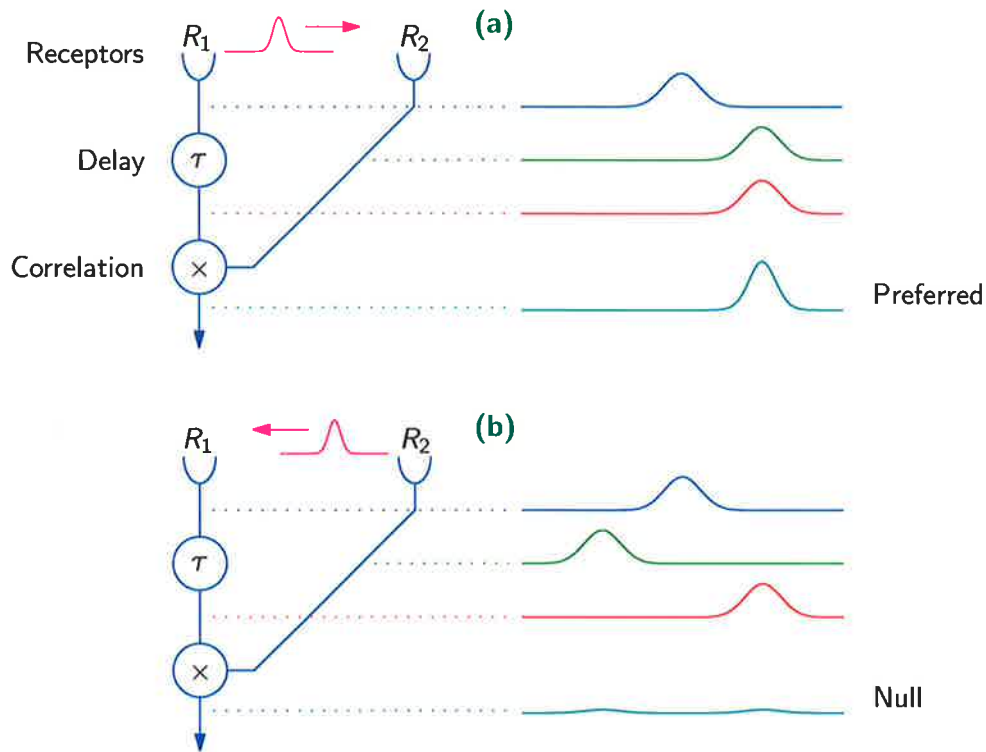


Figure 5.3. Preferred and null directions of a correlation detector. The top detector has the pulse moving in the preferred direction, which has a large response. The bottom detector has the pulse moving in the null direction, giving only a small response.

output. The spatial separation of the receptors is compensated by the delay in the left channel. This delay allows the detector to be tuned to a specific spike velocity.

However, if the pulse is moving in the opposite direction as shown in Figure 5.3b, a different response is generated. Instead, the delay stage in the left channel shifts the pulses further apart, thus producing a minimal response. The small response that is generated is only due to the background luminance. The two directions shown in Figure 5.3 are referred to as the *preferred* and *null direction*. This is common in many types of local motion detection systems.

A bidirectional detector is formed by placing two EMDs back-to-back and taking the difference from the subunits, as shown in Figure 5.4 [Reichardt, 1987; van Santena and Sperling, 1985; Borst and Egelhaaf, 1989]. Motion in either direction is detected by the two unidirectional EMD units tuned to opposite directions.

From Figure 5.4, the response is given by

$$r(t) = R_2(t)R_1(t - \tau) - R_1(t)R_2(t - \tau), \quad (5.1)$$

assuming the delay unit is ideal.

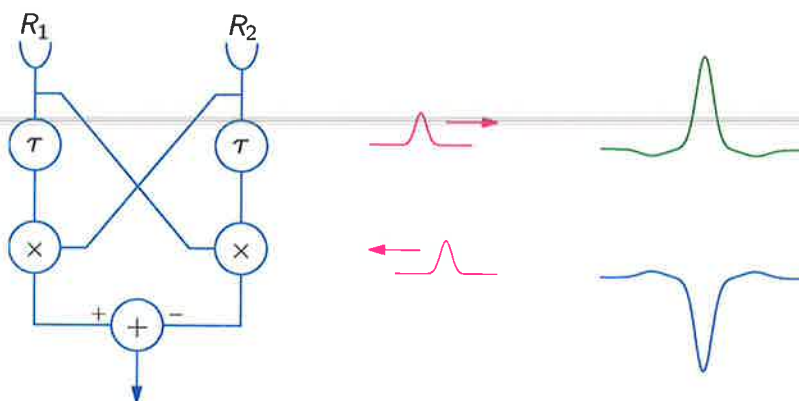


Figure 5.4. The Reichardt motion detector (RMD) consisting of two EMDs back-to-back.

The preferred direction for one arm is the null direction of the other, thus taking the difference of the two arms gives a response indicating the direction of the stimulus.

The diagrams shown in Figure 5.3 and Figure 5.4 have been designed to be illustrative, accordingly ideal inputs and filters were assumed. Since we wish to design an edge detector, thus the inputs need filtering to possess more desirable characteristics. That is, more like the pulses shown in Figure 5.4. Simply differentiating the signal (i.e. a highpass filter) is not appropriate for reasons mentioned later.

Next, we detail the test signal and the delay stage, taking into account the realisation of a physical system. The Reichardt detector is explained showing the advantages of filtering the input signals.

Test Signal

A standard test stimulus used for edges is given by

$$L(t) = L_0(1 + cu(t - \tau)), \quad (5.2)$$

where L_0 is the background luminance, c is the contrast between the edge and the background ($-1 \leq c \leq 1$), and $u(t)$ is the unit step function that gives 1 when $t \geq 0$, or 0 otherwise. A positive edge is generated when $c > 0$ (i.e. $L_0 \rightarrow L_0(1 + c)$) and a negative edge formed when $c < 0$ (i.e. $L_0 \rightarrow L_0(1 - c)$). The edge can be moving either left or right.

Delay Stage

Another difference from the illustrative diagrams is how the delay is implemented. Instead of being an ideal delay it is realised as a lowpass filter (LPF) with time constant

τ . The transfer function is

$$H(s) = \frac{A}{A + s}, \quad (5.3)$$

where A is the *cut-off frequency* – the inverse of the time constant. The time constant controls the decay rate, in other words, how quickly the output takes to reach the input. This implementation of the delay stage is easily realisable in analog electrical or biological circuits [Mead, 1989; Harrison and Koch, 2000; Moini et al., 1997; Torre and Poggio, 1978]. The variable $s = j\omega$ is commonly used to transform the frequency to the s -domain. Laplace transforms can be employed for artificial signals, but are not good for real broadband signals. For an iterative approach, the following relation can be used,

$$y_n = y_{n-1} + (1 - e^{-Ah})(x_n - y_{n-1}),$$

where $y(t)$ is the filtered output from the signal $x(t)$. This can be derived intuitively with some knowledge about LPFs, however a formal derivation is given in Appendix C. It is worth noting that Ae^{-At} is the inverse Laplace transform of $H(s)$ from (5.3).

Response of the Reichardt Detector

The response to a leftward moving edge from the Reichardt detector is shown in Figure 5.5a. We immediately notice two peaks generated for a single edge. This is because the interaction between the channels is a multiplication operation, i.e. *excitatory*.




Figure 5.5. Response of the Reichardt detector to a step input. The responses are to a rightward moving edge with $L_0 = 10$, $c = 0.25$ and $A = 30$. The response in (a) is from the detector shown in Figure 5.4 and (b) has the inclusion of bandpass filters after the receptors as described later in the text.

5.2 Biological Processing Schemes

Whenever an edge passes a receptor, the system is perturbed and generates a response. Since each output in the local Reichardt detector is sourced from two receptors, there are two pulses generated. Furthermore, the output pulses are of different magnitudes and signs due to the nonlinearity of the detector.

The peaks of the pulses occur at the times when the edge passes the receptors and no signal has passed through the delay filter. Thus we are able to find the magnitude of these peaks. For a single step moving rightward with a background luminance of L and a step luminance of $L(1 + c)$, the first peak is given by

$$r_{\text{peak1}} = -cL^2 \quad (5.4a)$$

and the second by

$$r_{\text{peak2}} = cL^2(1 + c). \quad (5.4b)$$

These expressions are valid for both positive and negative c . If the edge is travelling in the opposite direction, the peaks are interchanged and the polarity reversed. This makes sense as the two channels that form the RMD are identical.

For this detector to satisfy the criteria set out in Sec. 5.1.3, we need to consider it in the widefield configuration. By simply spatially adding the outputs of each local detector the criteria can be satisfied. In the case of a rightward step the second peak of the n th RMD corresponds in time to the first peak of the $(n + 1)$ th RMD, thus giving an output of

$$r_{\text{total}(\rightarrow)} = r_{\text{peak1}} + r_{\text{peak2}} = c^2L^2 > 0.$$

For a leftward pulse,

$$r_{\text{total}(\leftarrow)} = -c^2L^2 < 0,$$

which correctly detects motion irrespective of the contrast change. Nevertheless, the problem of the response at the detector limits remain. Since we must spatially sum all the outputs we cannot determine where the edge is in the widefield array. A solution is to group the detectors to form *detection windows*, although this only minimises the problem and the undesirable spikes of the wrong polarity appear at either edge of the detection window [Harmer and Abbott, 2000b; Beare, 1997].

RMD with Bandpass Filters

To improve the response of the Reichardt detector, extra filters are added. Incorporating a bandpass filter (BPF) after the receptors converts the edge into a pulse which is

desirable. An alternative is to simply differentiate the incoming signal, but this has poor noise performance. Hence, the BPF is used, which removes the high frequency components.

Usually a BPF is made from a lowpass and highpass filter in series [Harrison and Koch, 2000]. However, a very close approximation is to take the difference of two LPFs with different time constants as shown in Figure 5.6. See Appendix C for a comparison of magnitude responses. Tuning the parameters A_1 and A_2 in the LPFs controls the width and height of the BPF response. Assuming that $A_1 < A_2$, then A_1 determines the low frequencies that are attenuated and A_2 determines the high frequencies attenuated. In terms of frequency response, A_1 and A_2 are the 3 dB frequencies that define the bandwidth of the BPF. The response to a step edge using two LPFs to implement the BPF is shown in Figure 5.6b. This produces the preferred input for the Reichardt detector, as shown in Figure 5.5b.




Figure 5.6. The bandpass filter and response to step edges. The equivalent circuit for the BPF consisting of two LPFs with $A_1 < A_2$. The response to stimulus with luminance levels 10 and 12.5.

The RMD is far simpler than the original schematic shown in Reichardt (1961), which contains extra spatial and temporal filters to accurately model physiological properties. The basic behaviour of the RMD does not change greatly when extra linear, time-invariant or temporal filters are added [van Santena and Sperling, 1985]. Adding extra filters forms an elaborated Reichardt detector (ERD), which can be used to focus on different characteristics.

5.2.2 Barlow and Levick's Inhibitory Mechanism

Shortly after the work of Hassentein and Reichardt, Barlow and Levick (1965) found evidence of a complimentary *inhibitory* directionally selective mechanism in rabbit retinal ganglion cells. Instead of relying on passing activity from one receptor to the next in the preferred direction, the mechanism prevents activity passing through the receptors in the null direction. The two complementary schemes are shown in Figure 5.7, where only one direction is shown in each for clarity [Barlow and Levick, 1965; Ullman, 1983].




Figure 5.7. Two methods for directionally sensitive motion detectors. The preferred direction is towards the right for both schemes. (a) Reichardt's excitatory scheme uses the correlation of two signals, one of which is a delayed. A response is only given when the two signals coincide, that is, for an object moving rightward at the tuned velocity. (b) Barlow and Levick's inhibitory scheme functions by complementary mechanics. Instead of passing activity to the right, it passes an inhibitory signal to the left. The signal moving in the null direction will diminish the response.

The complementary aspects are depicted by the inclusion of the NOT operation in the inhibitory path and the lateral flow of information is in opposite directions – the preferred direction for the excitatory and the null direction for the inhibitory mechanism. However, both still possess the essential requirements of being asymmetrical and non-linear. The mechanisms proposed by Barlow and Levick (1965) were in terms of digital logic, though a pure delay for the inhibitory mechanism is not completely necessary since the inhibition only needs to exceed the excitation [Barlow and Levick, 1965; Torre and Poggio, 1978]. The other components can be transformed to an analog system using the appropriate conversions.

Shunting Inhibition

A type of lateral inhibition developed by Pinter (1983) describes the neurochemistry of visual cells. This is referred to as *shunting inhibition* and it is implemented by the shunting neuron [Bouzerdoux and Pinter, 1989; Bouzerdoux and Pinter, 1990], which is governed by

$$\frac{dm}{dt} = L - am - m \sum_i k_i f_i(X_i), \quad (5.5)$$

where $m(t)$ is the response (membrane voltage), $L(t)$ is the excitatory input, a is the self decay of the neuron, k_i are the weights, f_i are the activation functions and $X_i(t)$ are the inhibitory inputs. There is a clear nonlinear interaction between the inputs.

Following from previous designs of the RMD, a version of the local feedforward and feedback configurations are shown in Figure 5.8 [Beare, 1997]. The shunting neuron is denoted by an M, the excitatory input by a small circle (●) and the inhibitory input by a short bar (■). The feedback arrangement eliminates the need for a delay filter due to the implicit self decay of the neuron, however it is more difficult to determine basic characteristics.

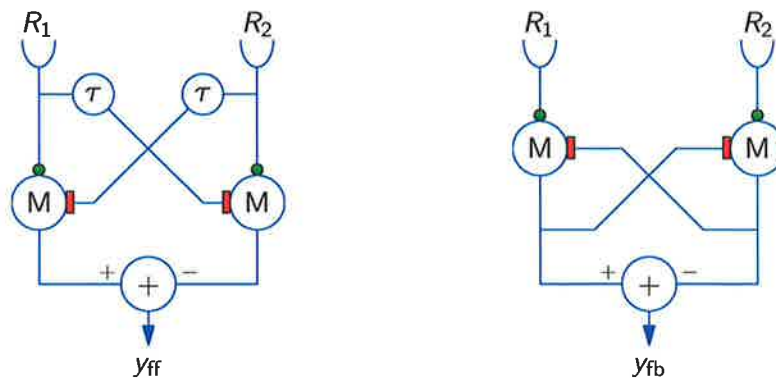


Figure 5.8. Feedforward and feedback local inhibitory motion detectors (LIMD). The feedforward configuration is a direct conversion of Barlow and Levick's digital scheme in Figure 5.7b, using the shunting inhibitory neuron as the interaction. The feedback configuration uses the self decay of the neuron in place of the low pass filters.

The responses for the two arrangements in Figure 5.8 are shown in Figure 5.9, which are quite similar. As they stand, the detectors suffer from two problems. The sign of the response is dependent on the direction of contrast change of the edge, and the response to a stationary edge is nonzero. The solution to these problems is presented in Sec. 5.2.3. We briefly present the steady state and adaptive properties of the shunting neuron.




Figure 5.9. Response of the LIMDs to a step input. The response of the feedforward (y_{ff}) and feedback (y_{fb}) LIMD to a rightward moving edge passing the receptors R_1 and R_2 as shown with $L_0 = 10$, $c = 0.25$, $a = 10$ and $A = 8$.

Steady State Response

We can simplify the interaction within the neuron by ignoring all the internal dynamics, which is achieved by setting $\dot{m} = 0$. Using (5.5), we let $k = 1$ and $f(X) = X$ to give the steady state response as

$$m = \frac{L_{exc}}{a + L_{inh}}, \quad (5.6)$$

where the L_{exc} and L_{inh} represent the excitatory and inhibitory inputs respectively. This is also the response gained when an edge has been sitting between two adjacent receptors for a sufficiently long time. Essentially we have reduced the neuron to a division operation which retains its adaptive properties while simplifying the mechanics.

The response of the feedforward arrangement shown in Figure 5.8a is given by

$$m = \frac{L_1}{a + L_2} - \frac{L_2}{a + L_1}. \quad (5.7)$$

This shows, due to the nonlinearity of the neuron, that a zero response is only given if $L_1 = L_2$. This is the undesirable result of a nonzero response when an edge is held between the receptors.

The steady state response of the feedforward inhibitory detector is shown in Figure 5.10 for a rightward moving edge. This highlights the importance of the delay filter, which controls how quickly the output reaches its final value after the initial jump




Figure 5.10. Steady state response of the feedforward LIMD. Response to the same edge shown in Figure 5.9. Omitting the internal dynamics of the neuron causes vertical jumps in the response. The curved sections are due to the LPF.

when the edge passes a receptor. The self decay a of the neuron is present in the denominator of (5.6) and influences the magnitude of the response.

Adaptive Properties

The adaptive properties of a sensor are important as they can compress the dynamic range of the output, making the post-processing easier. For example, standard outdoor luminance levels can vary as much as 12 orders of magnitude, while a photosensor operates over only half of this range [Gruss et al., 1991].

The Reichardt detector adapts poorly due to its multiplicative interaction, which leads to a much larger output than input range. This can be seen from (5.4) where the peak responses are proportional to L^2 .

In contrast, the neuron uses a division interaction to compress the dynamic range, i.e. it is adaptive. The peak response occurs once the LPF has decayed and has the value

$$m_{\text{peak}} = \frac{Lc(a + L(2 + c))}{(a + L)(a + L(1 + c))} \quad (5.8)$$

where the edge has a luminance of $L(1 + c)$ on a background of L . This is shown by the dashed line in Figure 5.9 and Figure 5.10.

In Figure 5.11, the peak response in (5.8) is plotted against luminance for several values of contrast. At the extremes of luminance the response asymptotically levels. In fact,

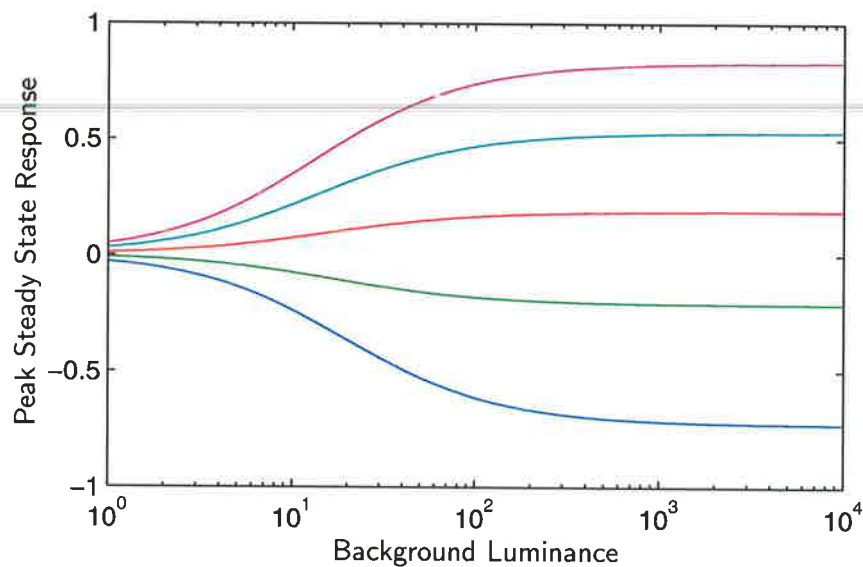


Figure 5.11. Adaptive properties of the feedforward LIMD. The peak response of the feedforward LIMD given by (5.8) for contrast values of 0.5, 0.3, 0.1, -0.1 and -0.3 from top to bottom ($a = 10$).

we can show that as $L \rightarrow \infty$ then $m_{\text{peak}} \rightarrow c(2+c)/(1+c)$ and as $L \rightarrow 0$ then $m_{\text{peak}} \rightarrow 0$. Thus, at high luminance levels the contrast becomes important.

Adaptive Bandpass Filter

Using the adaptive properties of the shunting neuron, it can be configured to form an adaptive bandpass filter [Beare, 1997]. The filter configuration and response to edges are shown in Figure 5.12. It can be explained by considering the neuron as a low-pass filter with a time constant controlled by the inhibitory input. Rewriting (5.5) as $\dot{m} = -(a+X)m + L$ gives an easy comparison to a LPF that has the form $\dot{y} = -\alpha y + \beta x$. Then the bandpass filter is formed by taking the difference of two LPFs with different time constants. Since one of the neurons in Figure 5.12 has a delay element in its feedback path, the time constants of each arm are different, hence a bandpass filter is formed. The response shown in Figure 5.12 shows the filter functioning correctly. This configuration could be customised further by placing delay stages in both arms, though they still need to be different. By plotting the peak response against luminance, the adaptive properties are displayed [Beare, 1997].




Figure 5.12. Adaptive bandpass filter architecture and response. The LPF configuration is shown on the left. The response on the right is to the dashed input shown with luminance of 12.5 and background 10, neurons have $a = 10$ and $A = 15$. This resembles the desired output of the BPF as shown in Figure 5.6.

5.2.3 Directionally Selective Local Inhibitory Motion Detector

Although each arm is asymmetrical in the previous LIMD detectors, the local detector as a whole is symmetrical, which presents problems. Ideally we require the properties of an adaptive nonlinear interaction that gives zero response to stationary inputs. An arrangement referred to as a directionally selective local inhibitory motion detector (DSLIMD) achieves this goal [Beare, 1997].

To provide the asymmetry in the local detector the excitatory and inhibitory input to both neurons are made to originate from the same receptor as shown in Figure 5.13. The delay element inserted before the left neurons inhibitory input provides the critical asymmetry, otherwise the response would always be zero.

The response to an edge moving in opposite directions is shown in Figure 5.13 where some useful properties are exhibited. The DSLIMD structure exhibits a preferred and null direction, observed by the different heights of the response pulses. This is similar to the EMDs in Figure 5.3. The EMD was converted to a useful detector by placing two of them back-to-back. As we will see later, the same can be achieved here.

Another property that may seem peculiar at first is that the pulses originate at different times even though the times the edge passes the receptors are simply interchanged. This is because it is an inhibitory detector.

5.2 Biological Processing Schemes

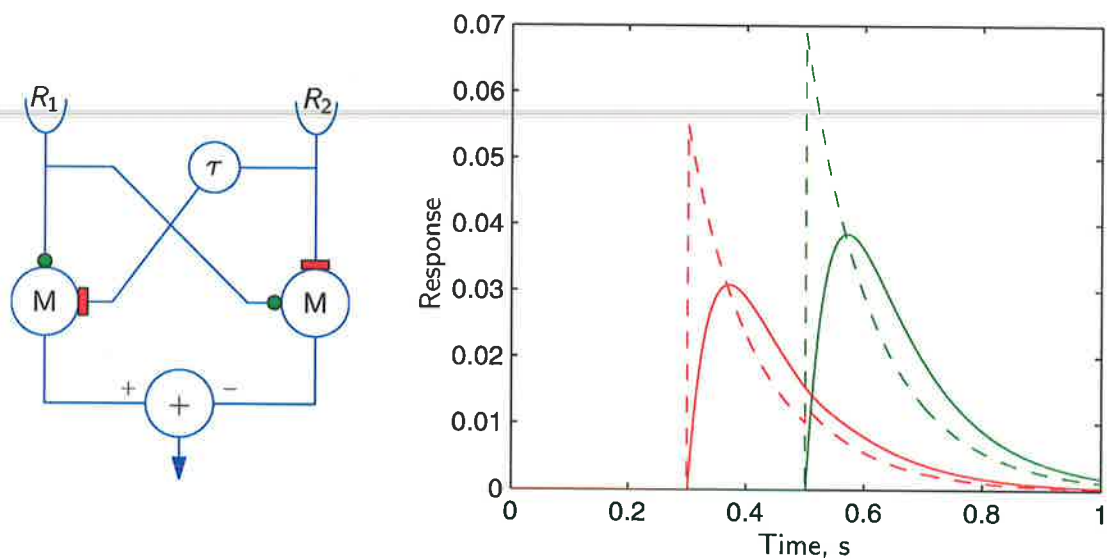


Figure 5.13. Asymmetrical LIND and responses to a step edge. The excitatory and inhibitory inputs are each derived from the same receptors. The delay element provides the asymmetry. The larger pulse in the response is for a rightward moving edge with $L_0 = 10$ and $c = 0.25$ passing the receptors (R_1, R_2) at times (0.3, 0.5). This is the preferred direction. The smaller pulse is for an identical leftward moving edge ($c = -0.25$), passing the receptors at (0.5, 0.3). This is the null direction. The dashed lines are the respective steady state responses. The neurons have $a = 10$ and $A = 8$.

We will consider the two cases. When the edge moves rightward, R_1 jumps from L to $L(1 + c)$. Assume that $c > 0$ for now. This signal passes identically to the excitatory input of the two neurons. The delay element is unchanged, and thus the output from both neurons increases equally and is cancelled by the final differencing operation. When the edge passes R_2 , the inhibitory input of the right neuron changes from L to $L(1 + c)$, but the change to the left neurons inhibitory input is delayed. Thus, the output from the right neuron is diminished (it is being inhibited), and a net response is generated. For a leftward moving edge, the process is very similar. As soon as the edge passes R_2 a response is generated – even though R_1 has no knowledge of it. The reason for the reduced magnitude is the rightward edge, R_1 has the value $L(1 + c)$ whereas for the leftward edge it is only L , hence there is not as much excitation to the neurons.

The arrangement in Figure 5.13 is not yet capable of bidirectional motion detection. The problem is that the responses are dependent on the direction of change of contrast, and there exist null and preferred directions. This can be seen by looking at the steady

state response for a rightward moving pulse

$$m_{\text{peak}} = \frac{c(1+c)L^2}{(a+L)(a+L(1+c))}, \quad (5.9)$$

and for a leftward moving pulse

$$m_{\text{peak}} = \frac{cL^2}{(a+L)(a+L(1+c))}. \quad (5.10)$$

These are clearly dependent on the sign of c . Moreover, even for an edge that is 'reversed', where it moves right with $c = 0.25$ and $L = 10$, and then left with $c = 0.2$ and $L = 12.5$, the magnitudes are different.

To rectify this as suggested earlier, two asymmetrical LIMDS can be placed back-to-back. The slight difference from the EMD case is that the local detectors need to be spatially displaced so that the responses from the inhibitory inputs align. The two asymmetrical LIMDs are shown as part of the widefield configuration of the DSLIMD in Figure 5.14.

Using the steady state simplification, the peak response occurs when an edge passes the inhibitory receptor, that is, before the delay element has kicked in. For a rightward edge, the peak is

$$m_{\text{peak}} = m_{\text{peak}(\rightarrow)} - m_{\text{peak}(\leftarrow)} = \frac{c^2L^2}{(a+L)(a+L(1+c))}, \quad (5.11)$$

which is always greater than zero for $-1 \leq c \leq 1$. For a leftward pulse, the peak response is simply the negative of (5.11). The adaptive properties are clearly shown by plotting m_{peak} against L .

The response of the DSLIMD to left and right moving edges, going to and from light and dark are shown in Figure 5.15. The different magnitudes of the pulses do not indicate preferred and null direction, but rather the different levels in luminance. If the edges only swapped between two values of luminance, the same peak magnitudes are observed.

5.2.4 Horridge's Template Model

The template model, originally developed by Horridge to model insect vision, shares both biological and engineering heritage [Horridge, 1990; Horridge and Sobey, 1991].




Figure 5.14. The DSLIMD in the widefield configuration. The section enclosed by the dashed line is the DSLIMD. It spans across three receptors, but only gives a response when an edge passes the central one, the inhibitory input.

Motion information is stored in *templates*, which contain current and delayed spatiotemporal data. This is loosely classified as a correlation type of detector, though the templates have no specific physical interpretation.

The mechanics of the template model are relatively simple, a key advantage of the system. The architecture is shown in Figure 5.16. The channels are first bandpass filtered to indicate the change of intensity. The signal is then thresholded and sampled into one of three quantised states, increased (\uparrow), decreased (\downarrow) or no-change ($-$) in intensity depending on whether the signal was above, below or between the threshold limits. The templates are then formed by joining adjacent spatial channels at successive sampling instants. This forms the 2×2 template that contains the spatiotemporal information.

Given that each element can have three values and there are four elements in each template, there are $3^4 = 81$ total possible templates. Obviously, not all the templates are directionally sensitive, only those that have a 3 : 1 diagonally symmetric structure shown in Figure 5.17. However, not all the templates are equally reliable. The




Figure 5.15. Response of the DSLIMD to all possible edges. The rightward edge passes the receptors (R_1, R_2, R_3) at times (0.1, 0.3, 0.5), and the leftward edges at (0.5, 0.3, 0.1). The edge and neuron properties are as for Figure 5.13b.

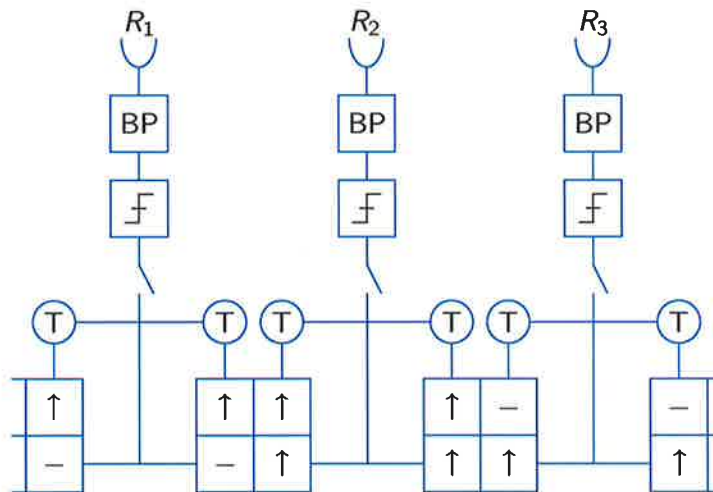


Figure 5.16. Horridge template model in a widefield configuration. The signal from the receptors is bandpass filtered, thresholded, sampled and then stored in a 2×2 template containing adjacent spatial and temporal channels (spaced by T seconds). The template is then used with the look-up table in Figure 5.17. The configuration of arrows and dashes shown are arbitrary and are drawn for illustrative purposes.

5.2 Biological Processing Schemes

templates that contain three no-change states are possibly due to flicker noise or object occlusions, whereas the templates without any no-change states are possibly due to coherent motion on a varying noise intensity background. This leaves eight templates that reliably detect coherent motion [Nguyen et al., 1995]. These are the coloured templates shown in Figure 5.17.

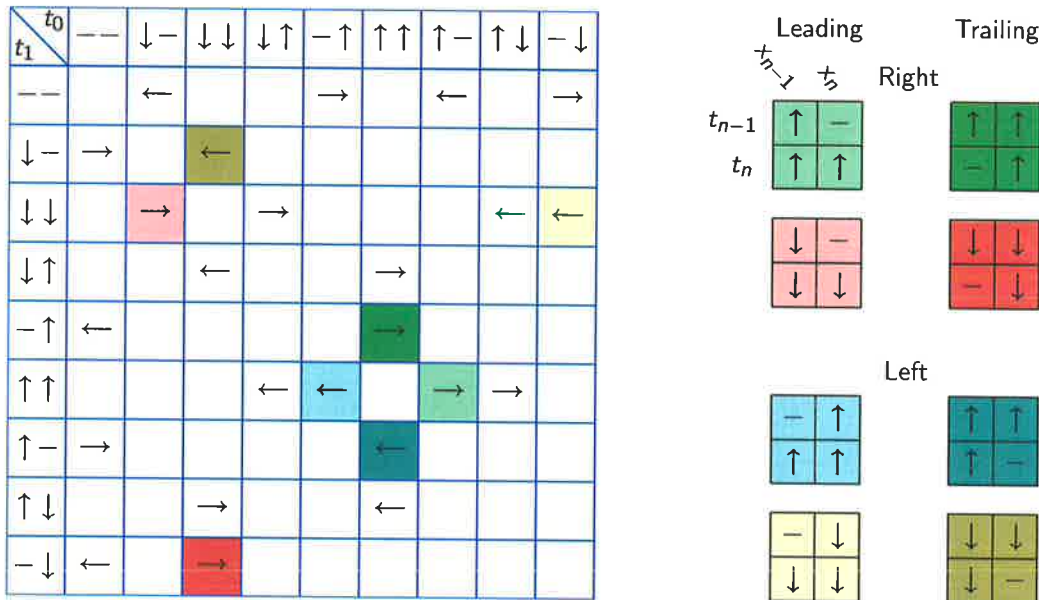


Figure 5.17. Look-up table of the templates. Only eight of the possible 81 templates indicate coherent motion, these are the coloured ones. The colours correspond to the conjugate templates.

Furthermore, these remaining eight templates have special properties. Whereas the previous schemes detect the actual edge, the template model tracks the leading and trailing parts of the edge. The two templates are generated for a single edge and referred to as *conjugate* templates. Therefore, there are only four conjugate pairs that form the directionally sensitive templates of coherent motion. These are shown to the right of the look-up table in Figure 5.17 with the conjugate templates shaded light and dark of the same colour. There is a template pair for each edge moving either left or right, on either a lighter or darker background. To visualise this, the templates shown in the response in Figure 5.18 are colour coded with the respective templates in the look-up table in Figure 5.17. The lighter shaded colour represents the leading edge and the darker shading the trailing edge.

The spacing between conjugate pairs of templates is dependent on the threshold in conjunction with the BPF cut-off frequencies as shown in Figure 5.19. Placing the threshold near the top of the peak of the response from the BPF generates templates that are close




Figure 5.18. Response of the template model to moving objects. The response to an object moving left and right. The symbols \triangleleft and \triangleright denote leftward and rightward motion respectively. The colours correspond to the templates in Figure 5.17.

together. But lowering the threshold, the distance between the templates increases due to the increasing width of the BPF response. This is relevant in the next chapter where the threshold has a range of implications.




Figure 5.19. Spacing of templates with the threshold placement. The two different values of threshold, θ_1 and θ_2 , generate conjugate templates with different separations.

Once these template pairs are known to exist, it makes tracking algorithms possible. This allows higher level of processing such as velocity estimation, image segmentation and time-to-impact to be found [Nguyen, 1996; Yakovleff, 1995; Beare, 1997; Moini, 1997; Yakovleff and Moini, 1996; Beare et al., 1995].

5.2 Biological Processing Schemes

Another major difference between the template model and the previous schemes is that the outputs are quantised, there is no sense of confidence levels. With the previous schemes, the higher the contrast of the edge or the closer it is to the tuned velocity for example, the larger the response. This gives some idea of how genuine the response is. The template model identifies edges without indicating confidence levels.

5.2.5 Gradient Schemes

Gradient schemes originated in computer science fields with the intention of being used to detect velocity and edges in artificial visual systems [Fennema and Thompson, 1979; Ullman, 1981; Poggio et al., 1985; Marr and Hildreth, 1980].

The acquisition of images, generally through a digitised video camera, is usually of high quality which makes gradient schemes viable. The gradient approach uses spatial and temporal derivatives and consequentially is quite susceptible to noise since the differential operation tends to accentuate noise.

These systems assume the global change in luminance is zero, or known. Then, by applying the chain rule to $dL/dt = 0$ we get [Horn and Schunck, 1981]

$$\frac{\partial L}{\partial x} \frac{dx}{dt} + \frac{\partial L}{\partial y} \frac{dy}{dt} + \frac{\partial L}{\partial t} = 0,$$

where $L(x, y, t)$ is the luminance at the point (x, y) at time t . Alternatively we can write this as $\nabla L \cdot v = -L_t$, where ∇L is the spatial gradient and $v = [v_x, v_y]$ is the velocity. Reducing this to one dimension gives the velocity as

$$v_x = -L_t/L_x$$

as also shown by Borst and Egelhaaf (1989).

However, Nakayama (1985) points out that the scheme does not appear applicable to biological image processing. This is because spatial gradients are poorly sensed by the visual system and there is difficulty when dealing with variations in global luminance.

Although the gradient scheme may not be utilised in biological systems, its mathematical origins make it suitable for processing computer generated scenes, which can then be used as a *benchmark* for later work. Since the images are computer generated, the effect of noise can be eliminated.

We can predict where the motion is by again employing the temporal and spatial gradients by

$$M(x, t) = -\frac{\partial L}{\partial t} \frac{\partial L}{\partial x} = -L_{xt}. \quad (5.12)$$

This provides a change of contrast independent measure for *motion intensity* of edges. The minus sign simply orientates positive M as being rightward in accordance to the previous models discussed. An examples of (5.12) is shown in Figure 5.20.




Figure 5.20. Gradient operations applied to a test image. The original image of a light spot moving back and forth with spatial and temporal gradients. The colour bar on the far right represents magnitude; red is rightward and blue is leftward motion. Image intensities have been normalised to full scale.

Although M is suitable for determining where the edges are, it cannot be applied to the template model as it tracks the leading and trailing parts of an edge. A double spatial and temporal derivative is used, but this gives no directional information. Thus, with use of the sign of M we can form

$$\hat{M}(x, t) = \text{sign}(M)|L_{xx}L_{tt}|, \quad (5.13)$$

which can be used for the template model. An example of this is shown in Figure 5.21. The abrupt change in direction that occurs at the tuning points of \hat{M} can be eliminated by multiplying (5.13) by $|M|$, but this biases the weighting of the ‘tracks’.

To calculate the derivative in practice, forward (or backward) differences are taken at the edges, and centred differences for the interior points. This provides an accurate representation of where the motion is, but due to the use of future quantities, it is not practically realisable. Thus, we expect a delay to be present in the output of the biological schemes, but this can be handled with the same method used when generating the correlation coefficients for aperiodic stochastic resonance.

5.3 Rudimentary Noise Analysis




Figure 5.21. Gradient operations for the template model. The same original image as in Figure 5.20 with double spatial and temporal gradients. Image intensities have been normalised.

5.3 Rudimentary Noise Analysis

In order to gain a quick, simple assessment of the relative robustness of the schemes, noise is added to the test signal. The simulation results in Figure 5.22 show the signals and responses when small and large amounts of noise are added. The bandpass filter is included simply to see how it fares as a filter – it is not a motion detector.

The clear result from Figure 5.22 is that removing the dynamical components from a system, namely the steady state DSLIMD, produces poor results. The Reichardt detector and DSLIMD show strong robustness against noise with the DSLIMD suppressing the noise slightly more.

5.4 Summary

We have described three biologically inspired motion detection schemes that satisfy the criteria outlined. They are the Reichardt detector, the DSLIMD and the template model.

The RMD uses an excitatory mechanism to compare adjacent channels, whereas the DSLIMD uses an inhibitive mechanism. Though there is not a great deal of difference for normal operation in good conditions, once noise is added to the signal the DSLIMD appears to suppress the noise more. This is dealt with quantitatively in the next chapter.

The template model, which has a mixed biological and engineering background, uses a look-up table to determine if there is any motion. The response, in the form of motion

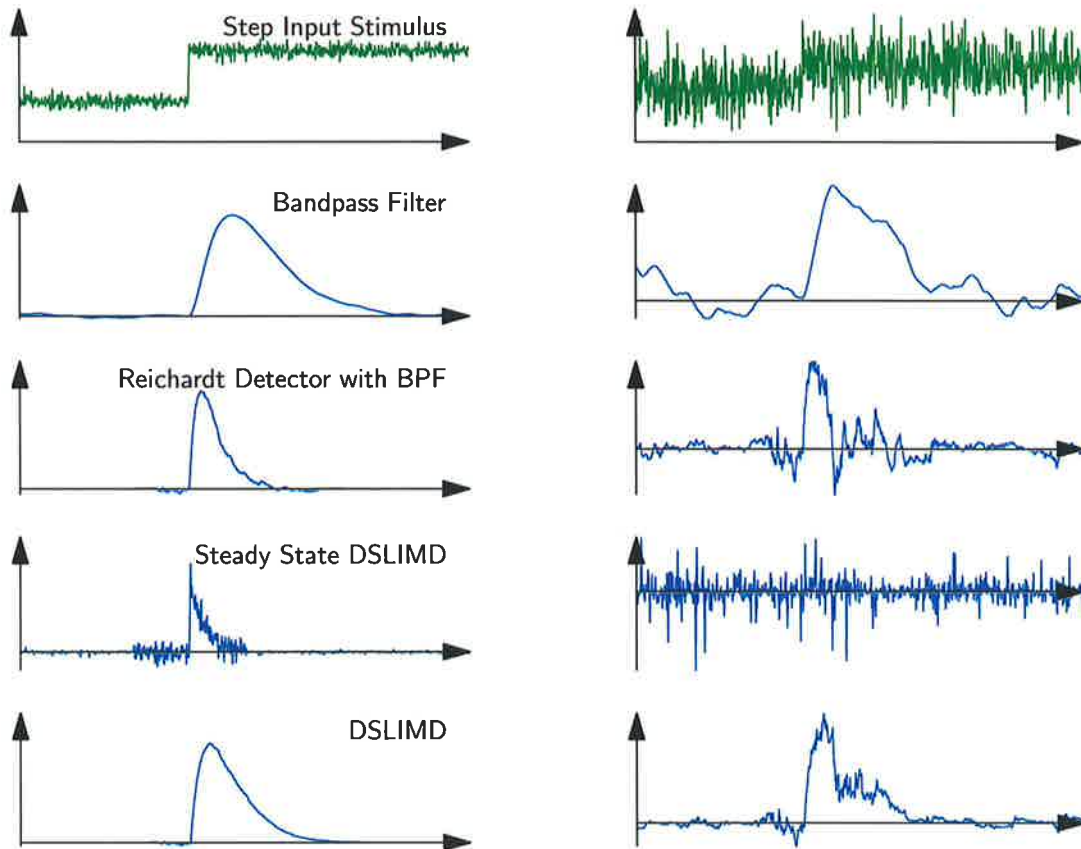


Figure 5.22. Simple noise comparison between several motion detection schemes. The left column has a small amount of added noise ($\sigma = 0.2$) whereas the right column has considerable added noise ($\sigma = 2.0$). The responses have been normalised to full scale.

vectors, only indicates if the presence or absence of motion. However, the threshold can be tuned to give some control over the probability of detection and false alarm rates. The threshold is also discussed in more detail in the next chapter.

We now have both sets of tools that are required. The concepts of SR given in Chapter 4, and the motion detection schemes just described. In the next chapter we will look at the possible ways that noise can be used to benefit edge detection schemes.

Stochastic Motion Processing

THIS chapter serves to merge the motion detection schemes described Chapter 5 with stochastic resonance techniques from Chapter 4. The three schemes, namely the Reichardt detector, DSLIMD and template model were each chosen due to their differing internal mechanisms and directional selectivity.

Under certain conditions we show that it is possible to benefit from noise using SR and SSR techniques in the template model. However, within the other two schemes, the addition of noise only degrades the response.

6.1 Schemes Investigated

The three motion detection schemes discussed in Chapter 5 are investigated in more depth. Each of the schemes has a particular reason for being studied in terms of its detection mechanism. The Reichardt detector utilises an excitatory mechanism in the form of a correlation unit, the DSLIMD employs an inhibitory mechanism that is implemented by the shunting neuron, and the template model uses a mixed biological-engineering approach.

We set the scene by describing a generalised framework in which to test and compare the schemes, and explore the possible configurations of the template model. The trends of individual motion detection schemes are given, and then their performance in the network configuration. We then consider the template model with suboptimal threshold settings. We finish by briefly looking at distributed thresholds and oversampling.

6.2 Architecture of Detectors

To keep the configuration of detectors as general as possible, we use the parallel network configuration shown in Figure 4.27 (on page 134). A visualisation in the context of motion detectors is shown in Figure 6.1. The standard motion detection system is formed by setting $N = 1$.

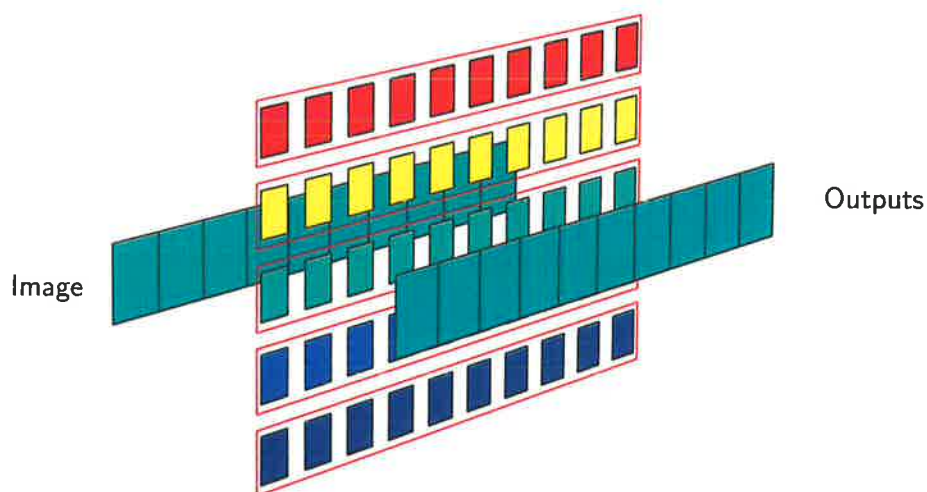


Figure 6.1. Network of widefield detectors. The rows in the centre represent a widefield array of detectors.

6.2.1 Template Model Configurations

The architecture of the template model becomes interesting when applied to the parallel network configuration. This is the configuration used for SSR, where the increased number of detectors along with the addition of noise, improve performance. This could be beneficial to the template model. The parallelism is due to some of the functional units being separable, namely the bandpass filter, thresholding and template formation units. The different configurations are shown in Figure 6.2, the dashed box indicating the parts that are placed in parallel in Figure 6.1. The outputs of the dashed regions are then appropriately summed or averaged.

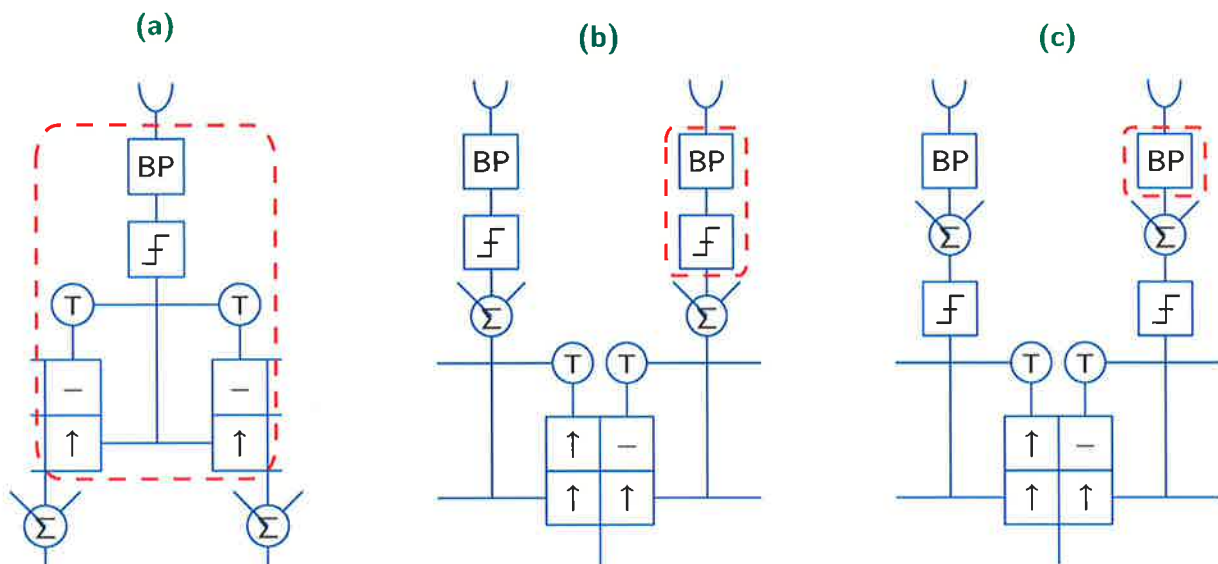


Figure 6.2. Network configurations for the template model. The dashed boxes are the parts to be parallelised. (a) The whole template model. (b) Bandpass and thresholding units. (c) Bandpass units. The outputs from the parallelised components are then summed or averaged, marked by Σ .

Figure 6.2a shows the whole template model as a parallel device, summing the motion vectors to produce the output. Motion vectors in opposite directions add destructively, resulting in no motion, hence with pure noise and enough elements in the network, there is very little or no net motion. In Figure 6.2b only the BPF and thresholding units are part of the parallel network. This also gives a discrete output, in the form of intensity changes (increase, decrease or none), which interact in similar fashion to the motion vectors. The output is then used to form the templates. In Figure 6.2c only the BPFs are contained in the network. The difference of this configuration is that

6.2 Architecture of Detectors

the outputs to be summed are continuous. Since BPFs are fairly generic units, similar techniques can be applied to other schemes.

Moving from the left to the right in Figure 6.2, the complexity decreases as does the computational time – compare a whole network of template detectors to simply filtering all the inputs. Another trend is that the more components placed in the network, the greater the information lost. Thresholding reduces the amount of information from a continuous range to one of three discrete values. Again the template formation reduces this information to abstract templates. This poses the question; why not simply add noise to each of the signals in the network, and then average before passing to the template model? The averaging effect of the noise (i.e. the central limit theorem) recovers the signal. However, this would form a linear system as far as the network is concerned, and any decrease in SNR at the input correspondingly decreases the SNR at the output.

To gauge qualitatively the performance of each configuration a simple noise analysis is shown in Figure 6.3. A reasonably large network size of $N = 32$ was used to see the effect of the network.

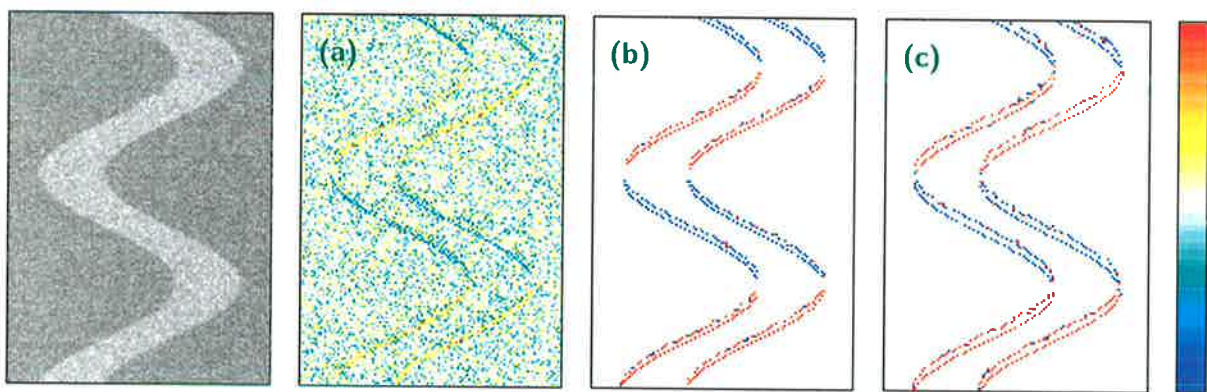


Figure 6.3. Simple noise analysis of template configurations from Figure 6.2. The original noisy image on the left ($\sigma = 1$) with the responses corresponding to the parallel configurations in Figure 6.2. In (a) there are N rows (in parallel) of motion vectors, whereas for (b) and (c) there is only a single row of outputs. The colour bar on the right indicates the direction and strength of the response. The top half is rightward, and the bottom half is leftward motion and the responses are normalised to full scale.

The response in Figure 6.3a appears comparatively worse due to the large amount of background noise. This is because every template that is due to noise alone is shown – note that most of the background noise is small, consisting only of a couple of templates. For the other two schemes, (b) and (c), most of the stray templates are effectively

cancelled due to the summing or averaging of signals before the template formation stage.

6.3 Effectiveness of Additive Noise

Considering the system, the most logical choice of metric is the correlation coefficient. Using this measure, we investigate the three schemes. Initially we only consider single widefield configurations with a simple stimulus. Parallel networks and special cases are studied thereafter.

Test Patterns

The test sequences are based on what could be expected to be captured from a standard PC. A sample rate of 16 frames/second ($h = 1/16$) and an image width of about 120 pixels were used. As mentioned earlier, the schemes process the horizontal lines of a 2-D image independently, which means we only need to consider a vector image to gauge performance. Two test sequences were used. The first is a spot (intensities 10 and 12.5) moving back and forth with slightly blurred edges, see Figure 6.4a. The second consists of a square wave grating (slightly blurred) of the same intensities moving in one direction for the first half of the time and the other direction for the last half of the time, see Figure 6.4b. From the test images, the benchmarks M from (5.12) and \hat{M} from (5.13) are generated.

The intensities mentioned and shown in Figure 6.4 are the luminance levels. For simulation and presentation purposes it is convenient to work with 8-bit grey scale images. This gives values of 0 to 255. By choosing reasonably low background luminance levels the dynamic range can be kept within the grey scale image range. When discussing the adaptive properties in Chapter 5 (on page 163), it is mentioned that the luminance levels span several orders of magnitude. However, photodetectors in electrical circuits [Harrison and Koch, 2000] and receptors in biological visual systems logarithmically compress the incoming (light) signals. This means that we can use the limited range of luminance levels within grey scale images without compromise.

It is common to use sinusoidal gratings to evaluate the performance of motion detection schemes. They can provide information such as the magnitude and phase response against frequency. However, this is more useful when using the detector to measure velocity [Dror et al., 2001; O'Carroll et al., 1997]. In this work we are only




Figure 6.4. Test patterns for widefield detectors. Light and dark intensities are 12.5 and 10 respectively. (a) A single spot moving back and forth. (b) Multiple spots moving left and then right.

interested in the detection of edges, thus a square wave grating is more appropriate. The edges are slightly blurred by convolving the test signal with a very short Hanning window, the length of which controls the shape of the pulse generated by M or \hat{M} . This is similar to the commonly used $G * I$ convolution for smoothing, G is a 2-d Gaussian and $I(x, y)$ is a 2-d intensity image. As the length of the Hanning window approaches zero the pulse produced from the differentiation becomes sharper. The blurring is only introduced along the spatial dimensions to mimic the real (optical) blurring that occurs in physical capture devices. If the blurring is too great, the square grating starts to resemble a sinusoidal grating, which is of little use.

6.3.1 Individual Widefield Characteristics

The following subsections show the correlations of the detector responses in their widefield configuration to the images shown in Figure 6.4. The cross-correlations are taken between M (or \hat{M}) and the response using correlation coefficients (i.e. relations (4.27) and (4.28)), taking into account the extra spatial dimension.

The accompanying plots for each of the detectors (Figures 6.5-6.7) also show images of the responses for several noise values to visually show the effect of the added noise, and to easily compare the schemes.

Each of the schemes was tuned for the small range of velocities present in the two test sequences.

Reichardt Detector

The correlations for the Reichardt detector with the bandpass units are shown in Figure 6.5. The plot of C_0 gives the perception that the detector is immune to noise, which is

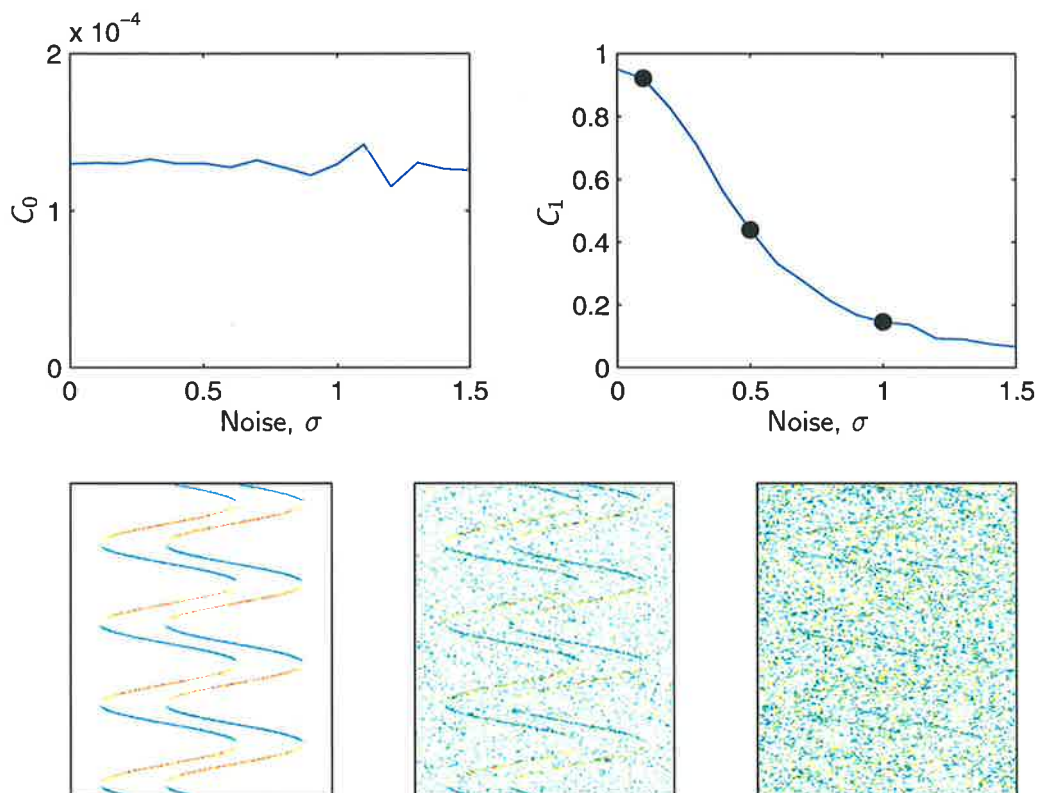


Figure 6.5. Correlation coefficient for Reichardt detector. The top two plots show the absolute and normalised correlation coefficients. The bottom three images are the detector outputs for the three noise intensities shown in the C_1 plot.

clearly not true as shown in the images along the bottom. This is partly due to the test images chosen – the benchmark is mostly zero. Thus any extra noise in the response in this region does not decrease the correlation C_0 . However, this is taken into account in C_1 , the normalised power norm. Here, the extra noise in the response increases the standard deviation in the denominator, which in turn decreases C_1 as expected.

6.3 Effectiveness of Additive Noise

Accordingly, C_1 should be used for any quantitative measure. In fact, it is also possible for C_1 to exhibit SR whereas C_0 does not [Stocks, 2001a]. Nevertheless, in this case C_1 also decreases with increased noise.

From the bottom three plots in Figure 6.5, the noise in the response is distributed reasonably independently throughout the image. This is due to the multiplication of the noisy inputs. The main noise inhibitor is the bandpass filter.

DSLIMD

The correlations for the DSLIMD are shown in Figure 6.6. The characteristics are similar to those of the Reichardt detector – the addition of noise only degrades the performance. The only noticeable departure is that C_0 also falls with increasing noise. The DSLIMD also sustains a higher value of C_1 for larger noise intensities.

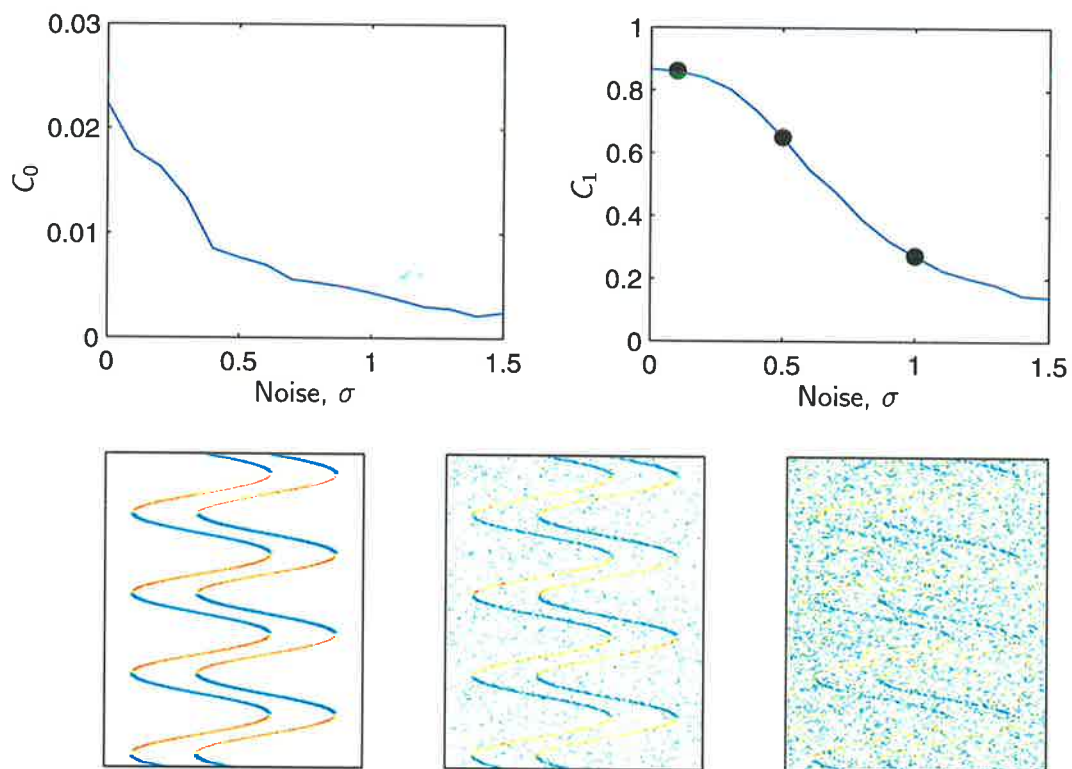


Figure 6.6. Correlation coefficient for DSLIMD. The top two plots show the absolute and normalised correlation coefficients. The bottom three images are the detector outputs for the three noise intensities shown in the C_1 plot.

Template Model

Since the outputs of the template model are spiky motion vectors, a short 2-D Hanning window is passed over the response. This is essentially the same technique used for finding the instantaneous firing rate from action potentials of neurons [Collins et al., 1995a]. It simply smooths the response to allow better matching to the motion benchmark, \hat{M} in this case.

To tune the template model, the sampling time implicitly determines the velocity. The threshold, which is set in conjunction with the BPF, is tuned to the contrast level of the images. For the correlations shown in Figure 6.7 the parameters have been adjusted so the double peaks for each of the edges in \hat{M} best match the distance between conjugate templates in the response. Alternatively, the threshold could have been fixed and the required properties of \hat{M} changed by varying the length of the Hanning window.

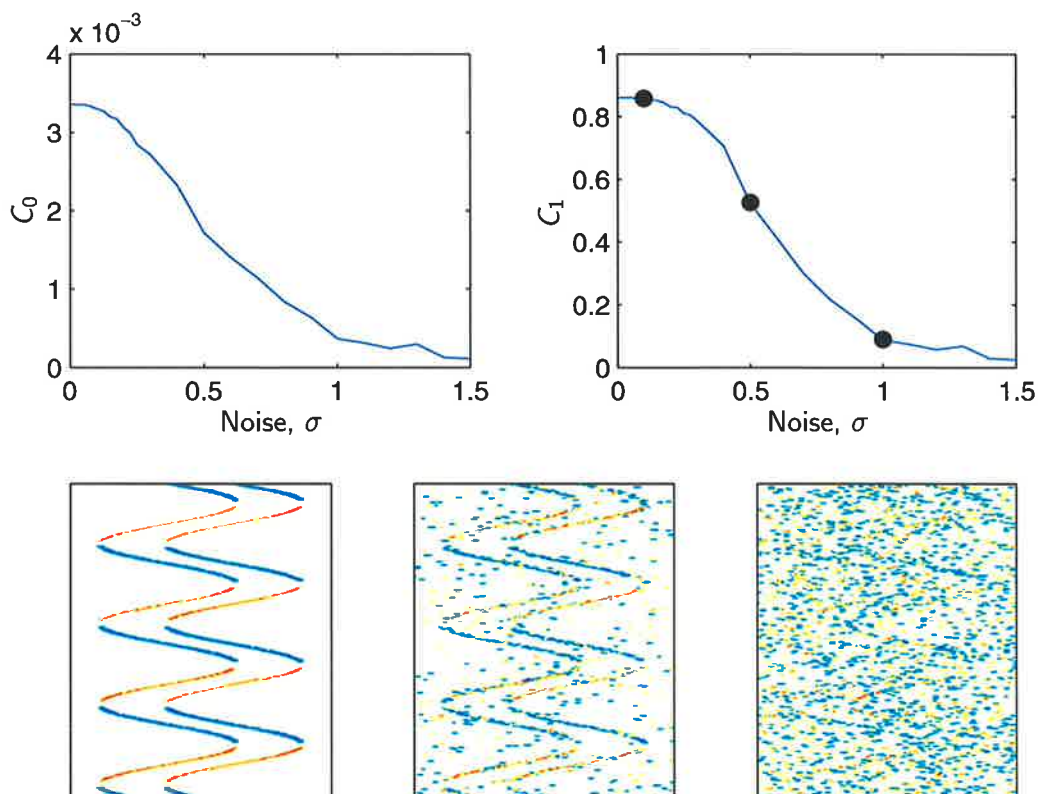


Figure 6.7. Correlation coefficient for Template Model. The top two plots show the absolute and normalised correlation coefficients. The bottom three images are the detector outputs for the three noise intensities shown in the C_1 plot.

6.3 Effectiveness of Additive Noise

Again, the correlations in Figure 6.7 show that additive noise does not assist the performance. For very small noise levels the correlations remain high, but then they fall rapidly.

In summary, additive noise only proves to degrade the performance of the detectors when correctly tuned to the input signals. Incorrectly tuned detectors will be considered in a subsequent section.

Comparing the noise robustness between the schemes, it is easy to see the DSLIMD is the most robust to noise with a noticeably higher correlation coefficient for larger noise intensities. The initial correlations when $\sigma = 0$ are highly dependent on the statistics of the test signal and the tuning of the detector. Thus, the relative trends from that point are more important than its absolute value.

The astute observer may note that the qualitative results shown in Figure 5.22 for $\sigma = 2.0$ appear better than the results in Figure 6.5 and 6.6 for $N = 1$ at $\sigma = 1.5$. The reason is due to different parameter settings. Figure 5.22 uses a much higher sampling rate ($h = 0.001$ s), and the generated pulse from the detector is spread over a longer time (larger bandwidth) for demonstration purposes. Here, we are only using $h = 1/16$ s and we want the pulses to die down quickly before the next edge arrives. The trade-off is that a shorter pulse width is more susceptible to noise since we need to allow more higher frequency components through the filters. These two factors dramatically diminish the robustness to noise, though the general trends remain the same.

Tuning Parameters

In the absence of noise, the value of the correlation coefficient mainly depends on the parameter values. These include the cut-off frequencies for the low pass and bandpass filters, spatial separation of receptors, neuron self decay and threshold values. Their purpose is to tune the detector for a particular edge velocity and control the shape of the response.

In the previous simulations, the parameters were tuned so that the output response matched the benchmark M or \hat{M} fairly accurately. The more highly tuned the detector, the larger the variance in C_1 between different types of stimulus when $\sigma = 0$. Thus, the tuning of the detector depends on the desired application.

6.3.2 Parallel Network of Detectors

Though adding noise degrades performance, by increasing the number of devices in the network the correlations are improved for the Reichardt and DSLIMD as shown in Figure 6.8.

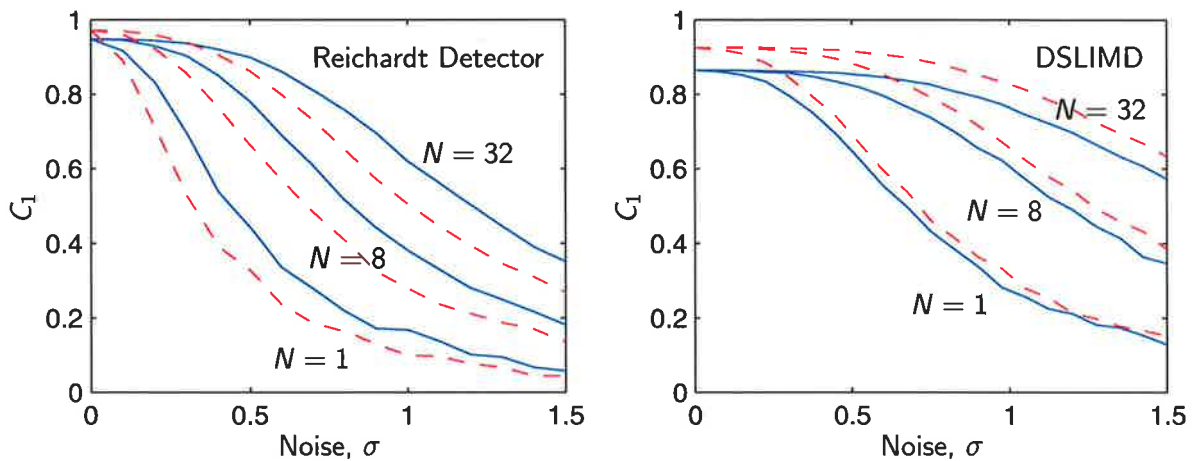


Figure 6.8. Correlations for a network of Reichardt and DSLIMDs. The solid and dashed lines are for a single spot and square wave grating respectively.

A possible explanation for the degradation of performance with noise may be that the system is near optimal. In the absence of noise, the correlations are very close to unity with either high or low contrast images. The addition of noise does not add any more degrees of freedom to the system in the same way as the threshold network [Stocks and Mannella, 2000], thus no extra information is gained.

Alternatively, the system as a whole may only be weakly nonlinear. We know that the interactions within the local detectors are nonlinear, but this does not characterise the overall detector behaviour. Detector linearity is considered next.

6.3.3 Detector Linearity

The linearity of the response may assist in explaining some of the trends in the correlations. This is investigated by plotting the peak response of the three detectors against the contrast, as shown in Figure 6.9. The peak response is taken from actual simulations of an edge moving at a constant velocity, as shown in Figure 6.4b. The height of the motion vectors for the template model are chosen arbitrarily at 0.4.




Figure 6.9. Linearity of the detectors against contrast. The peak response of individual detectors from simulations using the square wave grating. The two template model responses are for different threshold values. The response is given in the relevant units for each of the schemes.

The responses in Figure 6.9 from the Reichardt detector and DSLIMD do not contain any discontinuous nonlinearities in the same way that the template model, or other systems studied earlier do. Though they are nonlinear, over a small contrast range the response is approximately linear. Given that the Reichardt and DSLIMD are only weakly nonlinear, it should be of no surprise that the correlations monotonically decrease with the addition of noise. The step nonlinearity of the template model, which is on par with the LCC and LCD systems in Chapter 4, suggests it should be able to display SR and SSR under the appropriate conditions.

6.3.4 Suboptimal Parameter Settings

So far we have shown that neither SR or SSR are exhibited for a tuned detector. However, once the parameters are moved from their 'optimal' settings (with respect to matching the benchmark), SR and SSR characteristics can be observed.

Due to the linearity characteristics of the detectors shown in Figure 6.9, only the template model is able to readily display any type of SR. The particular parameter of interest is the threshold. Its adjustment allows the system to be either subthreshold to exhibit classical SR, or suprathreshold to exhibit SSR. We will consider the configurations shown in Figure 6.2a and 6.2b.

SR Behaviour

If the thresholds are set high, in the absence of noise, no motion vectors will be generated. This is typical practice when configuring a standard system for classical SR. Since the signal is subthreshold, noise is required to generate a response, thus in the absence of noise a zero correlation coefficient is given. When there is excess noise, the response is completely governed by it, causing a low correlation. Therefore, there must exist some value of nonzero noise intensity that maximises the correlations. The correlations for the template model of Figure 6.2a, which display SR are shown in Figure 6.10,




Figure 6.10. SR in the template model. The threshold ($\{\theta_i\} = 0.31$) is placed slightly higher than the peak response from the BPF of the square wave grating. The different line styles refer to correlating the response with M and \hat{M} respectively.

Recalling that the threshold level determines the separation between conjugate templates (see Figure 5.19), hence how well the response fits either M or \hat{M} . If the threshold is high, only the top of the peak of the signal (from the BPF) is crossed and the conjugate templates are close together, in which case M is a better benchmark. Conversely, if the threshold is low, the bottom of the peak is crossed where the width is greater and the templates are further apart, in which case \hat{M} is a better measure. This is the ‘normal’ operating condition. Therefore, since we are dealing with high values of threshold in Figure 6.10, correlations with M have the higher values.

6.3 Effectiveness of Additive Noise

SSR Behaviour

For a low threshold setting the system is suprathreshold, that is, there is always a strong response even in the absence of noise. As expected from the results of SRR and as shown in Figure 6.7, noise degrades the performance for a *single widefield* detector. For more detectors in parallel, the performance actually improves until an optimal noise intensity is reached, whence the performance continues to degrade upon the addition of further noise, as shown in Figure 6.11. These characteristics are very similar to Stocks' work with threshold networks [Stocks, 2000b; Stocks and Mannella, 2000].




Figure 6.11. SSR in the template model. The threshold ($\{\theta_i\} = 0.03$) is placed much lower than the peak BPF response of the square wave grating.

The reason for the improvement in correlations is similar to that of the threshold network. In the absence of noise, all of the detectors act identically, thus all giving the same outputs – motion to the left, right or none. When a little noise is added, the motion vectors deviate from their noiseless positions a little. Essentially, the outputs are able to explore more degrees of freedom. In terms of the correlations, the noise allows the motion vectors to spread and provide more information, hence allow better matching to the benchmark. Once excessive noise is added, the spread becomes large and templates start to appear due to the noise alone, thus decreasing the correlations. Increasing the number of detectors in the network improves the correlations, though there is a diminishing rate of return for the number of detectors added.

Since there is a low threshold value, the conjugate templates are distanced further apart. Thus, as mentioned earlier, for low noise levels \hat{M} is the better measure, which is consistent with Figure 6.11. However, once the noise increases, the difference in correlations between M and \hat{M} diminishes.

The results shown in Figures 6.10 and 6.11 use the network configuration with the full template model parallelised, i.e. as in Figure 6.2a. Figure 6.12 shows SR and SSR using the configuration of Figure 6.2b. The notable difference being that the noise induced maximums occur at higher noise values.

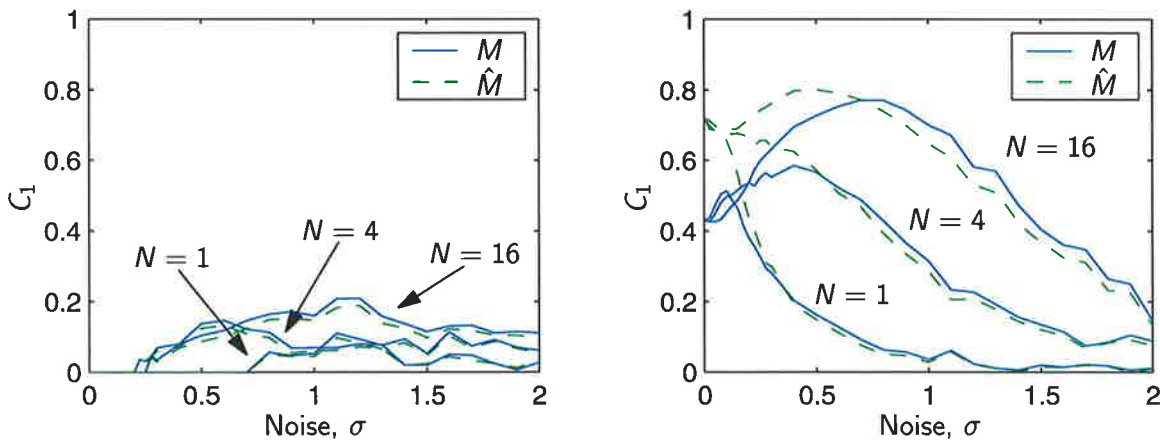


Figure 6.12. SR and SSR in the template model. The simulation from Figures 6.10 and 6.11 are repeated, but using the template model configuration of Figure 6.2b.

Changing Conditions

The SR and SSR behaviour of the template model could be useful if the lighting conditions (or temperature for mm-waves) change without the ability to modify the threshold. The important property is the input image contrast. This is equivalent to the energy of the image. The absolute value of the luminance is less important since the dc component is removed by the BPF.

Consider an image of average luminance and contrast for which the template model is tuned, particularly the BPF and threshold. Decreasing the luminance and contrast simulates low light conditions while increasing them simulates bright conditions. The correlations for the three conditions are shown in Figure 6.13.

When discussing SSR in Sec. 4.4.4, it was shown to be more robust because there are less fragile settings to consider. More importantly though, SSR is still observed if the signal amplitude is increased – as long as the noise is also increased proportionately.

6.3 Effectiveness of Additive Noise




Figure 6.13. Correlations for the template model under changing conditions. The solid and dashed curves represent correlations with M and \hat{M} respectively, with $N = 32$ and a threshold of 0.1. The luminance for low and high light conditions vary by $0.3\times$ and $4\times$ respectively.

This remains true for the template system. Under the bright conditions when the contrast increases (i.e. bigger signal), more noise than normal must be added in order to achieve affective SSR gain. For the case shown in Figure 6.13 where the intensity has increased by 4 times, the amount of noise must also be increased by 4 times compared to the noise that the other cases are receive.

When conditions change and the system becomes suboptimal, SR is able to assist for small signal amplification, while SSR for large signal enhancement. Thus, we can use noise induced assistance when the system becomes suboptimal. The observation that stochastic resonance only benefits systems that are suboptimal agrees with our earlier experience in Chapter 4. A more elegant and efficient solution to the problem of changing contrast levels would be to employ some type of automatic threshold adjustment. Besides giving better results, it would probably be easier to implement than adding noise of the correct intensity.

6.3.5 Distributed Threshold Settings

As pointed out in Stocks' SSR work [Stocks, 2000b] with threshold networks, having the same threshold settings is an inefficient use of N devices – it is suboptimal. An

optimal setting is obtained by choosing the thresholds so that each possible outcome is equally likely to maximise the entropy of the output. By distributing the thresholds of the template model in the network, an optimal, or near optimal configuration may be found.

The response when the thresholds are uniformly placed with respect to the height of the pulse generated by the BPF (in the absence of noise) is shown in Figure 6.14. For an optimal setting, C_1 would be unity in the absence of noise, however there are many other parameters so we expect C_1 to be slightly less than unity.




Figure 6.14. Distributed threshold settings in the template model. The thresholds were distributed evenly between the central 90% (approx.) of the height of the BPF output.

Note there is no small rise as seen for the 'standard' curve in Figure 6.13, even though inputs are the same. This shows that the distributed threshold gives a more optimal performance.

6.3.6 Oversampling

At the start of Chapter 4 we discussed dithering as an example of SR and how the process of oversampling improved the accuracy of an ADC. Given that the responses shown in Figure 5.22 appear more immune to noise than those shown in Sec. 6.3.1 (due to their higher sampling rate), oversampling may be able to improve the correlations. This is indeed the case, as shown by the template responses in Figure 6.15 with 8 and 16 times oversampling. This only improves the performance, it is not intended to display

6.4 Limitations of SR in Motion Detection




Figure 6.15. Oversampling in the template detector. The labels refer to the number of times oversampled, that is the number of internal template formations taken during each sampling period. Simulations used $N = 8$ and $\{\theta_i\} = 0.1$.

SR. The implementation used is very basic and further improvement may be possible with a more efficient algorithm.

6.4 Limitations of SR in Motion Detection

In the previous section, we show that simply adding noise to images to be processed by a nonlinear system does not always help. To determine if SR is going to take place, we need to investigate the detector linearity. This is not the interaction between adjacent channels (which is always very nonlinear), but how the output response varies with edge contrast. This is shown for the three detectors considered in Figure 6.9. Though we require a nonlinear response to exhibit SR, when used as a velocity detector a linear response is often desired. This means the velocity is directly proportional to the response. Thus, we have conflicting requirements, which means compromises are inevitable.

In accordance with previous observations of stochastic resonance, if the system is operating optimally, no improvement can be gained by adding noise to the system. By distributing the thresholds in the template model, an almost optimal response is attained, where we demonstrate that performance declines when noise is added to the system.

Also, the noise intensities where the SR and SSR peaks occur are not very high. In most cases they occur well before $\sigma = 1$, which is the value of noise shown in Figure 6.3. Even at this noise intensity, it is easy to see with the naked eye where the edges are located. Thus, the assistance that can be offered by stochastic resonance is limited to low noise intensities.

6.5 Summary

We give a comparison of the performance in noisy environments for three motion detectors. A comprehensive investigation of the robustness to noisy signals for the two newer detectors, the DSLIMD and template model, has not previously been undertaken. We compare their performance to the widely used and well known Reichardt detector. The results show that the DSLIMD performs the best – though at the expense of being a little more complex, but still only using local information.

Secondly, we applied the techniques of SR to motion detection. We found that with certain conditions, and utilising a network structure, that SR and SSR are exhibited. This requires a strong nonlinearity of the detector response as demonstrated by the template model, and more than one detector (for SSR) in the network. In low contrast conditions, SR assists by increasing the dynamic range of the system, and in high contrast, SSR provides enhancement of signals above threshold.

Chapter 7

Conclusions and Directions

IN this chapter, we give an extended summary of the work undertaken in this thesis. We also recommend future work and directions that could be undertaken and finish with some conclusive remarks.

7.1 Summary

This thesis addresses the problem of motion detection in noisy environments. We consider two fields of noise induced enhancement, namely, Brownian ratchets and stochastic resonance. The principle in common with these two phenomena is that noise plays a constructive role.

Chapter 2 reviews the origins of Brownian motion and some of the flawed attempts that have been made in the past to rectify this motion to generate perpetual motion. We show how one such device, the ratchet and pawl machine, could be linearised and used to transport Brownian particles. In order for the linearised version of the Brownian ratchet to work, it must be kept in nonequilibrium. That is, energy must be supplied to the device. This type of device has been shown to exist in biology (e.g. molecular motors), and can be artificially constructed to work with colloidal solutions.

In Chapter 3, we expound a mathematical abstraction of the Brownian ratchet, referred to as Parrondo's games. Originally devised by Prof. Juan Parrondo, they consist of two simple gambling games that have the unique property of losing when played individually, but have a winning expectation when the games are randomly (or periodically) switched. There are several explanations on how the games function, but in essence it is due to their similarity to the flashing Brownian ratchet. That is, the random mixing leads to the counter-intuitive result. If there is insufficient mixing (i.e. noise), or if there is too much, the desired result does not occur – we need an optimal amount of noise to gain the best result. This is exactly the principle of stochastic resonance, which is studied in Chapter 4. Even though the games have no direct application to motion detection, the principle that noise or randomness can become useful, in the presence of a nonlinearity or asymmetry, is important.

Chapter 4 begins the second part of the thesis dealing with stochastic resonance. We start by describing the basics of SR in an example black box system, where we show the signature response of SR with the signal-to-noise ratio (SNR). We also consider the signal-to-threshold distance, where we show that the performance of a system can be more efficiently improved by lowering the threshold to eliminate the signal-to-threshold distance. In this case noise no longer offers any assistance to the system. This is supported by a number of papers. This means SR is appropriate for systems that are constrained, as internal parameters cannot be changed, and doing so may compromise the system. This is reasonably common for many real-world systems.

We describe a number of nonlinear static and dynamic systems that are commonly used in SR. This includes excitable systems such as thresholds and neurons, and bistable systems such as the double well potential and Schmitt trigger. White and coloured noise algorithms are also given.

We finish off Chapter 4 by describing the measures that are used to quantify SR. When the SR phenomenon was first described by Benzi in the early 1980s, SR was characterised by the SNR. Though this is not a real frequency resonance, the SNR is accepted as a suitable measure of the coherence between the input and output. It is also used widely by engineers and is well understood. A real resonance is described in terms of residence times, which quantify the synchronisation of the switching mechanism due to the forcing signal. However, the SNR and residence times are only appropriate for periodic signals. To deal with aperiodic signals, cross correlation coefficients are employed. This can be further generalised by measuring the information content which makes no assumptions *a priori* about the type of input signal. Nevertheless, all of the aforementioned measures effectively measure the resolution of the signal. Though this may be appropriate in some cases (suprathreshold SR for example), in most cases the effect of noise on the bandwidth must be taken into account. The bandwidth may not be independent of the noise intensity. This is because the useful quantity is the information transfer rate of the system, which is the resolution (i.e. information content) times the bandwidth. Both of these factors are taken into account in the discussion of channel capacity, which is based on Shannon's well known channel capacity formula.

Chapter 5 simply gives the background of several motion detection schemes in which SR will be investigated in Chapter 6. We utilise edge detection algorithms that have been biologically inspired by the study of the insect visual systems and the neurochemistry of neurons. The three schemes studied are the Reichardt detector, directionally sensitive local inhibitory motion detector (DSLIMD) and the Horridge template model. Each of the schemes function by different mechanisms, which influences their ability to suppress noise and their applicability to adapt SR. We detail the mechanisms used in each of the schemes and give a brief comparison to a noisy signal.

In Chapter 6 we apply some of the techniques of SR to motion detection, and in doing so give a relative comparison between the three schemes. Firstly, we consider each of the schemes and plot the correlation coefficients between the responses to noise signals, and a benchmark generated from a clean signal using a gradient scheme. Though the Reichardt detector has been known for a long time (1950s), the template model and

7.2 Future Directions

DSLIMD are relatively new schemes (early and late 1990s respectively). The quantitative analysis shows the relative performance between each of the schemes for increasing noise intensity. For low noise values, all of the schemes maintain reasonably high correlations. For larger amounts of noise the DSLIMD has higher coefficients, thus is the most robust, as also shown by the brief noise analysis in Chapter 5. This is due to the inhibitory mechanism this detector employs.

Secondly, we merge SR and motion detection. Since all the detectors can detect motion in the absence of noise (i.e. suprathreshold), we utilise the structure of SSR as used in the threshold network. We show the intuitive result that the performance improves as the number of widefield detectors in the network are increased. However, in the Reichardt detector and DSLIMD there is no improvement as the noise intensity increases. This is possibly due to the detectors being optimally configured (they have very high correlations with no noise), and they are only weakly nonlinear. Though the local detector interactions may be highly nonlinear, the detector as a whole appears only weakly nonlinear.

To clarify this the template model contains a threshold parameter, so it possesses a strong nonlinearity. That is, a small variation in the threshold value (or alternatively signal strength) produces either none or all of the desired templates. In contrast to the other schemes, a small variation in signal strength produces approximately a proportional change in the response. We show that for the appropriate conditions the template model can display SR or SSR when used in the network configuration. Thus, if the conditions change suddenly, or there is some other trade-off in the system, stochastic resonance can help.

7.2 Future Directions

The directions for future work that could continue from this thesis are on several fronts; the work relating to Parrondo's games, SR, and motion detection.

Parrondo's Games

In regards to Parrondo's games, we concluded that they were an instructive example of noise-induced cooperative behaviour, though there was no apparent link to motion detection. However, Parrondo's games are nevertheless an interesting body of work and deserve future development in their own right. In terms of applications, other

researchers are already investigating directions in biogenesis [Davies, 2001] and the stock market [Boman et al., 2001] for example. In terms of the mathematics of the games, there are opportunities for the development of new games with different rules and state dependencies. This may allow analogies to be found with other types of physical processes. Also, another direction would be the use of genetic algorithms to evolve the rules of the games in order to maximise (i) parameter space and (ii) rate of return.

Stochastic Resonance

The problem of how to detect weak signals in a noisy environment is difficult, and an effective algorithm would be useful and profitable. The work by Dawe using ISHTAR (page 104) to detect stealthy targets appears to be such a solution. However, the signals and principles remain unclear. More work is necessary to identify the exact mechanisms at work.

Though Dawe's ISHTAR is a sonar application, it seems plausible that this algorithm could be generalised to other media and applications. A possible approach would be to clearly identify the types of signals used, and to accurately model the medium in which the signals traverse or radiate. This would allow the received signal statistics to be determined, and hence conclude whether SR could take place. However, modelling all but the simplest media and boundary conditions is a challenging problem, which ultimately affects the prediction of SR.

In Chapter 4, we describe the application of Shannon's channel capacity to the filtered output of a single threshold system. We also describe Stocks' work using a network of thresholds, which can be quantified in terms of the information content since the output is not filtered. However, the situation changes when the output is filtered – we effectively increase the number of output states and the AMI is no longer proportional to the channel capacity. It would be useful to know how this filtering (and its specifications) modifies the channel capacity for a multithreshold system. Initially simulations or experiments may prove as a useful tool.

Motion Detection

For the motion detection schemes, the next logical step is to implement them in (analog) VLSI and interface this with the millimetre wave antenna array, which is currently

7.3 Closing Comments

in the design and testing phase. The template model has been successfully implemented in the visible spectrum. Though the addition of noise does not always help, the improved robustness to noise will be advantageous.

We have investigated three key motion detection schemes, which is a subset of all the types of motion detectors. There are new schemes that have only recently been devised (e.g. [Harris et al., 1999]), that may be also appropriate for investigation.

A worthwhile approach would be to determine the linearity of each type of detector (in terms of peak response to an edge) against the contrast of the input. Although the internal interactions of the local detector mechanism may contain highly nonlinear operations, the peak response against contrast may be only weakly nonlinear. This will help to readily determine if SR is able to take place.

The investigation of the detector in Chapter 6 only considers a small range of tuning parameters. A formal analysis of the performance with respect to variations of tuning parameters was beyond the scope of this thesis. For example, the effects of tuning bandpass filter parameters for different ranges of velocities would most likely modify the robustness in some way. We suspect, the smaller the velocity range, the more noise is attenuated, at the expense of being a more general purpose detector. Other factors, such as noise distributions and image statistics (i.e. the types of images), also warrant investigation.

7.3 Closing Comments

Biological systems process noisy signals very effectively – neurons have an output SNR in the order of 0 dB. Neurons have also been shown to utilise SR, which they have done after millions of years of evolution. We have employed biological motion detection schemes based on neurons and visual systems, and applied these to SR. Under certain conditions we show that noise can indeed be of assistance, though there are limitations to its application. It is hoped that the work presented in this thesis enhances the understanding of SR and motion detection of noisy signals and can be used as a stepping stone for future work.

Parrondo's Games Analysis

A.1 Constraints for the Generalised modulo M Game

The parameters of Parrondo's games can be chosen such that individually each game is losing but a randomisation between the games is winning. In this section we present the mathematical analysis for the generalised modulo M games that establishes this [Harmer et al., 2000b; Harmer et al., 2001]. We do this by establishing conditions for recurrence of the corresponding discrete-time Markov chains.

The analysis of game A is elementary and can be found in many textbooks, for example see Karlin and Taylor (1975), but we present it here in the interest of motivating our analysis of game B .

We win a single round of game A with probability p and lose with probability $1 - p$. Assuming that they bet one unit on each round of the game, we wish to calculate the probability f_j that the player's capital ever reaches zero given that they start with a capital of j units. It is a consequence of Markov chain theory [Karlin and Taylor, 1975, p. 93] that either

1. $f_j = 1$ for all $j \geq 0$, in which case the game is either fair or losing, or
2. $f_j < 1$ for all $j > 0$, in which case there is some probability that our capital will grow indefinitely and so the game is winning.

For $j \geq 1$, let $f_j^{(n)}$ be the probability that our capital reaches zero within the first n games, given that it starts at j . It is easy to see that $f_0^{(n)} = 1$ for all n . For each j , the sequence $\{f_j^{(n)}\}$ is increasing and thus must have a limit which is f_j , as defined above. By conditioning on what happens at the first time point, we derive the equation

$$f_j^{(n+1)} = pf_{j+1}^{(n)} + qf_{j-1}^{(n)}. \quad (\text{A.1})$$

A.1 Constraints for the Generalised modulo M Game

It follows that f_j is the minimal nonnegative solution to the equation

$$f_j = pf_{j+1} + qf_{j-1}, \quad (\text{A.2})$$

subject to the boundary condition

$$f_0 = 1. \quad (\text{A.3})$$

We can see how this makes sense by considering the DTMC in Figure A.1. In state j , the probability of going back to zero is f_j . If we move up a state, it becomes pf_{j+1} , that is, the probability of moving up times the probability of going back to zero from that state. Similarly, when moving down a state it is $(1-p)f_{j-1}$.




Figure A.1. Portion of a DTMC. Finding a relation for the probability f_j of reaching zero given we start in state j .

The general solution to (A.2) has the form

$$f_j = A \left(\frac{1-p}{p} \right)^j + B, \quad (\text{A.4})$$

where A and B are constants. Invoking the boundary condition (A.3) gives $B = 1 - A$ and (A.4) becomes

$$f_j = A \left[\left(\frac{1-p}{p} \right)^j - 1 \right] + 1. \quad (\text{A.5})$$

If $(1-p)/p \geq 1$, the minimal nonnegative solution to (A.2) occurs when $A = 0$ and so $f_j = 1$ for all $j \geq 0$. If $(1-p)/p < 1$, the minimal nonnegative solution to (A.2) occurs when $A = 1$ and so

$$f_j = \left(\frac{1-p}{p} \right)^j$$

for all $j > 0$. Thus we can write

$$f_j = \min \left(1, \left(\frac{1-p}{p} \right)^j \right) \quad (\text{A.6})$$

and we observe that the game is winning if

$$\frac{1-p}{p} < 1, \quad (\text{A.7})$$

that is if $p > 1/2$. However, this only provides half of the information as we cannot distinguish between losing and fair games according to the conditions previously stated. Namely, if $f_j = 1$ the game is either losing or fair. The solution is to switch the players around, then by symmetry, we deduce the game is losing if

$$\frac{1-p}{p} > 1, \quad (\text{A.8})$$

and is fair if

$$\frac{1-p}{p} = 1. \quad (\text{A.9})$$

This result, of course, accords with our intuition.

Now let us turn to game *B*. Here the probability that the player wins a single round depends on the value of their current capital. If the capital is a multiple of M , the probability of winning is p_1 , whereas if the current capital is not a multiple of M , the probability of winning is p_2 . The corresponding losing probabilities are $1 - p_1$ and $1 - p_2$ respectively. Let g_j be the probability that our capital ever reaches zero given that we start with a capital of j units. As with game *A*, Markov chain theory tells us that either

1. $g_j = 1$ for all $j \geq 0$, in which case the game is either fair or losing, or
2. $g_j < 1$ for all $j > 0$, in which case there is some probability that the player's capital will grow indefinitely and so the game is winning.

Following the derivation of game *A*, for $i \geq 1$ and $j \in \{1, \dots, M-1\}$, the set of numbers $\{g_k\}$ satisfies the equations

$$g_{Mi} = p_1 g_{Mi+1} + (1 - p_1) g_{Mi-1} \quad (\text{A.10})$$

and

$$g_{Mi+j} = p_2 g_{Mi+j+1} + (1 - p_2) g_{Mi+j-1} \quad (\text{A.11})$$

subject to the boundary condition

$$g_0 = 1. \quad (\text{A.12})$$

For $j \in \{1, \dots, M-1\}$, the general solution to equation (A.11) is

$$g_{Mi+j} = C_i \left(\frac{1-p_2}{p_2} \right)^j + D_i, \quad (\text{A.13})$$

A.1 Constraints for the Generalised modulo M Game

with

$$C_i = \frac{g_{Mi} - g_{M(i+1)}}{1 - \left(\frac{1-p_2}{p_2}\right)^M} \quad (\text{A.14})$$

and

$$D_i = \frac{g_{M(i+1)} - g_{Mi} \left(\frac{1-p_2}{p_2}\right)^M}{1 - \left(\frac{1-p_2}{p_2}\right)^M}. \quad (\text{A.15})$$

Substituting this into equation (A.10), we derive the equation

$$\begin{aligned} & \left[1 - \left(\frac{1-p_2}{p_2}\right)^M\right] g_{Mi} \\ &= p_1 \left\{ g_{M(i+1)} \left[1 - \left(\frac{1-p_2}{p_2}\right)\right] + g_{Mi} \left[\left(\frac{1-p_2}{p_2}\right) - \left(\frac{1-p_2}{p_2}\right)^M \right] \right\} \\ &+ (1-p_1) \left\{ g_{Mi} \left[1 - \left(\frac{1-p_2}{p_2}\right)^{M-1}\right] + g_{M(i-1)} \left[\left(\frac{1-p_2}{p_2}\right)^{M-1} - \left(\frac{1-p_2}{p_2}\right)^M \right] \right\} \end{aligned}$$

for $i \geq 1$. Grouping the $g_{M(i-1)}$, g_{Mi} and $g_{M(i+1)}$ terms and after some tedious manipulation, for $i \geq 1$, this reduces to

$$\begin{aligned} \left[p_1 p_2^{M-1} + (1-p_1)(1-p_2)^{M-1} \right] g_{Mi} &= \left[p_1 p_2^{M-1} \right] g_{M(i+1)} \\ &+ \left[(1-p_1)(1-p_2)^{M-1} \right] g_{M(i-1)} \end{aligned}$$

For $i \geq 0$, the general solution to this is

$$g_{Mi} = C_i \left(\frac{(1-p_1)(1-p_2)^{M-1}}{p_1 p_2^{M-1}} \right)^i + D_i.$$

Use of the boundary condition $g_0 = 1$ yields

$$g_{Mi} = C_i \left[\left(\frac{(1-p_1)(1-p_2)^{M-1}}{p_1 p_2^{M-1}} \right)^i - 1 \right] + 1,$$

and we similarly deduce that the minimal nonnegative solution is

$$g_{Mi} = \min \left(1, \left(\frac{(1-p_1)(1-p_2)^{M-1}}{p_1 p_2^{M-1}} \right)^i \right).$$

As for game A, we deduce that game B is winning if

$$\frac{(1-p_1)(1-p_2)^{M-1}}{p_1 p_2^{M-1}} < 1, \quad (\text{A.16})$$

losing if

$$\frac{(1-p_1)(1-p_2)^{M-1}}{p_1 p_2^{M-1}} > 1 \quad (\text{A.17})$$

and fair if

$$\frac{(1-p_1)(1-p_2)^{M-1}}{p_1 p_2^{M-1}} = 1. \quad (\text{A.18})$$

Now consider the randomised game, where the probabilities q_1 and q_2 are defined in (3.18a) (3.18b) respectively. We observe that this is identical to game B except that the probabilities have changed. It follows from (A.16-A.18) that the randomised game is winning, losing and fair if

$$\frac{(1-q_1)(1-q_2)^{M-1}}{q_1 q_2^{M-1}} < 1, \quad (\text{A.19})$$

losing if

$$\frac{(1-q_1)(1-q_2)^{M-1}}{q_1 q_2^{M-1}} > 1 \quad (\text{A.20})$$

and fair if

$$\frac{(1-q_1)(1-q_2)^{M-1}}{q_1 q_2^{M-1}} = 1. \quad (\text{A.21})$$

The existence of the paradox of Parrondo's games will be established if we can find parameters p, p_1, p_2, γ and M for which

$$\frac{1-p}{p} > 1, \quad \frac{(1-p_1)(1-p_2)^{M-1}}{p_1 p_2^{M-1}} > 1 \quad \text{and} \quad \frac{(1-q_1)(1-q_2)^{M-1}}{q_1 q_2^{M-1}} < 1$$

are satisfied, as stated in equations (3.31).

A.2 Calculating the Equilibrium Distribution

This section outlines a method for determining the equilibrium distributions, which was originally due to Mihoc and Fréchet, and expanded for the periodic case by Pyke (2001). It states the equilibrium distribution is proportional to the diagonal cofactors of the transition matrix \mathbb{P} ,

$$\boldsymbol{\pi} = \frac{1}{D} \text{diag}(\text{cofac}(\mathbb{I} - \mathbb{P})), \quad (\text{A.22})$$

where D is the normalisation constant.

Strictly speaking the cofactors of an $n \times n$ matrix with entries a_{jk} are

$$C_{jk} = (-1)^{j+k} M_{jk},$$

A.2 Calculating the Equilibrium Distribution

where M_{jk} is the minor of a_{jk} , which is equivalent to the determinant of order $n - 1$ [Kreyszig, 1953, p. 373]. This is easiest illustrated by way of example of a 3×3 matrix.

Given the matrix

$$\mathbb{A} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix},$$

selected minors are given by

$$M_{11} = \begin{vmatrix} a_{22} & a_{23} \\ a_{32} & a_{33} \end{vmatrix}, \quad M_{12} = \begin{vmatrix} a_{21} & a_{23} \\ a_{31} & a_{33} \end{vmatrix}, \quad M_{13} = \begin{vmatrix} a_{21} & a_{22} \\ a_{31} & a_{32} \end{vmatrix}.$$

The cofactors are then

$$\begin{aligned} C_{11} &= +M_{11} & C_{12} &= -M_{12} & C_{13} &= +M_{13} \\ C_{21} &= -M_{21} & C_{22} &= +M_{22} & C_{23} &= -M_{23} \\ C_{31} &= +M_{31} & C_{32} &= -M_{32} & C_{33} &= +M_{33} \end{aligned},$$

where the signs form a checkerboard pattern

$$\begin{array}{ccc} + & - & + \\ - & + & - \\ + & - & + \end{array}.$$

The diagonal is simply the main diagonal, hence for \mathbb{A}

$$\text{diag}(\mathbb{A}) = \begin{bmatrix} a_{11} \\ a_{22} \\ a_{33} \end{bmatrix}.$$

Now, finding the cofactors of $\mathbb{I} - \mathbb{P}$ using \mathbb{P} from (3.17) with $M = 3$ gives

$$\mathbb{I} - \mathbb{P} = \begin{bmatrix} 1 & p_2 - 1 & -p_2 \\ -p_1 & 1 & p_2 - 1 \\ p_1 - 1 & -p_2 & 1 \end{bmatrix}, \quad (\text{A.23})$$

and the minors of interest are

$$\begin{aligned} M_{11} &= \begin{vmatrix} 1 & p_2 - 1 \\ -p_2 & 1 \end{vmatrix} = 1 - p_2 + p_2^2, \\ M_{22} &= \begin{vmatrix} 1 & -p_2 \\ p_1 - 1 & 1 \end{vmatrix} = 1 - p_2 + p_1 p_2, \\ M_{33} &= \begin{vmatrix} 1 & p_2 - 1 \\ -p_1 & 1 \end{vmatrix} = 1 - p_1 + p_1 p_2. \end{aligned}$$

These are equivalent to the cofactors since the sign of all the minors along the main diagonal are positive. Thus the stationary distribution from (A.22) is

$$\pi^B = \frac{1}{3 - p_1 - 2p_2 + 2p_1p_2 + p_2^2} \begin{bmatrix} 1 - p_2 + p_2^2 \\ 1 - p_2 + p_1p_2 \\ 1 - p_1 + p_1p_2 \end{bmatrix}, \quad (\text{A.24})$$

where the normalisation constant is the sum of the diagonal cofactors.

A.3 Rate of Winning

Denoting the capital at time n by X_n , the rate of winning as a function of the number of games played intuitively given by $E[X_{n+1} - X_n]$. This is always equivalent to $E[X_{n+1}] - E[X_n]$ no matter the dependencies of the arguments of E [Yates and Goodman, 1999, p. 232]. This type of relation for the variance is only true if the random variables are mutually independent. Thus, the rate of winning is

$$\begin{aligned} r(n) &= E[X_{n+1}] - E[X_n] \\ &= \sum_{j=-\infty}^{\infty} j[\pi_j(n+1) - \pi_j(n)], \end{aligned} \quad (\text{A.25})$$

where π_j is the stationary probability of state j at time n .

We can write the global balance equation as

$$\pi_j(n+1) = p_{j-1,j}\pi_{j-1}(n) + p_{j+1,j}\pi_{j+1}(n),$$

where $p_{j,k}$ is the transition probability of moving from the j th to the k th state. Using the previous two equations gives

$$r(n) = \sum_{j=-\infty}^{\infty} j[p_{j-1,j}\pi_{j-1}(n) + p_{j+1,j}\pi_{j+1}(n) - \pi_j(n)]. \quad (\text{A.26})$$

Then by expanding the $j = k$ and $k \pm 1$ indices of the summation,

$$\begin{aligned} j = k - 1 &: (k - 1)[p_{k-2,k-1}\pi_{k-2}(n) + p_{k,k-1}\pi_k(n) - \pi_{k-1}(n)] \\ j = k &: (k) [p_{k-1,k}\pi_{k-1}(n) + p_{k+1,k}\pi_{k+1}(n) - \pi_k(n)] \\ j = k + 1 &: (k + 1)[p_{k,k+1}\pi_k(n) + p_{k+2,k+1}\pi_{k+2}(n) - \pi_{k+1}(n)] \end{aligned}$$

and collecting the $\pi_k(n)$ terms gives

$$\pi_k(n)[(k - 1)p_{k,k-1} - k + (k + 1)p_{k,k+1}] = \pi_k(n)[2p_{k,k+1} - 1].$$

and A_2 contains the closing probabilities given by

$$A_2 = \begin{bmatrix} 1-p_1 & 0 & 1-p_3 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 1-p_2 & 0 & 1-p_4 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (\text{A.31})$$

and A_1 contains the probabilities of staying in the same state, which is not possible, hence $A_1 = 0$. We note that the sum of the sub-matrices gives the transition matrix (3.50), i.e. $\mathbb{P}_{B'} = A_0 + A_1 + A_2$.

From Costa et al. (2002) we can determine if the QBD process is positive recurrent (losing), null recurrent (fair) or transient (winning) according to if

$$f = eA_0\pi - eA_2\pi \quad (\text{A.32})$$

is less than, equal to or greater than zero. The vector e is a row of ones, which simply serves as a summation. However, one still needs to calculate the stationary distribution π . Using the stationary distribution (3.52), which can be used for both games A and B' leads to same results of equations (3.55).

With a little manipulation the expression f can be written as

$$\begin{aligned} f &= \sum_i p_i \pi_i - \sum_i (1-p_i) \pi_i \\ &= \sum_i 2p_i \pi_i - 1, \end{aligned}$$

which is equivalent to setting p_{win} from (3.53) to a half.

SR Calculations

B.1 Random Number Generators

At the heart of any random number generator (RNG) is an algorithm that produces uniform numbers, usually in the range $[0, 1]$. The uniform numbers can then be transformed to any other distributions via an appropriate mapping function. Thus, to set a good foundation, an accurate uniform RNG is required. Of course, not all mapping functions are created equal, with some being better than others in terms of accuracy, speed or space/time complexity.

Most of the time the built-in RNGs provided with compilers are adequate. However, as stated in the main text, in circumstances where a very large quantity or a high throughput is required, or the quality of random numbers is uncertain, custom designed algorithms need to be used.

In the remainder of this section, we give a uniform RNG (RCARRY) and two commonly used versions of the Box-Müller algorithm to generate standard normal variates. Other mapping algorithms, such as the numerical inversion or wedge-tail methods, that produce normal distributions are also common [Torralba and Chakrabarti, 1993; Knuth, 1969]. A brief description of the ziggurat algorithm is given, which claims to give fast and accurate normal variates.

RCARRY

This is known as a *subtract-and-borrow* algorithm, called RCARRY [James, 1990]. The basic formula is

$$x_n = (x_{n-r} \pm x_{n-s} \pm c) \bmod b,$$

where $r > s$ are the lags, c is a carry bit and b is the word size. In James (1990), it is proposed to use $b = 2^{24}$ with $r = 24$ and $s = 10$. The carry bit is equal to 0 unless the sum is greater than b , in which case it is 1 (in the least significant bit position).

B.1 Random Number Generators

Therefore, we need to store r seeds, two indices and the carry bit, which is the state of the generator. The state needs to be initialised to use the generator.

Box-Müller

The algorithm is as follows [Mannella, 1989; Gillespie, 1996].

1. Generate two random variables U_1 and U_2 that are uniformly distributed between 0 and 1.
2. Compute the two variables $r = \sqrt{-2 \ln U_1}$ and $\theta = 2\pi U_2$.
3. The normal deviates are then formed by $n_1 = r \sin U_1$ and $n_2 = r \cos U_2$.

Each iteration provides two standard normal random numbers. Though this is a simple algorithm, but it can be slow to implement due to the sin, cos, and ln functions.

Transformed Box-Müller

The algorithm is as follows [Knuth, 1969; Mannella, 1989].

1. Generate U_1 and U_2 as above and transform them according to $V_1 = 2U_1 - 1$ and $V_2 = 2U_2 - 1$, which makes them uniformly distributed between -1 and 1 .
2. Compute $S = V_1^2 + V_2^2$.
3. If $S \geq 1$, then go to step 1 and start over, otherwise continue.
4. Compute $S = \sqrt{(-2 \ln S)/S}$ and then form the normal deviates by $n_1 = SV_1$ and $n_2 = SV_2$.

This algorithm eliminates the trigonometric functions, but the trade-off is that steps 1 to 3 are executed 1.27 times on average [Knuth, 1969]. For each completed iteration two normally distributed variables are produced.

Ziggurat

This method is named after the shape of a single, convenient density – a ziggurat – and is due to Marsaglia and Tsang (1984). To explain how it functions, refer to the diagram in Figure B.1. We want to generate a random number variable for the decreasing function $f(x)$.




Figure B.1. A decreasing symmetric density function. The regions are formed by choosing an appropriate unit area rectangle.

We choose a unit rectangle that crosses $f(x)$ twice as shown in Figure B.1, which gives five regions. Region 3 is a rotated and translated version of region 4, thus region 2 and 5 have the same area.

To generate the random deviate, choose a point $P(X, Y)$ uniformly from the rectangle. If P is in region 1, return X , or if in region 3, return $a - X$ (this is the fast part), otherwise further action is required (this is the slow part). Clearly, if region 1 dominates the rectangle, which can be made so by using the ziggurat density, the algorithm is fast.

The ziggurat is a step function that is made up of n layers, each with width of $1/n$. The structure is similar to the result of quantising a signal using a zero-order hold sampling method. Further details and code can be found in Marsaglia and Tsang (1984).

In summary, these algorithms are only necessary if the supplied RNG is substandard for the purposes required. For example, in MATLAB 4.x the RNG has a short period that could repeat within a hour on even a standard computer. Fortunately, MATLAB 5.x has an improved RNG, which is used for most of the simulations in this thesis.

B.2 Numerically Integrating SDEs

We wish to numerically integrate the FHN neuronal model of (4.12) using the RK-4 algorithm. The following code uses the standard RK-4 to solve the deterministic part, while a separate algorithm is used to solve the stochastic part. This is referred to as *coupling* the two algorithms. The algorithm is given as a non-specific pseudo-code.

```

procedure RK-4
  %% find solution to  $y' = f(t, y)$  given initial conditions  $y(t_0) = y_0$ 
  %% from  $t_0$  to  $t_N$  using  $N$  increments of size  $h$ .

```

B.2 Numerically Integrating SDEs

```
require: signal s
variables: t, v, w, η, k's

for n in 0..N-1 loop

    %% find the required constants
    kv1 = h * fhn1(tn, vn, wn, sn, ηn)
    kw1 = h * fhn2(tn, vn, wn)

    kv2 = h * fhn1(tn + ½h, vn + ½kv1, wn + ½kw1, sn+1/2, ηn+1/2)
    kw2 = h * fhn2(tn + ½h, vn + ½kv1, wn + ½kw1)

    kv3 = h * fhn1(tn + ½h, vn + ½kv2, wn + ½kw2, sn+1/2, ηn+1/2)
    kw3 = h * fhn2(tn + ½h, vn + ½kv2, wn + ½kw2)

    kv4 = h * fhn1(tn + h, vn + kv3, wn + kw3, sn+1, ηn+1)
    kw4 = h * fhn2(tn + h, vn + kv3, wn + kw3)

    %% update the variables
    tn+1 = tn + h
    vn+1 = vn + ⅙(kv1 + 2kv2 + 2kv3 + kv4)
    wn+1 = wn + ⅙(kw1 + 2kw2 + 2kw3 + kw4)
    ηn+1 = ηn * exp(-α) + randn * sqrt(D(1 - exp(-2α))/τc)

end loop;

end RK-4;
```

This requires the following functions for the FHN neuron.

```
function fhn1(t, v, w, s, η)
    return (v(v - a)(1 - v) - w + A + s + η)/ε
end fhn1;

function fhn2(t, v, w)
    return v - w - b
end fhn2;
```

The problem is how to deal with values of s and η , that are both essentially random signals (on different time scales), between the sampling times. Since the mean frequency of the signal is much smaller than the integration size, it can be held constant without any problems. However, it is not immediately clear how to represent the noise process that may have a correlation time the same order as the step size. A number of interpolation schemes are explored in Davis (2000). The scheme with the smallest error uses a type of rectangular window.

The technique employed is to hold the noise constant at all times within the numerical integration step, Δ . When the step finishes at $t = i\Delta$, the noise value for the next

integration step is updated and will be different. This means the noise will have two different values at $t = i\Delta$, just before and just after $i\Delta$. Though this may be difficult (or impossible) to program using a built-in integration algorithm, it lends itself well for the algorithm given above. This is a natural approach when using varying time signals that cannot be known or predicted in the future. Thus, for the algorithm given above, we can take $s_{n+1} = s_{n+1/2} = s_n$ and $\eta_{n+1} = \eta_{n+1/2} = \eta_n$ for each interval Δ .

The algorithm given is only a relatively simple one that is adequate for our purposes. However, there are other techniques that use Runge-Kutta or Euler algorithms with coloured noise [Fox, 1989a] that are more complicated and accurate [Honeycutt, 1992a; Honeycutt, 1992b].

B.3 AMI Expressions for Gaussian Signals and Noise

In this section the calculations used to get from (4.32) to (4.34) are given. Preliminary identities that are required are first derived. The notation for the network given in Figure 4.27 is used along with the Heaviside function (4.30). We assume that the signal and noise have Gaussian distributions of the form given in (4.35) and all the thresholds are identically set to θ . For convenience, the original AMI relation to be simplified is reproduced,

$$AMI = - \sum_{n=0}^N P_y(n) \lg P_y(n) - \left(- \int_{-\infty}^{\infty} \sum_{n=0}^{\infty} P_x(x) P(n|x) \lg P(n|x) dx \right). \quad (\text{B.1})$$

At the output of the network, the probability the output is in state n is given by

$$P_y(n) = \int_{-\infty}^{\infty} P_y(n|x) P_x(x) dx. \quad (\text{B.2})$$

B.3.1 Preliminary Calculations

We first derive some identities that will be useful for later calculations.

B.3 AMI Expressions for Gaussian Signals and Noise

Finding $P_{1|x}$

$P_{1|x}$ is the conditional probability of a device being in state 1 given the signal value x . It is given by the cumulative distribution

$$P_{1|x} = \int_{\theta-x}^{\infty} P_{\eta}(\eta) d\eta \quad (\text{B.3})$$

$$= \frac{1}{\sqrt{\pi}} \int_{\theta-x}^{\infty} \frac{1}{\sqrt{2\sigma_{\eta}^2}} \exp\left(\frac{-\eta^2}{2\sigma_{\eta}^2}\right) d\eta. \quad (\text{B.4})$$

This is an integral of the normal function and hence can be written in terms of the complementary error function given by $\text{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^{\infty} \exp(-t^2) dt$. Using the substitution of $k = \eta/\sqrt{2\sigma_{\eta}^2}$, $P_{1|x}$ is reduced to

$$\begin{aligned} P_{1|x} &= \frac{1}{\sqrt{\pi}} \int_{\frac{\theta-x}{\sqrt{2\sigma_{\eta}^2}}}^{\infty} \exp(-k^2) dk \\ &= \frac{1}{2} \text{erfc}\left(\frac{\theta-x}{\sqrt{2\sigma_{\eta}^2}}\right). \end{aligned} \quad (\text{B.5})$$

Correspondingly, $P_{0|x}$ is the conditional probability a device is in state 0 given x , and since a device must be in either state 0 or 1, we have

$$P_{0|x} = 1 - P_{1|x}. \quad (\text{B.6})$$

Demonstrating that $\sum_{n=0}^N P(n|x) = 1$

$P(n|x)$ is the probability of the output being in state n given the input signal x . Thus, the probability that n devices are triggered is given by the binomial distribution as

$$P(n|x) = C_n^N P_{1|x}^n P_{0|x}^{N-n}. \quad (\text{B.7})$$

The summation of (B.7) then becomes

$$\begin{aligned} \sum_{n=0}^N P(n|x) &= \sum_{n=0}^N C_n^N P_{1|x}^n P_{0|x}^{N-n} \\ &= (P_{1|x} + P_{0|x})^N = 1, \end{aligned} \quad (\text{B.8})$$

by use of (B.6). Intuitively we see this must be true since we are summing over the complete probability space, which must be unity.

Demonstrating that $\sum_{n=0}^N nP(n|x) = NP_{1|x}$

Starting with $\sum_0^N nP(n|x)$ we substitute (B.7) to give

$$\sum_{n=0}^N nP(n|x) = \sum_{n=0}^N nC_n^N P_{1|x}^n P_{0|x}^{N-n}. \quad (\text{B.9})$$

With a little manipulation we can show that $nC_n^N = C_{n-1}^{N-1}N$, and rewriting (B.9) gives

$$\begin{aligned} \sum_{n=0}^N nP(n|x) &= \sum_{n=0}^N NP_{1|x} C_{n-1}^{N-1} P_{1|x}^{n-1} P_{0|x}^{N-n} \\ &= NP_{1|x}, \end{aligned} \quad (\text{B.10})$$

using the result of (B.8).

Simplification of $\sum_{n=0}^N P(n|x) \lg P(n|x)$

We wish to massage $\sum_0^N P(n|x) \lg P(n|x)$, which forms part of (B.1), into a form that will be beneficial in later calculations. The log term is expanded by multiplying its argument by $C_n^N (C_n^N)^{-1}$ to give

$$\sum_{n=0}^N P(n|x) \lg P(n|x) = \sum_{n=0}^N P(n|x) \lg C_n^N + \sum_{n=0}^N P(n|x) \lg P(n|x) (C_n^N)^{-1}. \quad (\text{B.11})$$

Simplifying the notation by setting $p = P_{1|x}$, $q = P_{0|x}$, and then using (B.7) to give

$$\begin{aligned} \sum_{n=0}^N P(n|x) \lg P(n|x) &= \sum_{n=0}^N P(n|x) \lg C_n^N + \sum_{n=0}^N P(n|x) \lg p^n q^{N-n} \\ &= \sum_{n=0}^N P(n|x) \lg C_n^N \\ &\quad + \lg p \sum_{n=0}^N nP(n|x) + N \lg q \sum_{n=0}^N P(n|x) - \lg q \sum_{n=0}^N nP(n|x). \end{aligned}$$

Using the relations (B.8) and (B.10) yields

$$\begin{aligned} \sum_{n=0}^N P(n|x) \lg P(n|x) &= \sum_{n=0}^N P(n|x) \lg C_n^N + Np \lg p + N \lg q - Np \lg q \\ &= \sum_{n=0}^N P(n|x) \lg C_n^N + N [p \lg p + q \lg q], \end{aligned} \quad (\text{B.12})$$

which is the desired result. It should be noted that the corresponding equation given in [Stocks, 2001b; Stocks, 2001c] (which is equation (3) in both), contains a typesetting error, though the end result is of course correct.

B.3.2 Simplifying the average mutual information

Substituting the simplification (B.12) into the original expression for AMI given by (B.1) gives

$$AMI = - \sum_{n=0}^N P_y(n) \lg P_y(n) + \int_{-\infty}^{\infty} P_x(x) \sum_{n=0}^N P(n|x) \lg C_n^N dx + N \int_{-\infty}^{\infty} P_x(x) [p \lg p + q \lg q] dx.$$

Assuming that we can interchange the integral and summation, the middle term becomes

$$\sum_{n=0}^N \int_{-\infty}^{\infty} P_x(x) P(n|x) \log C_n^N dx = \sum_{n=0}^N P_y(n) \log C_n^N \quad (B.13)$$

with the aid of (B.2). Further simple manipulation yields

$$AMI = - \sum_{n=0}^N P_y(n) \log P'(n) - \left(- N \int_{-\infty}^{\infty} P_x(x) [p \log p + q \log q] dx \right), \quad (B.14)$$

where $P_y(n) = C_n^N P'(n)$, which is the desired result.

Numerical Approach

Using the notation of (B.14), the AMI can easily be evaluated numerically from the PDFs of the signal and noise. Figure B.2 shows the relations between several key variables. We only require the variance of the noise and signal, and the threshold. To find AMI from experimental data, histograms can be used to find $P_x(x)$ and $P_\eta(\eta)$ from which $P_{1|x}$ can be determined.




Figure B.2. Relationships between variables for the threshold network. $P_x(x)$ and $P_\eta(\eta)$ are the PDFs of the signal and noise respectively. The cumulative distributions shown in the middle plot are related to the areas under $P_\eta(\eta)$ for a particular value of x .

B.4 Shannon's Channel Capacity

We start by giving the noise power as $P_n = S_n B_{n,\text{eff}}$, where S_n is the PSD of the noise (which is flat for white noise) and $B_{n,\text{eff}}$ is the effective bandwidth of the noise. This means, for a signal within a certain bandwidth, we can limit the PSD of the noise by employing a simple linear filter with a cut-off frequency equal to the signal bandwidth. Thus, the noise bandwidth is equal to the signal bandwidth B_s , to give

$$P_n = S_n B_s. \quad (\text{B.15})$$

Substituting (B.15) into the Shannon formula of (4.39) gives the information transfer rate as

$$C = B_s \log_2 \left(1 + \frac{P_s}{S_n B_s} \right). \quad (\text{B.16})$$

Taking the channel capacity at the output of the STR means P_s becomes $P_{s,\text{out}}$ and S_n becomes $S_{n,\text{out}}$. Using the signal-to-noise ratio given in (4.24) we have

$$C = B_s \log_2 \left(1 + \frac{\text{SNR}_{\text{out}}}{B_s} \right), \quad (\text{B.17})$$

which gives the channel capacity in terms of the output SNR and signal bandwidth.

To find the maximal bandwidth of the signal B_s , we need to consider Shannon's sampling theorem and the mean level crossing frequency $\nu(U_t)$, given in [Kiss, 1996] as

$$\nu(U_t) = \frac{2}{\sigma} \exp\left(\frac{-U_t^2}{2\sigma^2}\right) \left(\int_0^\infty f^2 S(f) df \right)^{1/2}, \quad (\text{B.18})$$

where the noise power $\sigma^2 = \int_0^\infty S(f) df = B_{n,\text{in}} S_{n,\text{in}}$. From the sampling theorem, B_s is approximately equal to half the mean spike frequency, which is half the mean level crossing frequency due to noise in any direction, thus $B_s = \nu(U_t)/4$. By direct integration we have

$$\frac{2}{\sigma} \left(\int_0^\infty f^2 S(f) df \right)^{1/2} = \frac{2}{\sqrt{3}} B_{n,\text{in}},$$

then combining with (B.18) we find the signal bandwidth as

$$B_s = \frac{B_{n,\text{in}}}{2\sqrt{3}} \exp\left(\frac{-U_t^2}{2B_{n,\text{in}} S_{n,\text{in}}}\right). \quad (\text{B.19})$$

From (A.3.4) in Kiss (1996), the output SNR is given as

$$\begin{aligned} \text{SNR}_{\text{out}} &= \frac{\nu(0)(AU_t)^2}{\sigma^4} \exp\left(\frac{-U_t^2}{2\sigma^2}\right) \\ &= \frac{2}{\sqrt{3}} B_{n,\text{in}} \frac{(AU_t)^2}{(B_{n,\text{in}} S_{n,\text{in}})^2} \exp\left(\frac{-U_t^2}{2B_{n,\text{in}} S_{n,\text{in}}}\right), \end{aligned} \quad (\text{B.20})$$

B.4 Shannon's Channel Capacity

where $\nu(0)$ is found from (B.18). This is given in terms of B_s by

$$SNR_{\text{out}} = \frac{(2AU_t)^2}{(B_{n,\text{in}}S_{n,\text{in}})^2} B_s. \quad (\text{B.21})$$

Substituting (B.21) and (B.19) back into (B.17) gives the channel capacity as

$$C = \frac{B_{n,\text{in}}}{2\sqrt{3}} \exp\left(\frac{-U_t^2}{2B_{n,\text{in}}S_{n,\text{in}}}\right) \log_2\left(1 + \frac{(2AU_t)^2}{(B_{n,\text{in}}S_{n,\text{in}})^2}\right), \quad (\text{B.22})$$

which is the desired result.

Appendix C

Lowpass & Highpass Filters

This section briefly deals with simple RC filters and their transfer functions. We require to find expressions for certain components of biological and digital systems which can be easily simulated or implemented in analog VLSI [Mead, 1989].

Firstly, we consider the delay stage. To approximate an ideal delay, a lowpass filter is employed. This is easily realisable and simple to simulate, and we derive an updating formula for it.

Secondly, we consider a bandpass filter, which typically consist of a LPF and a HPF in series. In terms of hardware implementation, this is a straightforward exercise. However, for simulations the form of the transfer function of the HPF does not lend itself to producing a simple updating formula. We consider an alternative implementation of the BPF and give the respective updating formula.

Circuits, Transfer Functions and Differential Relations

The standard RC lowpass and highpass filters are shown in Figure C.1. In the frequency domain the impedances are $Z_R = R$ and $Z_C = \frac{1}{j\omega C}$ for the resistor and capacitor




Figure C.1. Simple RC circuits for lowpass and highpass filters. The input voltage is represented as V_{in} and $v_{in}(t)$ in the frequency and time domain respectively. The components have resistance R and capacitance C with voltage drops of V_R and V_C respectively.

respectively. Using the substitution $s = j\omega$, the transfer function can be written as

$$H_L(s) = \frac{V_C}{V_{in}} = \frac{\frac{1}{RC}}{s + \frac{1}{RC}} = \frac{\omega_c}{s + \omega_c}, \quad (\text{C.1})$$

where $\omega_c = 1/RC$ is the cut-off frequency. This is the frequency where the output power is half of the input, that is $|H(s)|^2 = \frac{1}{2}$. This is referred to as the 3 dB point. The reason for transforming into the s -domain is that it is amenable to several types of mathematical analysis, such as Laplace transforms for example.

In the time domain, the current-voltage relations are $i = v_R/R$ and $i = Cdv_C/dt$ for the resistor and capacitor respectively. Equating the two currents and using $v_{in} = v_R + v_C$ leads to the relation

$$\dot{v}_C = -\frac{1}{\tau}v_C + \frac{1}{\tau}v_{in}, \quad (\text{C.2})$$

where $\tau = 1/\omega_c$ is the time constant.

Repeating this procedure for the highpass filter shown in Figure C.1 gives the transfer function to be

$$H_H(s) = \frac{s}{s + \omega_c}, \quad (\text{C.3})$$

and the differential equation as

$$\dot{v}_R = -\frac{1}{\tau}v_R + \dot{v}_{in}. \quad (\text{C.4})$$

However, this has several problems. For the transfer function, H_H is not readily solvable by Laplace transforms and the \dot{v}_{in} in differential form complicates matters.

As an aside, one can switch between the time and s -domain directly by substituting s for the operator d/dt . This was first discovered by Oliver Heaviside [Josephs, 1950], who developed the theory to enable easy manipulation of circuit equations. Although he did not understand the deep mathematical basis of it, he was able to generalise the theory. This enabled him to predict the behaviour of many circuits. It was not until about 10 years later when mathematicians started to take notice, where it was revealed they were actually Laplace transforms – formalised by Laplace some 100 years earlier [Mead, 1989, p. 136].

Intuitively a bandpass filter can be formed by cascading the two circuits shown in Figure C.1 to give the transfer function,

$$H_B(s) = \frac{\omega_c s}{s^2 + 3\omega_c s + \omega_c^2}. \quad (\text{C.5})$$

The order in which the filters are applied is irrelevant. Recall that $\mathcal{L}\{h_L(t)h_H(t)\} \neq H_L(s)H_H(s)$, a convolution is required (\mathcal{L} is the Laplace transform operator).

When dealing with stochastic signals it is useful to have an updating formula that only depends on current and previous values. We start by looking at the form of the LPF given in (C.2). We can find an exact solution to (C.2) in the following manner.

First define $\lambda = 1/\tau$ and multiply through by $e^{\lambda t}$ to give

$$e^{\lambda t} \frac{dv_C}{dt} + e^{\lambda t} \lambda v_C = e^{\lambda t} \lambda v_{in}.$$

Noting that the LHS is equivalent to $\frac{d}{dt}(e^{\lambda t} v_C(t))$, we integrate both sides with respect to t yielding

$$e^{\lambda t} v_C(t) = K + \int_0^t e^{\lambda s} \lambda v_{in}(s) ds,$$

where K is a constant. Rearranging gives the final form of

$$v_C = Ke^{-\lambda t} + \lambda e^{-\lambda t} \int_0^t e^{\lambda s} v_{in}(s) ds.$$

Updating Formula

For the sake of further analysis, we relabel $v_C \rightarrow y$ for the output and $v_{in} \rightarrow x$ for the input. To produce an updating formula we use time increments $t = h(n-1), hn, h(n+1), \dots$ for step size h . Then at the n th step we have $t_n = nh$ and thus

$$y_n = Ke^{-\lambda nh} + \lambda e^{-\lambda nh} \int_0^{nh} e^{\lambda s} x(s) ds, \quad (C.6)$$

where y_n is shorthand for $y(nh)$ and similarly for x_n . At the next time step, $t_{n+1} = (n+1)h$ and we have

$$y_{n+1} = Ke^{-\lambda h(n+1)} + \lambda e^{-\lambda h(n+1)} \int_0^{h(n+1)} e^{\lambda s} x(s) ds. \quad (C.7)$$

Breaking the integral into $\int_0^{nh} + \int_{nh}^{h(n+1)}$ and grouping selected terms gives

$$y_n = e^{-\lambda h} \left[Ke^{-\lambda nh} + \lambda e^{-\lambda nh} \int_0^{nh} e^{\lambda s} x(s) ds \right] + \lambda e^{-\lambda h(n+1)} \int_{nh}^{h(n+1)} e^{\lambda s} x(s) ds.$$

Notice that within the square brackets we actually have y_n from (C.6). Also since $x(t)$ remains constant in the range $t = nh$ to $h(n+1)$, we can substitute x_{n+1} for $x(s)$ in the integrand of the second integral. Thus we can complete the second integral to give

$$y_{n+1} = e^{-\lambda h} y_n + (1 - e^{-\lambda h}) x_{n+1}, \quad (C.8)$$

and with some simple manipulation we have

$$y_{n+1} = y_n + (1 - e^{-\lambda h})(x_{n+1} - y_n). \quad (\text{C.9})$$

This is the desired updating formula for the lowpass filter. It only depends on the previous value of the output (y_n) and the current value of the signal (x_{n+1}). Since the $1 - e^{-\lambda h}$ term does not change during the algorithm, it can be precomputed which leaves only elementary operations to update y_n .

Alternatively, by using a numerical integration algorithm, the Runge-Kutta (second or fourth order) or Euler's method for example, we arrive at

$$y_{n+1} = y_n + h\lambda(x_{n+1} - y_n).$$

This can be seen by taking the expansion of the exponential in (C.9), which gives $1 - e^{-\lambda h} = h\lambda + O((h\lambda)^2)$. This is only valid for small $h\lambda$, with the error being negligible when $h\lambda < 1$.

One may notice that the OU process described by (4.21) on page 121 is very similar to the LPF described by (C.2). Ignoring the constant in front of the input term, why are the derived formulas (4.22) and (C.8) so different? The reason is the assumptions made during the analysis. For the LPF we assumed the input signal was constant during the integration interval, while for the OU process this assumption cannot be made, and a formal SDE approach is required.

Bandpass Filter

Creating an updating formula for the highpass filter is not as simple. Fortunately, to produce the desired bandpass filter, only lowpass filters are necessary. By taking the difference between two LPFs that have different time constants, a BPF is formed. The difference in frequency response using this method as opposed to a 'real' BPF is negligible, as shown in Figure C.2, where several filter responses have been plotted.

Given the two LPF outputs of y and z with $\lambda_1 < \lambda_2$ respectively, then the bandpass filtered output can be given by

$$b_{n+1} = z_{n+1} - y_{n+1}.$$




Figure C.2. Magnitude responses of lowpass, bandpass and highpass filters. The two lowpass filters have cut-off frequencies of 20 and 800 rad/s, as shown by the dashed lines at the -3 dB points. The HPF has $\omega = 20$ rad/s. The response of the BPFs formed by taking the difference between the two LPFs and cascading a LPF and HPF is almost identical.

Bibliography

- ABBOTT-D (2001). Overview: Unsolved problems of noise and fluctuations, *Chaos*, **11**(3), pp. 526–538.
- ABBOTT-D AND PARFITT-A (1997). Extension of the vision paradigm to millimeter waves, *Proceedings of SPIE*, Vol. 3207, pp. 103–106.
- ABBOTT-D, DAVIS-B. R AND PARRONDO-J. M. R (2000). The problem of detailed balance for the Feynman-Smoluchowski Engine (FSE) and the multiple pawl paradox, in D. Abbott and L. B. Kish (eds.), *Second International Conference on Unsolved Problems of Noise and Fluctuations*, American Institute of Physics, Adelaide, Australia, pp. 213–220.
- ABBOTT-D, DAVIS-B. R, PHILLIPS-N. J AND ESHRAGHIAN-K (1996). Simple derivation of the thermal noise formula using window-limited Fourier transforms, *IEEE Transactions in Education*, **39**(1), pp. 1–13.
- ABBOTT-D, MOINI-A, YAKOVLEFF-A, NGUYEN-X. T, BLANKSBY-A, KIM-G, BOUZERDOUM-A, BOGNER-R. E AND ESHRAGHIAN-K (1994). A new VLSI smart sensor for collision avoidance inspired by insect vision, *Proceedings of SPIE*, Vol. 2344, pp. 105–115.
- ABBOTT-D, YAKOVLEFF-A, MOINI-A, NGUYEN-X. T, BLANKSBY-A, BEARE-R, BEAUMONT-SMITH-A, KIM-G, BOUZERDOUM-A, BOGNER-R. E AND ESHRAGHIAN-K (1995). Biologically inspired obstacle avoidance – a technology independent paradigm, *Proceedings of SPIE*, Vol. 2591, pp. 2–12.
- ADJARI-A AND PROST-J (1992). Drift induced by a periodic potential of low symmetry: pulsed dielectrophoresis, *C. R. Academy of Science Paris, Série II*, **315**, pp. 1635–1639.
- AJDARI-A, LEWINER-J, PROST-J AND VIOVY-J (1997). *Devices for separating particles contained in a fluid*, number 5 593 565, United States Patent.
- ALLISON-A AND ABBOTT-D (2001). Control systems with stochastic feedback, *Chaos*, **11**(3), pp. 715–724.
- ALLISON-A, PEARCE-C. E. M AND ABBOTT-D (2002). State-space visualisation and fractal properties of Parrondo's games, in A. Nowack (ed.), *Annals of the International Society on Dynamic Games*, Birkhäuser. Submitted.
- ANDÒ-B AND GRAZIANI-S (2000). *Stochastic Resonance: Theory and Applications*, Kluwer Academic Publishers.
- ARENA-P, FAZZINO-S, FORTUNA-L AND MANISCALCO-P (2001). Non linear dynamics and the Parrondo paradox, *XIV Italian Meeting on Game Theory and Applications*, Ischia, Italy. Unpublished.
- ASDI-A. S AND TEWFIK-A. H (1995). Detection of weak signals using adaptive stochastic resonance, *1995 International Conference on Acoustics, Speech, and Signal Processing*, IEEE Conference Proceedings, IEEE, New York, NY, USA, pp. 1332–1335.
- ASTUMIAN-R. D (1997). Thermodynamics and kinetics of a Brownian motor, *Science*, **276**, pp. 917–922.

Bibliography

- ASTUMIAN-R. D (2001). Making molecules into motors, *Scientific American*, **285**(1), pp. 57–64.
- ASTUMIAN-R. D AND BIER-M (1994). Fluctuation driven ratchets: Molecular motors, *Physical Review Letters*, **72**(11), pp. 1766–1769.
- ASTUMIAN-R. D AND MOSS-F (1998). Overview: The constructive role of noise in fluctuation driven transport and stochastic resonance, *Chaos*, **8**(3), pp. 533–538.
- BADER-J. S, ROTHBERG-J. M, DEEM-M. W, MULHERN-G. T AND WENT-G. T (1999). *Separation of charged particles by a spatially and temporally varying electric field*, number 5 938 904, United States Patent.
- BALÁZSI-G, KISS-L. B AND MOSS-F. E (2000). Spatiotemporal stochastic resonance and its consequences in a neuron system, in D. Abbott and L. B. Kish (eds.), *Second International Conference on Unsolved Problems of Noise and Fluctuations*, Vol. 511, American Institute of Physics, Adelaide, Australia, pp. 159–168.
- BARLOW-H. B AND LEVICK-W. R (1965). The mechanism of directionally selective units in rabbit's retina, *Journal of Physiology*, **178**, pp. 477–504.
- BARTUSSEK-R, HÄNGGI-P AND KISSNER-J. G (1994). Periodically rocked thermal ratchets, *Europhysics Letters*, **28**(7), pp. 459–464.
- BEARE-R (1997). *Image segmentation based on local motion detection*, PhD thesis, University of Adelaide.
- BEARE-R, BLANKSBY-A AND BOUZERDOUM-A (1995). Low level visual motion processing using local motion detectors, *IEEE International Conference on Neural Networks Proceeding*, Vol. 1, Perth, WA, Australia, pp. 1–6.
- BENNETT-C. H (1982). The thermodynamics of computation—a review, *International Journal of Theoretical Physics*, **21**(12), pp. 905–940.
- BENNETT-C. H (1987). Demons, engines and the second law, *Scientific American*, **257**(5), pp. 88–96.
- BENZI-R (1980). Short contributions, NATO International School of Climatology. Unpublished Short Communication.
- BENZI-R, PARISI-G, SUTERA-A AND VULPIANI-A (1982). Stochastic resonance on climate change, *Tellus*, **34**(1), pp. 10–16.
- BENZI-R, SUTERA-A AND VULPIANI-A (1981). The mechanism of stochastic resonance, *Journal of Physics A*, **14**, pp. L453–457.
- BEZRUKOV-S. M (1998). Stochastic resonance as an inherent property of rate-modulated random series of events, *Physics Letters A*, **248**(1), pp. 29–36.
- BIALEK-W AND DEWEESE-M (1995). Random switching and optimal processing in the perception of ambiguous figures, *Physical Review Letters*, **74**(15), pp. 3077–3080.
- BIALEK-W, DEWEESE-M, RIEKE-F AND WARLAND-D (1993). Bits and brains: information flow in the nervous system, *Physica A*, **200**(1-4), pp. 581–593.

- BIER-M (1997). Brownian ratchets in physics and biology, *Contemporary Physics*, **38**(6), pp. 371–379.
- BIER-M AND ASTUMIAN-R. D (1996). Biasing Brownian motion in different directions in a 3-state fluctuating potential and an application for the separation of small particles, *Physical Review Letters*, **76**(22), pp. 4277–4280.
- BISHOP-C. M (1996). *Neural Networks for Pattern Recognition*, Oxford Press, chapter 9, pp. 346–349.
- BLACKWELL-D AND GIRSHICK-M. A (1954). *Theory of Games and Statistical Decisions*, John Wiley & Sons, New York, NY, USA.
- BOMAN-M, JOHANSSON-S AND LYBACK-D (2001). Parrondo strategies for artificial traders, *The 2001 International Conference of Intelligent Agent Technology*, World Scientific, Maebashi Cite, Japan.
- BORST-A AND EGELHAAF-M (1989). Principles of visual motion detection, *Trends in neurosciences*, **12**(8), pp. 297–306.
- BOUZERDOUM-A AND PINTER-R. B (1989). Image processing in biological and computer vision systems, *Proceedings of SPIE*, Vol. 1199, pp. 1229–1240.
- BOUZERDOUM-A AND PINTER-R. B (1990). A shunting inhibition motion detector that can account for the functional characteristics of fly motion-sensitive interneurons, *IEEE International Joint Conference on Neural Networks*, Vol. 1, New Yorke, NY, USA, pp. 149–153.
- BRILLOUIN-L (1950). Can the rectifier become a thermodynamical demon?, *Physical Review*, **78**(5), pp. 627–628.
- BRUCE-I. C, IRLICHT-L. S, WHITE-M. W, O'LEARY-S. J, DYNES-S, JAVEL-E AND CLARK-G. M (1999a). A stochastic model of the electrically stimulated auditory nerve: Pulse-train response, *IEEE Transactions on Biomedical Engineering*, **46**(6), pp. 630–637.
- BRUCE-I. C, WHITE-M. W, IRLICHT-L. S, O'LEARY-S. J AND CLARK-G. M (1999b). The effects of stochastic neural activity in a model predicting intensity perception with cochlear implants: Low-rate stimulation, *IEEE Transactions on Biomedical Engineering*, **46**(12), pp. 1393–1404.
- BRUCE-I. C, WHITE-M. W, IRLICHT-L. S, O'LEARY-S. J, DYNES-S, JAVEL-E AND CLARK-G. M (1999c). A stochastic model of the electrically stimulated auditory nerve: Single-pulse response, *IEEE Transactions on Biomedical Engineering*, **46**(6), pp. 617–29. See also [Bruce et al., 1999a; Bruce et al., 1999b].
- BULSARA-A. R AND GAMMAITONI-L (1996). Tuning into noise, *Physics Today*, **49**(3), pp. 39–45.
- BULSARA-A. R AND ZADOR-A (1996). Threshold detection of wideband signals: A noise-induced maximum in the mutual information, *Physical Review E*, **54**(3), pp. 2185–2188.
- CAPURRO-A, PAKDAMAN-K, NOMURA-T AND SATO-S (1998). Aperiodic stochastic resonance with correlated noise, *Physical Review E*, **58**(4), pp. 4820–4827.
- CECCATTO-H. A AND SCHRECKENBERG-M (1989). Motion detection in noisy environments: Phase transition on a model system, *Physical Review Letters*, **63**(19), pp. 2152–2155.
- CHAITIN-G (1999). *The Unknowable*, Springer-Verlag.

Bibliography

- CHAPEAU-BLONDEAU-F (1996). Stochastic resonance in the Heaviside nonlinearity with white noise and arbitrary periodic signal, *Physical Review E*, **53**(5), pp. 5469–5472.
- CHAPEAU-BLONDEAU-F (1997a). Input–output gains for signal in noise in stochastic resonance, *Physics Letters A*, **232**, pp. 41–48.
- CHAPEAU-BLONDEAU-F (1997b). Noise-enhanced capacity via stochastic resonance in an asymmetric binary channel, *Physical Review E*, **55**(2), pp. 2016–2019.
- CHAPEAU-BLONDEAU-F AND GODIVIER-X (1997). Theory of stochastic resonance in signal transmission by static nonlinear systems, *Physical Review E*, **55**(2), pp. 1478–1495.
- CHIALVO-D. R, LONGTIN-A AND MÜLLER-GERKING-J (1997). Stochastic resonance in models of neuronal ensembles, *Physical Review E*, **55**(2), pp. 1798–1808.
- CHOW-C. C, IMHOFF-T. T AND COLLINS-J. J (1998). Enhancing aperiodic stochastic resonance through noise modulation, *Chaos*, **8**(3), pp. 616–620.
- COLLINS-J. J, CHOW-C. C AND IMHOFF-T. T (1995a). Aperiodic stochastic resonance in excitable systems, *Physical Review E*, **52**(4), pp. 3321–3324.
- COLLINS-J. J, CHOW-C. C AND IMHOFF-T. T (1995b). Stochastic resonance without tuning, *Nature*, **376**, pp. 236–238.
- COLLINS-J. J, CHOW-C. C, CAPELA-A. C AND IMHOFF-T. T (1996a). Aperiodic stochastic resonance, *Physical Review E*, **54**(5), pp. 5575–5584.
- COLLINS-J. J, IMHOFF-T. T AND GRIGG-P (1996b). Noise-enhanced information transmission in rat SA1 cutaneous mechanoreceptors via aperiodic stochastic resonance, *Journal of Neurophysiology*, **76**(1), pp. 642–645.
- COSTA-A, FACKRELL-M AND TAYLOR-P. G (2002). Two issues surrounding Parrondo’s paradox, in A. Nowack (ed.), *Annals of the International Society on Dynamic Games*, Birkhäuser. Submitted.
- CURRIE-N. C AND BROWN-C. E (1987). *Principles and applications of millimeter-wave radar*, Artech House, Norwood, MA, USA.
- DAVIES-P. C. W (2001). Physics and life: The Abdus Salam Memorial Lecture, in J. Chela-Flores, T. Tobias and F. Raulin (eds.), *Sixth Trieste Conference on Chemical Evolution*, Kluwer Academic Publishers, Trieste, Italy. In press.
- DAVIS-B. R (2000). Numerical methods for systems excited by white noise, in D. Abbott and L. B. Kiss (eds.), *Second International Conference on Unsolved Problems of Noise and Fluctuations*, American Institute of Physics, Adelaide, Australia, pp. 533–538.
- DAWE-R (2000a). *Improved sonar detector using stochastic resonance*, number 64137/00, Australian Patent Application.
- DAWE-R. L (2000b). ISHTAR: A counter stealth sonar, *UDT Pacific 2000 Conference*, pp. 318–322.
- DAWE-R. L AND GALBREATH-E. R (2000). Stochastic resonance applications in underwater acoustic signal detection, *Australian Acoustical Society Annual Conference*, pp. 47–53.

- DE-CHUN-G, GUNG-H, XIAO-DONG-W, CHUN-YAN-Y, GUANG-RONG-Q, RONG-L AND DA-FU-D (1992). Experimental study of the signal-to-noise ratio of stochastic resonance systems, *Physical Review A*, **46**(6), pp. 3243–3249.
- DEWEESE-M AND BIALEK-W (1995). Information flow in sensory neurons, *Il Nuovo Cimento D*, **17D**(7-8), pp. 733–741.
- DIJKSTRA-E. W (1990). Making a fair roulette from a possibly biased coin, *Information Processing Letters*, **36**, p. 193.
- DOOB-J. L (1953). *Stochastic Processes*, John Wiley & Sons, Inc., New York.
- DOUGLAS-J. K, WILKENS-L, PANTAZELOU-E AND MOSS-F (1993). Noise enhancement of information transfer in crayfish mechanoreceptors by stochastic resonance, *Nature*, **365**, pp. 337–339.
- DROR-R. O, O'CARROLL-D. C AND LAUGHLIN-S. B (2001). Accuracy of velocity estimation by Reichardt correlators, *Journal of the Optical Society of America A*, **18**(2), pp. 241–252.
- DUKE-T. A. J AND AUSTIN-R. H (1998). Microfabricated sieve for the continuous sorting of macromolecules, *Physical Review Letters*, **80**(7), pp. 1552–1555.
- DURRETT-R (1991). *Probability: Theory and Examples*, Duxbury Press, Belmont, California.
- DURRETT-R, KESTEN-H AND LAWLER-G (1991). Making money from fair games, in R. Durrett and H. Kesten (eds.), *Random Walks, Brownian motion, and Interacting particle systems*, Vol. 28, Birkhäuser: Boston, pp. 255–267.
- DYKMAN-M. I, LUCHINSKY-D. G, MANNELLA-R, MCCLINTOCK-P. V. E, SHORT-H. E, STEIN-N. D AND STOCKS-N. G (1994). Noise-induced linearisation, *Physics Letters A*, **193**, pp. 61–66.
- EICHWALD-C AND WALLECZEK-J (1997). Aperiodic stochastic resonance with chaotic input signals in excitable systems, *Physical Review E*, **55**(6 A), pp. 6315–6318.
- EINSTEIN-A (1905). On the movement of small particles suspended in stationary liquids required by the molecular-kinetic theory of heat, *Annalen der Physik*, **17**, pp. 549–560. Appearing in The Collection of Papers of Albert Einstein, vol. 2, pp. 123–134.
- ERTAS-D (1998). Lateral separation of macromolecules and polyelectrolytes in microlithographic arrays, *Physical Review Letters*, **80**(7), pp. 1548–1551.
- FAUCHEUX-L. P, BOURDIEU-L. S, KAPLAN-P. D AND LIBCHABER-A. J (1995). Optical thermal ratchet, *Physical Review Letters*, **74**(9), pp. 1504–1509.
- FAUVE-S AND HESLOT-F (1983). Stochastic resonance in a bistable system, *Physics Letters*, **97A**(1), pp. 5–7.
- FELDMAN-D, IMPAGLIAZZO-R, NAOR-M, NISAN-N, RUDICH-S AND SHAMIR-A (1993). On dice and coins: Models of computation for random generation, *Information and Computation*, **104**(2), pp. 159–174.
- FENNEMA-C. L AND THOMPSON-W. B (1979). Velocity determination in scenes containing several moving objects, *Computer Graphics and Image Processing*, **9**, pp. 301–315.

Bibliography

- FEYNMAN-R. P, LEIGHTON-R. B AND SANDS-M (1963). *The Feynman Lectures on Physics*, Vol. 1, Addison-Wesley, Reading, MA, pp. 46.1–46.9.
- FITZHUGH-R (1961). Impulses and physiological states in theoretical models of nerve membrane, *Biophysical Journal*, **1**, pp. 445–466.
- FOX-R. F (1989a). Numerical simulations of stochastic differential equations, *Journal of Statistical Physics*, **54**(5-6), pp. 1353–1366.
- FOX-R. F (1989b). Stochastic resonance in a double well, *Physical Review A*, **39**(8), pp. 4148–4153.
- FOX-R. F (1991). Second-order algorithm for the numerical integration of color-noise problems, *Physical Review A*, **43**(6), pp. 2649–2654.
- FOX-R. F AND LU-Y (1993). Analytical and numerical study of stochastic resonance, *Physical Review E*, **48**(5), pp. 3390–3398.
- FRANCESCHINI-N, PICHON-J. M AND BLANES-C (1992). From insect vision to robot vision, *Philosophical transactions of the Royal Society of London B*, **337**, pp. 283–294.
- FUH-C AND YEH-Y (2001). Random perturbation in games of chance, *Studies in Applied Mathematics*, **107**, pp. 207–215.
- GAMMAITONI-L (1995). Stochastic resonance and the dithering effect in threshold physical systems, *Physical Review E*, **52**(5), pp. 4691–4698.
- GAMMAITONI-L, HÄNGGI-P, JUNG-P AND MARCHESONI-F (1998). Stochastic resonance, *Reviews of Modern Physics*, **70**(1), pp. 223–287.
- GAMMAITONI-L, MARCHESONI-F, MENICHELLA-SAETTA-E AND SANTUCCI-S (1989a). Stochastic resonance in bistable systems, *Physics Review Letters*, **62**(4), pp. 349–351.
- GAMMAITONI-L, MENICHELLA-SAETTA-E, SANTUCCI-S, MARCHESONI-F AND PRESILLA-C (1989b). Periodically time-modulated bistable systems: Stochastic resonance, *Physical Review A*, **40**(4), pp. 2114–2119.
- GARCÍA-OJALVO-J AND SANCHO-J. M (1999). *Noise in Spatially Extended Systems*, Springer-Verlag.
- GARGAMO-L AND VACCARO-U (1999). Efficient generation of fair dice with few biased coins, *IEEE Transactions on Information Theory*, **45**(5), pp. 1600–1606.
- GILLESPIE-D. T (1996). Exact numerical simulation of the Ornstein-Uhlenbeck process and its integral, *Physical Review E*, **54**(2), pp. 2084–2091.
- GINGL-Z, KISS-L. B AND MOSS-F (1995). Non-dynamical stochastic resonance: Theory and experiments with white and arbitrarily coloured noise, *Europhysics Letters*, **29**(3), pp. 191–196.
- GODIVIER-X AND CHAPEAU-BLONDEAU-F (1997). Noise-assisted signal transmission in a nonlinear electronic comparator: Experiment and theory, *Signal Processing*, **56**(3), pp. 293–303.
- GOYCHUK-I AND HÄNGGI-P (1999). Quantum stochastic resonance in parallel, *New Journal of Physics*, **1**, pp. 14.1–14.14.

- GOYCHUK-I AND HÄNGGI-P (2000). Stochastic resonance in ion channels characterized by information theory, *Physical Review E*, **61**(4), pp. 4272–4280.
- GRUSS-A, CARELY-L. R AND KANADE-T (1991). Integrated sensor and range-finding analog signal processor, *IEEE Journal of Solid-State Circuits*, **26**(3), pp. 184–191.
- HALL-L. T (2001). Monolithic beamforming arrays for passive mm-wave sensing applications, *Ph.D. upgrade report*, Adelaide University, Dept. of Electrical & Electronic Engineering, Adelaide, Australia.
- HÄNGGI-P AND BARTUSSEK-R (1996). Brownian rectifiers: How to convert Brownian motion into directed transport, *Nonlinear Physics of Complex Systems – Current Status and Future Trends, Lecture notes in physics*, Vol. 476, Springer, Berlin, pp. 294–308.
- HARMER-G. P, ABBOTT-D AND TAYLOR-P. G (2000a). The paradox of Parrondo's games, *Proceedings of the Royal Society A*, **456**(1994), pp. 247–260.
- HARMER-G. P, ABBOTT-D, TAYLOR-P. G AND PARRONDO-J. M. R (2000b). Parrondo's paradoxical games and the discrete Brownian ratchet, in D. Abbott and L. B. Kish (eds.), *Second International Conference on Unsolved Problems of Noise and Fluctuations*, Vol. 511, American Institute of Physics, Adelaide, Australia, pp. 189–200.
- HARMER-G. P, ABBOTT-D, TAYLOR-P. G AND PARRONDO-J. M. R (2001). Parrondo's games and Brownian ratchets, *Chaos*, **11**(3), pp. 705–714.
- HARMER-G. P, ABBOTT-D, TAYLOR-P. G, PEARCE-C. E. M AND PARRONDO-J. M. R (2000c). Information entropy and Parrondo's discrete-time ratchet, in D. S. Broomhead, E. Luchinskaya, P. McClintock and T. Mullin (eds.), *Stochastic and Chaotic Dynamics in the Lakes*, Vol. 502, American Institute of Physics, Melville, NY, USA, pp. 544–549.
- HARMER-G. P AND ABBOTT-D (1999a). Parrondo's paradox, *Statistical Science*, **14**(2), pp. 206–213.
- HARMER-G. P AND ABBOTT-D (1999b). Parrondo's paradox: losing strategies cooperate to win, *Nature*, **402**, p. 864.
- HARMER-G. P AND ABBOTT-D (1999c). Simulation of circuits demonstrating stochastic resonance, *Proceedings of SPIE*, Vol. 3893, pp. 195–203. Also appeared in Harmer and Abbott (2000a).
- HARMER-G. P AND ABBOTT-D (2000a). Simulation of circuits demonstrating stochastic resonance, *Microelectronics Journal*, **31**(7), pp. 553–559.
- HARMER-G. P AND ABBOTT-D (2000b). Smart sensor motion detection schemes in a noisy environment, *Proceedings of SPIE*, Vol. 4236, pp. 25–35.
- HARMER-G. P AND ABBOTT-D (2001). Motion detection and stochastic resonance in noisy environments, *Microelectronics Journal*, **32**(12), pp. 959–967.
- HARMER-G. P, DAVIS-B. R AND ABBOTT-D (2002). A review of stochastic resonance: Circuits and measurement, *IEEE Transactions on Instrumentation and Measurement*. Accepted.
- HARRISON-R. R AND KOCH-C (2000). A robust analog VLSI Reichardt motion sensor, *Analog Integrated Circuits and Signal Processing*, **24**, pp. 213–229.

Bibliography

- HARRIS-R. A, O'CARROLL-D. C AND LAUGHLIN-S. B (1999). Adaptation and the temporal delay filter of fly motion detectors, *Vision-Research*, **39**(16), pp. 2603–2613.
- HASSENSTEIN-B AND REICHARDT-W (1956). Systemtheoretische Analyse der Zeit-, Reihenfolgen- und Vorzeichenauswertung bei der Bewegungs-Perzeption des Rüsselkäfers *chlorophanus*, *Z. Naturforsch*, **11b**, pp. 513–524.
- HAUSER-M. W (1991). Principles of oversampling A/D conversion, *Journal of the Audio Engineering Society*, **39**(1/2), pp. 3–26.
- HAYAKAWA-H AND HONG-D. C (1997). Thermodynamical theory of weakly excited granular systems, *Physical Review Letters*, **78**(14), pp. 2764–2767.
- HELSTROM-C. W (1968). *Statistical Theory of Signal Detection*, Pergamon Press.
- HENEGHAN-C, CHOW-C. C, COLLINS-J. J, IMHOFF-T. T, LOWEN-S. B AND TEICH-M. C (1996). Information measures quantifying aperiodic stochastic resonance, *Physical Review E*, **54**(3), pp. 2228–2231.
- HILDRETH-E. C AND KOCH-C (1987). The analysis of visual motion: From computational theory to neuronal mechanics, *Annual Review of Neuroscience*, **10**, pp. 477–533.
- HILGERS-A, GREMM-M AND SCHNAKENBERG-J (1995). A criterion for stochastic resonance, *Physics Letters A*, **209**, pp. 313–316.
- HODGKIN-A. L AND HUXLEY-A. F (1952). A qualitative description of membrane current and its application to conduction and excitation in nerve, *Journal Physics*, **177**, pp. 500–544.
- HOEFFDING-W AND SIMONS-G (1970). Unbiased coin tossing with a biased coin, *Ann. Math. Statist.*, **41**, pp. 341–352.
- HONEYCUTT-R. L (1992a). Stochastic Runge-Kutta algorithms. I. White noise, *Physical Review A*, **45**(2), pp. 600–603.
- HONEYCUTT-R. L (1992b). Stochastic Runge-Kutta algorithms. II. Colored noise, *Physical Review A*, **45**(2), pp. 604–610.
- HORN-B. K. P AND SCHUNCK-B. G (1981). Determining optical flow, *Artificial Intelligence*, **17**, pp. 185–203.
- HORRIDGE-G. A (1990). A template theory to relate visual processing to digital circuitry, *Proceedings of the Royal Society of London B*, **239**, pp. 17–33.
- HORRIDGE-G. A (1992). What can engineers learn from insect vision?, *Philosophical transactions of the Royal Society of London B*, **337**, pp. 271–282.
- HORRIDGE-G. A AND SOBEY-P (1991). An artificial seeing system copying insect vision, *International Journal of Optoelectronics*, **6**(1/2), pp. 177–193.
- HOYLE-F (1957). *The Black Cloud*, Willian Heinemann Ltd, Great Britain.
- INCHIOSA-M. E AND BULSARA-A. R (1995). Nonlinear dynamics elements with noisy sinusoidal forcing: Enhancing response via nonlinear coupling, *Physics Letters E*, **52**(1), pp. 327–339.

- INCHIOSA-M. E AND BULSARA-A. R (1996). Signal detection statistics of stochastic resonators, *Physical Review E*, **53**(3), pp. R2021–R2024.
- ISHIOKA-S AND FUCHIKAMI-N (2001). Thermodynamics of computing: Entropy of nonergodic systems, *Chaos*, **11**(3), pp. 734–746.
- ITOH-T (1996). Simulating fair dice with biased coins, *Information and Computation*, **126**, pp. 78–82.
- JAEGER-H. M, NAGEL-S. R AND BEHRINGER-R. P (1996). Granular solids, liquids and gases, *Reviews of Modern Physics*, **68**(4), pp. 1259–1273.
- JAMES-F (1990). A review of pseudorandom number generators, *Computer Physics Communications*, **60**, pp. 329–344.
- JARAMILLO-F AND WIESENFELD-K (1998). Mechanoelectrical transduction assisted by Brownian motion: a role for noise in the auditory system, *Nature Neuroscience*, **1**(5), pp. 384–388.
- JOSEPHS-H. J (1950). *Heaviside's Electric Circuit Theory*, John Wiley & Sons.
- JUNG-P (1995a). Stochastic resonance and optimal design of threshold detectors, *Physics Letters A*, **207**, pp. 93–104.
- JUNG-P (1995b). Stochastic resonance in threshold devices, *Il Nuovo Cimento*, **17D**(7-8), pp. 827–834.
- JUNG-P AND HÄNGGI-P (1991). Amplification of small signals via stochastic resonance, *Physical Review A*, **44**(12), pp. 8032–8042.
- KÁDÁR-S, WANG-J AND SHOWALTER-K (1998). Noise-supported travelling waves in sub-excitable media, *Nature*, **391**, pp. 770–772.
- KAMEDA-H, ALTMAN-E, KOZAWA-T AND HOSOKAWA-Y (2000). Braess-like paradoxes in distributed computer systems, *IEEE Transactions on Automatic Control*, **45**(9), pp. 1687–1691.
- KARLIN-S AND TAYLOR-H. M (1975). *A First Course in Stochastic Processes*, Academic Press, New York.
- KEENER-J. P (1991). Analog circuitry for the van der Pol and FitzHugh-Nagumo equations, *IEEE Transactions on Systems, Man and Cybernetics*, **13**(5), pp. 1010–1014.
- KEIRSTEAD-W. P AND HUBERMAN-B. A (1986). Collective detection of motion in the presence of noise, *Kybernetik*, **12**, pp. 64–73.
- KELLY-T. R, TELLITU-I AND SESTELO-J. P (1998). New molecular devices: in search of molecular ratchets, *Journal of Organic Chemistry*, **63**, pp. 3655–3665.
- KESTENBAUM-D (1997). Sand castles and cocktail nuts, *New Scientist*, **154**(2083), pp. 25–28.
- KEY-E. S, KŁOSEK-M. M AND ABBOTT-D (2001). On Parrondo's paradox: how to construct unfair games by composing fair games, *Submitted to ANZIAM*.
- KIKKERT-C. J AND BIDDELI-A (1997). Hardware additive dither for analogue to digital converters, *Proceedings of the IREE 14th Australian Microelectronics Conference*, Melbourne, Victoria, pp. 156–161.

Bibliography

- KISH-L. B, HARMER-G. P AND ABBOTT-D (2001). Information transfer rate of neurons: stochastic resonance of Shannon's information channel capacity, *Fluctuation and Noise Letters*, 1(1), pp. L13–L19.
- KISS-L. B (1996). Possible breakthrough: significant improvement of signal to noise ratio by stochastic resonance, in R. Katz (ed.), *Chaotic, Fractal, and Nonlinear Signal Processing*, Vol. 375, American Institute of Physics, Mystic, Connecticut, USA, pp. 382–396.
- KITAMURA-K, TOKUNAGA-M, HIKIKOSHI-A AND YANAGIDA-T (1999). A single myosin head moves along an actin filament with regular steps of 5.3 nanometres, *Nature*, 397, pp. 129–134.
- KLAFTER-J, SHLESINGER-M. F AND ZUMOFEN-G (1996). Beyond Brownian motion, *Physics Today*, 49(2), pp. 33–39.
- KLARREICH-E (2001). Playing both sides, *The Sciences*, 41(1), pp. 25–29.
- KLAY-M. P AND FOULIS-D. J (1990). Maximum likelihood estimation on generalized sample spaces: an alternative resolution of Simpson's paradox, *Foundations of Physics*, 20(7), pp. 777–799.
- KNUTH-D. E (1969). *The Art of Computer Programming*, Vol. 2, second edn, Addison-Wesley Publishing Company, Reading, Massachusetts.
- KOCH-C (1998). *Biophysics of Computation: Information Processing in Single Neurons*, Oxford University Press.
- KOIVUNEN-T (1992). A noise-insensitive motion detector, *IEEE Transactions on Consumer Electronics*, 38(3), pp. 168–174.
- KORILIS-Y. A, LAZAR-A. A AND ORDA-A (1999). Avoiding the Braess paradox in non-cooperative networks, *Journal of Applied Probability*, 36(1), pp. 211–222.
- KREYSZIG-E (1953). *Advanced Mathematics for Engineers*, John Wiley & Sons, Inc., New York.
- LANDAUER-R (1961). Irreversibility and heat generation in the computing process, *IBM Journal of Research and Development*, 5, pp. 183–91. Reprinted in Leff and Rex (1990).
- LANZARA-E, MANTEGNA-R. N, SPAGNOLO-B AND ZANGARA-R (1997). Experimental study of a non-linear system in the presence of noise: The stochastic resonance, *American Journal of Physics*, 65(4), pp. 341–349.
- LAVENDA-B. H (1985). Brownian motion, *Scientific American*, 252, pp. 56–67.
- LAVITAN-I. B AND KACZMAREK-L. K (1997). *The Neuron: Cell and Molecular Biology*, second edn, Oxford University Press, New York.
- LAWRENCE-A (2001). Millimeter-wave ICs open up new spectrum, *Compound Semiconductor*, 7(4), pp. 1–5. http://www.compoundsemiconductor.net/Archive_7-4.htm.
- LEE-D. N (1976). A theory of visual control of braking based on information about time-to-collision, *Perception*, 5, pp. 437–459.
- LEE-D. N (1980). The optic flow field: the foundation of vision, *Philosophical Transactions of the Royal Society of London B*, 290, pp. 169–179.

- LEFF-H. S AND REX-A. F (1990). *Maxwell's Demon: Entropy, Information and Computing*, IOP Publishing Ltd, Great Britain.
- LEMONS-D. S AND GYTHIEL-A (1997). Paul Langevin's 1908 paper "On the theory of Brownian motion" ["Sur la théorie du mouvement brownien," C. R. Acad. Sci. (Paris) 146, 530–533 (1908)], *American Journal of Physics*, **65**(11), pp. 1079–1081.
- LEONARD-D. S (1992). Stochastic resonance in a random walk, *Physical Review A*, **46**(10), pp. 337–339.
- LINKE-H, HUMPHREY-T. E, L OFGREN-A, SUSHKOV-A. O, NEWBURY-R, TAYLOR-R. P AND OMLING-P (1999). Experimental tunneling ratchets, *Science*, **286**, pp. 2314–2317.
- LOERINCZ-K, GINGL-Z AND KISS-L. B (1996). A stochastic resonator is able to greatly improve signal-to-noise ratio, *Physics Letters A*, **224**(1-2), pp. 63–67.
- LONGTIN-A (1993). Stochastic resonance in neuron models, *Journal of Statistical Physics*, **70**(1/2), pp. 309–327.
- LONGTIN-A (1995). Synchronization of the stochastic Fitzhugh-Nagumo equations to periodic forcing, *Il Nuovo Cimento*, **17D**(7-8), pp. 835–846.
- LONGTIN-A, BULSARA-A AND MOSS-F (1991). Time-interval sequences in bistable systems and the noise-induced transmission on information by sensory neurons, *Physical Review Letters*, **67**(5), pp. 656–659.
- LUCHINSKY-D. G, MANNELLA-R, MCCLINTOCK-P. V. E AND STOCKS-N. G (1999a). Stochastic resonance in electrical circuits—I: Conventional stochastic resonance, *IEEE Transactions on Circuits and Systems—II: Analog and Digital Signal Processing*, **46**(9), pp. 1205–1214.
- LUCHINSKY-D. G, MANNELLA-R, MCCLINTOCK-P. V. E AND STOCKS-N. G (1999b). Stochastic resonance in electrical circuits—II: Nonconventional stochastic resonance, *IEEE Transactions on Circuits and Systems—II: Analog and Digital Signal Processing*, **46**(9), pp. 1215–1224.
- MAGNASCO-M. O (1993). Forced thermal ratchets, *Physical Review Letters*, **71**(10), pp. 1477–1481.
- MAGNASCO-M. O (1994). Molecular combustion motors, *Physical Review Letters*, **72**(16), pp. 2656–2659.
- MAGNASCO-M. O (1996). Szilard's heat engine, *Europhysics Letters*, **33**(8), pp. 583–588.
- MAKSE-H. A, HAVLIN-S, KING-P. R AND STANLEY-H. E (1997). Spontaneous stratification in granular mixtures, *Nature*, **386**, pp. 379–382.
- MALNIKOV-V. I (1993). Schmitt trigger: A solvable model of stochastic resonance, *Physical Review E*, **48**(4), pp. 2481–2489.
- MANNELLA-R (1989). Computer experiments in non-linear stochastic physics, in F. Moss and P. V. E. McClintock (eds.), *Noise in nonlinear systems: Experiments and simulations*, Vol. 3, Cambridge University Press, pp. 189–221.
- MANNELLA-R AND PALLESCHI-V (1989). Fast and precise algorithm for computer simulation of stochastic differential equations, *Physical Review A*, **40**(6), pp. 3381–3386.

Bibliography

- MANTEGNA-R. N AND SPAGNOLO-B (1994). Stochastic resonance in a tunnel diode, *Physical Review E*, **49**(3), pp. R1792–1795.
- MANTEGNA-R. N AND SPAGNOLO-B (1995). Stochastic resonance in a tunnel diode in the presence of white noise or coloured noise, *Nuovo Cimento*, **17D**(7-8), pp. 873–881.
- MANTEGNA-R. N AND SPAGNOLO-B (1996). Noise enhanced stability in an unstable system, *Physical Review Letters*, **76**(4), pp. 563–566.
- MAREK-A (1959). A note to recent theories of Brownian motion in non-linear systems, *Physica*, **25**, pp. 1358–1367.
- MARR-D AND HILDRETH-E (1980). Theory of edge detection, *Proceedings of the Royal Society of London B*, **207**, pp. 187–217.
- MARSAGLIA-G AND TSANG-W. W (1984). A fast, easily implemented method for sampling from decreasing or symmetric unimodal density functions, *SIAM Journal Science and Statistical Computation*, **5**(2), pp. 349–359.
- MAXWELL-J. C (1885). *Theory of Heat*, eighth edn, Longmans, Green, and Co., chapter 12, pp. 328–329.
- MCFEE-R (1971). Self-rectification in diodes and the second law of thermodynamics, *American Journal of Physics*, **39**(7), pp. 841–820.
- MCNAMARA-B AND WIESENFELD-K (1989). Theory of stochastic resonance, *Physical Review A*, **39**(9), pp. 4854–4869.
- MCNAMARA-B, WIESENFELD-K AND ROY-R (1988). Observation of stochastic resonance in a ring laser, *Physical Review Letters*, **60**(25), pp. 223–287.
- MEAD-C (1989). *Analog VLSI and Neural Systems*, Addison-Wesley Publishing Company.
- MEYER-D. A AND BLUMER-H (2001). Parrondo games as a lattice gas automata, *Ninth International Conference on Discrete Models in Fluid Mechanics*, Santa Fe, NM, USA. Preprint quant-ph/00110028, <http://xxx.lanl.gov>.
- MITAIM-S AND KOSKO-B (1998). Adaptive stochastic resonance, *Proceedings of the IEEE*, **86**(11), pp. 2152–2183.
- MOINI-A (1997). *Design of a VLSI motion detector based upon the insect visual system*, Masters thesis, University of Adelaide.
- MOINI-A, BOUZERDOUM-A, ESHRAGHIAN-K, YAKOVLEFF-A, NGUYEN-X. T, BLANKSBY-A, BEARE-R, ABBOTT-D AND BOGNER-R. E (1997). An insect vision-based motion detection chip, *IEEE Journal of Solid State Circuits*, **32**(2), pp. 279–283.
- MORAAL-H (2000). Counterintuitive behaviour in games based on spin models, *Journal of Physics A*, **33**(23), pp. L203–L206.
- MOSS-F AND WIESENFELD-K (1995). The benefits of background noise, *Scientific American*, pp. 50–53.
- MOSS-F, CHIOU-TAN-F AND KLINKE-R (1996). Will there be noise in their ears?, *Nature Medicine*, **2**(8), pp. 860–862.

- MOSTELLER-F (1965). *Fifty Challenging Problems in Probability*, Addison-Wesley, Reading, MA, USA.
- NAFIE-M AND TEWFIK-A. H (1998). Low power detection using stochastic resonance, *Thirty-Second Asilomar Conference on Signals, Systems and Computers*, IEEE, Piscataway, NJ, USA, pp. 1461–1465.
- NAHIN-P. J (2000). *Duelling Idiots and other Probability Puzzlers*, Princeton University Press.
- NAKAYAMA-K (1985). Biological image motion processing: A review, *Vision Research*, **25**(5), pp. 625–660.
- NEIMAN-A, SCHIMANSKY-GEIER-L AND MOSS-F (1997). Linear response theory applied to stochastic resonance in models of ensembles of oscillators, *Physical Review E*, **56**(1), pp. 9–12.
- NETRAVALI-A. N AND HASKELL-B. G (1988). *Digital pictures: Representation and compression*, Plenum Press, New York, USA.
- NEUFELD-E (1995). Simpson’s paradox in artificial intelligence and in real life, *Computational Intelligence*, **11**(1), pp. 1–10.
- NEUMANN-J. V (1944). *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, USA.
- NEUTS-M. F (1981). *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*, The John Hopkins University Press, USA.
- NG-J AND ABBOTT-D (2002). Introduction to quantum games and a quantum Parrondo game, in A. Nowack (ed.), *Annals of the International Society on Dynamic Games*, Birkhäuser. Submitted.
- NGUYEN-X. T (1996). *Smart VLSI micro-sensors for velocity estimation inspired by insect vision*, PhD thesis, University of Adelaide.
- NGUYEN-X. T, BOUZERDOUM-A, BOGNER-R. E, MOINI-A AND ESHRAGHIAN-K (1995). Feature representation of motion trajectories, *Proceedings of the IEEE International Conference of Neural Networks*, Vol. 6, Perth, WA, Australia, pp. 2922–2927.
- NICOLIS-C (1982). Stochastic aspects of climate transitions response to a periodic forcing, *Tellus*, **1**(34), pp. 1–9.
- NIELSEN-M. A AND CHUANG-I. L (2000). *Quantum Computation and Quantum Information*, Cambridge University Press, Cambridge, United Kingdom.
- NORRIS-J. R (1997). *Markov Chains*, Cambridge University Press.
- NUTTALL-A. N (2001). Private communication, Naval Underwater Warfare Centre (NUWC), Rhode Island, USA.
- O’CARROLL-D. C, LAUGHLIN-S. B, BIDWELL-N. J AND HARRIS-R. A (1997). Spatio-temporal properties of motion detectors matched to low image velocities in hovering insects, *Vision-Research*, **37**(23), pp. 3427–3439.
- ONSAGER-L (1931a). Reciprocal relations in irreversible processes I, *Physical Review*, **37**, pp. 405–426.
- ONSAGER-L (1931b). Reciprocal relations in irreversible processes II, *Physical Review*, **38**, p. 2265.

Bibliography

- PARRONDO-J. M. R (2001). The Szilard engine revisited: Entropy, macroscopic randomness, and symmetry breaking phase transitions, *Chaos*, **11**(3), pp. 725–733.
- PARRONDO-J. M. R AND ESPAÑOL-P (1996). Criticism of Feynman's analysis of the ratchet as an engine, *American Journal of Physics*, **64**(9), pp. 1125–1130.
- PARRONDO-J. M. R, HARMER-G. P AND ABBOTT-D (2000). New paradoxical games based on Brownian ratchets, *Physical Review Letters*, **85**(24), pp. 5226–5229.
- PEARCE-C. E. M (2000a). Entropy, Markov information sources and Parrondo games, in D. Abbott and L. B. Kish (eds.), *Second International Conference on Unsolved Problems of Noise and Fluctuations*, American Institute of Physics, Adelaide, Australia, pp. 207–212.
- PEARCE-C. E. M (2000b). On Parrondo paradoxical games, in D. Abbott and L. B. Kish (eds.), *Second International Conference on Unsolved Problems of Noise and Fluctuations*, American Institute of Physics, Adelaide, Australia, pp. 201–206.
- PERELMAN-Y. I (1967). *Zhivaya Matematika*, Nauka, Moscow. Reissue of the 1934 edition.
- PESKIN-C. S AND OSTER-G (1995). Coordinated hydrolysis explains the mechanical behavior of kinesin, *Biophysical Journal*, **68**(4), pp. 202–211.
- PINSKY-R AND SCHEUTZOW-M (1992). Some remarks and examples concerning the transient and recurrence of random diffusions, *Annales de l'Institut Henri Poincaré - Probabilités et Statistiques*, **28**(4), pp. 519–536.
- PINTER-R. B (1983). The electrophysiological bases for linear and for nonlinear product term lateral inhibition and the consequences for wide field textured stimuli, *Journal of Theoretical Biology*, **105**, pp. 233–243.
- PLANAT-M (2001). $1/f$ noise, the measurement of time and number theory, *Fluctuation and Noise Letters*, **1**(1), pp. R65–R80.
- POGGIO-T, TORRE-V AND KOCH-C (1985). Computational vision and regularization theory, *Nature*, **317**, pp. 314–319.
- PYKE-R (2001). On random walks related to Parrondo's games, *University of Washington, Seattle*. Unpublished.
- REICHARDT-W (1961). Autocorrelation, a principle for the evaluation of sensory information by the central nervous system, in W. A. Rosenblith (ed.), *Sensory Communication*, Contributions to the Symposium on principles of sensory communication, The MIT Press, Endicott House, MIT, pp. 303–317.
- REICHARDT-W (1987). Evaluation of optical motion information by movement detectors, *Journal of Comparative Physiology A*, **161**, pp. 533–547.
- REIF-R (1985). *Fundamentals of statistical and thermal physics*, McGraw-Hill Book Company, Singapore.
- REIMANN-P, GRIFONI-M AND HÄNGGI-P (1997). Quantum ratchets, *Physical Review Letters*, **79**(1), pp. 10–13.

- RIANI-M AND SIMONOTTO-E (1994). Stochastic resonance in the perceptual interpretation of ambiguous figures: A neural network model, *Physical Review Letters*, **72**(19), pp. 3120–3123.
- RIANI-M AND SIMONOTTO-E (1995). Periodic perturbation of ambiguous figure: a neural-network model and a nonsimulated experiment, *Nuovo Cimento*, **17D**(7-8), pp. 903–913.
- ROCHE-J. R (1992). Efficient generation of random variables from biased coins, *Bell Tech. Rep., AT&T Lab., File Case 20878*.
- RONCAGLIA-R AND TSIRONIS-G. P (1998). Discrete quantum motors, *Physical Review Letters*, **81**(1), pp. 10–13.
- ROSATO-A, STRANDBURG-K. J, PRINZ-F AND SWENDSEN-R. H (1987). Why the Brazil nuts are on top: Size segregation of particulate matter by shaking, *Physical Review Letters*, **58**(10), pp. 1038–1040.
- SANCHO-J. M, MIGUEL-M. S, KATZ-S. L AND GUNTON-J. D (1982). Analytical and numerical studies of multiplicative noise, *Physical Review A*, **26**(3), pp. 1589–1609.
- SEZGIN-M, BIRECIK-S, DEMIR-D, BUCAK-İ. Ö, ÇETIN-S AND KURUGÖLLÜ-F (1995). A comparison of visual target tracking methods in noisy environments, *Proceedings of the 1995 IEEE IECON: 21st International Conference on Industrial Electronics, Control, and Instrumentation*, Vol. 2, New York, NY, USA, pp. 1360–1365.
- SHANNON-C. E (1949). Communication in the presence of noise, *Proceedings of the IRE*, **37**(1), pp. 10–21.
- SHANNON-C. E (1951). Prediction and entropy of printed English, *The Bell System Technical Journal*, **30**, pp. 50–64.
- SHANNON-C. E AND WEAVER-W (1949). *The mathematical theory of communication*, The University of Illinois Press.
- SHLESINGER-M. F (1996). A brief history of random processes, in C. R. Doering, L. B. Kiss and M. F. Shlesinger (eds.), *Proceedings of the First International Conference on Unsolved Problems of Noise*, World Scientific, Szeged, Hungary, pp. 3–10.
- SHLESINGER-M. F, KLEFTER-J AND ZUMOFEN-G (1999). Above, below and beyond Brownian motion, *American Journal of Physics*, **67**(12), pp. 1253–1259.
- SIMONOTTO-E, RIANI-M, CHARLES-S, ROBERTS-M, TWITTY-J AND MOSS-F (1997). Visual perception of stochastic resonance, *Physical Review Letters*, **78**(6), pp. 1186–1189.
- SOANE-D. S AND SOANE-Z. M (1992). *Method and device for moving molecules by the application of a plurality of electrical fields*, number 5 126 022, United States Patent.
- SOKOLOV-I. M (1998). On the energetics of a nonlinear system rectifying thermal fluctuations, *Europhysics Letters*, **44**(3), pp. 278–283.
- SRINIVASAN-M. V, CHAHL-J. S, NAGLE-M. G AND ZHANG-S. W (1997a). Embodying natural vision into machines, in M. V. Srinivasan and S. Venkatesh (eds.), *From living eyes to seeing machines*, Oxford University Press, Oxford, UK, pp. 249–263.

Bibliography

- SRINIVASAN-M. V, ZHANG-S. W AND BIDWELL-N. J (1997b). Visually mediated odometry in honeybees, *The Journal of Experimental Biology*, **200**, pp. 2513–2522.
- SRINIVASAN-M. V, ZHANG-S. W, LEHRER-M AND COLLETT-T. S (1996). Honeybee navigation *en route* to the goal: Visual flight control and odometry, *The Journal of Experimental Biology*, **199**, pp. 237–244.
- STEWART-R. W AND PFANN-E (1998). Oversampling and sigma-delta strategies for data conversion, *Electronics & Communication Engineering Journal*, **10**(1), pp. 37–47.
- STOCKS-N. G (2000a). Private Communication.
- STOCKS-N. G (2000b). Optimising information transmission in model neuronal ensembles, in J. A. Freund and T. Poschel (eds.), *Stochastic Processing in Physics, chemistry and Biology*, Vol. 557 of *Lecture Notes in Physics*, Springer-Verlag, Berlin, pp. 150–159.
- STOCKS-N. G (2000c). Suprathreshold stochastic resonance, in D. S. Broomhead, E. Luchinskaya, P. McClintock and T. Mullin (eds.), *Stochastic and Chaotic Dynamics in the Lakes*, Vol. 502, American Institute of Physics, Ambleside, UK, pp. 415–421.
- STOCKS-N. G (2000d). Suprathreshold stochastic resonance in multilevel threshold systems, *Physics Review Letters*, **84**(11), pp. 2310–2313.
- STOCKS-N. G (2001a). Private Communication.
- STOCKS-N. G (2001b). Information transmission in parallel threshold arrays: Suprathreshold stochastic resonance, *Physical Review E*, **63**(041114), pp. 1–9.
- STOCKS-N. G (2001c). Suprathreshold stochastic resonance: an exact result for uniformly distributed signal and noise, *Physics Letters A*, **279**, pp. 308–312.
- STOCKS-N. G AND MANNELLA-R (2000). Suprathreshold stochastic resonance in a neuronal network model: A possible strategy for sensory coding, in N. Kosabov (ed.), *Future Directions for Intelligent Systems and Information Sciences*, Physica-Verlag, pp. 236–247.
- STOCKS-N. G, STEIN-N. D AND MCCLINTOCK-P. V. E (1993). Stochastic resonance in monostable systems, *Journal of Physics A*, **26**, pp. L385–L390.
- STOCKS-N. G, STEIN-N. D, SHORT-H. E, MANNELLA-R, LUCHINSKY-D. G, MCCLINTOCK-P. V. E AND DYKMAN-M. I (1996). Noise-induced linearization and delinearization, in M. Millonas (ed.), *Fluctuations and Order: The New Synthesis*, Springer-Verlag, New York, NY, USA, pp. 51–67.
- STRONG-S. P, KOBERLE-R, DE RUYTER VAN STEVENINCK-R. R AND BIALEK-W (1998). Entropy and information in neural spike trains, *Physical Review Letters*, **80**(1), pp. 197–200.
- SVOBODA-K, SCHMIDT-C. F, SCHNAPP-B. J AND BLOCK-S. M (1993). Direct observation of kinesin stepping by optical trapping interferometry, *Nature*, **365**, pp. 721–727.
- SZILARD-L (1929). Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen, *Zeitschrift für Physik*, **53**, pp. 840–856. English translation reprinted in Leff and Rex (1990).
- THE MCGRAW-HILL COMPANIES (2001). <http://www.mhhe.com/socscience/intro/ibank/set1.htm>.

- TIPLER-P. A (1991). *Physics for Scientists and Engineers*, third edn, Worth Publishers Inc., New York, NY, USA.
- TORAL-R (2001). Cooperative Parrondo's games, *Fluctuation and Noise Letters*, 1(1), pp. L7-L12.
- TORAL-R AND CHAKRABARTI-A (1993). Generation of Gaussian distributed random numbers by using a numerical inversion method, *Computer Physics Communications*, 74, pp. 327-334.
- TORRE-V AND POGGIO-T (1978). A synaptic mechanism possibly underlying directional selectivity to motion, *Proceedings of the Royal Society of London B*, 202, pp. 409-416.
- TRAVIS-J (1995). Making light work of Brownian motion, *Science*, 267, pp. 1593-1594.
- ULLMAN-S (1981). Analysis of visual motion by biological and computer systems, *Computer*, 14(8), pp. 57-68.
- ULLMAN-S (1983). The measurement of visual motion, *Trends in neurosciences*, 6, pp. 177-179.
- URICK (1967). *Sonar Signal Processing*, Cambridge University Press.
- VAN DEN BROECK-C, REIMANN-P, KAWAI-R AND HÄNGGI-P (1999). Coupled Brownian motors, in D. Reguera, M. Rubí and J. M. G. Vilar (eds.), *Lecture Notes in Physics: Statistical Mechanics of Biocomplexity*, Vol. 527, Springer-Verlag: Berlin, Heidelberg, New York, pp. 93-111.
- VANDERKOOY-J AND LIPSHITZ-S. P (1984). Resolution below the least significant bit in digital systems with dither, *Journal of the Audio Engineering Society*, 32(3), pp. 106-113.
- VAN SANTENA-J. P. H AND SPERLING-G (1985). Elaborated Reichardt detectors, *Journal of the Optical Society of America A*, 2(2), pp. 300-321.
- VILAR-J. M. G AND RUBÍ-J. M (1997). Stochastic multiresonance, *Physical Review Letters*, 78(15), pp. 2882-2885.
- VON NEUMANN-J (1951). Various techniques used in connection with random digits, *Applied Math Ser.*, 12, pp. 36-38.
- VON SMOLUCHOWSKI-M (1906). Zur kinetischen Theorie der Brownschen Molecularbewegung und der Suspensionen, *Annalen der Physik*, 21, pp. 756-780.
- VON SMOLUCHOWSKI-M (1912). Experimentall nachweisbare, der üblichen Thermodynamic widersprechende Molekularphänomene, *Physikalische Zeitschrift*, XIII(21-22), pp. 1069-1080.
- WACKER-A AND SCHÖLL-E (1995). Criteria for stability in bistable electrical devices with S- or Z-shaped current voltage characteristic, *Journal of Applied Physics*, 78(12), pp. 7352-7357.
- WHITESIDE-T. C. D AND SAMUEL-G. D (1970). Blur zone, *Nature*, 225, pp. 94-95.
- WIESENFELD-K AND JARAMILLO-F (1998). Minireview of stochastic resonance, *Chaos*, 8(3), pp. 539-548.
- WIESENFELD-K AND MOSS-F (1995). Stochastic resonance and the benefits of noise: From ice ages to crayfish and SQUIDS, *Nature*, 373, pp. 33-36.
- YAKOVLEFF-A (1995). *Architectural studies for visual processing*, PhD thesis, University of Adelaide.

Bibliography

- YAKOVLEFF-A AND MOINI-A (1996). From sensing to perceiving: An overview of the bugeye project, *Advanced Focal Plane Arrays and Electronic Cameras*, Berlin, Germany, pp. 86–96.
- YATES-R. D AND GOODMAN-D. J (1999). *Probability and Stochastic Processes: A Friendly Introduction for Electrical and Computer Engineers*, John Wiley & Sons.
- ZHOU-T, MOSS-F AND JUNG-P (1990). Escape-time distributions of a periodically modulated bistable system with noise, *Physical Review A*, **42**(6), pp. 3161–3169.
- ZOZOR-S AND AMBLARD-P (1999). Stochastic resonance in discrete time nonlinear AR(1) models, *IEEE Transactions on Signal Processing*, **47**(1), pp. 108–122.
- ZOZOR-S AND AMBLARD-P (2001). Erratum: Stochastic resonance in discrete time nonlinear AR(1) models, *IEEE Transactions on Signal Processing*, **49**(5), pp. 1107–1109.

Glossary

ADC	Analog to Digital Converter, 92
AMI	Average Mutual Information, 134
ASR	Aperiodic Stochastic Resonance, 129
BPF	BandPass Filter, 158
CCD	Charge Coupled Device, 4
CMOS	Complimentary Metal On Silicon, 4
dB	decibel, 125
DSLIMD	Directionally Selective Local Inhibitory Motion Detector, 165
DTMC	Discrete-Time Markov-Chain, 54
DWP	Double Well Potential, 116
EMD	Elementary Motion Detector, 154
FHN	FitzHugh-Nagumo neuron, 114
HPF	HighPass Filter, 223
ISHTAR	Improved Sonar Harnessing Temporal Acoustic Resonances, 104
LCC	Level Crossing Circuit, 110
LCD	Level Crossing Detector, 111
LIMD	Local Inhibitory Motion Detector, 161
lofargram	LOW Frequency And Ranging gram – a time-frequency plot used by the sonar community, 108
LPF	LowPass Filter, 156
LR	Likelihood Ratio, 107
OU	Ornstein-Uhlenbeck, 121
PDF	Probability Density Function, 45
PMMW	Passive MilliMetre Wave, 8
PSD	Power Spectral Density, 95
QBD	Quasi-Birth-and-Death, 210
RMD	Reichardt Motion Detector, 156
RNG	Random Number Generator, 213
ROC	Receiver-Operating-Characteristic, 104
SDE	Stochastic Differential Equation, 121
SNR	Signal-to-Noise Ratio, 93
SR	Stochastic Resonance, 91
SSR	Suprathreshold Stochastic Resonance, 135
STR	Stochastic Resonator, 94

Index

- action potential, 115
- ADC, 101, 135
- aperiodic stochastic resonance, 128, 130
- average mutual information, 132

- bandpass filter, 158, 184, 226
 - adaptive, 164
- Barlow and Levick's inhibitory mechanism, 160
- benchmark signal, 172
- biasing parameter, 40
- birth-and-death process, 55
- bistable systems, 115
- bona fide resonance, 96, 127
- Box-Müller algorithm, 214
- Brown, Robert, 18
- Brownian Motion, 18
- Brownian ratchets, 12, 17
 - linearised, 28
- bugeye project, 4

- channel capacity, 139
 - Shannon's formula, 140
- Co-operative Parrondo's games, 83
- coherence, 94
- coloured noise, 121
- crayfish mechanoreceptors, 98
- cross correlation coefficient, 128, 182
- cut-off frequency, 157

- delay element, 156, 223
- detailed balance, 60
- detection probability, 104
- detection window, 158
- discrete-time Markov chains, 54
- dithering, 100
 - additive, 103
 - c.f. stochastic resonance, 104
- double well potential, 116, 144
- DSLIMD, 165
 - noise robustness, 184
- dynamic range, 102

- Einstein, Albert, 19
- electromagnetic spectrum, 7
- elementary motion detector, 154
- entropy, 134
- equilibrium distribution, 56, 62, 79
- escape rate, 127
- excitatory mechanism, 157, 160

- fairness, 41
- filters, **223**
- FitzHugh-Nagumo neuron, **114**, 130, 146
- flashing ratchet, 30
- fractal properties, 83

- Gaussian noise, 120
- gradient schemes, 172

- Harmer, Gregory, 253
- Heaviside function, 110, 133
- Heaviside, Oliver, 224
- highpass filter, 223
- history dependent games, 74
- Horridge template model, 167
 - look-up table, 170
 - noise robustness, 185
- Horridge's template model
 - parallel configurations, 179
- hysteresis, 118

- inhibitory mechanism, 160
 - adaptive properties, 163
- interspike interval, 127
- interspike interval times, 127
- ISHTAR, 104
 - discussion, 107

- Langevin, Paul, 19
- laws of thermodynamics, 20

- level crossing circuit, 110
- level crossing detector, 111
- lowpass filter, 157, 223
- martingale, 43
- Maxwell's demon, 24
 - pressure, 26
 - temperature, 25
- Maxwell, James, 24
- millimetre waves, 8
 - antenna array, 9
 - atmospheric absorption, 10
 - passive imaging, 9
- mixing parameter, 48
- molecular motor, 33
- motion detection, 149
 - biological schemes, 154
 - criteria, 153
 - noise performance, 175
 - requirements, 153
- motion intensity, 173
- neuron models, 111
 - FitzHugh-Nagumo, 114
 - Hodgkin-Huxley, 114
 - integrate-and-fire, 113
- noise, 120
- noise induced linearisation, 144
- nonlinear systems, 109
- null direction, 155
- Ornstein-Uhlenbeck process, 121
- oversampling, 103, 193
- Parrondo's games, 37
 - analogy to ratchet, 53
 - analysis, 54
 - constraints, 63
 - construction, 39
 - distributions, 45, 62
 - entropy, 69
 - explanation, 49
 - fairness, 41
 - model as DTMCs, 57
 - playing, 41, 61
- Parrondo, Juan, 38
- perception, 150
- phase space, 145
- power spectral density, 95
- preferred direction, 155
- probability of a false alarm, 104
- probability space, 65, 81
- quartic potential, 117
- radiometry, 8
- random numbers, 122, 213
- randomised game, 47, 60
- range, 5
- ratchet and pawl machine, 21
 - engine, 24
- ratchet examples, 12, 31
- rate of return, 67
- RCARRY algorithm, 213
- receiver-operating-characteristic, 104
- recurrence, 57
- Reichardt detector, 154
 - noise robustness, 183
- residence time, 127
- Runge-Kutta algorithm, 215
- S-type characteristic, 119
- Schmitt trigger, 118
- second law of thermodynamics, 20
- sensing, 150
- shunting inhibition, 161
- shunting neuron, 161
- signal-to-noise ratio, 95, 118, 125, 130
- signal-to-threshold distance, 96
- skip-free, 55
- smart sensors, 4, 150
- state classification, 57
- state probabilities, 56
- state space, 85
- stationary probabilities, 56
- stochastic resonance, 13, 91
 - basics, 93

- bona fide resonance, 127
 - examples, 98–106
 - quantifying, 124
 - systems, 109
 - weak signal detection, 104
- stochastic resonator, 94
- subthreshold, 93
- suprathreshold, 96, 135
- suprathreshold stochastic resonance, 135
 - threshold placement, 137

- templates, 168
- test signal, 156, 181
- thermal noise, 19
- thesis layout, 2
- threshold devices, 110
- time constant, 157
- time-to-collision, 5, 171
- transience, 57
- transition matrix, 55
- transition probability, 55

- visual perception, 100

- white noise, 120

Resume



Greg Harmer graduated from the University of Adelaide with the B.Sc. (Applied Maths & Computer Science) degree in 1996 and the B.E. (Electrical and Electronic, Hons I) degree in 1997. He is currently undertaking a Ph.D. in the Department of Electrical and Electronic Engineering at the University of Adelaide. He has written a number of journal and conference papers, including a book chapter and was an invited speaker at UPoN'99, Adelaide, Australia.



"Mr. Osborne, may I be excused? My brain is full."

© Gary Larson

Erratum

- In the caption for Figure 4.9, the words 'detection' and 'false alarm' need to be switched.
- At the end of the paragraph following (4.5), $\sigma_s^2 = 4\sigma_n^2$ should be $\sigma_s^2 = 3\sigma_n^2$.
- At the end of the caption for Figure 4.10, add the following sentence. 'However, there is limited value of increasing $p(D)$ to 20% (which is generally unacceptable), at the expense of degrading the capability in the higher $p(D)$ range.'
- Add the following sentence to the end of the caption of Figure 4.11. 'Note that this is inconsistent with Dawe's assumptions in the theoretical treatment.'
- The last word in the paragraph at the top of page 109 should be 'tonals' instead of 'narrowband'.