

Phase Transitions in the Random Stirring Process on Infinite Graphs

by

Owen Daniel


Dissertation

Submitted for the degree of

Master of Science

Mathematics Institute
The University of Warwick

August 2012



Contents

Acknowledgements	iii
Declarations	iv
Abstract	v
Introduction	vi
List of Notation	x
Chapter 1 Quantum Mechanics of the Bose Gas	1
1.1 Mathematical Formulation of Quantum Mechanics	1
1.1.1 Feynman–Kac Formulae	2
1.1.2 The Ideal Bose Gas	5
1.2 The Spin- $\frac{1}{2}$ Heisenberg Ferromagnet	7
1.2.1 Equivalence to the Bose Gas	9
1.2.2 A Return to Feynman–Kac and the Random Stirring Representation . . .	12
1.2.3 Thermodynamic Limit and Phases	17
Chapter 2 The Random Stirring Process	19
2.1 Random Stirring Representations	19
2.1.1 The Interchange Process	19
2.1.2 Random Walk on the Symmetric Group	20
2.1.3 Coagulation-Fragmentation of Permutations	21
2.1.4 The Cyclic Time Random Walk	23
2.1.5 The RSP on an Infinite Graph	26
2.2 General Techniques and Results	27
2.2.1 Cycle Sizes and the Projection Graph	28
2.2.2 Transience and Recurrence of the CTRW	35
Chapter 3 The Critical Phase for Trees	39
3.1 T_c^I and a Critical Window	39
3.2 T_c^S and the Sharp-Threshold Conjecture	47

3.2.1	Analysis for Large T	47
3.2.2	Monotonicity around T_c^I	49
Chapter 4 Conclusion		52
4.1	Random Stirring on Finite Graphs	52
4.1.1	The Complete Graph and Poisson–Dirichlet Distribution	53
4.1.2	Representation Theory for the RSP	54
4.1.3	The Spectral Gap of the RSP (Aldous’s Conjecture)	55
4.2	Questions, Conjectures and the Future	56
Appendix A Hazard Functions		58

Acknowledgements

Once again I would like to thank Stefan Adams for his continuing support of my work and for his enthusiasm and knowledge. Stefan also deserves acknowledgement for recognising that every once in a while it is a good idea to go for a goodnight beer. I wish to thank Omer Angel, Alan Hammond, Daniel Ueltschi, and Peter Windridge for insightful conversations, and for sharing their thoughts on where this project could lead.

I thank Simon Bignold for driving me home in the snow, without which this dissertation would not have been possible.

Declarations

The work in this dissertation is largely a review of previous publications. To the best of the author's knowledge all results are correctly referenced. Many results are proven which are implicitly assumed in other papers, in these cases the proofs are the work of the author. For such results the proofs are not credited, though we often reference a paper where the result is used. Where we have reworked proofs from previous publications we write (for instance): "we follow [Ang03], p.13".

The major sources for this work are [Ang03, Tót93].

Abstract

We survey the random stirring process on infinite graphs, motivated by the study of the Heisenberg spin $1/2$ ferromagnet. We explain how infinite cycles in the random stirring process relate to thermodynamic functions for the Heisenberg model, and in turn how these relate to the study of Bose–Einstein condensation. We go on to study specific properties of random stirring and to derive alternative descriptions for the model, including Angel’s cyclic time random walk. This is used to show Angel’s proof that the random stirring process has a critical phase for regular trees.

Introduction

Before embarking on our journey, we take a few lines to describe both the physical motivation behind our work, and the purely probabilistic appeal of the random stirring process.

Physical Motivation

The aim of statistical mechanics is to study mathematical models of atomic systems. Having established a model, one looks to describe collective phenomena including phase transitions, critical scaling and fluctuation limit laws. The study of *classical* statistical mechanics has proven fruitful, and a powerful arsenal of broadly applicable tools have been developed, including Pirogov–Sinai theory, percolation techniques, and the random walk representation (to name but a few). On the other hand, quantum models are still ill-understood, and methods that work well in a classical setting prove less applicable in a quantum environment.

Phase Transition in the Bose Gas

One particularly interesting problem from quantum mechanics is the occurrence of Bose–Einstein condensation. In the 1920s, Satyendra Nath Bose and Albert Einstein conjectured that if a gas of weakly interacting bosons is cooled to a sufficiently low temperature, then a macroscopic proportion of the particles occupy the lowest possible energy state. This phenomenon has come to be known as Bose–Einstein condensation (BEC), and the collection of lowest energy particles is the *condensate*. Until the end of the twentieth century, BEC was a purely theoretical construct, though it was conjectured that for any (fixed) density of particles there is a critical temperature below which the condensate exists, and acts like a single quantum particle. In 1995, BEC was first physically realised by Cornell, Wieman, and Ketterle who were awarded the 2001 Nobel Prize for Physics for “the achievement of Bose–Einstein condensation in dilute gases of alkali atoms, and for early fundamental studies of the properties of the condensates”, [Nob01].

Many mathematical models for the Bose gas have been developed. The simplest of these, the ideal gas, was shown by Einstein to exhibit a phase transition at which BEC occurs, [Ein24]. For more involved models, in which particles have a weak interaction, explicit confirmation of a critical phase is yet to be achieved. One approach was introduced by Feynman [Fey53], in which a random walk representation is used. Depending on how one chooses to model the particle-to–

particle interaction, Feynman’s representation often simplifies to an elegant probabilistic model, which can then be approached using purely probabilistic techniques. In this work we consider an interaction which was studied by Tóth [Tót93], and results in an elegant representation as a random permutation. In this case, the condensate manifests itself as an infinite cycle in the permutation.

The Quantum Heisenberg Model

The starting point for Tóth’s paper is not in fact the Bose gas, but rather a seemingly unrelated model for a ferromagnet. A ferromagnetic material is such that when placed in a magnetic field, it takes on magnetic properties of its own, and maintains these properties even after the field is removed. This retained magnetism only occurs at sufficiently low temperature, and we say that the material undergoes *spontaneous magnetization*. To return a ferromagnet to its non-magnetized state one has to heat the material to a sufficiently high temperature, known as the Curie temperature. Above this temperature the material becomes a paramagnet: it exhibits magnetic properties whilst in the presence of the external field, but does not have these properties when the field is removed. As with the Bose gas, there are many models for ferromagnetic materials, and there is a well established theory for studying the phase transition at which spontaneous magnetization occurs.

The best known of these is the classical Ising model, a lattice spin model regularly described on \mathbb{Z}^d . In the case $d = 2$ it was shown by Peierls [Pei36] that the Ising model exhibited a phase transition, and Onsager [Ons44] described an explicit formula for the Curie temperature. Whilst these results have become classics of statistical mechanics, mathematicians have sought to find more mature models for magnetic materials, one class of which are the quantum Heisenberg models. In the Ising model, each site of the lattice is endowed with either a positive or negative charge (or *spin*) and neighbouring sites try to align their spins. The image is the same for Heisenberg models, except now the charge at each vertex is given by a complex vector in \mathbb{C}^2 . What is of interest to us is the fact that in one case, the spin-1/2 Heisenberg ferromagnet, the mathematical description coincides with that of the interacting Bose gas with a simple particle interaction. Consequences of this relationship were explored by Conlon and Solovej [CS91a, CS91b] and Tóth [Tót93], though these are in the form of upper bounds on thermodynamic functions, and no concrete results are established. This remains an area of active research, and developing the theory of phase transitions for these models will be a major achievement with broad consequences.

A Note to Probabilists

For purists with little or no interest in physics we address the following comforting remark: whilst the first chapter continues in the vein of this motivation and is steeped in the language of physics, future chapters dispense of this setting and deal with the subject of random spatial permutations as purely probabilistic objects. Whilst we could not possibly recommend that you

overlook the first sections, we reassure that there is light at the end of the tunnel. Even if it is some twenty pages away!

On a more serious note, we take this space to comment on the mathematical motivation for our work. Random permutations have an important place in probability and provide a complex playground for mathematicians to work in. A classical problem in the area is to study the properties of a uniformly chosen permutation of a finite set: how many cycles do we expect? what is the mean cycle length? etc. Such questions have been resolved for centuries and often rely on little more than a clever treatment of relevant generating functions, [FS09]. Of more interest to us is the study of random permutations on infinite sets. Whilst there is no natural way to choose a uniform permutation, we can consider permutation valued processes. These have a rich structure and can be identified with more classical probabilistic problems including coagulation–fragmentation processes and random walks on groups. They also provide examples of highly correlated percolation models, and the problems of interest are analogous to those of classical percolation: do infinite cycles occur? if so, are they unique? What critical phases describe the appearance of such cycles?

Whilst the primary motivation for studying the random stirring process is undeniably its physical relevance, it has the added probabilistic appeal that it is identified with a non-Markovian random walk: the cyclic time random walk, [Ang03]. Whilst we save the technical construction until Chapter 2, we whet the appetite by providing a family friendly description.

Consider the following lazy (random) tourist who rather than walking at random around a city (or graph perhaps) chooses to take buses. He waits at a bus stop until the first bus arrives at which point he boards, staying on the bus until it reaches its first stop. He then dismounts and awaits the next bus. Unlike the random walker who has no memory of where he has been, the random bus rider has some memory since buses arrive according to a timetable. We make the (somewhat unrealistic) assumption that buses arrive at the same time past the hour for every hour of the day. A second assumption (one that is even more unrealistic) is that if there is a bus going from A to B which leaves at some time t , then at this time there is also a bus leaving B going towards A . The questions of interest here are of course the same as those we ask of a standard random walk: what notions of recurrence and transience make sense? and, what conditions can be put on the graph to ensure transience?

For now we leave our discussion there, and leave the reader curious about these questions, and how they relate to random permutations.

Outline

The bulk of this work is divided into two sections. The first chapter introduces the mathematical background within which we view the physical problems mentioned in the motivation: namely, we provide a whistle-stop tour through quantum statistical mechanics and the relationship between lattice gas models and spin systems. The central results of this section concern Tóth’s representation of the partition function of the Heisenberg spin-1/2 ferromagnet in terms of the

random stirring process. The remaining chapters are then dedicated to probabilistic aspects of random stirring. Chapter 2 introduces the process in a general setting, and provides different representations for the random stirring dynamics, in particular detailing the relationship with the cyclic time random walk. We describe some of the questions of interest and prove several basic results. Chapter 3 then provides partial answers to some of the deeper questions, focusing on critical phases for the cyclic time random walk on infinite trees. We conclude in Chapter 4 by remarking on random stirring on finite graphs, and identify some unsolved problems that we are interested in working on in the future.

List of Notation

Mathematics

- $[N]$ The integers $1, 2, \dots, N$.
 $|\Lambda|$ Lebesgue d -volume for $\Lambda \subset \mathbb{R}^d$.
 Cardinality of a countable set.
 C_b^∞ Space of smooth bounded functions.

Probability

- $\mathcal{M}^1(X)$ Probability measures on (X, \mathcal{F}) .
 $\mathbf{1}_{\{A\}}$ Indicator variable of the event A .
 $\mu_{x,y}^\beta$ Brownian bridge measure.
 $\omega_{x,y}^\beta$ Continuous time random walk bridge measure.

Statistical Mechanics

- Z_Λ Canonical partition function.
 Ξ_Λ Grand canonical partition function.
 $\langle O \rangle$ Expectation value of the quantum observable O .
 $m^{(\text{th})}$ Thermodynamic magnetisation.
 $m^{(\text{sp})}$ Spontaneous magnetisation.

Graph Theory

- \mathbb{G} Infinite, locally finite, connected graph.
 $\text{deg}_{\mathbb{G}}(v)$ Degree of vertex v .
 $E_G(V_1, V_2)$ Set of edges between V_1 and V_2 .
 $C_x(\mathbb{G})$ Connected component containing x .
 $d_{\mathbb{G}}(u, v)$ Graph distance between u and v .

Permutations

- \mathcal{S}_P Symmetric group on the set P .
 \mathcal{T}_G Transpositions induced by G .
 \mathcal{P}_A Partitions of the set A .
 Id_P Identity permutation in \mathcal{S}_P .
 $\langle x \rangle_\sigma$ Orbit of x in σ ; also denoted $\text{Orb}_\sigma(x)$.
 $\lambda_\sigma(x)$ Length of the orbit of x in σ .
 $l_\sigma(m)$ Number of cycles of length m in σ .

Random Stirring Process

- SSEP Simple symmetric exclusion process.
 CTRW Cyclic time random walk.
 RSP Random stirring process.
 \mathcal{B}_e Ringing times of the edge e .
 $\mathcal{B}_{\mathbb{G}}$ Bar set of \mathbb{G} .
 $\mathcal{B}_{\mathbb{G}}^T$ Cyclic Bar set of \mathbb{G} .
 $\nu^{(\text{meso})}$ Density of mesoscopic cycles.
 $\nu^{(\text{macro})}$ Density of macroscopic cycles.
 O_∞ The event that an infinite cycle exists.
 $\mathcal{T}_{\mathbb{G}}$ Set of times at which there are infinite cycles.
 T_c^I First time at which infinite cycles occur.
 T_c^S Last time at which there are finite cycles.

Chapter 1

Quantum Mechanics of the Bose Gas

The aim of this section is to introduce the mathematical model for the Bose gas, and describe its stochastic representation in terms of random permutations. To do this we relate it to the Heisenberg spin-1/2 ferromagnet model. Since most of the results of this section are well documented, we omit proofs of general statements, and focus on those which will be of relevance for future sections.

1.1 Mathematical Formulation of Quantum Mechanics

Before entering into the quantum world we recall some classical mechanics. The classical canonical ensemble is described as a collection of $N \gg 1$ particles in a box $\Lambda \subset \mathbb{R}^d$, $|\Lambda| < \infty$, $d \geq 1$. The state of the particles is an element of the phase space $\Gamma_\Lambda = (\Lambda \times \mathbb{R}^d)^N$ of pairs (x, p) where $x \in \Lambda^N$ correspond to particle positions, and $p \in \mathbb{R}^{dN}$ describe the momenta. The energy of a configuration $(x, p) \in \Gamma_\Lambda$ is described by a Hamiltonian

$$H_N(x, p) := \sum_{i=1}^N \frac{p_i^2}{2m} + \sum_{1 \leq i < j \leq N} V(|x_i - x_j|), \quad (x, p) \in \Gamma_\Lambda, \quad (1.1)$$

where m is the mass of a particle, and the potential $V : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ describes how pairs of particles interact. We have assumed that the particles are not acted on by an external force. The canonical partition function is defined as

$$Z_\Lambda(\beta, N) := \frac{1}{N!} \int_{\Gamma_\Lambda} \exp(-\beta H_N(x, p)) dx dp. \quad (1.2)$$

Remark 1.1. The unexpected factorial term arises as a resolution of the Gibbs paradox, and is justified by the heuristic assumption that we cannot distinguish between particles, [Ada06]. This is better justified in quantum mechanics, where bosons are known to be indistinguishable.

In quantum mechanics, the Heisenberg uncertainty principle dictates that it is not possible to simultaneously measure the position and momentum of a particle; resultingly it no longer makes sense to consider particle configurations as elements of the phase space Γ_Λ . Instead, the state of a

system is described by a wave function $\psi(x)$, a complex valued function describing a distribution of particle locations in Λ : in the case of a single particle without spin, this simply means that the probability that the particle is in a measurable set $A \subset \Lambda$ is given by $\int_A |\psi(x)|^2 dx$. The wave function itself is defined to be a solution to the time-independent Schrödinger equation

$$E\psi(x) = \widehat{H}_N\psi(x),$$

where E is the *separation constant* describing the energy of ψ . The Hamiltonian \widehat{H}_N is given by a *Schrödinger operator* acting on $L^2(\Lambda^N)$

$$\widehat{H}_N\psi(x) := -\frac{\hbar^2}{2m} \sum_{i=1}^N \Delta_i^{(\text{bc})} \psi(x) + \sum_{1 \leq i < j \leq N} V(|x_i - x_j|) \psi(x), \quad \psi \in L^2(\Lambda^N), \quad (1.3)$$

where $\hbar \approx 1.05 \times 10^{-34}$ is Planck's constant, and $\Delta_i^{(\text{bc})}$ is the Laplacian associated to particle i under some boundary conditions 'bc', which henceforth we assume to be periodic. The definition of \widehat{H}_N is analogous to (1.1), where we replace the momenta with momentum operators $p_i \mapsto -i\hbar\nabla p_i$. For interacting bosons the Hamiltonian, \widehat{H}_N , acts on the space of symmetric wave functions (a consequence of bosons having integer spins): that is we only consider those wave functions that are equivalent under permutation of the particle indices, [LSSY05]. The eigenvalues of the operator \widehat{H}_N describe the possible energy states that the system can occupy, so the natural analogue of (1.2) is the quantum canonical partition function

$$Z_\Lambda(\beta, N) := \text{Tr}_{L^2_+(\Lambda^N)} \left(\exp(-\beta \widehat{H}_N) \right), \quad (1.4)$$

where $L^2_+(\Lambda^N) \subset L^2(\Lambda^N)$ is the sub-Hilbert space of symmetric L^2 -functions, and Tr denotes the trace of an operator.

So far we have described 'classical' quantum mechanics: whilst the language of probability is used, at this level we need little probabilistic machinery. We continue by introducing the Feynman–Kac representation of the partition function (1.4), the true starting point for the relationship between quantum mechanics and probability.

1.1.1 Feynman–Kac Formulae

Feynman-Kac formulae were introduced by Feynman [Fey48, Fey53] as a tool to make rigorous his abstract path integral. In the latter of these papers, Feynman derived a formula for the partition function of a Bose gas as an integral over a collection of particle trajectories. When the particles are modeled as being in \mathbb{R}^d , these trajectories are described by the paths of Brownian bridges, whilst in discrete spaces these are replaced with random walk bridges.

In one form, Feynman–Kac formulae allow us to derive stochastic representations for kernels of exponential operators. Given the operator $\exp(-t\widehat{H}_N)$, we wish to find a function

$K_t(x, y)$ such that

$$\exp\left(-t\widehat{H}_N\right)f(x) = \int_{\Lambda^N} K_t(x, y)f(y)dy, \quad f \in L^2(\Lambda^N). \quad (1.5)$$

To simplify notation we replace the Hamiltonian \widehat{H}_N with $H := H_0 + V$, where $H_0 := -\frac{1}{2}\Delta$ and V is a suitably regular potential. We assume that H acts on $L^2(\mathbb{R}^d)$. In the simple case $H = H_0$, it is well known that the kernel $K_t(x, y) = p_t(x, y)$ satisfying (1.5) is the heat-kernel

$$p_t(x, y) := (2\pi t)^{-\frac{d}{2}} \exp\left(-\frac{|x-y|^2}{2t}\right).$$

On realising that this is the transition kernel of a d -dimensional Brownian motion, the relationship between Hamiltonian operators and stochastic processes is less mystical.

Since exponentials of operators do not adhere to the multiplicative rules of the exponential function, obtaining the kernel for general Hamiltonians is somewhat more demanding. Fortunately we have at our disposal the Trotter product formula, which provides the equivalent multiplication rule for suitably ‘nice’ operators

$$\exp\left(-t(A_1 + A_2)\right) = \lim_{n \rightarrow \infty} \left(\exp\left(-\frac{t}{n}A_1\right) \exp\left(-\frac{t}{n}A_2\right) \right)^n, \quad (1.6)$$

see [Nel64] p.343 for a short proof in the special case that the operators A_1, A_2 are self-adjoint. Combined with the fact that the kernel of $\exp(-tH_0)\exp(-tV)$ is

$$K_t(x, y) = p_t(x, y)e^{-tV(y)}, \quad (1.7)$$

equation (1.6) enables us to find a representation for the kernel of $\exp(-tH)$. We sketch a basic Feynman–Kac result for smooth potentials, and direct readers to [LHB11] for a detailed proof, as well as an in depth treatment of the general topic of Feynman–Kac formulae.

Claim 1.2. *For the Hamiltonian $H = -\frac{1}{2}\Delta + V$, with $V \in C_b^\infty(\mathbb{R}^d)$, bounded and smooth*

$$\exp(-tH)g(x) = \mathbb{E}_x \left[\exp\left(-\int_0^t V(B_s)ds\right)g(B_t) \right], \quad g \in L^2(\mathbb{R}^d). \quad (1.8)$$

where \mathbb{E}_x is the expectation with respect to the Wiener measure \mathbb{P}_x of a Brownian motion started at x , $B_0 = x$.

Sketch proof. Recall that since Brownian motion is a Markov process with transition kernel $p_t(x, y)$ then for Borel sets $\{E_i\}_{i=1}^n$, and times $t_1 \leq t_2 \leq \dots \leq t_n$

$$\mathbb{P}_x[B_{t_1} \in E_1, \dots, B_{t_n} \in E_n] = \int_{E_1} \dots \int_{E_n} p_{t_1}(x, x_1) \dots p_{t_n - t_{n-1}}(x_{n-1}, x_n) dx_1 \dots dx_n,$$

Consequently, for bounded Borel-measurable f_1, \dots, f_n

$$\mathbb{E}_x[f(B_{t_1}) \cdots f(B_{t_n})] = \int_{\mathbb{R}^d} \cdots \int_{\mathbb{R}^d} f_1(x_1) p_{t_1}(x, x_1) \cdots f_n(x_n) p_{t_n}(x_{n-1}, x_n) dx_1 \cdots dx_n.$$

Using Fubini's theorem and the fact that the transition kernel coincides with the integral kernel of $\exp(-\frac{1}{2}\Delta)$ one obtains

$$= \exp(-tH_0) f_1(x) \cdots \exp(-tH_0) f_n(x).$$

To confirm the claim we apply the Trotter product formula (1.6)

$$\begin{aligned} \exp(-tH)g(x) &= \lim_{n \rightarrow \infty} \left(\exp\left(-\frac{t}{n}H_0\right) \exp\left(-\frac{t}{n}V\right) \right)^n g(x) \\ &= \lim_{n \rightarrow \infty} \int_{\mathbb{R}^d} K_{\frac{t}{n}}(x, x_1) dx_1 \cdots \int_{\mathbb{R}^d} K_{\frac{t}{n}}(x_{n-2}, x_{n-1}) dx_{n-1} \int_{\mathbb{R}^d} K_{\frac{t}{n}}(x_{n-1}, x_n) g(x_n) dx_n, \end{aligned}$$

where $K_t(x, y)$ is as in (1.7). Since the potential V only contributes a multiplicative factor to the kernel K , we set $f_i(y) = \exp\left(-\frac{t}{n}V(y)\right)$, $i = 1, \dots, n-1$, and $f_n(y) = \exp\left(-\frac{t}{n}V(y)\right)g(y)$; then from the expression for the Brownian expectation

$$= \lim_{n \rightarrow \infty} \mathbb{E}_x \left[\exp\left(-\frac{t}{n}V(B_{t_1})\right) \cdots \exp\left(-\frac{t}{n}V(B_{t_n})\right) g(t_n) \right],$$

with $t_i = i\frac{t}{n}$, $i = 0, \dots, n$. Which simplifies to

$$= \lim_{n \rightarrow \infty} \mathbb{E}_x \left[\exp\left(-\frac{t}{n} \sum_{i=0}^n V(B_{t_i})\right) g(t) \right].$$

Applying the dominated convergence theorem, we take the limit inside the expectation. Finally since V is smooth and $(B_t)_{t \geq 0}$ almost surely continuous, $\frac{t}{n} \sum_{i=0}^n V(B_{t_i}) \rightarrow \int_0^t V(B_s) ds$. \square

Whilst we have outlined Feynman–Kac representations in the continuous setting, discrete analogues hold; in fact it will be a discrete formula that we will ultimately be interested in. We delay introducing this until Section 1.2 and conclude this discussion with Feynman's original representation of the Bose gas. In the following we let \mathcal{S}_N denote the symmetric group on N elements, i.e. the set of permutations $\sigma : [N] \rightarrow [N]$, $[N] := \{1, 2, \dots, N\}$. We denote $\mu_{x,y}^\beta$ for the Brownian bridge measure from x to y over time horizon β , which is a regular Borel measure on the space $C([0, \beta]; \mathbb{R}^d)$ of continuous functions $f : [0, \beta] \rightarrow \mathbb{R}^d$

$$\mu_{x,y}^\beta(A) := \frac{\mathbb{P}_x[B \in A, B_\beta \in dy]}{dy}, \quad (1.9)$$

where A is measurable with respect to the topology of uniform convergence on $C([0, \beta]; \mathbb{R}^d)$. A more careful construction is given in [ACK11] which takes into account possible boundary

conditions on Λ .

Theorem 1.3 (Feynman–Kac Representation of the Bose Gas). *Let $\widehat{H}_N = H_0 + V$, where V decays sufficiently fast. The partition function has the representation*

$$Z_\Lambda(\beta, N) = \frac{1}{N!} \sum_{\sigma \in S_N} \int_\Lambda dx_1 \cdots \int_\Lambda dx_N \times \bigotimes_{i=1}^N \mu_{x_i, \sigma(x_i)}^\beta \left[\exp \left(- \sum_{1 \leq i < j \leq N} \int_0^\beta V(|B_s^{(i)} - B_s^{(j)}|) ds \right) \right] \quad (1.10)$$

See [Fey53] for the classical reference, or [Gin71] for a rigorous account.

1.1.2 The Ideal Bose Gas

For general potentials V the representation (1.10) remains intractable, though for simple choices of V progress can be made. The simplest of these is the *ideal gas*, for which $V \equiv 0$. We outline some results for this case, both folklore and new, and restrict our attention to $d = 3$.

Rather than working in the canonical ensemble, it is often more convenient to work in the grand canonical setting, where we no longer fix the number of particles, but rather control the expected number via a chemical potential $\mu \in \mathbb{R}$. The grand canonical partition function is defined as the generating function for canonical partition functions

$$\Xi_\Lambda(\beta, \mu) := \sum_{N \geq 0} Z_\Lambda(\beta, N) e^{-\beta \mu N}. \quad (1.11)$$

The following result dates back to Einstein [Ein24].

Proposition 1.4. *For the ideal gas, $V \equiv 0$, the grand-canonical partition function is given by*

$$\Xi_\Lambda(\beta, \mu) = \prod_{i \geq 0} \left(1 - \exp(-\beta(\epsilon_i - \mu)) \right)^{-1},$$

for $\mu < \epsilon_0$, where $0 < \epsilon_0 \leq \epsilon_1 \leq \cdots$ are the eigenvalues of the Laplace operator, and are explicitly known.

See [LSSY05] for a discussion. The eigenvalues ϵ_i describe the different energy levels that particles can occupy, and we define the occupation number η_i as the fraction of sites occupying energy state ϵ_i . Denoting $\langle \cdot \rangle$ for the expectation value of a quantum observable, the expected total number of particles is

$$\langle N \rangle = z \frac{\partial}{\partial z} \log \Xi_\Lambda(\beta, \mu), \quad z = \exp(\beta \mu),$$

and elementary differentiation enables us to evaluate this expression for $\mu < \epsilon_0$

$$\langle N \rangle = \sum_{i \geq 0} \left(\exp(\beta(\epsilon_i - \mu)) - 1 \right)^{-1}.$$

From this one sees how the chemical potential controls the expected number of particles since as $\mu \rightarrow \epsilon_0$, $\langle N \rangle$ is strictly increasing. In addition, it is apparent that μ cannot exceed ϵ_0 (or more generally any ϵ_i) since then $\eta_i < 0$, and a negative proportion of the particles would occupy the i -th energy level, a notion which is ill-defined. Since we work on a bounded box $\Lambda \subset \mathbb{R}^3$ it is reasonable to consider the density of particles in the box $\rho_\Lambda(\beta, \mu) := \langle N \rangle / |\Lambda|$. In taking the thermodynamic limit $|\Lambda| \rightarrow \infty$, we find $\epsilon_0 \rightarrow 0$ and the density $\rho_\Lambda(\beta, \mu)$ is bounded above as $\mu \rightarrow \epsilon_0$ by the critical density [LSSY05]

$$\rho_c(\beta) := \left(\frac{2\pi\hbar^2\beta}{m} \right)^{-\frac{3}{2}} \zeta\left(\frac{3}{2}\right), \quad (1.12)$$

where ζ denotes the Riemann-zeta function

$$\zeta(z) := \sum_{k=1}^{\infty} k^{-z}, \quad z \in \mathbb{C}.$$

What we see is that as we increase μ the total number of atoms does not continue growing, but rather is capped at a certain value, described by the critical density. The way in which we interpret this capping effect is to say that in fact there is no bound on the total number of particles, but rather on the number of particles which do not occupy the lowest energy state, ϵ_0 (the *ground state*): that is $\sum_{i \geq 1} \eta_i$ converges, whilst $\eta_0 \rightarrow \infty$. In this regime, we say that those particles that are in the ground state are the Bose–Einstein condensate.

In [Süt93], Sütő proposes a model for studying Bose interactions in the canonical ensemble, in which he expresses $Z_\Lambda(\beta, N)$ as a sum over the collection $\mathcal{P}_{[N]}$ of all partitions of the set $[N]$. Interpreting such a partition as defining the equivalence classes of a permutation, one obtains a permutation representation of the Bose gas where particles arrange into cycles depending on their energy state. Sütő argues that for general potentials, V , there exists a critical density $\rho_c^V(\beta)$ such that in the thermodynamic limit those particles occupying state ϵ_0 arrange themselves into a collection of infinite cycles, and *cycle percolation* occurs. Conversely, he notes that in the sub-critical regime, $\rho < \rho_c^V(\beta)$, the cycles are all finite and the η_i , $i \geq 0$ are all of order 1. In the case of the ideal gas Sütő confirms that this phase transition coincides with the occurrence of Bose–Einstein condensation. By defining suitable probability measures $\mu_N^\rho \in \mathcal{M}^1(\mathcal{S}_N)$, he demonstrates that in the thermodynamic limit probability mass is lost on finite cycles; in particular if $\langle 1 \rangle$ denotes the (random) cycle containing the ‘particle’ 1, then

$$\lim_{N \rightarrow \infty} \sum_{k \geq 1} \mu_N^\rho(\#\langle 1 \rangle = k) < 1 \quad \text{if } \rho > \rho_c(\beta),$$

where the critical density is exactly (1.12), [Süt93] pp.4705–6. The sub-critical claim is similarly confirmed, though what happens at criticality is unknown. In a later paper Sütő embellished the picture by confirming that the existence of infinite cycles exactly coincides with the existence of BEC in the case of a mean field Bose gas [Süt02].

Another approach to studying critical behaviour is by analysing the canonical *specific free energy* or the grand canonical *pressure*, defined respectively as

$$\begin{aligned} f_\rho(\beta) &:= \lim_{|\Lambda| \rightarrow \infty} -\frac{1}{\beta|\Lambda|} \log Z_\Lambda(\beta, N), \\ p(\beta, \mu) &:= \lim_{|\Lambda| \rightarrow \infty} \frac{1}{\beta|\Lambda|} \log \Xi_\Lambda(\beta, \mu), \end{aligned} \tag{1.13}$$

where in the case of the specific free energy it is assumed that the limit is taken such that it preserves the density $N/|\Lambda| \rightarrow \rho$. In either case, understanding the singularities of the thermodynamic function can be used to indicate the existence of a phase transition. As of yet, no explicit formula is known for the free energy under a fixed particle density and temperature. Recently progress was made by Adams, Collecchio and König who derived a variational formula for $f_\rho(\beta)$ under general (weak) requirements on the potential V , [ACK11]. By working with the Feynman–Kac representation (1.10) they derive an interpretation of this formula as an expectation with respect to a marked Poisson point process in \mathbb{R}^d , where marks are in the form of looped trajectories from the point back to itself (these correspond to the loops over which we integrate in the Feynman–Kac representation). As in the work of Sütő, the measures introduced are concentrated only on finite cycles and BEC manifests itself as a loss of probability mass, with the magnitude of this lost mass corresponding to the density of particles in the ground state. The analysis here generalises a variational formula derived in [Ada07] for the ideal gas, where Adams shows existence of a phase transition for dimensions $d \geq 3$, and that this occurs at the expected critical density. These calculations are described in detail in [Dan11], as well as similar results for phase transitions under different measures on spaces of permutations.

1.2 The Spin- $\frac{1}{2}$ Heisenberg Ferromagnet

As hinted in the introduction, the study of the Bose gas is closely related to a quantum magnetism model known as the spin-1/2 Heisenberg ferromagnet. In this section we derive this relationship, before stating the discrete Feynman–Kac representation which was promised in the previous pages. After reviewing some known results for the model, we will be in a strong position to move away from the motivational aspects of this work, and onto the main object of our interest: the random stirring process. A detailed introduction to quantum lattice models is [Sim93], whilst a thorough survey of the spin-1/2 Heisenberg ferro- and antiferromagnet is [GUW11].

A quantum spin system describes a collection of electrons each of which is endowed with a *spin* taking a value in a finite dimensional Hilbert space. For our purposes the electrons occupy the vertices of a lattice box $\Lambda \subset \mathbb{Z}^d$, and take spins in \mathbb{C}^2 ; we denote \mathbb{C}_x^2 for the Hilbert space at site $x \in \Lambda$. The Hilbert space of spin configurations is given by the tensor product

$$\mathcal{H}_\Lambda = \bigotimes_{x \in \Lambda} \mathbb{C}_x^2.$$

The space \mathcal{H}_Λ is itself a Hilbert space and one can define its inner product for generating elements $\otimes_{x \in \Lambda} f_x$, $f_x \in \mathbb{C}_x^2$, by

$$\langle \otimes_{x \in \Lambda} f_x | \otimes_{x \in \Lambda} g_x \rangle = \prod_{x \in \Lambda} \langle f_x | g_x \rangle_x,$$

where $\langle \cdot | \cdot \rangle_x$ is the inner product on \mathbb{C}_x^2 . The Hamiltonian of a spin-1/2 model is described by the set of Pauli matrices

$$\sigma^{(1)} := \frac{1}{2} \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma^{(2)} := \frac{1}{2} \begin{pmatrix} 0 & -i \\ -i & 0 \end{pmatrix}, \quad \sigma^{(3)} := \frac{1}{2} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

which act on the spaces \mathbb{C}_x^2 , and describe the spin of the electron in the x -, y -, and z -direction. The Pauli matrices are used to form local operators $\{\sigma_x^{(i)}\}_{x \in \Lambda}$, $i = 1, 2, 3$, which act on the whole of \mathcal{H}_Λ

$$\sigma_x^{(i)} = I_2 \otimes \cdots \otimes I_2 \otimes \sigma^{(i)} \otimes I_2 \otimes \cdots \otimes I_2,$$

the tensor product of $(\#\Lambda - 1)$ -copies of the 2×2 identity matrix, and the Pauli matrix appearing in the ' x -th' position. Roughly speaking the operators $\sigma_x^{(i)}$ correspond to having independent spins at each site $x \in \Lambda$ in the direction i ; this notion of independence is due to the operators satisfying the commutation relations

$$\begin{aligned} [\sigma_x^{(i)}, \sigma_y^{(j)}] &:= \sigma_x^{(i)} \sigma_y^{(j)} - \sigma_y^{(j)} \sigma_x^{(i)} \\ &= i \epsilon_{i,j} \delta_{x,y} \sigma_x^{(k)} \end{aligned}$$

where $i, j, k \in \{1, 2, 3\}$, $i \neq j \neq k$, and $\delta_{x,y}$ is the Kronecker delta; the term $\epsilon_{i,j}$ depends on the specific pair of directions i, j that are chosen. The important point being that if $x \neq y$ then the commutator is zero, with the interpretation that the operator acts independently on these sites. Detailed constructions of $\{\sigma_x^{(i)}\}_{x \in \Lambda}$, and a description of the commutation relations can be found in [GUW11], pp.180-1. To each local operator we define $\vec{\sigma}_x := (\sigma_x^{(1)}, \sigma_x^{(2)}, \sigma_x^{(3)})$ so that interactions between spin operators can be conveniently written in terms of a dot product

$$\vec{\sigma}_x \cdot \vec{\sigma}_y := \sigma_x^{(1)} \sigma_y^{(1)} + \sigma_x^{(2)} \sigma_y^{(2)} + \sigma_x^{(3)} \sigma_y^{(3)},$$

which is also an operator on \mathcal{H}_Λ .

For a constant $J \in \mathbb{R}$ and external magnetism $h \in \mathbb{R}$ we define the spin-1/2 Heisenberg operator

$$H_\Lambda^{J,h} := -J \sum_{\substack{x,y \in \Lambda \\ |x-y|=1}} \vec{\sigma}_x \cdot \vec{\sigma}_y - h \sum_{x \in \Lambda} \sigma_x^{(3)}. \quad (1.14)$$

The second term of the definition corresponds to an external magnetic force, which adds a bias for those sites which have spin in the z -direction. The spin-1/2 Heisenberg *ferromagnet* is the special case $J = 1$, whilst $J = -1$ corresponds to the *antiferromagnet*. Whilst we

will not discuss the antiferromagnet, a nice theory has been developed for it, analogous to Tóth's stochastic representation for the ferromagnet (which we will presently discuss), this is the Aizenmann–Nachtergaele loop-process. Once again the reader is directed to [GUW11] for a thorough comparison of these models. Henceforth we denote $H_\Lambda^h := H_\Lambda^{1,h}$ for the ferromagnetic Hamiltonian, and often refer to this as the Heisenberg model (since we will not be concerned with the more general form). Since we have endowed the lattice with an electron at every site we are naturally in a canonical-type ensemble, and define the partition function accordingly

$$\Xi_\Lambda(\beta, h) := \text{Tr} \left(\exp \left(-\beta H_\Lambda^h \right) \right). \quad (1.15)$$

Remark 1.5. The observant reader may be questioning the choice of the Greek letter Xi (Ξ) for this partition function since we have previously reserved this for the grand-canonical case; the motivation for this notation is provided in Theorem 1.11.

1.2.1 Equivalence to the Bose Gas

Whilst the Hamiltonian H_Λ^h may seem to require a considerably more complicated (or at least more abstract) definition, we continue by showing that it is in fact related to the Hamiltonian of the Bose gas; moreover, we will see that the canonical partition function (1.15) is no more than the grand canonical partition function (1.11) for a particular choice of pair interaction, a correspondence that dates back to Matsubara and Matsuda [MM56], and is pre-dated in the case of interacting fermions by the work of Jordan and Wigner [JW28].

First, define the so called *raising*, *lowering* and *magnon* operators

$$S_x^+ := \sigma_x^{(1)} + i\sigma_x^{(2)}, \quad S_x^- := \sigma_x^{(1)} - i\sigma_x^{(2)}, \quad m_x := S_x^+ S_x^-.$$

It is well known that these operators satisfy the following commutation relations

$$\begin{aligned} [S_x^+, S_y^+] &= [S_x^-, S_y^-] = 0, \\ [S_x^-, S_y^+] &= \delta_{x,y}(1 - 2m_x), \\ S_x^+ S_x^- &= m_x, \end{aligned}$$

which allow us to rewrite (1.14) as

$$H_\Lambda^h = \frac{1}{2} \sum_{\substack{x,y \in \Lambda \\ |x-y|=1}} (S_x^+ - S_y^+) (S_x^- - S_y^-) - \sum_{\substack{x,y \in \Lambda \\ |x-y|=1}} m_x m_y - h \sum_{x \in \Lambda} m_x + \frac{\#\Lambda}{2} \left(h - \frac{z}{4} \right) I_{\mathcal{H}}, \quad (1.16)$$

where z is the number of unordered nearest neighbour pairs, and $I_{\mathcal{H}}$ is the identity operator,

[MM56]. Under periodic boundary conditions, this simplifies further by noting that

$$\begin{aligned} \sum_{\substack{x,y \in \Lambda \\ |x-y|=1}} (S_x^+ - S_y^+) (S_x^- - S_y^-) &= \sum_{\substack{x,y \in \Lambda \\ |x-y|=1}} (S_x^+ S_x^- - S_x^+ S_y^- - S_y^+ S_x^- + S_y^+ S_y^-) \\ &= 2d \sum_{x \in \Lambda} S_x^+ S_x^- + \sum_{\substack{\langle x,y \rangle \in \Lambda \times \Lambda \\ |x-y|=1}} S_x^+ S_y^+ \end{aligned}$$

where $\langle x, y \rangle$ is the ordered pair of sites; and in terms of the discrete Laplacian

$$= \sum_{x,y \in \Lambda} S_x^+ \Delta_{x,y} S_y^-,$$

where

$$\Delta_{x,y} = \begin{cases} 2d & \text{if } |x - y| = 0, \\ 1 & \text{if } |x - y| = 1, \\ 0 & \text{if } |x - y| > 1. \end{cases} \quad (1.17)$$

In this form the Hamiltonian is comparable to that of a lattice Bose gas, in particular it can be related to a lattice gas in the grand canonical ensemble, defined in terms of *creation* and *annihilation* operators.

For each $x \in \Lambda$ we can define operators a_x^\dagger, a_x which act on the bosonic-Fock space, and correspond to adding and removing (respectively) particles from the Bose gas; these are the creation and annihilation operators. Related to these is the local particle number operator $n_x := a_x^\dagger a_x$. The detailed construction of these operators is not important to us, what is significant is that the Hamiltonian of a lattice Bose gas can be written as

$$H_\Lambda^{(\text{bose})} = - \sum_{x,y \in \Lambda} a_x^\dagger \Delta_{x,y} a_y + \frac{1}{2} \sum_{x,y \in \Lambda} V(x,y) n(x) n(y), \quad (1.18)$$

this is known as the method of *second quantization*, see for example [Sch08]. Whilst the connection between (1.18) and (1.16) currently appears weak, the link is clarified by the observation that the creation, annihilation, and particle number operators share similar commutation relations as the raising, lowering and magnon operators

$$\begin{aligned} [a_x^\dagger, a_y^\dagger] &= [a_x, a_y] = 0, \\ [a_x, a_y^\dagger] &= \delta_{x,y}, \\ a_x^\dagger a_x &= n_x. \end{aligned}$$

Choosing the potential

$$V(x,y) = \begin{cases} \infty & \text{if } |x - y| = 0, \\ -2 & \text{if } |x - y| = 1, \\ 0 & \text{if } |x - y| > 1, \end{cases} \quad (1.19)$$




Figure 1.1: Comparison of the potential (1.19) and the Lennard-Jones potential; both display hard-core repulsion for small inter-particle distance, and low interaction for moderate distances.

it is clear that (upto inclusion of the external field) the Heisenberg Hamiltonian is equivalent to that of the Bose lattice gas. Resultingly, Tóth expresses the canonical Heisenberg partition function in terms of the grand canonical partition function of the Bose lattice gas, [Tót93] p.77

$$\Xi_{\Lambda}(\beta, h) = \sum_{N=0}^{\infty} Z_{\Lambda}^{(\text{bose})}(\beta, N) e^{\beta h N}, \quad (1.20)$$

where for $N = \sum_{x \in \Lambda} n_x$, $Z_{\Lambda}^{(\text{bose})}(\beta, N)$ is the canonical partition function for the lattice gas (1.18). This correspondence between a lattice gas and a system of spins is deep, and allows us to express observables for each model in terms of the another, a comprehensive discussion of such links is contained in [MM56]. We simply observe the equality

$$N = \frac{\#\Lambda}{2} + \left\langle \sum_{x \in \Lambda} \sigma_x^{(3)} \right\rangle_{\beta, h},$$

where the term in brackets is the observable describing the average magnitude of the spin in the z -component. With this we see that the exponential term in the grand-canonical partition function (the right hand side of (1.20)) corresponds to the external field in (1.16), so that the chemical potential and the external magnetism have analogous roles in the two models.

Remark 1.6. Whilst the potential (1.19) may seem only of interest as a toy model it is sufficiently close to a realistic potential to pick up salient features. In particular it exhibits the same ‘well-like’ structure as the Lennard Jones potential (a more realistic model), Fig. 1.1; as Baxter observes [Bax08]: “the qualitative features of the liquid–gas transition are not expected to depend on the details of the potential: it should be sufficient that it have short-range repulsion and an attractive well.”

1.2.2 A Return to Feynman–Kac and the Random Stirring Representation

As long promised, we are now in a position to state the Feynman–Kac representation which enables us to describe the Heisenberg model within a probabilistic framework. In analogy to Theorem 1.3, we define the measure $\omega_{x,y}^\beta$, $x, y \in \Lambda$, to be the measure for continuous time random walks on Λ with jump rate $2d$, which form bridges between x and y over a period β . In keeping with the above, we assume that Λ has periodic boundary conditions.

Theorem 1.7. *For a Hamiltonian of the form (1.18), the N -particle canonical partition function, $Z_\Lambda(\beta, N) = Z_\Lambda^{(\text{bose})}(\beta, N)$, has a Feynman–Kac representation*

$$Z_\Lambda(\beta, N) = \frac{1}{N!} \sum_{\sigma \in \mathcal{S}_N} \sum_{x_1, \dots, x_N \in \Lambda} \bigotimes_{i=1}^N \omega_{x_i, \sigma(x_i)}^\beta \left[\exp \left(- \sum_{1 \leq i < j \leq N} \int_0^\beta V(X_s^{(i)}, X_s^{(j)}) ds \right) \right]. \quad (1.21)$$

Our aim for the remainder of this section will be to derive an expression for $\Xi_{\beta, h}$, throughout we follow [Tót93]. First we analyse the form that $Z_\Lambda(\beta, N)$ takes when we choose the potential (1.19). Consider the integral term inside the exponential; we observe that if up to time β no walks coincide (i.e. $X_s^{(i)} \neq X_s^{(j)}$, $\forall i \neq j$) then the integral takes the simple form

$$\begin{aligned} \sum_{1 \leq i < j \leq N} \int_0^\beta V(X_s^{(i)}, X_s^{(j)}) ds &= - \int_0^\beta \sum_{1 \leq i < j \leq N} 2 \mathbf{1}_{\{|X_s^{(i)} - X_s^{(j)}| = 1\}}, \\ &= - \int_0^\beta \sum_{1 \leq i, j \leq N} \mathbf{1}_{\{|X_s^{(i)} - X_s^{(j)}| = 1\}}, \end{aligned}$$

that is, at each time $0 \leq s \leq t$ the argument is the total number of ordered pairs $\langle i, j \rangle \in [N] \times [N]$ such that $X_s^{(i)}$ and $X_s^{(j)}$ are adjacent in Λ . For general sets $A \subset \Lambda$ let $\mathcal{N}(A) = \mathcal{N}_\Lambda(A)$ denote the number of ordered pairs of adjacent elements in A

$$\mathcal{N}(A) := \# \{ \langle a, b \rangle : a, b \in A, |a - b| = 1 \}.$$

Also define the first collision time

$$\tau := \inf \left\{ t > 0 : \exists i, j \in [N], i \neq j, \text{ st. } X_t^{(i)} = X_t^{(j)} \right\}. \quad (1.22)$$

Then for any time $\beta > 0$

$$\exp \left(- \sum_{1 \leq i < j \leq N} \int_0^\beta V(X_s^{(i)}, X_s^{(j)}) ds \right) = \exp \left(\int_0^\beta \mathcal{N} \left(X_s^{(1)}, \dots, X_s^{(N)} \right) ds \right) \mathbf{1}_{\{\beta < \tau\}}.$$

To simplify the partition function further, we drop the summation over \mathcal{S}_N replacing it with a sum over all $A \subset \Lambda$, $\#A = N$, and an indicator that the walks all start and end in A . This indicator does not specify the exact initial and final location of any individual walk, and hence

accounts for all possible permutations on A

$$Z_\Lambda(\beta, N) = \sum_{\substack{A \subset \Lambda \\ \#A=N}} \mathbb{E} \left[\exp \left(\int_0^\beta \mathcal{N} \left(X_s^{(1)}, \dots, X_s^{(N)} \right) ds \right) \mathbf{1}_{\{\beta < \tau\}} \right. \\ \left. \times \mathbf{1}_{\{X_\beta^{(1)}, \dots, X_\beta^{(N)}\} = A} \mid \left\{ X_0^{(1)}, \dots, X_0^{(N)} \right\} = A \right], \quad (1.23)$$

where the expectation is with respect to the joint distribution of independent continuous time random walks. In this form we see that the partition function does not depend so much on the dynamics of the individual walks, but on the walks as a collection. With this in mind we define the process $\eta: \Omega \times \mathbb{R}_{\geq 0} \rightarrow \in^\Lambda$, on the power set of $V(\Lambda)$ (where Ω denotes a generic measurable space). Given an initial configuration $\eta_0 = A$ we define

$$\eta_t := \left\{ X_t^{(1)}, \dots, X_t^{(N)} \right\},$$

and observe that up to time τ , $\eta_t \in \{A' \subset \Lambda: \#A' = N\}$. The distribution of the process up to time τ is in fact identical to that of a well known process, the *simple symmetric exclusion process* or SSEP. The finite state space SSEP with N particles on the lattice box Λ , is defined as follows: at time 0 the process starts from an initial configuration of N particles placed on different sites of Λ . Independently, each particle waits an exponential- $2d$ period of time (assuming periodic boundaries) and then chooses a neighbouring site uniformly at random. If the chosen site is unoccupied, the particle moves to this new site; else it remains where it is. For a rigorous construction of this process, as well as a thorough study of its properties, see [Lig85]. The fact that sites can only hold at most a single particle imposes a hard-core interaction which is not shared by the simple random walks $X^{(1)}, \dots, X^{(N)}$; however, up to the first time at which a jump is blocked, the distributions are the same. The random blocking time for the SSEP is governed by the value of $\mathcal{N}(\eta_t)$, since more neighbouring particles implies a greater chance of blocked jumps. We can use this to describe a coupling between the random walks and a killed-SSEP, in which we extend the state space of the SSEP to include a cemetery state ∂ . Now the process $\xi: \mathbb{R}_{\geq 0} \rightarrow \{A' \subset \Lambda: \#A' = N\} \cup \partial$ evolves as a SSEP up to a time τ_∂ , after which it remains in the cemetery, i.e. $\xi_t = \partial, t > \tau_\partial$. The time τ_∂ is determined by its instantaneous rate, or *hazard function*

$$\lim_{\Delta t \rightarrow 0} \frac{\mathbb{P}[\tau_\partial < t + \Delta t \mid \tau_\partial > t]}{\Delta t} = \mathcal{N}(\xi_t).$$

We observe that this instantaneous rate coincides with that of the collision time τ . Since up to the collision time the random walk dynamics are the same as the exclusion process, we obtain the following claim.

Claim 1.8. *The ‘location set’ of N -continuous time nearest-neighbour random walks on Λ ran up to the time τ given by (1.22), and the N -particle killed SSEP have the same distribution if*

started from the same initial configuration $\{X_0^{(1)}, \dots, X_0^{(N)}\} = \xi_0 = A$

$$\left(\tau, \{X_t^{(1)}, \dots, X_t^{(N)}\} : t < \tau\right) \stackrel{(d)}{=} \left(\tau_\partial, \xi_t : t < \tau_\partial\right)$$

for any $A \subset \Lambda$, $\#A = N$.

We are now in a position to find a simple closed expression for the canonical partition function.

Proposition 1.9 (Tóth, [Tót93] pp.78-80). *Let $(\xi_t)_{t \geq 0}$ be a SSEP of N -particles on $\Lambda \subset \mathbb{Z}^d$ with periodic boundary conditions, and exponential rates $(2d)^{-1}$. The canonical partition function of the lattice Bose gas with potential V given by (1.19) is given by*

$$Z_\Lambda(\beta, N) = \sum_{\substack{A \subset \Lambda \\ \#A=N}} \mathbb{P}[\xi_\beta = A \mid \xi_0 = A]. \quad (1.24)$$

Proof. By the preceding claim we can rewrite (1.23) in terms of the SSEP

$$Z_\Lambda(\beta, N) = \sum_{\substack{A \subset \Lambda \\ \#A=N}} \mathbb{E} \left[\exp \left(\int_0^\beta \mathcal{N}(\xi_s) ds \right) \mathbf{1}_{\{\beta < \tau_\partial\}} \mathbf{1}_{\{\xi_\beta = A\}} \mid \xi_0 = A \right],$$

where the expectation is with respect to the distribution of the killed-SSEP, (τ_∂, ξ) . Using the tower property for conditional expectations applied to the killing time and the SSEP, the indicator term can be written as a conditional probability. That is

$$= \sum_{\substack{A \subset \Lambda \\ \#A=N}} \mathbb{E} \left[\exp \left(\int_0^\beta \mathcal{N}(\xi_s) ds \right) \mathbb{P}[\beta < \tau_\partial \mid (\xi_t)_{0 \leq t \leq \beta}] \mathbf{1}_{\{\xi_\beta = A\}} \mid \xi_0 = A \right].$$

Considering the conditional probability term, it is a basic result of survival analysis that given the hazard function of a variable, its survival probability can be computed; since we are conditioning on the behaviour of $(\xi_t)_{0 \leq t \leq \beta}$ our hazard function $\mathcal{N}(\xi_t)$ is determined, and we see that the survival probability is given by

$$\mathbb{P}[\beta < \tau_\partial \mid (\xi_t)_{0 \leq t \leq \beta}] = \exp \left(- \int_0^\beta \mathcal{N}(\xi_t) dt \right),$$

(a derivation of this formula is given in Appendix A). Consequently we see that the two exponential terms in the partition function cancel, and we obtain the desired formula. \square

We can now tackle the real result, which is deriving an expression for the grand canonical partition function $\Xi_\Lambda(\beta, h)$; to do this we introduce the random stirring process, which will be the object of our consideration for the remaining chapters of this work.

The random stirring process, or RSP, on the box Λ is a permutation valued random process on the sites $V(\Lambda)$, $\sigma : \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathcal{S}_\Lambda$. As originally described by Tóth, the process is initialised as the identity permutation $\sigma_0 = \text{Id}_\Lambda$, then independently and with exponential rate 1 the labels of adjacent sites are swapped; that is if at time t the adjacent pair (x, y) ‘rings’

then we set $\sigma_t(x) = \sigma_{t^-}(y)$, $\sigma_t(y) = \sigma_{t^-}(x)$, where t^- is understood to be the time immediately before the ring. The resulting process defines a bijection (since by the dynamics no two sites ever have the same label), and hence is a permutation.

Claim 1.10. *Let $(\sigma_t)_{t \geq 0}$ be an RSP on Λ and for a subset $A \subset \Lambda$, define the process*

$$\xi_t := \{\sigma_t(x) : x \in A\}.$$

Then $\xi : \Omega \times \mathbb{R}_{\geq 0} \rightarrow \{A' \subset \Lambda : \#A' = \#A\}$ is a SSEP with initial configuration A .

To see this, we note that ξ_0 corresponds to starting with a collection of indistinguishable particles at each site of A ; each time the labels of x and y are exchanged in the RSP, if there is a particle at either x or y (or both) then it swaps sites (if there are particles at both x and y they swap positions). Following the trajectory of a single particle we see that if it always jumps to sites that were not previously occupied it performs a continuous time random walk. When we reach a time at which a particle exchanges sites with another, we follow this new particle instead. Since the particles are indistinguishable we do not ‘see’ the swap, and it appears as though the original particle has just sat at its site (i.e. had its jump repressed). Consequently, each trajectory looks like that of a single particle in the SSEP.

We are now ready for Tóth’s simple representation of the spin-1/2 ferromagnet. For a permutation $\sigma \in \mathcal{S}_P$ (for some discrete set P) we denote

$$\langle x \rangle_\sigma = \text{Orb}_\sigma(x) := \left\{ \sigma^k(x) : k \in \mathbb{N} \right\}, \quad x \in P,$$

for the cycle (or, *orbit*) of σ containing x . The length of the cycle is $\lambda_\sigma(x) := \#\langle x \rangle_\sigma$, and we write $l_\sigma(m) := m^{-1} \#\{x : \lambda_\sigma(x) = m\}$, for the number of cycles of length m (i.e. m -cycles). When referring to the RSP $(\sigma_t)_{t \geq 0}$ we write $\langle x \rangle_t$, $\lambda_t(x)$, $l_t(m)$ in place of $\langle x \rangle_{\sigma_t}$, $\lambda_{\sigma_t}(x)$, $l_{\sigma_t}(m)$.

Theorem 1.11 (Tóth, [Tót93] pp.80-1). *Let $\Xi_\Lambda(\beta, h)$ be the Heisenberg spin-1/2 partition function, then*

$$\Xi_\Lambda(\beta, h) = \mathbb{E} \left[\prod_{n \geq 1} (1 + e^{\beta h n})^{l_\beta(n)} \right]. \quad (1.25)$$

Proof. We consider first the following auxiliary problem: independently and with probability $p \in [0, 1]$ we place a particle at each site $x \in \Lambda$, and use this as the initial condition ξ_0 for an SSEP on Λ . We run the SSEP dynamics until time $t > 0$ and are interested to know the probability that the configuration $\xi_t = \xi_0$. We see that

$$\begin{aligned} \mathbb{P}[\xi_t = \xi_0] &= \sum_{A \subset \Lambda} \mathbb{P}[\xi_0 = A] \mathbb{P}[\xi_t = A \mid \xi_0 = A] \\ &= \sum_{A \subset \Lambda} p^{|A|} (1-p)^{|\Lambda \setminus A|} \mathbb{P}[\sigma_t(A) = A], \end{aligned}$$

where A is interpreted as the set of initial particle locations, and the second line follows from the preceding claim. Interpreting A as a collection of marked points, the condition $\sigma_t(A) = A$

specifies that the permutation σ_t maps marked points to marked points; consequently any orbit of $\langle x \rangle_t$ is either contained entirely in A or in $\Lambda \setminus A$; we say that the cycles are *monochrome*

$$\begin{aligned} \mathbb{P}[\sigma_t \text{ has monochrome cycles}] &= \sum_{A \subset \Lambda} \mathbb{P}[\{\text{marks}\} = A] \mathbb{P}[\text{monochrome cycles} \mid \{\text{marks}\} = A] \\ &= \sum_{A \subset \Lambda} p^{|A|} (1-p)^{|\Lambda \setminus A|} \mathbb{P}[\sigma_t(A) = A], \end{aligned}$$

so an equivalent problem is to study monochrome cycles. Using the fact that the marks in Λ are distributed independently of σ_t

$$\begin{aligned} \mathbb{P}[\xi_t = \xi_0] &= \mathbb{P}[\sigma_t \text{ has monochrome cycles}] \\ &= \mathbb{E}[\mathbf{1}_{\{\text{all cycles monochrome}\}}] \\ &= \mathbb{E}\left[\prod_{m \geq 1} \left(\mathbf{1}_{\{m\text{-cycle monochrome}\}}\right)^{l_t(m)}\right], \end{aligned}$$

where the final indicator is for a single cycle of length m . Using the independence of the marks again, the indicator can be written explicitly since the probability that an m -cycle is monochrome is exactly $p^m + (1-p)^m$ (either all sites marked or all sites not-marked). Hence

$$= \mathbb{E}\left[\prod_{m \geq 1} \left(p^m + (1-p)^m\right)^{l_t(m)}\right].$$

We return to the problem at hand. From Proposition 1.24, we express the grand-canonical partition function as

$$\Xi_{\Lambda}(\beta, h) = \sum_{A \in \Lambda} e^{\beta h |A|} \mathbb{P}[\xi_{\beta} = A \mid \xi_0 = A],$$

where the conditional probability is with respect to a SSEP measure on $\#A$ -particles. Choosing $p = e^{\beta h} / (1 + e^{\beta h}) \in [0, 1]$, we obtain

$$\begin{aligned} \Xi_{\Lambda}(\beta, h) &= \left(1 + e^{\beta h}\right)^{|\Lambda|} \sum_{A \subset \Lambda} p^{|A|} (1-p)^{|\Lambda \setminus A|} \mathbb{P}[\xi_{\beta} = A \mid \xi_0 = A] \\ &= \left(1 + e^{\beta h}\right)^{|\Lambda|} \mathbb{E}\left[\prod_{m \geq 1} \left(p^m + (1-p)^m\right)^{l_{\beta}(m)}\right]. \end{aligned}$$

Given the identity $|\Lambda| = \sum_{m \geq 1} m l_t(m)$, we can bring the constant term inside the product

$$= \mathbb{E}\left[\prod_{m \geq 1} \left(1 + e^{\beta h m}\right)^{l_{\beta}(m)} \left(p^m + (1-p)^m\right)^{l_{\beta}(m)}\right].$$

Rearranging terms we obtain the desired result. \square

1.2.3 Thermodynamic Limit and Phases

We have already seen that the lattice Bose gas can be interpreted in terms of the Heisenberg model, and of course this relationship can be reverted. In particular we expect that at low temperature and zero external field ($h = 0$) the spin-1/2 model behaves like the ideal Bose lattice gas, [CS91b]. Recalling the definition of the pressure, equation (1.13), it is expected that at low temperatures the Heisenberg pressure converges to that of the ideal Bose lattice gas.

Conjecture 1.12. *Let $p(\beta, h)$ be the Heisenberg spin-1/2 ferromagnetic pressure, at inverse temperature $\beta > 0$ and external field $h \in \mathbb{R}$. For $d \geq 1$*

$$\lim_{\beta \rightarrow \infty} \beta^{(d+2)/2} p(\beta, 0) = C^{(\text{bose})} := -\frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \log(1 - e^{-k^2}) dk.$$

It was in this context that Tóth introduced the RSP, and the main result of [Tót93] provides a lower bound for the pressure.

Theorem 1.13 (Tóth, [Tót93]). *For $d \geq 3$*

$$C^{(\text{bose})} \log 2 \leq \liminf_{\beta \rightarrow \infty} \beta^{(d+2)/2} p(\beta, 0).$$

The RSP can also be applied to identifying phase transitions in the Heisenberg model, which are seen through the behaviour of the order parameters

$$m^{(\text{th})}(\beta) := -\lim_{h \searrow 0} \frac{p(\beta, h) - p(\beta, 0)}{h},$$

$$m^{(\text{sp})}(\beta, h) := \liminf_{|\Lambda| \rightarrow \infty} \frac{1}{|\Lambda|} \left\langle \left| \sum_{x \in \Lambda} \sigma_x^{(3)} \right| \right\rangle_{\beta, h},$$

The first of these is the *thermodynamic magnetisation* which describes the residual magnetic effect which occurs when we gradually remove an external magnetic field, whilst the *spontaneous magnetisation* describes the average magnitude of the spin in the z -direction. For both order parameters, at high temperature one expects them to be 0, whilst at low temperatures ordering amongst the spins occurs (*symmetry breaking*), and they become strictly positive. Manipulating (1.25) Tóth obtains the formula

$$m^{(\text{sp})}(\beta) := m^{(\text{sp})}(\beta, 0) = \frac{1}{2} \lim_{n \rightarrow \infty} \lim_{\Lambda \rightarrow \mathbb{Z}^d} \frac{\mathbb{E} \left[\mathbf{1}_{\{\lambda_\beta(0) > n\}} 2^{\sum_{k \geq 1} l_\beta(k)} \right]}{\mathbb{E} \left[2^{\sum_{k \geq 1} l_\beta(k)} \right]}, \quad (1.26)$$

where the powers of 2 correspond to setting $h = 0$ in $(1 + e^{\beta h m})$. In this form it is apparent that a change of behaviour in the spontaneous magnetism is triggered by the expected cycle sizes in the RSP. This equation is essentially saying that in the limit we are interested in the occurrence of infinite cycles, but since (for the purposes of this section) we work on sequences of

finite graphs, we must define what we mean by an infinite cycle. We define two relevant notions

$$\nu^{(\text{meso})}(\beta, h) = \lim_{K \rightarrow \infty} \liminf_{|\Lambda| \rightarrow \infty} \frac{1}{|\Lambda|} \mathbb{E}_{\Lambda}^{\beta, h} \left[\# \left\{ x \in \Lambda : K < \lambda_{\beta}(x) < \frac{|\Lambda|}{K} \right\} \right], \quad (1.27)$$

$$\nu^{(\text{macro})}(\beta, h) = \lim_{\epsilon \rightarrow 0} \liminf_{|\Lambda| \rightarrow \infty} \frac{1}{|\Lambda|} \mathbb{E}_{\Lambda}^{\beta, h} \left[\# \{x \in \Lambda : \lambda_{\beta}(x) > \epsilon |\Lambda|\} \right]. \quad (1.28)$$

The *mesoscopic* cycles in the permutation are those that are infinite but make up a ‘negligible’ proportion of the total sites; *macroscopic* cycles, on the other hand, are those which are infinite and include a positive fraction of the sites. The total expected density of infinite cycles is $\nu^{\infty}(\beta, h) := \nu^{(\text{meso})}(\beta, h) + \nu^{(\text{macro})}(\beta, h)$.

Proposition 1.14. *For $d \geq 1$,*

$$\begin{aligned} \nu^{\infty}(\beta) := \lim_{h \rightarrow 0} \nu^{\infty}(\beta, h) > 0 &\implies m^{(\text{th})}(2\beta) > 0, \\ \nu^{(\text{macro})}(\beta, 0) > 0 &\iff m^{(\text{sp})}(2\beta) > 0. \end{aligned}$$

see [GUW11], pp.196-7. As it stands little is known about the value ν^{∞} . Although in general we study the RSP to shed light on the physical model, one exception is provided by the Mermin-Wagner theorem, which implies that symmetry breaking does not occur for the Heisenberg model in $d = 1, 2$.

Theorem 1.15 (Mermin-Wagner, [MW66]). *For $d = 1, 2$, $m^{(\text{th})}(\beta) = 0$ for all $\beta > 0$, and the RSP has almost surely no infinite cycles, $\nu^{\infty}(\beta) = 0$.*

For $d \geq 3$ symmetry breaking can occur, and it is conjectured by Tóth that there is a phase transition above which the spontaneous magnetisation is non-zero: so that in particular the RSP has macroscopic cycles. As a first approach to showing such a phase transition, Tóth suggests dropping the powers of 2 in (1.26), and consider the ‘toy’ thermodynamic limit

$$m(\beta) := \lim_{n \rightarrow \infty} \lim_{\Lambda \rightarrow \mathbb{Z}^d} \mathbb{P}[\lambda_{\beta}(0) > n].$$

Conjecture 1.16 (Tóth, [Tót93]). *For $d = 3$ both m and $m^{(\text{sp})}$ undergo a phase transition. i.e. there exist $\beta_c, \beta_c^{(\text{sp})} > 0$ such that*

$$m(\beta) \begin{cases} = 0 & \text{for } \beta < \beta_c \\ > 0 & \text{for } \beta > \beta_c \end{cases} \quad \text{and} \quad m^{(\text{sp})}(\beta) \begin{cases} = 0 & \text{for } \beta < \beta_c^{(\text{sp})} \\ > 0 & \text{for } \beta > \beta_c^{(\text{sp})}. \end{cases}$$

Both of these conjectures are far from being solved, though the case for $m(\beta)$ has received some attention. In the remaining sections we introduce the equivalent problem on an infinite graph, and highlight the results of Angel and Hammond for infinite trees.

Chapter 2

The Random Stirring Process

We introduced the random stirring process to provide a stochastic description for the phase transition expected in the Heisenberg spin-1/2 ferromagnet. We move away from this motivation, and now focus on the RSP itself. We first define the process in greater detail, before focusing on its basic properties.

2.1 Random Stirring Representations

In Section 1.2 the RSP was introduced as the following process: each site of a box $\Lambda \subset \mathbb{Z}^d$ is occupied by a labeled particle, then independently and at unit exponential rate particles at neighbouring sites are swapped. In this form the process is known as the *interchange* process, and can be defined on more general graphs. The aim of this section is to introduce several alternative descriptions of the RSP on finite graphs, before discussing the construction in the infinite setting.

2.1.1 The Interchange Process

Let $G = (V(G), E(G))$ be a finite graph, where we abbreviate $V = V(G)$, $E = E(G)$ when there is no ambiguity. Denote $\mathcal{S}_G = \mathcal{S}_V$ for the set of permutations on the vertices V . To each edge $e = uv \in E(G)$ we associate a transposition $\tau_e := (u, v) \in \mathcal{S}_G$. Conversely given the transposition $\tau = (u, v)$ we write $e_\tau := uv$. We denote $\mathcal{T}_G := \{\tau_e : e \in E\} \subset \mathcal{S}_G$, and observe that if G is connected then \mathcal{T}_G generates \mathcal{S}_G : that is, every element of \mathcal{S}_G can be obtained as a finite composition of elements of \mathcal{T}_G . We always assume that G is connected.

The RSP is a permutation valued random process $\sigma: \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathcal{S}_G$ described by the following coupling. To each $e \in E$ we associate an independent unit-rate Poisson process \mathcal{B}_e , which together induce a countable subset of $\mathbb{R}_{\geq 0} \times E$,

$$\mathcal{B}_G := \bigcup_{e \in E} \{(t, e) : t \in \mathcal{B}_e\}.$$

We call \mathcal{B}_G the *bar set*, and elements of \mathcal{B}_G are *bars*. If $(t, e) \in \mathcal{B}_G$ we say that the edge e

‘rings’ at time t . Since the composition of countably many Poisson processes almost surely do not coincide, $\cap_{e \in E} \mathcal{B}_e = \emptyset$, we can define an (almost sure) strict ordering on the bar set: $(t, e) \prec (t', e')$ if $t < t'$. Subsequently we write $\mathcal{B}_G = (t_i, e_i)_{i=1}^{\infty}$ as an ordered list. The RSP is started from the identity permutation, $\sigma_0 = \text{Id}_G$, and then evolves by composing the permutation with the transpositions associated to successive ringing times. More precisely, for $t \in (t_{i-1}, t_i]$, $i \geq 1$, $t_0 = 0$

$$\sigma_t := \begin{cases} \sigma_{t_{i-1}} & \text{for } t_{i-1} < t < t_i, \\ \tau_{e_i} \circ \sigma_{t_{i-1}} & \text{for } t = t_i. \end{cases}$$

The process is defined so that it is càdlàg, and is immediately seen to be a continuous time Markov process. Its embedded jump chain $Y_i := \sigma_{t_i}$, $i \geq 0$, is the homogeneous Markov chain on \mathcal{S}_G where successive jumps are determined by the uniform distribution over \mathcal{T}_G

$$P(\tau) = \begin{cases} \frac{1}{|E|} & \text{if } \tau \in \mathcal{T}_G, \\ 0 & \text{else.} \end{cases}$$

The k -step density is given by the convolution $P^{*k} = P * P^{*(k-1)}$, where

$$Q * R(\sigma) := \sum_{\rho \in \mathcal{S}_G} Q(\sigma \rho^{-1}) R(\rho), \quad Q, R \in \mathcal{M}^1(\mathcal{S}_G).$$

Connectivity of G ensures that the process σ_t (respectively Y_i) is irreducible, and so is ergodic (it is an irreducible finite state space Markov chain). In addition, the discrete chain has a symmetric transition matrix

$$\begin{aligned} P(\sigma, \rho) &= \frac{1}{|E|} \mathbf{1}_{\{\exists \tau \in \mathcal{T}_G \text{ st. } \rho = \tau \circ \sigma\}} \\ &= \frac{1}{|E|} \mathbf{1}_{\{\exists \tau^{-1} \in \mathcal{T}_G \text{ st. } \sigma = \tau^{-1} \circ \rho\}} \\ &= P(\rho, \sigma). \end{aligned}$$

Ergodicity and symmetry combined ensure that Y_i , and hence σ_t , has a unique stationary distribution $\mathcal{U} \in \mathcal{M}^1(\mathcal{S}_G)$, the uniform distribution, which is the uniform distribution on \mathcal{S}_G .

2.1.2 Random Walk on the Symmetric Group

Since the RSP is permutation valued it is apparent that there should be an algebraic description of the model, and this is provided through the *Cayley graph* associated to the group.

For a finite group (H, \circ) with generating set $J \subset G$ we define the graph $\Gamma = \Gamma_J^H$, a directed graph whose vertices are identified with the group elements, $V(\Gamma) = H$. The directed edge $\langle u, v \rangle \in E(\Gamma)$ if $g = j \circ h$ for some $j \in J$. We consider the Cayley graph $\Gamma_G = \Gamma_{\mathcal{T}_G}^{\mathcal{S}_G}$. If $\tau \in \mathcal{T}_G$ then $\tau^{-1} = \tau \in \mathcal{T}_G$, and it follows that $\langle \rho, \sigma \rangle \in E(\Gamma_G)$ if and only if $\langle \sigma, \rho \rangle \in E(\Gamma_G)$,




Figure 2.1: The graph $G = K_3$ and its generated Cayley graph.

so the Cayley graph can be considered to be undirected. Fig. 2.1 depicts the Cayley graph for $G = K_3$, the complete graph on 3 vertices.

The RSP is equivalent to the continuous time random walk on Γ_G in which the walk waits an exponential time at a site before jumping to a neighbouring site which is chosen uniformly amongst all neighbours. Group closure ensures that every $\rho \in \mathcal{S}_G$ has exactly $\#\mathcal{T}_G$ neighbours in Γ_G and

$$\#V(\Gamma_G) = (\#V(G))!, \quad \#E(\Gamma_G) = \frac{\#\mathcal{T}_G}{2} \#V(\Gamma_G).$$

Remark 2.1. We highlight why the random walk representation has its drawbacks, at least from a computation approach. Even for small graphs the associated Cayley graph is comparatively large: for the (rather small) lattice box $\Lambda = [-1, 1]^3 \subset \mathbb{Z}^3$, has Cayley graph Γ_Λ with over 10^{28} vertices and 2×10^{71} edges.

2.1.3 Coagulation-Fragmentation of Permutations

The third representation we consider stems entirely from an understanding of the composition of permutations, and the recognition that the act of composing a transposition with a permutation has a *coagulation-fragmentation* effect. Consider a permutation $\sigma \in \mathcal{S}_N$ written in cycle notation $\sigma = (x_{11}, \dots, x_{1n_1})(x_{21}, \dots, x_{2n_2}) \cdots (x_{k1}, \dots, x_{kn_k})$, where $x_{ij} \in [N]$ are distinct, k is the number of cycles, and $\sum_{i=1}^k n_i = N$. We have the following composition rules, which are justified in Fig. 2.2.

Claim 2.2. Let $\tau = (x_{ik}, x_{jl})$, $1 \leq k \leq n_i$, $1 \leq l \leq n_j$. For $i \neq j$

$$\tau \circ (x_{i1}, \dots, x_{in_i})(x_{j1}, \dots, x_{jn_j}) = (x_{i1}, \dots, x_{i(k-1)}, x_{jl}, \dots, x_{jn_j}, x_{j1}, \dots, x_{j(l-1)}, x_{ik}, \dots, x_{in_i}),$$

whilst for $i = j$ and $k < l$

$$\tau \circ (x_{i1}, \dots, x_{in_i}) = (x_{i1}, \dots, x_{ik}, x_{i(l+1)}, \dots, x_{in_i})(x_{i(k+1)}, \dots, x_{il}).$$

Since composing the permutation σ with $\tau = (x_{ik}, x_{jl})$ only affects the i - and j -cycles, the above claim provides the rule for composing any permutation with a transposition. We observe that

if x and y belong to different cycles of σ then in $(x, y) \circ \sigma$ these cycles are merged (coagulated); conversely, if they both lie in the same cycle of σ then in $(x, y) \circ \sigma$ the cycle is split (fragmented), with x and y in different cycles.




Figure 2.2: The possible effect of pre-composing a permutation with a transposition; if the entries of the transposition are in different cycles coagulation occurs (top); conversely if they are in the same orbit it is fragmented (bottom).

Before remarking on the RSP, we recall that any permutation in \mathcal{S}_N induces a partition in \mathcal{P}_N , the set of partitions of $[N]$

$$\mathcal{P}_N := \left\{ \{A_1, \dots, A_n\} : \cup_{i=1}^n A_i = [N], \emptyset \neq A_i \subseteq [N], A_i \cap A_j = \emptyset, \forall i \neq j \right\}.$$

Namely, if $\sigma \in \mathcal{S}_N$ then choosing one representative element, x_1, \dots, x_n , from each cycle we obtain a partition $\lambda(\sigma) = \{A_1, \dots, A_n\}$, with $A_i = \langle x_i \rangle_\sigma$, defining a surjection $\lambda: \mathcal{S}_N \rightarrow \mathcal{P}_N$. Elementary combinatorics confirms that this is not a bijection for $N \geq 2$, since

$$\#\lambda^{-1}(\{A_1, \dots, A_n\}) = \prod_{i=1}^n (\#A_i)!. \quad (2.1)$$

Let $\lambda_t = \lambda(\sigma_t) \in \mathcal{P}_G$, be the projection of the RSP, where \mathcal{P}_G is the set of partitions of V . In light of (2.1), the partition process captures the size of the cycles in the RSP, but loses information regarding the actual permutation structure.

Suppose we wanted to construct a random partition with the same law as λ_t , but without

reference to the RSP. Let $(\tilde{\lambda}_t)_{t \geq 0}$ be a coagulation-fragmentation process taking values in \mathcal{P}_G , for distinct blocks $A_t, B_t \in \tilde{\lambda}_t$, $A_t \neq B_t$, define their instantaneous rate of merging to be $\#E_G(A_t, B_t)$, where for $V_1, V_2 \subset V$

$$E_G(V_1, V_2) := \{e = uv \in E : u \in V_1, v \in V_2 \text{ or } v \in V_1, u \in V_2\},$$

is the set of edges between V_1 and V_2 . Since there are $\#E_G(\langle x \rangle_t, \langle y \rangle_t)$ possible transpositions which merge the cycles containing x and y , the rate of coagulation is the same as in $(\lambda_t)_{t \geq 0}$, and we remark that the coagulation is independent of which transposition is chosen: if $y \notin \langle x \rangle_\sigma$

$$\lambda\left((x, y) \circ \sigma\right) = \lambda\left((x', y') \circ \sigma\right), \quad \forall x' \in \langle x \rangle_\sigma, y' \in \langle y \rangle_\sigma.$$

Turning to fragmentation, we see that if the rate of splitting for a block $A_t \in \tilde{\lambda}_t$ is $\#E(A_t, A_t)$, then this matches the splitting rate for λ_t . Whilst we have identified the correct rate of fragmentation, to define the actual splitting dynamics is harder since the pair of cycles one obtains after applying a transposition (x, y) , $y \in \langle x \rangle_\sigma$ is dependent on the specific elements x, y that are chosen, and the order of the elements in the cycle. Subsequently, to construct a coagulation-fragmentation model which is equivalent to the RSP requires us to be in a set of *ordered partitions* in which blocks are ordered lists of vertices. It is not hard to see that such a set is equivalent to \mathcal{S}_G , so we gain little from such a representation.

Whilst referring to the mapping $\lambda_t = \lambda(\sigma_t)$ as a *representation* of the RSP is somewhat misleading, in as much as that we are unable to construct the process λ_t independently of the RSP, it highlights the manner in which the cycles of $(\sigma_t)_{t \geq 0}$ interact. In particular, after each ringing time the size of the largest cycle can be at most twice the size of the previous largest cycle,

$$\max_{x \in V} \lambda_t(x) \leq 2 \max_{x \in V} \lambda_{t-}(x).$$

For large graphs G and small times t this bound is sharp since most cycles formed will be pairs, and until fragmentation starts to occur the bound continues to perform well. However, once a macroscopic proportion of the sites are in large cycles, fragmentation occurs at a positive fraction of the ringing times. Ultimately we see that the critical behaviour of the RSP is directly related to the growth of blocks in $(\lambda_t)_{t \geq 0}$.

2.1.4 The Cyclic Time Random Walk

We finally turn our attention to the cyclic time random walk. Not only does this representation prove very useful in the study of random stirring, but it is also of interest in itself as an example of a non-Markovian random walk.

Fix $T > 0$, which we will refer to as the period, and define the cyclic bar set $\mathcal{B}_G^T = (t_i, e_i)_{i=1}^\infty$ by taking the bars of \mathcal{B}_G which ring in $[0, T)$ and then replicating them on the intervals $[kT, (k+1)T)$, $k \geq 0$

$$\mathcal{B}_G^T := \{(t, e) : (t \bmod T, e) \in \mathcal{B}_G\}. \quad (2.2)$$




Figure 2.3: The Harris diagram for a particular bar configuration, describing cyclic walks $\{X_v^T\}_{v \in [-3,3]}$, $T = 2$. The walks start at the bottom of the diagram at time 0 and then traverses up the vertical poles, crossing a horizontal bar whenever one is reached. The walk X_0^2 is shown in bold. This bar configuration defines the permutation $\tilde{\sigma} = (0, -2, 1, -3)(-1)(2)(3)$.

The cyclic time random walk (henceforth CTRW), $X_v^T : \mathbb{R}_{\geq 0} \rightarrow G$, is a deterministic process given the random environment \mathcal{B}_G^T . The walk starts at $X_v^T(0) = v$ and waits there until an adjacent edge rings, at which point it jumps to the adjacent vertex. For $t \in (t_{i-1}, t_i]$

$$X_v^T(t) = \begin{cases} X_v^T(t_{i-1}) & \text{for } t_{i-1} < t < t_i, \\ X_v^T(t_{i-1}) & \text{for } t = t_i, \text{ if } X_v^T(t_{i-1}) \notin e_i, \\ u & \text{for } t = t_i, \text{ if } X_v^T(t_{i-1}) \in e_i. \end{cases}$$

Here we adopted the notation $u \in e$ to mean that u is a vertex adjacent to the edge e .

Clearly the set \mathcal{B}_G^T is not Poisson distributed since the bars which ring after time T are deterministic, conditioned on those which ring in $[0, T]$. Consequently one sees that the random walk X_v^T does not have the Markov property: at time T we know that the walk will not jump until at least time $T + t_1$, so the waiting times are not independent of the past. Running the walk up to time $t \leq T$, however, the story is different since the set $\mathcal{B}_G^T \cap \{[0, T] \times E\}$ is Poisson distributed. Restricted to this time interval, $(X_v^T)_{0 \leq t \leq T}$ is the continuous time random walk

(i.e. the walk that waits an independent exponential waiting time at a vertex before jumping to a uniformly chosen neighbour). The relationship between the RSP and the CTRW is a consequence of the following observation.

Proposition 2.3. *The collection $\{X_v^T\}_{v \in V(G)}$ of CTRWs on the same cyclic bar set \mathcal{B}_G^T , almost surely defines a random permutation $\tilde{\sigma} \in \mathcal{S}_G$*

$$\tilde{\sigma}(v) := X_v^T(T). \quad (2.3)$$

We must check that each of the walks is at a distinct vertex at time T , in particular we must confirm that the walks never meet one another.

Proof. We show the stronger result that for $u, v \in V$, $u \neq v$, $X_u^T(t) \neq X_v^T(t)$ for all $t \geq 0$. Suppose for a contradiction that for some time $t > 0$ there are u, v , $u \neq v$, such that $X_u^T(t) = X_v^T(t)$ (the case $t = 0$ is covered by the initial conditions). Let t_* be the first time at which the walks meet, $t_* = \inf\{s \geq 0: X_u^T(s) = X_v^T(s)\}$. Since the walks are constant between ringing times, t_* must be a ringing time, and both of the particles must have arrived at this vertex at that time, since if one of them had arrived previous to t_* it would have left when that ring occurs. Since the probability that two edges ring at the same time is a null-event, almost surely t_* corresponds to a unique edge, and hence both X_u^T and X_v^T jumped to this site along the same bar, so $X_u^T(t_*^-) = X_v^T(t_*^-)$, which contradicts t_* being the first time that they meet. \square

This leads us to the following conclusion, which is the backbone of Angel's paper [Ang03].

Theorem 2.4. *The permutation $\tilde{\sigma}$ induced by CTRWs $\{X_v^T\}_{v \in V}$ has the same distribution as the marginal of the RSP at time T*

$$\left(\sigma_T; \mathcal{B}_G\right) \stackrel{(d)}{=} \left(\tilde{\sigma}; \mathcal{B}_G^T\right).$$

The following proof is included for completeness, though the technical details obscure the simplicity of the underlying statement.

Proof. Let \mathcal{B}_G^T be the cyclic bar set obtained from \mathcal{B}_G , so that the permutations σ_T , $\tilde{\sigma}$ are coupled; we have just to confirm that $\sigma_T = \tilde{\sigma}$. Up to time T the number of rings is almost surely finite $\mathcal{B}_G \cap \{[0, T] \times E\} = \{(t_i, e_i)\}_{i=1}^R$, for some $R < \infty$. The RSP is defined to be $\sigma_T = \tau_R \circ \tau_{R-1} \circ \dots \circ \tau_1$ where $\tau_i = \tau_{e_i}$. Consider the value of $\sigma_T(v)$. Applying Claim 2.2 we see that $\sigma_T(v)$ is obtained by locating the first of the transpositions that contain v , $\tau_{i_v} = (u_1, v)$, then finding the next transposition to contain u_1 , $\tau_{i_{u_1}} = (u_2, u_1)$ $i_{u_1} > i_v$, and so on until we arrive at u_k for which there is no $\tau_{i_{u_k}} > \tau_{i_{u_{k-1}}}$ in the list $(\tau_i)_{i=1}^R$, then $\sigma_T(v) = u_k$. Now considering instead the CTRWs we can order the edges in terms of their ringing times $(e_i)_{i=1}^R$. The final location of the CTRW started from v , $X_v^T(T) = \tilde{\sigma}(v)$ is obtained by crossing adjacent edges. That is, we locate the first edge $e_{i_v} = vu'_i$ adjacent to v , then the first edge adjacent to u'_i which rings after e_{i_v} , and so on until we reach u'_i for which there are no adjacent bars which

ring later; we set $\tilde{\sigma}(v) = u'_l$. Now by the mapping between edges and transpositions $\tau_i = \tau_{e_i}$, $u'_i = u_i$, and $k = l$, and in particular $\sigma_T(v) = u_k = u'_l = \tilde{\sigma}(v)$. \square

2.1.5 The RSP on an Infinite Graph

One problem that we face is how to describe the random stirring process in the context of an infinite graph. Whilst the representations described above all turned out to be equivalent for finite graphs, when we work in an infinite setting this equivalence is no longer certain. Throughout the remainder of this work we will denote $\mathbb{G} = (V(\mathbb{G}), E(\mathbb{G}))$ for an infinite but locally finite graph: that is $\#V(\mathbb{G}) = \infty$, $\deg_{\mathbb{G}}(v) < \infty$ for all $v \in V(\mathbb{G})$. In addition we assume that \mathbb{G} is connected.

We recall that an infinite permutation on a countable set V (which we take to be the vertices of \mathbb{G}) is a bijection $\sigma: V \rightarrow V$, and we denote \mathcal{S}_V for the set of all bijections. It is well known that \mathcal{S}_V is not countable. In the finite case, we considered the edges $e \in E(G)$ as generators of a group, and saw that so long as G was connected the generated group was $\mathcal{S}_G = \mathcal{S}_{V(G)}$, the symmetric group of G . In the case of an infinite connected graph the edge set $E = E(\mathbb{G})$ is countable, and so the group generated by E , \mathcal{S}_E , must also be countable; consequently $\mathcal{S}_E \neq \mathcal{S}_V$. In fact $\mathcal{S}_E \subset \mathcal{S}_V$ is a sub-group of finite support, that is any $\sigma \in \mathcal{S}_E$ contains only finitely many cycles of length > 1 , so that all but a finite set of elements are fixed points. Which of these two groups do we consider the symmetric group for the graph \mathbb{G} ? If we try to define the RSP in terms of a random walk on the Cayley graph we see that there are immediate problems since the Cayley graph of \mathcal{S}_V has uncountably many vertices each joined by uncountably many edges, whilst the Cayley graph of \mathcal{S}_E is countable, but not locally finite. Either graph obstructs our definition of the RSP.

If we consider the interchange dynamics instead we run into a different problem. A Poisson bar set $\mathcal{B}_{\mathbb{G}}$ on an infinite graph has the property that there are almost surely infinitely many ringing times in any finite interval. Consequently the interchange process applies infinitely many transpositions in any finite interval. Since the definition of a generating set only encompasses finite composition of generating elements then the countable composition of transpositions need no longer be in \mathcal{S}_E , however it is not in fact clear that it is in \mathcal{S}_V at all!

Example 1. Consider the graph $\mathbb{G} = \mathbb{Z}_{\geq 0}$ and the bars $(t_i, e_i)_{i=1}^{\infty}$, where $e_i = i(i-1) \in E(\mathbb{G})$ and $t_{i-1} < t_i$ for all $i \geq 1$. The resulting permutation σ is defined for $i \geq 1$ as $\sigma(i) = i-1$, and has inverse $\sigma^{-1}(i) = i+1$; however $\sigma(0)$ is not defined, and there is no element $z \in V(\mathbb{G})$ with $\sigma^{-1}(z) = 0$. Consequently $\sigma \notin \mathcal{S}_V$.

How can we salvage the random stirring process dynamics and describe the process in an infinite setting? Having ruled out the interchange process as a possible construction, we turn at last to the CTRW. Whilst it is true that the cyclic bar set $\mathcal{B}_{\mathbb{G}}^T$ also has infinitely many jumps in any interval, the CTRW only ‘sees’ finitely many of these jumps, and so the walk dynamics do not change in the infinite setting.

Proposition 2.5. *Upto any finite time $S \geq 0$, X_v^T sees only a finite sub-graph of \mathbb{G}*

$$\#\{X_v^T(t) : t \in [0, S]\} < \infty \quad a.s.$$

Proof. Consider first the case $0 \leq S < T$, and X_v^T is a standard continuous time random walk on the graph \mathbb{G} up to time S . The number of jumps performed by X_v^T in $[0, S]$ is then Poisson distributed (with intensity S), and hence is almost surely finite. Now suppose that $T \leq S < 2T$; we have just to check that the process makes only finitely many jumps in $[T, S]$. However the jumps of this process are exactly those of $X_{X_v^T(T)}^T = X_{\tilde{\sigma}(v)}^T$ on the interval $[0, S - T)$, of which there are finitely many as this is a continuous time random walk. For more general $S < \infty$ we can extend the above by induction. \square

Unlike the interchange process, being on an infinite graph makes no difference to the rate of jumps of the CTRW. In light of this we can define an infinite graph version of the RSP in terms of its marginals $\sigma_T : \Omega \rightarrow \mathcal{S}_V$

$$\sigma_T(v) := X_v^T(T), \quad v \in V(\mathbb{G}),$$

Remark 2.6. So far we have not defined random stirring on an infinite graph as a process, though for our purposes knowledge of the marginals is sufficient, and we are less interested in the evolution of the RSP over time. The reason for this goes back to our motivation, where we recall that the expected phase transition of the Heisenberg model is signaled by the existence of infinite cycles in the marginal σ_β , for fixed $\beta > 0$, see Corollary 1.16. So knowledge of the marginals is enough. We note the rather unfortunate convention that the period is denoted by T , since the inverse temperature is now identified as $\beta = T$, whereas in statistical mechanics T is the temperature and $\beta = T^{-1}$.

2.2 General Techniques and Results

Having established suitable definitions of the RSP and CTRW, we proceed to elicit basic properties of the processes, before approaching deeper questions in subsequent sections. The problems that we are interested in are summarised below.

How long does the RSP on a finite graph take to reach its stationary distribution?

What is the cycle structure of a typical sample from the marginal σ_T ? In particular what are the expected sizes of cycles? Are there infinite cycles?

How does the CTRW behave? What is required of an infinite graph to ensure that the walk is transient or recurrent?

Most of these questions still await answers, though there is a burgeoning stock of conjectures. We have already seen that on a finite graph the RSP converges to the uniform distribution on \mathcal{S}_G , and recently much progress has been made on identifying the convergence rate on the

complete graph [Sch05], as well as confirming Aldous’s conjecture that the spectral gap of the RSP is equal to that of the continuous time random walk [CLR10]. We return to these results in Chapter 4.

We concentrate on providing basic observations about the later two topics. In particular we take the time to prove several statements that are assumed in the paper [Ang03], and bring some of our own ideas to the table as well. Whilst few of the results in this section are deep, the proofs become somewhat cumbersome: we have attempted to find a balance between simplifying proofs and introducing auxiliary notation, though we fear that some proofs still obscure the true simplicity.

2.2.1 Cycle Sizes and the Projection Graph

One construction which was not introduced in the previous section is the projection graph associated to the RSP. For $T \geq 0$, define the graph $\Pi_T = (V, E_T)$ with $V = V(\mathbb{G})$, and with those edges of \mathbb{G} that have rung by time T

$$E_T := \{e \in E(\mathbb{G}) : \exists t \in [0, T] \text{ st } (t, e) \in \mathcal{B}_{\mathbb{G}}\}. \quad (2.4)$$

We recall that a subset of vertices $U \subset V$ is connected if every pair of vertices in U have a path joining them. A connected component (or cluster) is a maximal connected set; we denote $C_v(\mathbb{G})$ for the connected component of \mathbb{G} containing $v \in V$. By definition we have the inclusion $\Pi_s \subseteq \Pi_t$ for $0 \leq s \leq t$, where \subseteq is the sub-graph relation: for graphs G_1, G_2 on the same vertex set, $G_1 \subseteq G_2$ if $E_1 \subseteq E_2$.




Figure 2.4: Three possible configurations of 4 bars on the graph K_3 ; all of them have the same projection graph in both labeled and unlabeled form, but describe different permutations. In particular the cycle sizes are not constant.

Proposition 2.7. *Let Π_T be the projection graph of σ_T . For any $v \in V$, the orbit $\langle v \rangle_T$ is connected in Π_T .*

Proof. Suppose $\langle v \rangle_T$ is not connected, then there is a $k \geq 0$ such that $\sigma_T^k(v) \in V$ is not connected to $\sigma_T^{k+1}(v)$, and without loss of generality $k = 0$ (else we can redefine $v = \sigma_T^k(v)$). However $\sigma_T(v) = X_v^T(T)$, and so v and $\sigma_T(v)$ must be connected since X_v^T has traversed a set of bars between them. \square

Whilst the cycles of σ_T are connected subsets of Π_T , they need not be entire connected components, as demonstrated in Fig. 2.4 where only the middle configuration corresponds to a single cycle. One improvement to Π_t is to define it as a labeled graph, where the labeling indicates the number of times an edge has rung, i.e. define $\iota_T: E(\mathbb{G}) \rightarrow \mathbb{N}$, $\iota_T(e) := \#\{t \leq T: (t, e) \in \mathcal{B}_{\mathbb{G}}\}$. Whilst the pair (Π_T, ι_T) gives us more information about the RSP, it is still limited in its use: referring back to Fig. 2.4 we see that all three bar sets have the same labeling function.

The following observation is the sole case (excluding trivial ones) in which knowing the labeling proves fruitful.

Lemma 2.8. *Let $C \subset \Pi_T$ be a connected component of the projection graph. Suppose*

- (I) C is a tree.
- (II) $\iota_T(e) = 1$, for every $e \in C$.

Then C corresponds to a single cycle in σ_T : $\langle u \rangle_T = \langle v \rangle_T$, $\forall u, v \in C$.

Proof. Suppose that C satisfies conditions (I) and (II), i.e. it is a 1-tree, and that $\langle u \rangle_T \neq \langle v \rangle_T$ for some $u, v \in C$. Since C is a tree, it contains a unique shortest path $Q_{u,v} = z_0, \dots, z_n$, where $z_0 = u, z_n = v, z_i \neq z_j$, for $i \neq j$, and $z_{i-1}z_i \in E(C)$, $i = 1, \dots, n$. In addition, we can choose the representative elements to be such that $z_1, \dots, z_{n-1} \notin \langle u \rangle_T \cup \langle v \rangle_T$. Since the orbits of u and v are distinct this path must be non-trivial, in that it must have at least one edge. It is sufficient to prove the lemma in the case that the path Q_{uv} consists of exactly a single edge, $e = uv$. By condition (II) the edge $e = uv$ has a unique bar $(t, e) \in \mathcal{B}_{\mathbb{G}} \cap [0, T] \times E(\mathbb{G})$, and from Claim 2.7 either $\langle u \rangle_{t^-} \neq \langle v \rangle_{t^-}$, $\langle u \rangle_t = \langle v \rangle_t$ (coagulation) or $\langle u \rangle_{t^-} = \langle v \rangle_{t^-}$, $\langle u \rangle_t \neq \langle v \rangle_t$ (fragmentation).

Suppose first that (t, e) causes a coagulation, so that $\langle u \rangle_t = \langle v \rangle_t$. Since the cycles are distinct at time T , there is a first time $S \in (t, T]$ at which the cycles become distinct, which requires an edge $e' = u'v'$ to ring, where $u', v' \in \langle u \rangle_{S^-} = \langle v \rangle_{S^-}$ and $u' \in \langle u \rangle_S \neq \langle v \rangle_S \ni v'$. We note that u' and v' are in the same cycle of σ_{S^-} and so are connected in Π_{S^-} by Proposition 2.7. But then in Π_S the edge e' defines a path between u' and v' which was not in Π_{S^-} by uniqueness of the ringing time, so u' is in a cycle of Π_T . Since $u' \in C_u(\Pi_S)$ then $u' \in C_u(\Pi_T) = C$, contradicting (I) since C contains a cycle.

Now consider the fragmentation case: $\langle u \rangle_t \neq \langle v \rangle_t$. Since u and v were in the same cycle at t^- they must be connected in Π_{t^-} (again using Proposition 2.7), and so there must be a path joining them P_{uv} . As in the previous case this path cannot contain the edge $e = uv$, and we obtain the same contradiction as above since when (t, e) rings a cycle is formed.

Finally, recall that we assumed in the above that the path joining u and v was a single edge, $Q_{uv} = uv$. If instead it consists of n -edges $P_{u,v} = u, z_1, \dots, z_{n-1}, v$ then by our assumption that the representative elements are closest ($z_i \notin \langle u \rangle_T \cup \langle v \rangle_T$, $i = 1, \dots, n-1$) then z_1 must be in a distinct cycle, and $\langle u \rangle_T \neq \langle z_1 \rangle_T$. But then the unique path between the cycles of x

and z_1 is a single edge, and by the above contradiction we see $\langle u \rangle_T = \langle z_1 \rangle_T$, contradicting the minimality of the path $P_{x,y}$. \square

Remark 2.9. Implicit in this proof was the assumption that the path Q_{uv} must be finite. Whilst there are infinite paths in any locally finite graph, given any two fixed vertices it is always possible to construct a finite path between them. Contained in this proof is also the observation that if a connected component C is a 1-tree in Π_T then up to time this time no vertices $v \in C$ have been involved in a fragmentation: i.e. any transposition $\tau = (u, v)$ with $v \in C$ causes a coagulation.

Remark 2.10. The requirements (I) and (II) are the only conditions which ensure that the labeling uniquely determines the cycle sizes within a component: once multiple edges and cycles occur, one must know the order the edges ring to find the exact cycle sizes. The exception to this is the case that $C \subset \Pi_T$ is a connected component comprising of a single edge, $e \in E$; in this case if $\iota_t(e)$ is odd then this corresponds to a single cycle, whilst if it is even then the vertices are in separate 1-cycles.

As a consequence of this lemma we obtain our first interesting result for the RSP. We recall that the notation $\lambda_\sigma(v)$ denotes the length of the orbit containing v , and that $l_T(k)$ is the number of orbits of length k in σ_T . The graph distance $d_{\mathbb{G}}: V(\mathbb{G}) \times V(\mathbb{G}) \rightarrow \mathbb{N}$ gives the length of the shortest path between vertices

$$d_{\mathbb{G}}(u, v) := \begin{cases} \min_{P_{u,v} \subseteq \mathbb{G}} \#\{e: e \in P_{u,v}\} & \text{if } v \in C_u(\mathbb{G}), \\ \infty & \text{if } v \notin C_u(\mathbb{G}), \end{cases}$$

where the minimum is over all paths connecting u and v , and the length of a path is defined in terms of the number of edges so that $d(u, u) = 0$. The diameter of a set $V' \subset V$ is defined as $\text{diam}_{\mathbb{G}}(V') = \sup_{u,v \in V'} d_{\mathbb{G}}(u, v)$. The graph \mathbb{G} is said to be of uniformly bounded degree d if the degree of each vertex is at most d , $\deg_{\mathbb{G}}(v) \leq d$. Since we deal extensively with sub-graphs in the following, we sometimes abuse notation and write $v \in \mathbb{G}$, in place of $v \in V(\mathbb{G})$ (similarly for edges we write $e \in \mathbb{G}$). Henceforth we denote \mathbb{P}_T for the law of the marginal σ_T and the related CTRW.

Theorem 2.11. *Let \mathbb{G} be an infinite, connected graph of uniformly bounded degree d , $d < \infty$. For any time $T > 0$ and integer $k \geq 1$, the RSP on \mathbb{G} has almost surely infinitely many orbits of length k*

$$\mathbb{P}_T \left[l_T(k) = \infty \right] = 1. \tag{2.5}$$

Proof. Our proof relies on the observation that for any $k \geq 1$ there is a countable collection of sets $\{A_i\}_{i=1}^{\infty}$ satisfying

$$\text{diam}_{\mathbb{G}}(A_i) \geq k + 1, \quad \#E(A_i) \leq U, \quad V(A_i) \cap V(A_j) = \emptyset, \tag{2.6}$$

for all $i, j \geq 1$, $i \neq j$, where U is a universal constant. We construct such sets as follows; choose an arbitrary site $v_1 \in \mathbb{G}$ and let $A_1 = B_{k+1}(v_1)$ be the ball of radius $k + 1$ around v_1 in the graph metric, $d_{\mathbb{G}}$. Since \mathbb{G} is connected v_1 is connected to sites arbitrarily far away, and consequently there is at least one vertex, u , on the boundary of A_1 , i.e. with $d_{\mathbb{G}}(v_1, u) = k + 1$, and so $\text{diam}_{\mathbb{G}}(A_1) \geq k + 1$. Since \mathbb{G} is of uniformly bounded degree the ball has a bounded volume, $\#V(A_1) \leq d^{k+1} =: U'$ and subsequently has a bounded number of edges

$$\#E(A_1) \leq \frac{U'(U' - 1)}{2} =: U,$$

so that A_1 satisfies the second requirement of (2.6). Now consider the graph $\mathbb{G} \setminus A_1$ with $V(\mathbb{G} \setminus A_1) = V(\mathbb{G})$ and $E(\mathbb{G} \setminus A_1) = E(\mathbb{G}) \setminus E(A_1)$. Whilst $\mathbb{G} \setminus A_1$ need not be connected, it contains an infinite component, \mathbb{G}_1 . Now \mathbb{G}_1 is infinite, connected and of uniformly bounded degree (with bound $\leq d$), so we construct $A_2 \subset \mathbb{G}_1$ as we did for A_1 . This construction can now be continued inductively by letting \mathbb{G}_k be an infinite component from $\mathbb{G}_{k-1} \setminus A_k$. The A_i are clearly disjoint.

Using Lemma 2.8 we obtain a lower bound on the probability of the event $\{\lambda_T(v) = k\}$, to which we apply the Borel–Cantelli lemma. By construction each A_i contains at least one path from the central vertex v_i to the boundary of A_i ; let P_i be a subpath obtained by removing the first and last edges from such a path. Then P_i is a tree on $k - 1$ edges and k vertices. From Lemma 2.8, if P_i is a 1-tree in Π_T then it defines a k -cycle, so for $v \in P_i$

$$\begin{aligned} \mathbb{P}_T[\lambda_T(v) = k] &\geq \mathbb{P}_T[P_i \text{ is a 1-tree}] \\ &\geq \mathbb{P}_T[\iota_T(e) = 1, e \in P_i; \iota_T(e') = 0, e' \in A_i \setminus P_i] \\ &= \prod_{e \in P_i} \mathbb{P}_T[\iota_T(e) = 1] \prod_{e' \in A_i \setminus P_i} \mathbb{P}_T[\iota_T(e') = 0], \end{aligned}$$

where we used independence of the bar sets $\{\mathcal{B}_e\}_{e \in E}$ to split the probability. Using the fact that the number of rings up to time T is Poisson-T distributed

$$\begin{aligned} &= (Te^{-T})^{k-1} (e^{-T})^{\#E(A_i \setminus P_i)} \\ &= T^{k-1} e^{-T\#E(A_i)}, \end{aligned}$$

and using the bounded volume condition of (2.6)

$$\geq T^{k-1} e^{-TU}.$$

The probability that there are infinitely many k -cycles is greater than the probability that infinitely many of the A_i contain k -cycles, which in turn is greater than the probability that infinitely many of the P_i are 1-trees. Since the balls are disjoint the events $\{P_i \text{ is a 1-tree}\}$ are independent. Finally, since the lower bound $T^{k-1} e^{-TU}$ is not summable we apply the Borel–Cantelli lemma to deduce the result. \square

Remark 2.12. The proof above extends to the stronger statement that at any fixed time there

are almost surely infinitely many cycles of all finite sizes: $\mathbb{P}_T[\lambda_T(k) = \infty, \forall k \geq 1] = 1$. Since the countable union of countable sets is itself countable we can find collections $\{A_i^{(k)}\}_{i=1}^\infty, k \geq 1$ such that $A_i^{(k)} \neq A_j^{(l)}$ for $i, j, k, l \geq 1, i \neq j$, where $A_i^{(k)}$ is a k -ball as in the above. For each k we define the event $A^{(k)} = \{P_i^{(k)} \text{ is a 1-tree i.o.}\}$, where i.o. is read as infinitely often, and $P_i^{(k)}$ is a k -path in the interior of $A_i^{(k)}$. The events $\{A^{(k)}\}_{k \geq 1}$ are independent since the balls are disjoint, and they are not summable since $\mathbb{P}_T[A^{(k)}] = 1$, and Borel–Cantelli insures that there are infinitely many k -cycles in σ_T for all k simultaneously.

The technique employed above extends to show that any finite configuration of labels occurs infinitely often in the projection graph (Π_T, ι_T) ; in line with Remark 2.10, general cases tend not to provide significant additional insight.

Ultimately we are more interested in the possibility of infinite cycles in the RSP, and at what times such cycles occur. Let $O_\infty := \{l_T(\infty) \geq 1\}$ denote the event that the RSP contains at least one infinite cycle. In [Ang03], Angel provides a heuristic description of why O_∞ is a zero-one event, a statement that we return to shortly. Consequently we define the set of critical times

$$\mathcal{T}_\mathbb{G} := \left\{ T \geq 0 : \mathbb{P}_T[O_\infty] = 1 \right\},$$

from which we derive two notions of critical times

$$\begin{aligned} T_c^I &= T_c^I(\mathbb{G}) := \inf\{T : T \in \mathcal{T}_\mathbb{G}\}, \\ T_c^S &= T_c^S(\mathbb{G}) := \sup\{T : T \in [0, \infty) \setminus \mathcal{T}_\mathbb{G}\}. \end{aligned}$$

If $\mathcal{T}_\mathbb{G} = \emptyset$ we define $T_c^I = T_c^S = \infty$. T_c^I is the unique time such that for all smaller times the RSP almost surely does not contain infinite cycles, whilst T_c^S is the unique time such that at all later times there are infinite cycles. Directly from the definition one sees that $T_c^I \leq T_c^S$, and that if $\mathcal{T}_\mathbb{G}$ is a one-sided interval, $\mathcal{T}_\mathbb{G} = (T^*, \infty)$ or $\mathcal{T}_\mathbb{G} = [T^*, \infty)$, then $T_c^I = T_c^S = T^*$. It seems reasonable to believe that once infinite cycles occur in the RSP then they do not disappear, which is to say that $\mathcal{T}_\mathbb{G}$ is a one-sided interval, this was conjectured in [Ang03, Ha12a].

Conjecture 2.13 (Sharp-Threshold). *Either $\mathcal{T}_\mathbb{G}$ is empty, or it is a one-sided interval, and $T_c^I(\mathbb{G}) = T_c^S(\mathbb{G})$.*

We now return to the claim that O_∞ is a zero-one event, for which the following simple proposition will be useful.

Proposition 2.14. *O_∞ is invariant under composition with finitely many transpositions. i.e. if σ contains an infinite cycle then for any $\{\tau_i\}_{i=1}^N \in \mathcal{T}_\mathbb{G}$, $\tau_N \circ \dots \circ \tau_1 \circ \sigma$ also contains an infinite cycle.*

Proof. Consider an arbitrary permutation $\sigma \in S_V$ on a countable set, and suppose $u \in V$ is in an infinite cycle, $\lambda_\sigma(u) = \infty$. We consider the effect of composing σ with $\tau = (u, v)$. If $v \notin \langle u \rangle_\sigma$ then in $\tau \circ \sigma$ the infinite cycle has grown since a merge has occurred; in particular the cycle is still infinite. If instead the transposition causes a fragmentation then the cycle is partitioned

into two, however at least one of the resulting cycles $\langle u \rangle_{\tau \circ \sigma}$, $\langle v \rangle_{\tau \circ \sigma}$ must be infinite, and so the event is preserved, albeit with a different infinite cycle. Clearly this carries over to any finite sequence $\{\tau_i\}_{i=1}^N$. \square

One would be forgiven for thinking that this ensures that the sharp-threshold conjecture must hold; unfortunately the problem lies in the fact that since an infinite orbit necessarily contains infinitely many vertices, in any interval of time infinitely many edges that are adjacent to the orbit will ring, possibly decomposing the cycle into finite ones.

We return to the claim that O_∞ is a zero-one event, $\mathbb{P}_T[O_\infty] \in \{0, 1\}$. Without constructing the relevant σ -algebras it is not possible to confirm this rigorously, though we consider some heuristics. In a spatial or graphical setting, Kolmogorov's zero-one law says (heuristically) that if an event does not depend on the behaviour of the model on any finite set (for instance a box containing finitely many vertices) then it must be a zero-one event. For the RSP a statement of this form is seen to hold for the event O_∞ : if we consider the ringing times on any finite set of edges then in any finite time interval these edges ring only finitely often, so by Proposition 2.14 the existence of an infinite cycle is independent of these edges almost surely.

Claim 2.15. *For $T \geq 0$, the existence of an infinite orbit is a zero-one event: $\mathbb{P}_T[O_\infty] \in \{0, 1\}$.*

We proceed to derive a trivial lower bound for T_c^I . Before describing the result we briefly recap some required notions from bond-percolation, and direct readers to [BR06] for a thorough study.

Bond percolation on \mathbb{G} is a measure on the space of mappings $\{\omega: E \rightarrow \{0, 1\}\}$, which we call *configurations*. The edge $e \in E$ is said to be open in the configuration ω if $\omega(e) = 1$, else it is closed. Each configuration defines a unique subgraph $G_p(\omega) \subseteq \mathbb{G}$ with $e \in E(G_p(\omega))$ if $\omega(e) = 1$. The best studied measure is the Bernoulli- p measure, $p \in [0, 1]$, in which an edge is open with probability p independently of all other edges, $\mu_p \in \mathcal{M}^1(\{0, 1\}^{E(\mathbb{G})})$

$$\mu_p(\omega(e) = 1) = p, \quad \forall e \in E.$$

If $C_v(\omega) \subset G_p(\omega)$ is the connected component of the configuration containing $v \in V$, we define

$$\theta_v(p) := \mu_p(\#C_v(\omega) = \infty), \quad p \in [0, 1],$$

the probability that v is in an infinite cluster. It is easily confirmed that for any $p \in [0, 1]$ either $\theta_v(p) = 0, \forall v \in V$ or $\theta_v(p) > 0, \forall v \in V$. Consequently we define $p_c = p_c(\mathbb{G})$ such that for all $v \in V$

$$\theta_v(p) \begin{cases} = 0, & \text{for } p < p_c, \\ > 0, & \text{for } p > p_c. \end{cases} \quad (2.7)$$

Note that we do not specify what happens at criticality. A simple application of the Kolmogorov zero-one law tells us that for $p > p_c$ there is almost surely at least one infinite cluster (which is in fact known to be unique), and we say that percolation occurs. Much work has been done to identify values of $p_c(\mathbb{G})$ the highlight being the Harris-Kesten theorem: $p_c(\mathbb{Z}^2) = 1/2$. Explicit

results are also known for other lattices, and for infinite regular trees.

Remark 2.16. We observe that the monotonicity of $\theta_v(p)$ is readily proven using a simple coupling argument; resultingly the analogue of the sharp-threshold conjecture for bond percolation is trivial (and is implicitly contained in (2.7)). This is just one example of where the complicated dependency structure in the RSP makes heuristically simple claims hard to prove.

Through relating the projection graph, Π_T , to bond percolation on \mathbb{G} we obtain a lower bound on the time at which infinite cycles can occur in the RSP. Though the proof is simple, we state it as a theorem since it is the only result known for critical times of the RSP on a general graph. We follow [Ang03], p.10.

Theorem 2.17 (Angel, [Ang03]). *For an infinite, connected, locally finite \mathbb{G}*

$$T_c^I(\mathbb{G}) \geq \log \left(\frac{1}{1 - p_c(\mathbb{G})} \right). \quad (2.8)$$

Proof. Fix $T \geq 0$. Since cycles of σ_T are connected sets in Π_T the size of the largest cycle is bounded above by the size of the largest component, and there is an infinite cycle only if Π_T has an infinite connected component. Considering Π_T , we recall that the edge $e \in E(\mathbb{G})$ is in Π_T only if $\mathcal{B}_e \cap [0, T] \neq \emptyset$; since the bar sets are independent of one another and Poisson distributed $\mathbb{P}_T[e \notin \Pi_T] = e^{-T}$, the probability that a Poisson- T variable is equal to 0. Subsequently Π_T has the same distribution as the graph G_p , $p = 1 - e^{-T}$. Rearranging we see that Π_T almost surely contains an infinite component if $T > -\log(1 - p_c(\mathbb{G}))$. \square

Considering $\mathbb{G} = \mathbb{Z}$, it is easily shown that $p_c(\mathbb{Z}) = 1$ and subsequently $T_c^I(\mathbb{Z}) = T_c^S(\mathbb{Z}) = \infty$. In two dimensions, as a consequence of the Harris-Kesten theorem we obtain $T_c^I(\mathbb{Z}^2) \geq \log 2$. Recall that in Section 1.2 we described the phase transition of the Heisenberg spin-1/2 ferromagnet in terms of the density of infinite cycles in the thermodynamic limit of the RSP, $\nu^\infty(\beta)$, $\beta > 0$. It is important to make a distinction between the cycles that occur in the thermodynamic limit, and those that are present in the RSP on an infinite graph; however, for general statistical mechanical models, if one is able to define an infinite volume limit one would expect critical values in this setting to agree with those of the thermodynamic limit. With this in mind, we recall that the Mermin-Wagner theorem, Theorem 1.15, assured us that the thermodynamic limit of the RSP on $\Lambda \rightarrow \mathbb{Z}^d$, $d = 1, 2$ does not undergo a phase transition, and $\nu^\infty(\beta) = 0$ for all $\beta > 0$. This would suggest that $T_c^I(\mathbb{Z}^2) = \infty$, highlighting the poor performance of the lower bound obtained from Theorem 2.17 in the case $d = 2$. To make the link between the thermodynamic limit and infinite volume models stronger we recall that Tóth defined the toy order parameter

$$m(\beta) := \lim_{n \rightarrow \infty} \lim_{\Lambda \rightarrow \mathbb{Z}^d} \mathbb{P}[\lambda_\beta(0) > n],$$

and conjectured that in $d \geq 3$ there is a unique $\beta_c \in (0, \infty)$ at which this parameter has a

singularity (Conjecture 1.16, page 18)

$$m(\beta) \begin{cases} = 0 & \beta < \beta_c, \\ > 0 & \beta > \beta_c. \end{cases}$$

It is generally believed that β_c should coincide with the critical times T_c^I, T_c^S , so that combined with the sharp-threshold conjecture, Conjecture 2.13, there is a single critical time of interest.

2.2.2 Transience and Recurrence of the CTRW

Throughout we have emphasised that the CTRW provides a tractable setting in which to study the occurrence of infinite cycles in the RSP. We continue by studying the key results which confirm this relationship, and which provide the foundations for Angel's proof that the RSP on a d -regular tree ($d \geq 5$) has infinite cycles. We closely follow [Ang03].

Recall that the CTRW is determined by the bar set $\mathcal{B}_{\mathbb{G}}^T$ of cyclic ringing times. The walk X_v^T stays at its current site until an adjacent edge rings $\mathcal{B}_{\mathbb{G}}^T$, at which time the walk jumps across this bar. We refer to this bar as the *effective bar*, i.e. the effective bar for the CTRW from v at time t , $X_v^T(t)$, is the first bar $(t_i, e_i) \in \mathcal{B}_{\mathbb{G}}^T$ such that $t_i > t$ and $X_v^T(t)$ is adjacent to the edge e_i .

By construction the reversed walk is also defined, and is equivalent to the CTRW on the bar set $\bar{\mathcal{B}}_{\mathbb{G}}^T$, where $(\bar{t}, e) \in \bar{\mathcal{B}}_{\mathbb{G}}^T$ if and only if $(- \bar{t} \bmod T, e) \in \mathcal{B}_{\mathbb{G}}^T$. Observing that the ringing times of $\bar{\mathcal{B}}_{\mathbb{G}}^T$ in $[0, T)$ are exactly those of $\mathcal{B}_{\mathbb{G}}^T$ reflected in the line $T/2$, then the two bar sets are equal in distribution since the Poisson process is invariant under reflection.

Claim 2.18. *For $T > 0$ the cyclic bar sets $\mathcal{B}_{\mathbb{G}}^T$ and $\bar{\mathcal{B}}_{\mathbb{G}}^T$ are equal in distribution, $\bar{\mathcal{B}}_{\mathbb{G}}^T \stackrel{(d)}{=} \mathcal{B}_{\mathbb{G}}^T$, and consequently $\sigma_T \stackrel{(d)}{=} \sigma_T^{-1}$.*

The relationship between the RSP and CTRW is given by $\sigma_T(v) = X_v^T(T)$, from which it follows that $\sigma_T^k(v) = X_v^T(kT)$, so that the orbit of v is completely determined by X_v^T . If the orbit of v is of period k , $\lambda_T(v) = k$, then $X_v^T(kT) = v$ and more generally $X_v^T(mkT) = v$, for all $m \geq 0$. This is a special case of the following more general observation.

Proposition 2.19. *If for some time $t > 0$ and integer $k \geq 0$, $X_v^T(t) = X_v^T(t + kT)$ then*

$$X_v^T(s) = X_v^T(s \bmod kT), \quad \forall s \geq 0. \quad (2.9)$$

Proof. Let (t_i, e_i) be the effective bar of $X_v^T(t)$, and denoting $\tilde{t} = t + kT$ let (t_j, e_j) be the effective bar for $X_v^T(\tilde{t})$. If $t_j \neq t_i + kT$ then either $t_j < t_i + kT$ or $t_j > t_i + kT$. The later of these cannot be the case since $(t_i + kT, e_i) \in \mathcal{B}_{\mathbb{G}}^T$ since the bar set is T -cyclic, and e_i is adjacent to $X_v^T(\tilde{t})$, so this bar must be effective. Similarly, if we suppose $t_j < t_i + kT$ then again using the cyclic structure of the bar set $(t_j - kT, e_j) \in \mathcal{B}_{\mathbb{G}}^T$; furthermore since $t_j \in [\tilde{t}, t_i + kt)$ then $t_j - kT \in [t, t_i)$, and since e_j is adjacent to $X_v^T(t)$ then this contradicts (t_i, e_i) being effective. So $t_j = t_i + kT$ and $e_j = e_i$, so that $X_v(s) = X_v(s + kT)$ for all $s \in [t, t_i]$.

Since $X_v^T(t)$ almost surely makes only finitely many jumps in $[t, t+kT]$, then by induction $X_v(s) = X_v(s+kT)$ for all $s \in [t, t+kT]$, so when $s = t+kT$ we have $X_v^T(t+2kT) = X_v^T(t+kT) = X_v^T$, iterating this we see that $X_v^T(s) = X_v^T(s + mkT)$ for any $m \geq 0$ and $s \in [0, kT]$. So this confirms (2.9) for $s \geq t$. To extend the statement for $0 \leq s < t$ one can apply the same reasoning to the reversed process on the bar set $\overline{\mathcal{B}}_{\mathbb{G}}^T$. \square

Corollary 2.20. *A site $v \in V(\mathbb{G})$ is in an infinite cycle of the RSP if and only if for all $t \geq 0$*

$$X_v^T(t) \neq X_v^T(t + kT), \quad \forall k \geq 1. \quad (2.10)$$

If X_v^T satisfies (2.9) then the walk is at v at time kT and periodicity of the bar set ensures the walk repeats the same path on $[mkT, (m+1)kT]$, $m \geq 0$. This resembles the behaviour of a recurrent random walk in that it does not ‘escape to infinity’. Following the definition of Angel, we say that the CTRW is (cyclically) recurrent if it almost surely satisfies the periodicity equation (2.9). If the CTRW is not recurrent then it is (cyclically) transient. We have the following alternative characterisations of recurrence.

Proposition 2.21 (Angel, [Ang03]). *The following are equivalent,*

- (i) X_v^T is recurrent, that is X_v^T is almost surely periodic in the sense of (2.9).
- (ii) X_v^T almost surely returns to its origin infinitely often.
- (iii) X_v^T almost surely visits only finitely many vertices.
- (iv) X_v^T almost surely returns to its origin.

Proof. It is immediately seen that (i) implies all of (ii), (iii), and (iv).

(ii) \Rightarrow (i) Suppose that X_v^T returns to v at times $\{t_i\}_{i=1}^{\infty}$, necessarily these are all ringing times of bars adjacent to v . Since $\deg_{\mathbb{G}}(v) < \infty$ and $T < \infty$ there are almost surely only finitely many rings adjacent to v in each interval $[mT, (m+1)T)$. By the pigeonhole principle, there exist $i, j \geq 1$, $i \neq j$ for which $t_i \bmod T = t_j \bmod T$, and so the walk is periodic.

(iii) \Rightarrow (i) Since the walk visits only finitely many sites there is a $u \in \{X_v^T(t) : t \geq 0\}$ which is visited infinitely often. The argument is now the same as the above since there are only finitely many rings adjacent to u modulo T .

(ii) \Leftrightarrow (iv) Clearly if X_v^T almost surely returns to the origin infinitely often then it almost surely returns to the origin. Suppose $\mathcal{B}_{\mathbb{G}}^T$ is such that X_v^T returns to v only finitely often, and let t be the final arrival time of X_v^T at v . Consider the bar set $\widetilde{\mathcal{B}}_{\mathbb{G}}^T$ obtained by translating the ringing times of $\mathcal{B}_{\mathbb{G}}^T$ by a factor of $-t \bmod T$, so that $(s, e) \in \widetilde{\mathcal{B}}_{\mathbb{G}}^T$ only if $(s + t \bmod T, e) \in \mathcal{B}_{\mathbb{G}}^T$. This bar set has the property that the CTRW started at v never returns to the origin, and since the Poisson ringing times are invariant under translation $\mathcal{B}_{\mathbb{G}}^T \stackrel{(d)}{=} \widetilde{\mathcal{B}}_{\mathbb{G}}^T$. So if there is a positive probability that the walk returns only finitely often, then there is a positive probability that the walk never returns to the origin, which is the contrapositive of (iv) \Rightarrow (ii). \square

As it stands, we know that the RSP contains an infinite cycle if there is a vertex $v \in V$ for which the associated CTRW is not periodic. This is a somewhat unsatisfactory state of affairs:

for us to confirm that the RSP has (or does not have) infinite cycles we must check the paths of an infinite collection of CTRWs. Fortunately this burden is eased since we will see that if a CTRW started from one vertex is transient, then so are the CTRWs started at any other vertex. Before proving this we make an observation which will be needed in the proof.

For adjacent vertices $u, v \in V$, $e = uv \in E$, let $(t, e) \in \mathcal{B}_{\mathbb{G}}^T$ be a bar between them. Local finiteness of \mathbb{G} ensures that there is an interval $I = [s_-, s_+]$ in which (t, e) is the unique bar adjacent to either u or v . We claim that the exact ringing time of e in I is in fact immaterial, i.e. if we replace (t, e) with (t', e) , $t' \in I$ then the obtained RSP at time $T > s_+$, is unaffected: $\sigma'_T = \sigma_T$.

Theorem 2.22 (Angel, [Ang03]). *The CTRW, X_v^T , started from $v \in V$ is transient if and only if X_u^T is transient, for all $u \in V$.*

Proof. We follow [Ang03], p.13. From Corollary 2.20, X_v^T is not periodic if and only if v is in an infinite cycle, so we prove

$$\mathbb{P}_T[\lambda_T(v) = \infty] > 0 \iff \mathbb{P}_T[\lambda_T(u) = \infty] > 0, \quad \forall u \in V.$$

Since \mathbb{G} is connected it is sufficient to show this for vertices u adjacent to v , since any further sites are joined by a finite path, and an elementary induction argument follows. Our aim is to show that if v is in an infinite cycle at time T then there is a positive probability that any adjacent vertex u is in this cycle, and hence has a positive probability of being in an infinite cycle. Suppose for a contradiction that X_u^T is recurrent, $\mathbb{P}_T[\lambda_T(u) < \infty] = 1$.

Consider the event that v has an infinite orbit whilst u has a finite orbit, and also the event that no edge adjacent to either u or v rings in $[T - \epsilon, T]$

$$A := \{\lambda_T(u) < \infty, \lambda_T(v) = \infty\}, \quad B_\epsilon := \{\mathcal{B}_e^T \cap [T - \epsilon, T] = \emptyset, \forall e \sim u, v\}$$

Let $\delta := \mathbb{P}_T[\lambda_T(v) = \infty]$ then we observe the Fréchet inequality

$$\mathbb{P}_T[A] \geq \mathbb{P}_T[\lambda_T(u) < \infty] + \mathbb{P}_T[\lambda_T(v) = \infty] - 1 = \delta.$$

Since B_ϵ is the event that the union of $\deg_{\mathbb{G}}(u) + \deg_{\mathbb{G}}(v) - 2$ independent edges do not ring in $[T - \epsilon, T]$, we can choose $\epsilon = \epsilon(\delta)$ small enough that $\mathbb{P}_T[B_\epsilon] > 1 - \delta/2$. Using the Fréchet inequality again

$$\mathbb{P}_T[A \cap B_\epsilon] \geq \mathbb{P}_T[A] + \mathbb{P}_T[B_\epsilon] - 1 > \delta/2. \tag{2.11}$$

Consequently $\mathbb{P}_T[A \mid B_\epsilon] = \mathbb{P}_T[A \cap B_\epsilon] / \mathbb{P}_T[B_\epsilon] > \delta/2$.

Consider instead the event B_ϵ^{uv} that in $[T - \epsilon, T]$ the edge $uv \in E$ rings exactly once, and no other edge adjacent to u or v rings. From the discussion prior to this proof, the distribution of σ_T is independent of the actual ringing time of (u, v) in $[T - \epsilon, T]$. Consequence sampling from the conditional distribution $\mathbb{P}_T[\cdot \mid B_\epsilon^{uv}]$ is equivalent to sampling $\mathbb{P}_T[\cdot \mid B_\epsilon]$ and then composing

with the transposition (u, v)

$$\sigma_T \sim \mathbb{P}_T[\cdot \mid B_\epsilon^{uv}] \quad \Rightarrow \quad (u, v)^{-1} \circ \sigma_T = (u, v) \circ \sigma_T \sim \mathbb{P}_T[\cdot \mid B_\epsilon].$$

It follows from the previous paragraph that $(u, v) \circ \sigma_T \in A$ with strictly positive probability; but this implies that in σ_T the cycles containing u and v are merged, and $\lambda_T(u) = \infty$, contradicting our assumption that X_u^T is recurrent. \square

We say that \mathbb{G} is T -recurrent if the X_0^T (on \mathbb{G}) is recurrent, else it is T -transient.

Corollary 2.23. *If \mathbb{G} is T -transient then $\mathbb{P}_T[O_\infty] = 1$.*

Chapter 3

The Critical Phase for Trees

Our final observation in the previous chapter identified the critical times $T_{\mathbb{G}}^I, T_{\mathbb{G}}^S$, with periods of the CTRW at which the walk is transitive. To date, little progress has been made in studying the RSP on infinite graphs, with the only sources (known to the author) being the paper of Angel [Ang03], and the preprints of Hammond [Ha12a, Ha12b]. Undoubtedly the case of most interest to probabilists and statistical mechanists alike is the case $\mathbb{G} = \mathbb{Z}^d$, $d \geq 2$, but the techniques developed so far are only suitable for the case where the graph is an infinite tree.

The aim of this chapter is to summarise the results for trees. In the first section we provide a detailed description of Angel's proof that the RSP on a regular tree has a critical phase during which there are infinite cycles. The second section turns to the work of Hammond, summarising his results concerning the sharp-threshold conjecture for trees of high degree.

3.1 T_c^I and a Critical Window

Recall that T_c^I is the first time at which infinite cycles occur in the RSP

$$T_c^I = \inf \{T \geq 0: \mathbb{P}_T[O_\infty] = 1\}.$$

Let $\mathbb{G} = \mathbb{T}$ be an infinite, locally finite, tree. If each vertex has the same degree, $d \geq 2$, we say that the tree is d -regular, and denote it \mathbb{T}_d . Infinite trees provide a geometry in which the critical probability for percolation is known $p_c(\mathbb{T}_d) = (d-1)^{-1}$, see [BR06] pp.6–8 (or Corollary 3.5 below). Theorem 2.17 gives

$$T_c^I = T_c^I(\mathbb{T}_d) \geq \log \left(\frac{d-1}{d-2} \right) = d^{-1} + \frac{3}{2}d^{-2} + O(d^{-3}). \quad (3.1)$$

The central result of this section is Angel's proof that for dimensions $d \geq 5$ there is a critical window $I_d \subset \mathcal{I}_{\mathbb{T}_d}$ in which the RSP has infinite cycles almost surely.

For general infinite trees we designate a root vertex for the tree, $0 \in V(\mathbb{T})$, and refer to the root as the origin. We write $X^T = X_0^T$ for the CTRW on \mathbb{T} started from the origin. We say that v is a descendent of u if the shortest path joining u to v does not contain the origin, and if

$d_{\mathbb{T}}(0, u) \leq d_{\mathbb{T}}(0, v)$. If in addition u and v are adjacent we say that v is the child of u , and that u is the father of v ; we denote $\bar{v} = u$ for the father. Angel's identification of a critical window relies on finding a configuration of bars which limits the behaviour of a CTRW, and which can be used to bound the probability of returning to the origin away from zero. This configuration is described via the notion of *covering vertices*.

Following Angel we say that $v \in V$ is *good* if the edge connecting v to its father, $\bar{v}v \in E$, has a unique ringing time $t_v \in [0, T)$. Suppose that \bar{v} and v are both good, with unique ringing times $t_{\bar{v}}, t_v \in [0, T)$. Let J denote the interval between these times,

$$J = [t_v \vee t_{\bar{v}}, t_v \wedge t_{\bar{v}}),$$

and let $J^c = [0, T) \setminus J$. Henceforth we refer to J^c as an interval, viewing it modulo T . If u is a sibling of v (that is it has the same parent, $\bar{u} = \bar{v}$) we say that it *covers* v if $\bar{v}u \in E$ rings in both J and J^c , $\mathcal{B}_{\bar{v}u}^T \cap J \neq \emptyset$, $\mathcal{B}_{\bar{v}u}^T \cap J^c \neq \emptyset$. Suppose for simplicity that \bar{v} has two children u, v , and consider a CTRW which arrives at \bar{v} for the first time (so it cannot yet have visited u or v). The walk must arrive at time $t_{\bar{v}}$, and so arrives at the start of either J or J^c . If u covers v then regardless of which of J or J^c we are in, the CTRW necessarily visits u before v . We are really interested in the case that v is not covered by any of its siblings. In the following when we say $t \notin J$ we take this to mean that $t \bmod T \notin J$.

Lemma 3.1 (Angel, [Ang03]). *Suppose that v and \bar{v} are both good vertices, and v is not covered by any siblings. If $X^T(t) = \bar{v}$ at some $t \geq 0$, then one of the following holds.*

- (i) *There exists $s > t$ such that $X^T(s) = v$.*
- (ii) *X^T never returns to \bar{v} . i.e. if t^* is the effective time for $X^T(t)$ then $X^T(s) \neq \bar{v}$ for all $s > t^*$.*

Proof. We follow [Ang03], p.14. Suppose that (ii) does not hold so that if X^T visits a descendent of \bar{v} it must return. Suppose without loss of generality that $t_{\bar{v}} < t_v$, so that X^T jumps to \bar{v} at a time in J . Let $\{u_i\}_{i=0}^n$ denote the descendants of \bar{v} , with $u_0 = v$. Since v is uncovered, for each $1 \leq i \leq n$ the ringing times of $\bar{v}u_i \in E$ are either only in J or only in J^c . If none of these are in J then t_v is the effective time for $X^T(t)$, and so the walk visits v . Else a collection of siblings do ring in J ; abusing notation we denote these $\{u_i\}_{i=1}^n$. Since $n < \infty$, and the interval $J = [t_{\bar{v}}, t_v)$ is bounded, there are almost surely only finitely many bars adjacent to \bar{v} in J . Since the walk arrives at \bar{v} at the leftmost time in J it must visit a sibling, u_1 . Given that (ii) does not hold, however, the walk must return to \bar{v} (i.e. it performs an excursion), and returns via $\bar{v}u_1$. By assumption the ringing times of this edge are all in J , so any excursion from \bar{v} to a child u_1 which starts at a time in J must return at a time which is in J .

Since there are only finitely many bars in J either the walk only performs finitely many excursions from \bar{v} which visit $\{u_i\}_{i=1}^n$, or it performs some of these excursions infinitely often. In the first of these cases, let $s \in J$ denote the time at which X^T returns to \bar{v} from its final excursion via a sibling of v . Then no bars adjacent to \bar{v} ring in $[s \bmod T, t_v)$, and so the

next jump of X^T is to v . On the other hand, if the walk performs some excursions infinitely often, then it visits some particular sibling infinitely often, which from Proposition 2.21 implies periodicity. If X^T is periodic it must return to the origin infinitely often, and so in particular it requires the walk to re-cross the edge $\bar{v}\bar{v}$ connecting \bar{v} to its father. But this has a unique bar $t_{\bar{v}}$, and so to return the walk must be at \bar{v} at $t_{\bar{v}} \in J^c$. But if we have never visited v then this is not possible, since any other bar returns X^T to \bar{v} at a time in J . \square

Remark 3.2. The proof relied on the fact that between any two vertices in a tree there is a unique shortest path. This meant that for a walk to return to a previously visited vertex, v , it must do so via the same edge it originally crossed when leaving v . This was used both in showing that excursions from v return in the interval J , and also in arguing that for the walk to return to the origin it must cross $\bar{v}\bar{v}$. The notion of covering vertices would not have been strong enough to provide this result on graphs which contain complicated cyclic structure. It seems reasonable, however, that this method could be extended to simple cyclic graphs, in which between any two vertices there is at most one cycle, such as cactus graphs.

Lemma 3.1 provides us with a situation in which transience of X^T is ensured.

Proposition 3.3. *Let $\{v_i\}_{i=1}^\infty$ be an infinite path in \mathbb{T} , with $\bar{v}_i = v_{i-1}$ for $i \geq 2$. Suppose the v_i are all good, and uncovered for $i \geq 2$. If X^T visits v_0 , it does not return to the origin.*

Proof. It suffices to show that once X^T visits v_k it does not return to v_{k-1} . Suppose $X^T(t) = v_k$ for some $t > 0$; since v_{k+1} is uncovered by its siblings then either (i) or (ii) hold. In the later case, the walk never returns to v_k , and so cannot return to v_{k-1} . On the other hand, if (i) holds then the walk visits v_{k+1} before it can return to v_{k-1} . However hit now goes on indefinitely, since at each step we either leave the path $\{v_i\}_{i=1}^\infty$ and never return, or we go further up the sequence. Either way the walk does not return to v_{k-1} . \square

The purpose of this result is transparent: if we can bound the probability that $\mathcal{B}_{\mathbb{T}}^T$ contains a good uncovered path $\{v_i\}_{i=1}^\infty$ away from zero, then this is sufficient for the CTRW to be transient. It remains to find a bound on the probability that such a path exists. Before proceeding we provide a brief summary of Galton–Watson processes, a detailed introduction is given in [LP12].

The homogeneous Galton–Watson process is a Markov chain $W = (W_n)_{n \geq 1}$, described by an offspring distribution $\rho \in \mathcal{M}^1(\mathbb{N})$. For $i \geq 1$, $n \geq 1$ let $\zeta_i^{(n)} \sim \rho$ be i.i.d., initialising $W_0 = 1$ we define

$$W_{n+1} := \sum_{i=1}^{W_n} \zeta_i^{(n+1)}.$$

The process W can be visualised as an evolving tree; in the initial generation ($n = 0$) we have just a root. At the first generation we sample $\zeta_1^{(1)} \sim \rho$ and endow the root with that many children. In generation n there are W_n new children, and for each of these we sample $\zeta_i^{(n+1)} \sim \rho$, $1 \leq i \leq W_n$ and append this many children. The total size of the tree up to time n is $\sum_{k=1}^n W_k$. If $\rho_0 = \rho(0) > 0$ then there is a positive probability that a vertex added

in the n -th generation will have no children, and more over the probability that there are no vertices added in generation $(n + 1)$ is $\rho_0^{W_n}$. We define the event $\{\exists n \text{ st. } W_n = 0\}$, and observe that if this occurs then $W_m = 0$, $\sum_{k=1}^m W_k < \infty$ for all $m \geq n$, and we say that the process is extinct. If extinction does not occur then the tree has an infinite branch. Whilst there are many interesting results for branching processes, we are interested only in the probability that extinction does or does not occur, this is the content of the following classical theorem, see for example [LP12] pp.143–4.

Theorem 3.4. *Let W_n be a homogeneous Galton–Watson process with offspring distribution ρ , $\rho(1) \neq 1$. Then*

$$\mathbb{P}[\exists n \text{ st. } W_n = 0] < 1 \iff \sum_{k=1}^{\infty} k\rho(k) \leq 1.$$

Note that the sum on the right hand side is exactly the expected number of offspring of a vertex. As a first application we can confirm the critical probability for bond percolation on \mathbb{T}_d .

Corollary 3.5. *For a d -regular tree, $d \geq 2$, $p_c(\mathbb{T}_d) = (d - 1)^{-1}$.*

Proof. Consider \mathbb{T}_d to be rooted, so that every vertex $v \neq 0$ has a unique parent and $d - 1$ children. Let $p \in [0, 1]$ denote the percolation probability, so the probability that v is connected to any child is p independently of all other vertices. Consequently the total number of children of v in the percolation graph is binomial $B(d - 1, p)$, independently of all $u \in V \setminus \{v\}$. Subsequently the cluster containing the origin is a Galton–Watson tree with offspring distribution $B(d - 1, p)$, where we ignore, for clarity, the disparity of the parity at the origin. Since the mean of the binomial distribution is $(d - 1)p$, the result follows. \square

Our aim is to show that the sub-tree of good uncovered vertices which are connected to 0 forms a Galton–Watson tree, and subsequently X^T is transient if this tree is super-critical. For a good parent vertex $\bar{v} \in V(\mathbb{T})$ let $N_{\bar{v}}$ denote the number of good uncovered children of \bar{v} . Since we have not introduced the requisite σ -algebras to make the following result rigorous we state it as a claim.

Claim 3.6 (Angel, [Ang03]). *For any good parent $\bar{v} \in V$, $N_{\bar{v}}$ is independent of $t_{\bar{v}}$, and further is independent of N_u , for any other good parent $u \in V$.*

Proof. We expand on the sketch given by Angel, [Ang03] pp.14–5. The event that $v \in V$ is good and uncovered is a local property in that it depends only on those edges that are adjacent to v , and more over is not dependent on the specific ringing times of the edges, but only on their relative ordering. Consequently the event is invariant under a uniform shift of all of the ringing times. Since the distribution of the bar set is invariant under shifts it follows that v being uncovered is independent of the ringing time, $t_{\bar{v}}$, of its parent. Consequently $N_{\bar{v}}$ is independent of $t_{\bar{v}}$.

If u is a second good parent with child w , by the locality property the event that w is uncovered is independent of v being uncovered if they are not adjacent. It follows that if u and

\bar{v} are not adjacent then N_u and $N_{\bar{v}}$ are independent. Without loss of generality, the remaining case is that u is a good child of \bar{v} . For a child w of u , the only edge that has an effect on both $N_{\bar{v}}$ and on whether w is good and uncovered is $uw \in E$, which has the unique ringing time t_w . But we have already seen that w being uncovered is independent of t_w . It follows that N_u is independent of $N_{\bar{v}}$. \square

If $\mathbb{G} = \mathbb{T}_d$ then for any good uncovered vertices $\{v_i\}_{i=1}^\infty$, N_v is i.i.d. and we denote its law $\mathcal{N} = \mathcal{N}(T) \in \mathcal{M}^1(\{0, 1, \dots, d-1\})$; we also need the law $\mathcal{N}_0 = \mathcal{N}_0(T) = \mathcal{M}^1(\{0, 1, \dots, d\})$ which is the expected number of good children of the root. Since a vertex v is good if the edge $\bar{v}v$ has a unique ringing time in $[0, T]$ it is seen that $\mathcal{N}_0 = B(d, p)$ is binomially distributed with $p = Te^{-T}$.

Theorem 3.7 (Angel, [Ang03]). *For $\mathbb{G} = \mathbb{T}_d$, $d \geq 3$, if $N \sim \mathcal{N}$ is such that $\mathbb{E}_T[N] > 1$, then $\mathbb{P}_T[O_\infty] = 1$. Further*

$$\mathbb{E}_T[N] = (d-1)e^{-T} \int_0^T (e^{-c} + e^{c-T} - e^{-T})^{d-2} dc > 1$$

Proof. We follow [Ang03], p.15. Fix $T > 0$ and consider the random sub-graph of \mathbb{T}_d on the same vertex set $V(\mathbb{T}_d)$, but with the edge $\bar{v}v \in E(\mathbb{T}_d)$ present only if both \bar{v} and v are good and uncovered; for vertices v that are children of the root, we add $0v \in E(\mathbb{T}_d)$ if v is good. Let \mathbb{T} denote the connected component of this graph which contains the root. If \mathbb{T} is infinite then this means that there is an infinite sequence of good uncovered vertices, and by Lemma 3.1 the CTRW defined by $\mathcal{B}_{\mathbb{T}_d}^T$ is not periodic.

Let v be a child of 0 in the connected component \mathbb{T} , by construction the number of children of v has law \mathcal{N} . Moreover this is the case independently for any vertex $u \in V(\mathbb{T}) \setminus \{0\}$, and we see that the branch started from v is a Galton–Watson tree with offspring distribution \mathcal{N} . The tree \mathbb{T} itself is not a homogeneous Galton–Watson tree, since the offspring distribution of 0 is \mathcal{N}_0 . Ultimately the different distribution of the root is immaterial since

$$\begin{aligned} \mathbb{P}_T[\mathbb{T} \text{ is finite}] &= \sum_{k=1}^d \mathbb{P}_T[0 \text{ has } k \text{ children}] \mathbb{P}_T[\mathbb{T} \text{ is finite} \mid 0 \text{ has } k \text{ children}] \\ &= \sum_{k=1}^d \binom{d}{k} p^k (1-p)^{d-k} \prod_{i=1}^k \mathbb{P}_T[\text{Galton–Watson } \mathcal{N}\text{-tree is finite}] \end{aligned}$$

and consequently $\mathbb{P}_T[\mathbb{T} \text{ is finite}] = 1$ if and only if $\mathbb{P}_T[\text{Galton–Watson } \mathcal{N}\text{-tree is finite}] = 1$. That is to say, \mathbb{T} has an infinite branch if and only if $\mathbb{E}_T[N] > 1$, where $N \sim \mathcal{N}$. In short, the CTRW is transient (equivalently the RSP has infinite cycles) if $\mathbb{E}_T[N] > 1$.

To identify the formula for $\mathbb{E}_T[N]$, consider a good child v of the root 0, and let $\{u_i\}_{i=1}^{d-1}$

be the children of v in \mathbb{T}_d . Using the symmetry of the children

$$\begin{aligned}
\mathbb{E}_T[N_v] &= \sum_{i=1}^{d-1} \mathbb{P}_T[u_i \text{ is good and uncovered}] \\
&= (d-1)\mathbb{P}_T[u_1 \text{ is good and uncovered}] \\
&= (d-1)\mathbb{P}_T[u_1 \text{ good}] \mathbb{P}_T[u_1 \text{ uncovered} \mid u_1 \text{ good}] \\
&= (d-1)Te^{-T}\mathbb{P}_T[u_1 \text{ uncovered} \mid u_1 \text{ good}].
\end{aligned}$$

Considering the conditional probability term, without loss of generality we can assume that $t_v = 0$ (by shift invariance), so that $J = [0, t_{u_1}]$. Suppose that the unique ringing time of $vu_1 \in E$ is fixed $t_{u_1} = c \in [0, 1]$, the sibling u_j , $1 \leq j \leq d-1$ covers u_1 only if it rings in both the interval $[0, c)$ and $[c, T)$; since the intervals are disjoint this is the product of the probability that a rate c and a rate $T-c$ Poisson variable are both non-zero,

$$\mathbb{P}_T[u_j \text{ covers } u_1 \mid t_{u_1} = c] = (1 - e^{-c})(1 - e^{c-T}).$$

It follows that conditioned on $t_u = c$ the probability that u is uncovered is

$$\begin{aligned}
\mathbb{P}_T[u_1 \text{ uncovered} \mid t_{u_1} = c] &= \prod_{j=2}^{d-1} \mathbb{P}_T[u_j \text{ does not cover } u_1 \mid t_{u_1} = c] \\
&= \left(1 - (1 - e^{-c})(1 - e^{c-T})\right)^{d-2}
\end{aligned}$$

Conditioned on the edge vu ringing once, the actual ringing time t_u is uniform in $[0, T]$, hence

$$\begin{aligned}
\mathbb{E}_T[N_v] &= (d-1)Te^{-T}\mathbb{P}_T[u_1 \text{ uncovered} \mid u_1 \text{ good}] \\
&= (d-1)Te^{-T} \int_0^T \frac{1}{T} \left(1 - (1 - e^{-c})(1 - e^{c-T})\right)^{d-2} dc \\
&= (d-1)e^{-T} \int_0^T (e^{-c} + e^{c-T} - e^{-T})^{d-2} dc. \tag{3.2}
\end{aligned}$$

□

The ground work is done, and it remains for us to analyse when the expectation exceeds 1. Following [Ang03] p.15, we observe that the integrand is symmetric about $T/2$, and further is bounded from below by

$$e^{-c} + e^{c-T} - e^{-T} \geq e^{-c} \wedge e^{c-T}, \quad c \in [0, T],$$




Figure 3.1: Plots of $\mathbb{E}_T[N]$ for $d = 3$ (dotted), 5 (dashed), 10 (solid).

from which

$$\begin{aligned} \mathbb{E}_T[N] &\geq 2(d-1)e^{-T} \int_0^{T/2} e^{-c(d-2)} dc \\ &= \frac{2d-2}{d-2} (e^{-T} - e^{-Td/2}) \\ &\geq 2(e^{-T} - e^{-Td/2}). \end{aligned}$$

We can use this lower bound to obtain an initial critical window I_d for sufficiently large d . We look to see for which values of $T > 0$, $d \geq 3$ the following inequality holds

$$e^{-T} - e^{-Td/2} > \frac{1}{2}, \quad (3.3)$$

for any pair (T, d) for which it is true then $T \in \mathcal{I}_{\mathbb{T}_d}$. Rearranging the inequality for d we obtain

$$d > -\frac{2}{T} \log \left(e^{-T} - \frac{1}{2} \right) =: f(T),$$

the right hand of which has singularities at $T = 0$ and $\log 2$. Considering $f: (0, \log 2) \rightarrow \mathbb{R}$ it is seen to be continuous, positive, and has a unique minima $d_{\min} \approx 8.8$ (this could be shown rigorously, but we choose to believe Fig. 3.2). It follows by continuity that for any $d \geq 9$ there is an interval $I_d \subset (0, \log 2)$ for which $d_{\min} \leq f(T) < d$ for $T \in I_d$; that is to say for $d \geq 9$ inequality (3.3) holds, and so the RSP has infinite cycles for $T \in I_d$. Angel takes this further and identifies that for $d \geq 9$ a suitable interval is given by $I_d = [4/d, 1/2]$; it is easily confirmed that $f(T) < d$ at both end points of this interval, and since f has a unique minima it follows that $f(T) < d$ for all $T \in I_d$.

Via asymptotic analysis of (3.2), Angel claims that $\mathbb{E}_T[N] = 2(e^T - 1)^{-1} + O(d^{-1})$; he




Figure 3.2: Plot of $f : (0, \log 2) \rightarrow \mathbb{R}$ as defined in (3.3); we see that for any $T \in (0, \log 2)$ and sufficiently large enough d the RSP σ_T on \mathbb{T}_d has infinite cycles almost surely.

continues by remarking that this enables us to extend the upper bound of the interval I_d up to $\log 3 - \epsilon$, where $\epsilon = \epsilon(d)$ is arbitrarily small for high enough d . Further, by approximating the integrand via a cut-off Taylor expansion he lowers the rightmost bound

$$I_d = \left[d^{-1} + \left(\frac{13}{6} + \epsilon \right) d^{-2}, \log 3 - \epsilon \right], \quad d \geq 9. \quad (3.4)$$

We see that the lower bound (3.1) performs well, matching the first order term of the upper bound on T_c^I found above, and provides a bound on the second order term

$$\frac{3}{2} \leq \liminf_{d \rightarrow \infty} d^2 (T_c^I - d^{-1}) \leq \limsup_{d \rightarrow \infty} d^2 (T_c^I - d^{-1}) \leq \frac{13}{6}. \quad (3.5)$$

Remark 3.8. We briefly summarise the requirements that were put on the tree. For the proof of Theorem 3.7 we required that the tree is d -regular for $d \geq 3$. Here d -regularity was required to ensure that the distribution of N_v for any good vertex v was identically distributed: if the tree was not regular this would no longer have been true, and the tree would no longer have been a homogeneous Galton–Watson tree. If we replaced this with the requirement that the degree of the vertices is bounded above and below then with some small modifications a similar result would hold. Clearly the proof cannot work in the case of a 2-regular tree, since there is no notion of being an uncovered vertex; this is a trivial problem, however, since $\mathbb{T}_2 = \mathbb{Z}$ and we have already seen that $T_c^I(\mathbb{Z}) = \infty$.

The requirement that $d \geq 9$ came in during our attempt to bound the expectation $\mathbb{E}_T[N]$, and was required to ensure that the inequality (3.3) holds for at least some T . Rather than bounding $\mathbb{E}_T[N]$ to find the suitable interval one could numerically approximate the integral

(3.2) for fixed d . This was done in [Ang03] and intervals I_d were found for $5 \leq d \leq 8$

$$I_5 = [0.40, 1.11], \quad I_6 = [0.28, 1.23], \quad I_7 = [0.22, 1.28], \quad I_8 = [0.18, 1.30].$$

The cases $d = 3, 4$ are not resolvable using this method. The explanation of this is seen in Fig. 3.1 which shows the graphs of $\mathbb{E}_T[N]$ for three different values of d . The dotted line corresponds to the case $d = 3$ and it is clear that there is no value of $T > 0$ for which $\mathbb{E}_T[N] > 1$; a similar situation occurs for $d = 4$. This is not to say that there is not a critical phase for $d = 3, 4$, but just that aperiodic CTRWs do not arise as a result of an infinite sequence of uncovered vertices: this is just one of many configurations that could cause the walk to be aperiodic.

3.2 T_c^S and the Sharp-Threshold Conjecture

Having established that the RSP on \mathbb{T}_d , $d \geq 5$, undergoes a critical phase it is natural to ask whether trees satisfy the sharp-threshold conjecture, Conjecture 2.13. Recall that T_c^S is the last time that the RSP has only finite cycles

$$T_c^S := \sup \{T \geq 0: \mathbb{P}_T[O_\infty] = 0\}.$$

As already observed, if the sharp-threshold conjecture holds and $T_c^I = T_c^S =: T_c$ then the set of super-critical times takes the form of a single infinite interval $\mathcal{S}_{\mathbb{T}_d} = [T_c, \infty)$, or $\mathcal{S}_{\mathbb{T}_d} = (T_c, \infty)$. In two pre-prints [Ha12a, Ha12b], Hammond proves the sharp-threshold conjecture for regular trees of high enough degree. Hammond's result is obtained by a 'patchwork' approach involving the following steps.

- Analysis for large T .
Find an upper bound $T_0 \geq T_c^S$, with $T_0 = T_0(d)$, so that the RSP has infinite cycles for all sufficiently large times, $[T_0, \infty) \subseteq \mathcal{S}_{\mathbb{T}_d}$.
- Monotonicity around T_c^I .
Angel's critical window (3.4), provides an upper bound $I_d^- \geq T_c^I$, where I_d^-, I_d^+ are the left and right boundaries of I_d . Confirm that $T_c^S \notin (T_c^I, I_d^+]$.
- Patch together bounds.
Hold ones breath and hope that $T_0 \leq I_d^+$, ensuring that $T_c^S = T_c^I$.

The first of these steps is the content of Hammond's shorter article [Ha12a], whilst [Ha12b] is dedicated to the monotonicity result.

3.2.1 Analysis for Large T

We saw that Angel provided details about the nature of $\mathcal{S}_{\mathbb{T}_d} \cap [0, \log 3]$, and identifies an asymptotic bound (3.5) on the second order term of T_c^I . In [Ha12a], Hammond considers instead the

behaviour of the CTRW when the period T is large, obtaining both qualitative and quantitative results. The first of which is the following.

Theorem 3.9 (Hammond, [Ha12a]). *Let \mathbb{T} be of uniformly bounded degree, for some $d \geq 3$*

$$3 \leq \deg_{\mathbb{T}}(v) \leq d, \quad \forall v \in V(\mathbb{T}).$$

There exists $T_0 \geq 0$ such that X^T is transient for all $T \geq T_0$.

Similarly to Angel's proof of Lemma 3.1, Hammond relies on finding a configuration of bars which limit the freedom of the CTRW. Consider the trajectory of the walk X^T up to a time $S > 0$. Suppose that before time S the walk arrives at a vertex \bar{v} for the first time, and from there travels to a child v across the bar $(t, \bar{v}v) \in \mathcal{B}_{\mathbb{T}}^T$; suppose in addition that the time spent at \bar{v} before jumping to v is less than $T/2$. The bar $(t, \bar{v}v)$ is said to be *useful* if the next vertex visited after v is a child of v , and the walk does not return to v before time S . We denote U_S for the number of bars which are useful at time S .

Remark 3.10. Whilst this description is somewhat of a mouthful, the idea is simple: after the walk crosses the bar $(t, \bar{v}v)$ it visits a child of v , and is then contained in the sub-tree of strict descendants of v for all times before S .

We observe several properties afforded by useful bars. If there is a positive probability that the CTRW induced by $\mathcal{B}_{\mathbb{T}}^T$ has a useful bar (for some $S > 0$) which is never recrossed, then necessarily the CTRW is transient. This is simply because if the edge $e = \bar{v}v$ is crossed only once then the walk remains in the sub-tree of descendants of \bar{v} at all later times, and so cannot be periodic. For this reason when proving transience focus is placed on the case in which we condition that the CTRW returns to useful bars.

If the CTRW is recurrent then at all times S there are a bounded number of useful bars; once the walk has completed an orbit then $U_S = 0$, for all later S .

If $(t, \bar{v}v)$ is useful at time S and the CTRW eventually returns to v then the excursion performed between the two visits to v is independent of the history of the walk up to time t ; this is trivially true since the conditions of being useful ensure that any edges crossed during the excursion were not crossed before time t .

Finally we note that conditioned on X^T returning to v , there is a positive probability that the walk visits a previously undiscovered vertex. This is the content of Lemma 2.5 of [Ha12a] pp.8–10, where Hammond shows that there is a uniform bound on the probability that conditioned on X^T returning to v , the walk next visits a sibling of v , or it visits \bar{v} and then a sibling of \bar{v} . Due to our condition that $\deg_{\mathbb{T}}(\bar{v}), \deg_{\mathbb{T}}(v) \geq 3$ we know that they both have siblings, and further more by definition of $(t, \bar{v}v)$ being useful, neither of these siblings had been previously visited.

We provide a rough heuristic for the proof, for convenience let $\mathbb{T} = \mathbb{T}_d$, $d \geq 3$. Suppose that $T \gg 1$ and consider the path of X^T during $[0, T)$, which is that of a standard continuous time random walk. For large T , the walk has taken approximately T steps, each of which is away

from the root with probability $1 - d^{-1}$. The expected height at time T is then $T(1 - d^{-1})$, since the expected number of ‘up’ steps is binomially distributed. Since most of the steps are across previously unvisited bars, the trajectory of X_T^T is likely to contain many useful bars, and the probability that U_T exceeds a certain (large) value can be bounded arbitrarily close to 1 for large enough T , Lemma 2.11 [Ha12a] pp.12–3.

Consider the first time at which the walk first returns to a T -useful bar, at which point there is a positive probability that the walk next visits a previously unvisited vertex. If it does not, then the walk continues back down the tree, and will encounter another useful bar. Since there are many useful bars it is likely that we will often reach a useful bar, and from there visit a previously unvisited vertex.

If at time t we visit a previously unvisited vertex, u , then the walk at this point is (almost) independent of its past. Moreover, if the walk then goes higher in the tree, and does not return to u in $[t, t + T]$ then this walk is an independent continuous time random walk, and as before, scales a large height, and creates many useful bars. The exact number of bars, U_{t+T} , is equal to the number of useful bars at time T , minus those that are wasted before visiting u , plus the useful bars created in $[t, t + T]$. The hope is that on each excursion up a new branch of the tree the walk creates more useful bars than it wastes on its way back down the tree, and in this way never uses up all the useful bars: in particular this means the walk cannot return to the root. In Proposition 1.3 [Ha12a] pp.15–7, Hammond shows that this is the case by constructing a sequence of times $\{S_i\}_{i=1}^\infty$ for which $U_{S_i} \rightarrow \infty$ as $i \rightarrow \infty$, ensuring that X^T cannot be recurrent since the number of useful bars is not bounded.

In the concluding pages of [Ha12a], Hammond makes some quantitative statements that hold for trees of large enough dimension. Suppose that the degree of each vertex is bounded both from above and below: $d \leq \deg_{\mathbb{T}}(v) \leq d'$.

Theorem 3.11 (Hammond, [Ha12a]). *The RSP on \mathbb{T} has the following critical windows.*

(i) For $d \geq 40$, $[429d^{-1}, \infty) \subset \mathcal{I}_{\mathbb{T}}$.

(ii) If $\mathbb{T} = \mathbb{T}_d$ is d -regular, $d \geq 1287$ then $[d^{-1} + (13/6 + \epsilon)d^{-2}, \infty) \subset \mathcal{I}_{\mathbb{T}_d}$.

Here $\epsilon = \epsilon(d)$ can be made arbitrarily small for large enough d .

See [Ha12a] pp.21–22. To obtain (ii) Hammond already utilises a patchwork approach, by showing for $d \geq 1287$ the interval (3.4) obtained by Angel, overlaps with the one-sided interval obtained in (i) of the above.

3.2.2 Monotonicity around T_c^I

Theorem 3.11 makes the sharp-threshold conjecture seem legitimate (at least for large d), since we have the equivalent of (3.5) for T_c^S

$$\frac{3}{2} \leq \liminf_{d \rightarrow \infty} d^2(T_c^S - d^{-1}) \leq \limsup_{d \rightarrow \infty} d^2(T_c^S - d^{-1}) \leq \frac{13}{6}.$$

To confirm Conjecture 2.13 it remains to show $(T_c^I, d^{-1} + (13/6 + \epsilon)d^{-2}) \subset \mathcal{I}_{\mathbb{T}_d}$. Hammond proves the following.

Theorem 3.12 (Hammond, [Ha12b]). *Let $d \geq 1641$, and $0 < T < T' < d^{-1} + 3d^{-2}$. If $T \in \mathcal{T}_{\mathbb{T}_d}$ then $T' \in \mathcal{T}_{\mathbb{T}_d}$.*

Define $p_\infty: [0, \infty) \rightarrow [0, 1]$

$$p_\infty(T) := \mathbb{P}_T[\langle 0 \rangle_T = \infty],$$

the probability that the CTRW started from 0 (or any other vertex) is not periodic (so that recurrence of the CTRW, X^T , is equivalent to $p_\infty(T) = 0$). The proof of Theorem 3.12 is achieved by showing that p_∞ is monotone increasing on $[0, d^{-1} + 3d^{-2}]$ when $d \geq 1641$. To confirm this one wants to show that p_∞ is differentiable on $[0, d^{-1} + 3d^{-2}]$, and moreover that it has positive derivative. Whilst we make no attempt to prove Theorem 3.12 we describe the objects used, from which at least statements of the propositions can then be understood. Recalling our various descriptions of recurrence in Proposition 2.21, we note that $p_\infty(T)$ is the probability that X^T walks an unbounded distance away from the origin. Letting $p_N(T)$ denote the probability that the walk reaches a height N (from the origin), one sees that $p_N(T) \geq p_\infty(T)$ for all $T \geq 0$, and further $p_\infty(T) = \lim_{N \rightarrow \infty} p_N(T)$, where the convergence is pointwise. Considering the approximation p_N to p_∞ , Hammond obtains an expression for the derivative of $d/dT p_N(T)$ and proceeds to show that this is positive on the required interval. The equation for the derivative is analogous to Russo's formula of lattice percolation, [BR06] pp.46-7, and is written in terms of *pivotal bars*.

Independently of $\mathcal{B}_{\mathbb{T}}^T$ sample an additional bar (e, t) , where $t \in [0, T]$ is uniform and $e = \bar{v}v$ is chosen uniformly amongst those edges at distance at most N from the root, $d(0, v) \leq N$. We are interested in how the additional bar affects the path of the CTRW started from 0, and in particular whether it causes (or stops) the walk from reaching height N , in which case we would say that the bar is pivotal. We denote

$$\begin{aligned} \mathbf{P}_N^+ &:= X^T \text{ reaches height } N \text{ on the bar set } \mathcal{B}_{\mathbb{T}}^T \cup (e, t), \text{ but not on } \mathcal{B}_{\mathbb{T}}^T, \\ \mathbf{P}_N^- &:= X^T \text{ reaches height } N \text{ on the bar set } \mathcal{B}_{\mathbb{T}}^T, \text{ but not on } \mathcal{B}_{\mathbb{T}}^T \cup (e, t). \end{aligned}$$

Remark 3.13. Rigorous definitions of the events $\mathbf{P}_N^+, \mathbf{P}_N^-$ are simple to state, however rely on additional notation being introduced for the hitting time of height N , [Ha12b] p.7.

Lemma 3.14 (Hammond, [Ha12b]). *For a d -regular tree, $d \geq 3$, and $N \in \mathbb{N}$, $p_N: (0, \infty) \rightarrow [0, 1]$ is differentiable, and for $T > 0$*

$$\frac{d}{dT} p_N(T) = \#E_N \left(\mathbb{P}_T[\mathbf{P}_N^+] - \mathbb{P}_T[\mathbf{P}_N^-] \right), \quad (3.6)$$

where $\#E_N = d((d-1)^N - 1)/(d-2)$ is the number of edges in \mathbb{T}_d up to height N .

With this formula in hand it remains to show positivity of the derivative in $[0, d^{-1} + 3d^{-2}]$. Such a result is given for trees of high enough degree, $d \geq 1641$, Proposition 1.9 [Ha12b] p.8, from which Theorem 3.12 follows. Clearly at the heart of this positivity is the statement that

$\mathbb{P}_T[\mathbf{P}_N^+] \geq \mathbb{P}_T[\mathbf{P}_N^-]$, which Hammond is quick to point out is non-trivial, and is the main occupation of [Ha12b].

Ultimately, by combining the results of [Ang03, Ha12a, Ha12b] the following picture emerges.

Theorem 3.15. *For $d \geq 1641$, $T_c^I(\mathbb{T}_d) = T_c^S(\mathbb{T}_d) =: T_c(d)$. For $\epsilon > 0$ there exists $d_\epsilon \in \mathbb{N}$ such that for $d \geq d_\epsilon$*

$$T_c(d) \in \left[d^{-1} + \frac{3}{2}d^{-2}, d^{-1} + \left(\frac{13}{6} + \epsilon\right)d^{-2} \right].$$

We summarise what is known for the random stirring process on trees. Results are stated in order of decreasing strength, and are for d -regular trees; we abbreviate $\mathcal{T} = \mathcal{T}_{\mathbb{T}_d}$.

$d \geq d_\epsilon$	$[T_c, \infty) \in \mathcal{T}$ [Ha12b] $T_c \in [d^{-1} + \frac{3}{2}d^{-2}, d^{-1} + (\frac{13}{6} + \epsilon)d^{-2}]$ [Ang03]
$d \geq 1641$	$[T_c, \infty) \in \mathcal{T}$ [Ha12a] $T_c \in [d^{-1} + \frac{3}{2}d^{-2}, d^{-1} + 3d^{-2}]$ [Ang03, Ha12a]
$d \geq 1287$	$[d^{-1} + 2d^{-1}, \infty) \in \mathcal{T}$ [Ha12b]
$d \geq 40$	$[429d^{-1}, \infty) \in \mathcal{T}$ [Ha12a] $[d^{-1} + 3d^{-2}, 2d^{-1}] \in \mathcal{T}$ [Ang03, Ha12a]
$d \geq 9$	$[4d^{-1}, \frac{1}{2}] \in \mathcal{T}$ [Ang03]
$d = 8$	$[0.18, 1.30] \in \mathcal{T}$ [Ang03]
$d = 7$	$[0.22, 1.28] \in \mathcal{T}$ [Ang03]
$d = 6$	$[0.28, 1.23] \in \mathcal{T}$ [Ang03]
$d = 5$	$[0.40, 1.11] \in \mathcal{T}$ [Ang03]
$d \geq 3$	$[T_0(d), \infty) \in \mathcal{T}$ [Ha12a]

Chapter 4

Conclusion

Our aim has been to survey both the physical development of the RSP from its inception via the Feynman–Kac representation, through to the recent probabilistic studies of the process on infinite graphs. In this conclusion we finish with two outstanding sections. First we return to comment on the progress that has been made on studying the RSP on finite graphs, before concluding our study on infinite graphs by stating several open results which demand attention.

4.1 Random Stirring on Finite Graphs

In Section 1.2 we introduced Tóth’s formula for the spontaneous magnetisation of the Heisenberg spin-1/2 ferromagnet

$$m^{(\text{sp})}(\beta) := \frac{1}{2} \lim_{n \rightarrow \infty} \lim_{\Lambda \rightarrow \mathbb{Z}^d} \frac{\mathbb{E}_\beta \left[\mathbf{1}_{\{\lambda_\beta(0) > n\}} 2^{\sum_{k \geq 1} l_\beta(k)} \right]}{\mathbb{E}_\beta \left[2^{\sum_{k \geq 1} l_\beta(k)} \right]},$$

as well as the toy approximation

$$m(\beta) := \lim_{n \rightarrow \infty} \lim_{\Lambda \rightarrow \mathbb{Z}^d} \mathbb{P}_\beta [\lambda_\beta(0) > n].$$

Since these are defined in terms of the thermodynamic limit $\Lambda \rightarrow \infty$, it is of interest to study the RSP on finite graphs, and in particular on lattice boxes Λ . One would expect, however, that if one can show relevant phase transitions on infinite graphs then these should coincide with those obtained from taking thermodynamic limits.

As was observed in Section 2.1, the RSP on a finite connected graph is an irreducible finite state space Markov chain, with symmetric transition matrix. Consequently it is ergodic, and converges to the uniform distribution on \mathcal{S}_G . When we talk of convergence of distributions we mean in *total variation*, where for two measures μ, ν on the same measurable space (Ω, \mathcal{F})

the total variation distance is

$$\|\mu - \nu\|_{\mathbf{TV}} := \sup_{A \in \mathcal{F}} |\mu(A) - \nu(A)|,$$

and we say that a sequence of measures $(\mu_t)_{t=1}^{\infty}$ converge to μ in total variation, $\mu_t \xrightarrow{TV} \mu$, if $\lim_{t \rightarrow \infty} \|\mu_t - \mu\|_{\mathbf{TV}} = 0$.

In the knowledge that $\sigma_T \xrightarrow{TV} \pi$, where π is uniform on \mathcal{S}_G , it is natural to ask how fast this convergence is, and also what properties σ_T shares with π for large $T \geq 0$. In the following short sections we summarise several approaches to these questions.

4.1.1 The Complete Graph and Poisson–Dirichlet Distribution

When we consider $G = K_n$, the complete graph on n -vertices, the RSP becomes the continuous time process in which we wait an exponential period of time and then compose our permutation with a uniform transposition, and then repeat these dynamics. Dropping the exponential waiting times, we consider the discrete time analogue $(\sigma_t)_{t=1}^{\infty}$, $\sigma_t = \tau_t \circ \sigma_{t-1}$, where $\tau_k \in \mathcal{T}_n = \{(u, v) : u, v \in V(K_n), u \neq v\}$ are i.i.d. uniform for $t \geq 1$. We will refer to this as the interchange process.

Let $\pi_n \in \mathcal{S}_{K_n}$ be uniformly distributed. For fixed n , the rate at which σ_t converges to π_n as $t \rightarrow \infty$, is covered by the standard theory for Markov chain mixing, see [LPW09]. A deeper question is how the convergence is affected if we simultaneously take the limit $n \rightarrow \infty$. This is also of greater physical relevance, as it provides a mean field approximation to the thermodynamic limit of the RSP on a lattice. We consider what happens if $t = t(n)$, and $n, t(n) \rightarrow \infty$. As one might expect, the convergence properties in this regime depend on how exactly $t(n)$ grows with n . In [Sch05], Schramm makes the simple observation that if $t(n) = o(n \log n)$, then the distribution of $\sigma_{t(n)}$ is far from π_n : in particular, the interchange process is much more likely to have fixed points than π_n . On the other hand, in their seminal paper [DS81], Diaconis and Shahshahani show that for $t(n) = cn \log n$, $c > 1/2$ then convergence still occurs: $\sigma_{t(n)} \xrightarrow{TV} \pi_n$, as $n \rightarrow \infty$.

Schramm’s paper studies the convergence properties when $t(n) = cn$, and again finds that $c = 1/2$ plays an important role. In analogy to the percolation bound employed to bound the probability of infinite cycles, Theorem 2.17, Schramm bounds the size of the largest cycle in the interchange model $\sigma_{t(n)}$ by the size of the largest cluster in the Erdős-Renyi graph $G(n, p)$, $p = t(n)/\binom{n}{2}$. The Erdős-Renyi theorem asserts that for $c < 1/2$ the largest cycle of $\sigma_{t(n)}$ is of order $O(\log n)$, whilst for $c = 1/2$, it is of order greater than $n^{2/3}$, [Sch05] p.2. Again, comparison with a uniform sample π_n shows that convergence does not hold in this regime, since the probability that a given vertex is in a cycle of length $\log n$ (or $n^{2/3}$) in π_n goes to 0 as $n \rightarrow \infty$. The body of Schramm’s paper considers the case $c > 1/2$, and he proves that although the interchange process does not converge to the uniform distribution, that it has the same limiting cycle structure.

The cycle distribution of a uniform permutation $\pi_n \in \mathcal{S}_n$ is well understood. The theory

goes back to Shepp and Lloyd [SL66], who described the distribution of the largest cycle. Denoting $L_1(\pi_n)$ for the length of the largest cycle of π_n , they show that the law of $L_1(\pi_n)/n$ converges to the Dickman distribution, and that in particular $\lim_{n \rightarrow \infty} \mathbb{E}[L_1(\pi_n)/n] = 0.6243\dots$, the Golomb–Dickman constant. Letting $L_k(\pi_n)$ denote the length of the k -th largest cycle in π_n , it was subsequently shown by Kingman [Kin77], and Vershik and Schmidt [VS77], that the vector $n^{-1}(L_1(\pi_n), L_2(\pi_n), \dots) \xrightarrow{(d)} \text{PD}(1)$, where $\text{PD}(\theta)$, $\theta > 0$, is a class of distributions known as Poisson–Dirichlet distributions. Returning to [Sch05], Schramm shows for $t(n) > cn$, $c > 1/2$, the normalised cycle lengths of $\sigma_{t(n)}$ are also Poisson–Dirichlet: for all $\epsilon > 0$ there is an $n_0 = n_0(c, \epsilon)$ such that for all $n > n_0$ and $t(n) \geq cn$, $c > 1/2$

$$\mathbb{P}\left[\|Y - C_t^{-1}(L_1(\sigma_{t(n)}), L_2(\sigma_{t(n)}), \dots)\|_\infty < \epsilon\right] > 1 - \epsilon, \quad (4.1)$$

where $C_{t(n)}$ is the size of the largest component in the Erdős–Renyi graph $G(n, t(n))$, and $Y \sim \text{PD}(1)$.

The outcome of Schramm’s result is that although σ_t does not converge towards the uniform measure for $t(n) = o(n \log n)$, some of the structure of the limit of the uniform distribution is inherited when $t(n) > cn$, $c > 1/2$, namely the large cycle structure converges.

Remark 4.1. Returning to the RSP, σ_T , we note that the expected number of transpositions composed up to time T is $T \binom{n}{2}$, which grows with order n^2 . Consequently, for any $T > 0$, $t(n) = T \binom{n}{2}$ grows faster than cn for any $c > 0$; taking the limit $n \rightarrow \infty$, the cycle structure of σ_T converges to $\text{PD}(1)$, for any fixed $T > 0$. Since the $\text{PD}(1)$ distribution puts a (random) positive proportion of its mass on macroscopic cycles (where macroscopic is in the sense of (1.28), page 18), the RSP on K_n has macroscopic cycles in the thermodynamic limit at any $T > 0$. Moreover, since $T \binom{n}{2} > o(n \log n)$, Diaconis and Shahshahani’s result ensures the RSP converges with the uniform distribution on \mathcal{S}_n , as $n \rightarrow \infty$.

4.1.2 Representation Theory for the RSP

Diaconis and Shahshahani’s paper [DS81] makes use of a novel application of group representation theory to the study of random walks on groups, a technique introduced in the 1960s, see [Han65] and the references therein. This tool has recently been applied to the interchange process on finite graphs in the preprints [AK10a, AK10b, BK12].

Given a group (H, \circ) , a *representation* of H is a homomorphism $\rho: H \rightarrow \text{GL}(V)$, where V is a finite dimensional vector space, and $\text{GL}(V)$ is the space of linear maps on V . Of greatest importance in the study of group representations are the *irreducible* representations, that is those representations with no non-trivial invariant subspaces: if $W \subseteq V$ and $\rho(h)W \subset W$ for all $h \in H$, then $W = \{0\}$ or $W = V$. We are interested in the applications of representation theory to probabilistic models on the finite symmetric group \mathcal{S}_n . Recall that for $\sigma \in \mathcal{S}_n$, $l_\sigma(k)$ denotes the number of cycles in σ of length k . Alon and Kozma [AK10a] derive an explicit formula for

l_σ as a summation over irreducible representations of \mathcal{S}_n

$$l_\sigma(k) = \sum_{\rho \text{ irr. } \mathcal{S}_n} a_\rho \operatorname{Tr}(\rho(\sigma)), \quad (4.2)$$

where $a_\rho \in \mathbb{C}$ are known explicitly, and $\operatorname{Tr}(\rho(\sigma)) =: \chi_\rho(\sigma)$ is the *character* of $\rho(\sigma)$. Applying this formula to the RSP on a finite graph G , Alon and Kozma show that the probability that σ_T consists of a single cycle of size $n = \#V(G)$ is

$$\mathbb{P}_T[l_T(n) = 1] = \frac{1}{n} \prod_{i=1}^{n-1} (1 - e^{-\lambda_i t}),$$

where $(\lambda_i)_{i=1}^n$ are the eigenvalues of the discrete laplacian of the graph G , [AK10a] p.10. Equation (4.2) was subsequently employed by Berestycki and Kozma to find an expression for the expected cycle size of the RSP on the complete graph K_n

$$\mathbb{E}_T[l_T(k)] = \binom{n}{k} \left[\frac{1}{k} x \phi(x) + \int_x^1 \phi(y) dy \right], \quad 1 \leq k \leq n,$$

where $x = e^{-Tk}$, and $\phi(y) = y^{n-k}(1-y)^{k-1}$, [BK12] pp.7–12.

4.1.3 The Spectral Gap of the RSP (Aldous's Conjecture)

In Section 2.1 it was shown that the RSP on a finite graph has a representation as a Markov chain on a symmetric group. With this in mind, let $P \in \mathcal{M}^1(H)$ be a probability measure on the finite group H ; P defines a Markov chain on H with transition matrix M , $M(g, h) := P(hg^{-1})$, and k -step density $P^{*k} = P * P^{*(k-1)}$, where we recall $*$ denotes the convolution

$$(P * Q)(h) := \sum_{g \in H} P(hg^{-1})Q(g), \quad P, Q \in \mathcal{M}^1(H).$$

In the particular case of the RSP on G , $H = S_{V(G)}$ and the probability measure P is uniform on the transpositions in $\mathcal{T}_G = \{(u, v) : uv \in E(G)\}$. Many of the probabilistic results using representation theory are built around the group-theoretic Fourier transform, which converts convolution into multiplication. For a representation $\rho: H \rightarrow V$, it is defined as

$$\rho(P) := \sum_{h \in H} P(h)\rho(h) \in \operatorname{GL}(V), \quad (4.3)$$

(it is easily confirmed that $\rho(P * Q) = \rho(P)\rho(Q)$). One such application of this transform is the following description of the spectrum of the transition matrix M . Letting Spec_A denote the spectrum of a linear map A , Diaconis shows that

$$\operatorname{Spec}_M = \bigcup_{\rho \text{ irr. } H} \operatorname{Spec}_{\rho(P)}, \quad (4.4)$$

where the union is over all irreducible representations of H . See [DS81] p.175.

In the study of discrete state space Markov chains, knowledge of the spectrum of the transition matrix is an important aid to understanding the convergence (or *mixing*) rate of the process. Let $\lambda_1 \in \text{Spec}_M$ be the second largest eigenvalue of the transition matrix (necessarily $\lambda_1 \in [-1, 1)$), then the *spectral gap* $\gamma := 1 - \lambda_1$ can be used to provide both lower and upper bounds on the time taken till the Markov chain mixes

$$\left(\frac{1}{\gamma} - 1\right) \log\left(\frac{1}{2\epsilon}\right) \leq \inf\left\{t: \|P^t - \pi\|_{\mathbf{TV}} < \epsilon\right\} \leq \frac{1}{\gamma} \log\left(\frac{1}{\epsilon c}\right), \quad \epsilon > 0, \quad (4.5)$$

where π is the stationary distribution of the chain, and $c = c(\pi)$ is constant; see for example [LPW09] pp. 154–6. Using equation (4.4), Diaconis and Shahshahani find the spectrum of the RSP on K_n in terms of the characters of the irreducible representations, [DS81] p.175. Moreover they confirmed that the spectral gap of the RSP matches that of the continuous time random walk on K_n , a special case of a (recently solved) conjecture, often attributed to David Aldous.

Consider the RSP on a finite weighted graph $G = (G, \omega)$, where the labeling function of weights $\omega: E(G) \rightarrow \mathbb{R}_{\geq 0}$ describes the rate of the Poisson bar set on each edge: that is $\mathcal{B}_G = \cup_{e \in E} \mathcal{B}_e^\omega$, where \mathcal{B}_e^ω is a Poisson process on $[0, \infty)$ with intensity $\omega(e)$.

Remark 4.2. Whilst we have concentrated on the case in which the RSP is in a homogeneous environment, $\omega \equiv 1$, it is clear that the theory extends to this inhomogeneous setting.

Aldous’s conjecture states that on any weighted graph (G, ω) the continuous time random walk and the RSP have the same spectral gap. Following the path of just a single vertex $v \in V$ under the random stirring dynamics $\{\sigma_T(v): T \geq 0\}$, one sees that this trajectory has the same distribution as that of the continuous time random walk, from which it can be easily shown that $\lambda_1^{\text{RW}}(G) \geq \lambda_1^{\text{RSP}}(G)$, where $\lambda_1^{\text{RW}}, \lambda_1^{\text{RSP}}$ are the second largest eigenvalues of the continuous time random walk and RSP respectively. Using, amongst others, arguments from electrical network theory, Caputo, Liggett, and Richthammer confirmed the general form of the Aldous’s conjecture [CLR10]

$$\lambda_1^{\text{RW}}(G) = \lambda_1^{\text{RSP}}(G).$$

4.2 Questions, Conjectures and the Future

Although undoubtedly a lot of progress has been made on studying the RSP on finite graphs, there are still many open problems. A recurring theme in the previous section was that although general results are available, explicit formulae (with no unknown variables) can often only be identified in special cases such as when random stirring is done on a complete graph. With an eye to resolving Tóth’s conjecture, either in the toy- or full-form, it is important that we can find such rigorous results for lattice structures in which there is a greater geometry. As the authors of [AK10a] observe, for the time being the approximations of the relevant eigenvalues in the case that $G = \Lambda \subset \mathbb{Z}^d$ are not strong enough to provide any additional information for either conjecture.

In the infinite setting even less is known about the RSP, and we briefly summarise some questions that arise from the work of Angel and Hammond (many of which were stated in [Ang03]).

Formally we can extend the period of the CTRW to the case $T = \infty$, in which for any bounded time the walk is unaffected by the periodic nature of the bar set, and so is independent of its past. We therefore identify this case with the continuous time random walk. With this in mind, we note that the sharp-threshold conjecture can be restated as follows: for any graph \mathbb{G} there is a time $T_c = T_c(\mathbb{G})$ such that for $T \in [0, \infty)$, X^T is recurrent whenever $T < T_c$ and is transitive whenever $T > T_c$. Angel asks whether this can be extended to hold for $T \in [0, \infty) \cup \{\infty\}$, where X^∞ is the continuous time random walk.

Conjecture 4.3 (Angel, [Ang03]). *There exists $T_c = T_c(\mathbb{G})$ such that for all $T \in [0, \infty) \cup \{\infty\}$, \mathbb{G} is T -recurrent for $T < T_c$, and is T -transient for $T > T_c$. If the continuous time random walk is recurrent on \mathbb{G} , then $T_c(\mathbb{G}) = \infty$.*

A natural follow up question would be to ask whether the continuous time random walk being transient on \mathbb{G} implies that for high enough T the graph is T -transient. As Angel observes, [Ang03] p.12, this is not always true: there are graphs for which the continuous time random walk is transient, however their critical probability for bond percolation is $p_c(\mathbb{G}) = 1$, ensuring that the CTRW is always recurrent.

Should Conjecture 4.3 be true, it would provide a proof that the RSP has no phase transition on \mathbb{Z}^2 , agreeing with the statement of the Mermin-Wagner theorem (Theorem 1.15, page 18) which holds for the thermodynamic limit. If resolved, this conjecture would provide the first result for the RSP on a lattice. This brings us on to the most important of the open problems, resolving the behaviour of the RSP on \mathbb{Z}^d for $d \geq 2$, and in particular for $d = 3$. As we saw in Chapter 3, the techniques developed so far work for trees since we have control over the way in which the walk returns to previously visited vertices; for graphs with cycles such approaches are weakened, and so we need to find a new direction of attack.

Whilst undeniably lattices are of the greatest interest, it would also prove worthwhile to complete the picture for trees. This includes providing the proof of sharp-threshold for trees of low degree, and also removing the requirement that the trees are regular. Also of interest would be to identify the exact value of $T_c^I(\mathbb{T}_d)$ for regular trees; perhaps a more realistic problem of this type is to identify the second order term, improving on the bounds (3.5). The author is currently working on finding the critical temperature for the Heisenberg model on a tree, which would provide the thermodynamic phase transition which should agree with $T_c^I(\mathbb{T}_d)$.

Appendix A

Hazard Functions

We feel inclined to provide an introduction to hazard functions (also known as failure rates, or instantaneous rates) since though the material is elementary, it is contained in few texts intended for probabilists, and is normally introduced in the context of applied mathematics: in particular medical statistics, or economics. One (less applied) reference is [Ros09]. Let X be a continuous random variable taking values in $[0, \infty)$, so that X describes the time of occurrence of a random event. The hazard function $h : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ of X is

$$\begin{aligned} h(x) &:= \lim_{\Delta x \rightarrow 0} \frac{\mathbb{P}[X < x + \Delta x \mid X > x]}{\Delta x} \\ &= \lim_{\Delta x \rightarrow 0} \frac{1}{\mathbb{P}[X > x]} \frac{\mathbb{P}[x < X < x + \Delta x]}{\Delta x} \\ &= \frac{1}{1 - F(x)} \lim_{\Delta x \rightarrow 0} \frac{F(x + \Delta x) - F(x)}{\Delta x} \\ &= F'(x)/(1 - F(x)), \end{aligned}$$

where F is the cumulative distribution function. The first expression above motivates the term *instantaneous rate* since it describes the rate at which the random event falls within $[t, t + \Delta t]$.

Proposition A.1. *The distribution of X is uniquely determined by its hazard function.*

Proof. Observing that $-F'(x) = \frac{d}{dx}(1 - F(x))$ we can rewrite

$$h(x) = -\frac{d}{dx} \log(1 - F(x)).$$

Integrating both sides, and rearranging we obtain

$$F(x) = 1 - \exp\left(-\int_0^x h(y)dy\right).$$

Since the distribution of X is uniquely determined by F we are done. □

Corollary A.2.

$$\mathbb{P}[X < x] = \exp\left(-\int_0^x h(y)dy\right).$$

Bibliography

- [Ada06] S. ADAMS, *Lectures on Mathematical Statistical Mechanics*. Comm. Dublin Inst. Adv. Stud. **A** (30) (2006)
- [Ada07] S. ADAMS, *Large deviations for empirical path measures in cycles of integer partitions*. Preprint available at [arXiv:0702.053v2](https://arxiv.org/abs/0702.053v2) (2007)
- [ACK11] S. ADAMS, A. COLLEVECCHIO and W. KÖNIG, *A variational formula for the free energy of an interacting many-particle system*. Ann. Probab. **39** (2): 683–728 (2011)
- [AK10a] G. ALON and G. KOZMA, *Ordering the representations of S_n using the interchange process*. To appear in Can. Math. Bull. Preprint available at [arXiv:1003.1710v2](https://arxiv.org/abs/1003.1710v2) (2010)
- [AK10b] G. ALON and G. KOZMA, *The probability of long cycles in interchange processes*. Preprint available at [arXiv:1009.3723v3](https://arxiv.org/abs/1009.3723v3) (2010)
- [ALSS04] M. AIZENMAN, E. H. LIEB, R. SEIRINGER, J.P. SOLOVEJ and P.J. YNGVASON, *Bose-Einstein Quantum Phase Transition in an Optical Lattice Model*, Phys. Rev. A **70** (2) (2004)
- [Ang03] O. ANGEL, *Random Transpositions and the Cyclic Time Random Walk*. Discrete Random Walks (ed. Banderier, C., Krattenthaler, C.), Disc. Math. Theor. Comput. Sci. Proc. **AC**: 9–16 (2003)
- [Bax08] R. J. BAXTER, *Exactly Solved Models in Statistical Mechanics*. Dover Ed. (1st Ed. 1982), Dover, New York (2008)
- [BK12] N. BERESTYCKI and G. KOZMA, *Cycle structure of the interchange process and representation theory*. Preprint available at [arXiv:1205.4753v1](https://arxiv.org/abs/1205.4753v1) (2012)
- [BR06] B. BOLLOBÁS and O. RIORDAN, *Percolation*. CUP, Cambridge (2006)
- [CLR10] P. CAPUTO, T. M. LIGGETT, and T. RICHTHAMMER, *A recursive proof of Aldous’ spectral gap conjecture*. J. Am. Math. Soc. **23**: 831–851 (2010)
- [CS91a] J. CONLON and J.P. SOLOVEJ, *Random walk representations of the Heisenberg model*. J. Stat. Phys. **64** (1-2): 251–270 (1991)

- [CS91b] J. CONLON and J.P. SOLOVEJ, *Upper bound on the free energy of the spin 1/2 Heisenberg ferromagnet*. Lett. Math. Phys. **23**: 223–231 (1991)
- [Dan11] O. DANIEL, *Towards the formulation of an edge based random permutation model*. Masters Thesis. Available at: http://www2.warwick.ac.uk/fac/sci/masdoc/studentpages/students_2011/owendaniel/ (2011)
- [DS81] P. DIACONIS and M. SHAHSHAHANI, *Generating a random permutation with random transpositions*. Prob. Theory Rel. **57** (2): 159–179 (1981)
- [Dia88] P. DIACONIS, *Group Representations in Probability and Statistics*. IMS Monograph Series Volume 11 (1988)
- [Ein24] A. EINSTEIN, *Quantentheorie des einatomigen idealen Gases. 2. Abhandlung*. Sitzber. Preuss. Akad. **23** (3): 3–14 (1925)
- [Fey48] R. P. FEYNMAN, *Space-Time Approach to Non-Relativistic Quantum Mechanics*. Rev. Mod. Phys. **20** (2): 367–387 (1948)
- [Fey53] R. P. FEYNMAN, *Atomic theory of the λ -transition in helium*. Phys. Rev. **91** (6): 1291–1301 (1953)
- [FS09] P. FLAJOLET and R. SEDGEWICK, *Analytic Combinatorics*. CUP, Cambridge (2010)
- [Gin71] J. GINIBRE, *Some Applications of Functional Integration in Statistical Mechanics*. in Statistical Mechanics and Quantum Field Theory (ed. C DeWitt, R. Stora) pp. 327–427, Gordon and Breach, New York (1971)
- [GUW11] C. GOLDSCHMIDT, D. UELTSCHI and P. WINDRIDGE, *Quantum Heisenberg models and their probabilistic representations*. Entropy and the Quantum II (Contemp. Math.) (ed. Sims, R., Ueltschi, D.), AMS: 177–224 (2010)
- [Ha12a] A. HAMMOND, *Infinite cycles in the random stirring process on trees*. Preprint available at [arXiv:1202.1319v1](https://arxiv.org/abs/1202.1319v1) (2012)
- [Ha12b] A. HAMMOND, *Sharp phase transition in the random stirring model on trees*. Preprint available at [arXiv:1202.1322v1](https://arxiv.org/abs/1202.1322v1) (2012)
- [Han65] E. J. HANNAN, *Group representations and applied probability*. J. Appl. Probab. **2** (1): 1–68 (1965)
- [JW28] P. JORDAN and E. WIGNER, *Über das Paulische äquivalenzverbot*. Z. Phys. A - Hadron. Nucl. **47** (9-10): 631–651 (1928)
- [Kin77] J. F. C. KINGMAN, *The population structure associated with the Ewens sampling formula*. Theor. Popul. Biol. **11** (2): 274–283 (1977)
- [Kin93] J. F. C. KINGMAN, *Poisson Processes*. OUP, Oxford (1993)

- [LPW09] D. A. LEVIN, Y. PERES and E. L. WILMER, *Markov chains and mixing times*. AMS, Rhode Island (2009)
- [LSSY05] E. H. LIEB, R. SEIRINGER, J. P. SOLOVEJ and Y. YNGVASON, *The mathematics of the Bose gas and its condensation*. Oberwolfach Seminar **34**, Birkhäuser Verlag, Basel (2005)
- [Lig85] T. LIGGETT, *Interacting Particle Systems*. Springer, Berlin (1985)
- [LHB11] J. LÓRINCZI, F. HIROSHIMA and V. BETZ *Feynman-Kac-type theorems and Gibbs measures on path space: with applications to rigorous quantum field theory* De Gruyter, Berlin (2011)
- [LP12] R. LYONS and Y. PERES, *Probability on Trees and Networks*. CUP, in preparation. Current version available at <http://mypage.iu.edu/~rdlyons/>
- [MM56] T. MATSUBARA and H. MATSUDA, *A lattice model of liquid helium, I*. Prog. Theor. Phys. **16** (6): 569–582 (1956)
- [MW66] N. D. MERMIN and H. WAGNER, *Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic Heisenberg models*. Phys. Rev. Lett. **17** (22): 1133–1136.
- [Nel64] E. NELSON, *Feynman Integrals and the Schrödinger Equation*. J. Math. Phys. **5** (3): 332–343
- [Nel67] E. NELSON, *Dynamical Theories of Brownian Motion*. 2nd Ed. (1st Ed, 1967), PUP, Princeton (2001)
- [Nob01] *The Nobel Prize in Physics 2001*. URL: www.nobelprize.org/nobel_prizes/physics/laureates/2001 [accessed: 24 May 2012]
- [Ons44] L. ONSAGER, *Crystal statistics, I. A two-dimensional model with an order-disorder transition*. Phys. Rev. **65**: 117–149 (1944)
- [Pei36] R. PEIERLS, *On Ising’s model of ferromagnetism*. P. Camb. Philos. Soc. **32** (3): 477–481 (1936).
- [Ros09] S. ROSS, *Introduction to Probability Models*. 10th Ed. (1st Ed, 1980). Academic Press, Oxford (2010)
- [Sch05] O. SCHRAMM, *Compositions of random transpositions*. Isr. J. Math. **147**: 221–243 (2005)
- [Sch08] F. SCHWABL, *Advanced Quantum Mechanics*. Springer, Berlin (2008)
- [SL66] L. A. SHEPP and S. P. LLOYD, *Ordered cycle lengths in a random permutation*. T. Am. Math. Soc. **121** (2): 340–357 (1966)
- [Sim93] B. SIMON *The Statistical Mechanics of Lattice Gases. Volume I*. PUP, Princeton (1993)

- [Süt93] A. SÜTŐ, *Percolation transition in the Bose gas*. J. Phys. A: Math. Gen. **26**: 4689–4710 (1993)
- [Süt02] A. SÜTŐ, *Percolation transition in the Bose gas: II*. J. Phys. A: Math. Gen. **35**: 6995–7002 (2001)
- [Tót93] B. TÓTH, *Improved lower bound on the thermodynamic pressure of the spin 1/2 Heisenberg ferromagnet*. Lett. Math. Phys. **28** (1): 75–84 (1993)
- [VS77] A. M. VERSHIK and A. A. SCHMIDT *Limit measures arising in the asymptotic theory of symmetric groups. I*. Theor. Probab. Appl. **22** (1): 70–85 (1977)