The Voter Model

E. Schulpen

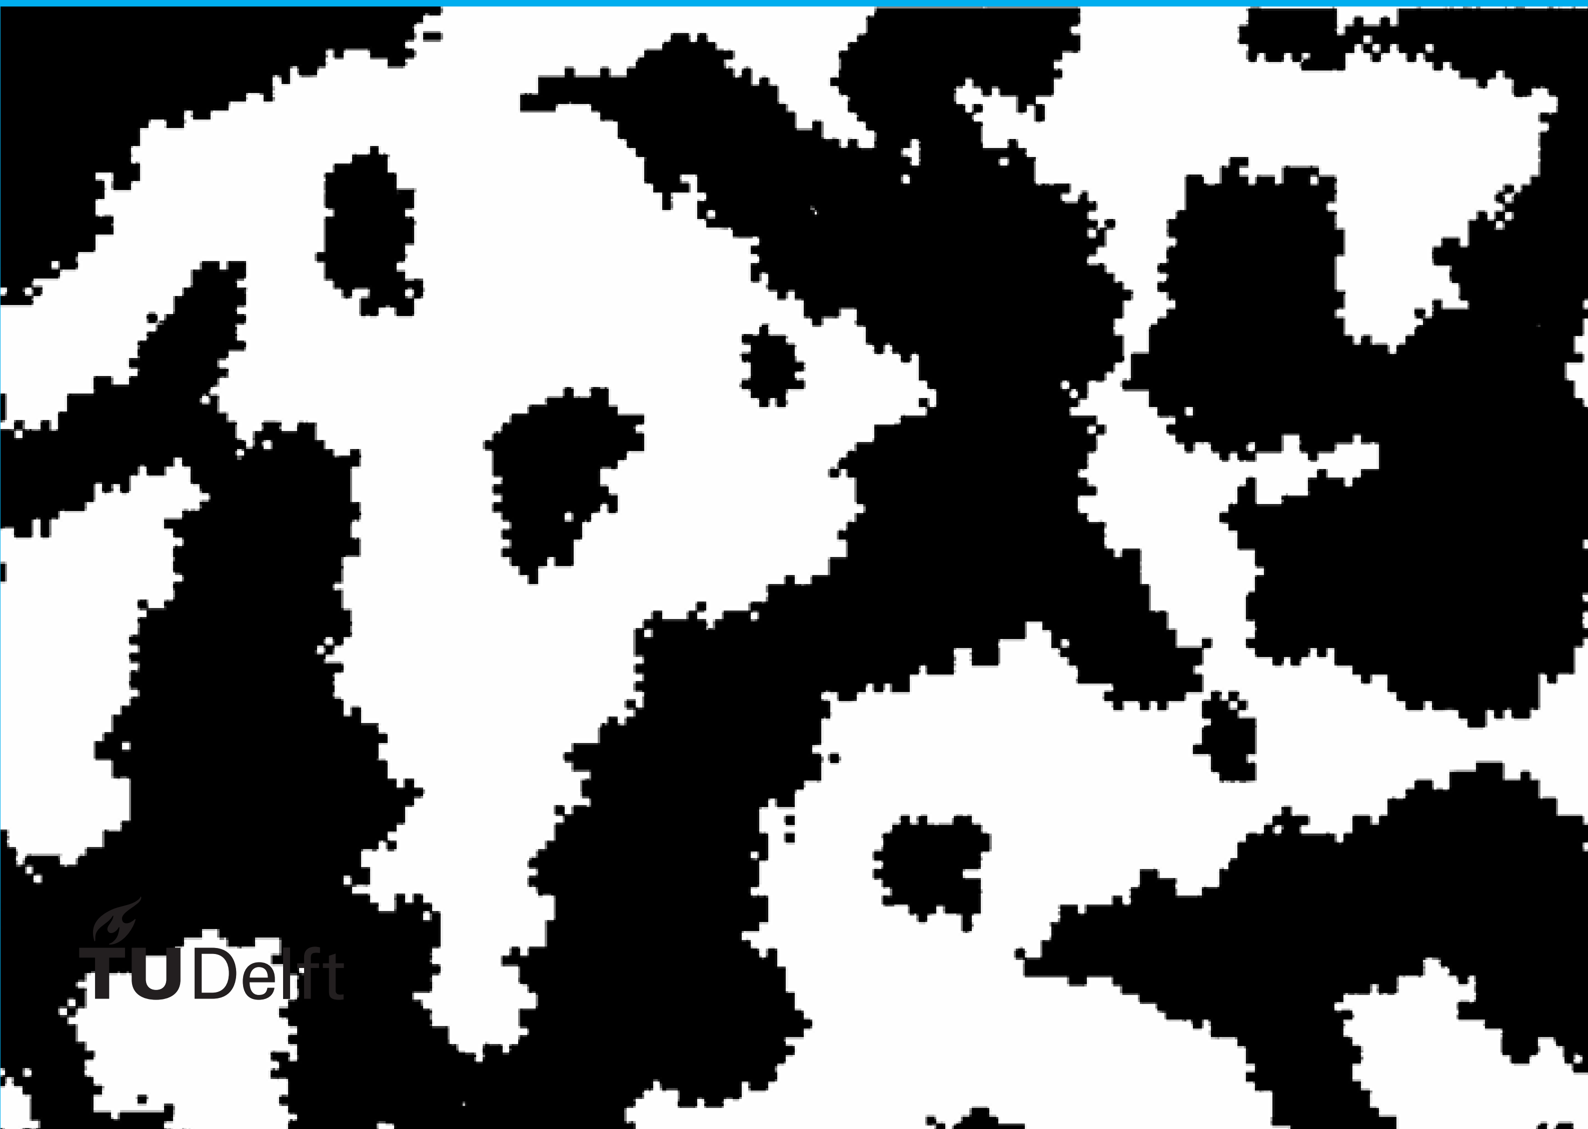

TU Delft

The Voter Model

by

E. Schulpen

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday August 27, 2025 at 3:30 PM.

Student number:	5344727	
Project duration:	April 24, 2025 – August 27, 2025	
Thesis committee:	Dr. R. Versendaal,	TU Delft, supervisor
	Dr. E. Lorist,	TU Delft

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

*Cover image: simulation of the threshold-*3 *voter model on a Moore neighborhood, adapted from Ten Lectures on Particle Systems by Thomas M. Liggett.*

TU Delft

Contents

Layman's Summary

How do opinions spread in a group of people? What happens if everyone occasionally changes their mind by copying a neighbor? Will the group eventually agree, or will different opinions persist? These are the kinds of questions the voter model helps us explore.

The voter model is a simple mathematical model where individuals are placed on a grid and randomly adopt the opinion of a neighbor. Despite its simplicity, the model captures essential features of real-world systems, from how diseases spread to how people form opinions in social networks. It belongs to a broader class of models known as interacting particle systems, which are used to study how local interactions can lead to complex global behavior.

In the voter model, each site can be in one of two states, often interpreted as two competing opinions. At random times, a site updates its state by copying the state of a randomly chosen neighbor. This rule is symmetric and neutral: no opinion is favored, and the dynamics are entirely driven by local imitation. Yet, the long-term behavior of the system depends crucially on the dimension of the space in which it evolves.

This thesis investigates the long-term behavior of the voter model, with a particular focus on the dichotomy between clustering and coexistence. In low-dimensional settings (such as a line or a plane), the system tends to reach consensus: eventually, all sites agree on the same opinion. In higher dimensions (such as a 3D grid), different opinions can coexist indefinitely, with no single opinion taking over the entire system.

To understand this phenomenon, the thesis develops a mathematical framework based on probability theory. Key tools include continuous-time Markov processes, random walks, and harmonic functions. A central concept is duality: a technique that allows us to study the evolution of the system by tracing the ancestry of opinions backward in time. This duality connects the voter model to coalescing random walks, which are easier to analyze and provide deep insights into the system's behavior.

The thesis also explores the structure of invariant measures, probability distributions that remain unchanged over time. In low dimensions, the only invariant measures are those corresponding to full consensus. In higher dimensions, however, there exists a whole family of invariant measures that represent stable mixtures of opinions. These measures are closely related to harmonic functions of the underlying random walk.

Finally, the thesis discusses several extensions of the model. These include the threshold voter model, where a site changes its state only if a sufficient number of neighbors disagree. These variations introduce asymmetry or nonlinearity into the dynamics and lead to qualitatively different outcomes, such as phase transitions or persistent coexistence even in low dimensions.

In summary, the voter model provides a powerful and elegant framework for understanding how simple local rules can give rise to rich and varied global behavior. Its analysis combines ideas from probability, statistical physics, and dynamical systems, and its insights are relevant to a wide range of applications in science and society.

Abstract

This thesis provides a rigorous and self-contained study of the linear voter model, a fundamental example of an interacting particle system. The model describes the evolution of binary states on a lattice, where each site updates its state by imitating a randomly chosen neighbor. Despite its simplicity, the model exhibits a rich dichotomy in long-term behavior: in low dimensions, the system clusters and reaches consensus, while in higher dimensions, it allows for coexistence of different states.

To analyze this behavior, we develop the necessary probabilistic and analytic framework, including continuous-time Markov processes, random walks, and potential theory. A central role is played by the duality between the voter model and coalescing random walks, which enables a detailed understanding of the system's dynamics. We also examine the classification of invariant measures and their connection to harmonic functions and martingale properties.

The thesis concludes with a discussion of extensions such as biased and threshold voter models, and outlines directions for further research. The results highlight the interplay between local interaction rules, spatial structure, and emergent global behavior in stochastic systems.

1

Introduction

Many phenomena in physics, biology, and the social sciences — such as the spread of diseases, opinion dynamics in social networks, or the behavior of magnetic materials — can be modeled as systems of simple, locally interacting agents. These systems are studied within the framework of interacting particle systems (IPS), a class of stochastic processes that includes models like the contact process, the exclusion process, and the voter model. Each of these models captures how local interaction rules can give rise to complex macroscopic behavior, such as phase transitions, clustering, or coexistence.

Among these, the voter model stands out for its simplicity and analytical tractability. In this model, each site on a lattice updates its state by imitating a randomly chosen neighbor. Despite this minimal rule, the model exhibits a rich dichotomy of behaviors: in low dimensions, it tends to cluster and reach consensus, while in higher dimensions, it allows for long-term coexistence of different states. One of the key reasons the voter model is so well understood is its duality with coalescing random walks, which enables powerful probabilistic techniques to analyze its long-term behavior. Additionally, the model conserves the average density of states and connects naturally to potential theory and harmonic functions.

In this thesis, we aim to provide a rigorous and self-contained analysis of the voter model, with a particular focus on its long-term behavior and invariant measures. We begin by developing the necessary mathematical background, including continuous-time Markov processes, random walks, and potential theory. We then define the voter model formally, explore its duality structure, and analyze its behavior in both finite and infinite settings. Special attention is given to the role of spatial dimension in determining whether the system clusters or coexists, and we conclude by discussing extensions such as biased and threshold voter models.

Chapter 2 introduces the framework of continuous-time Markov processes and interacting particle systems. We define generators, semigroups and invariant measures, and explain how spin-flip systems can be constructed using Harris' graphical representation. This chapter lays the probabilistic foundation for everything that follows.

Chapter 3 turns to random walks and potential theory. We discuss recurrence and transience of random walks, introduce Green's functions and harmonic functions, and explain the role of the Liouville property. These tools form the analytic backbone for the study of clustering and coexistence in the voter model.

Chapter 4 gives a detailed introduction of the voter model itself. We define the nearest-neighbor and k-th nearest-neighbor versions, and show how duality with coalescing random walks provides a powerful method of analysis. The chapter develops the central results on clustering in low dimensions and coexistence in high dimensions.

Chapter 5 investigates invariant measures and long-term behavior. We connect martingale properties and harmonic functions to the classification of extremal invariant measures, and describe convergence to equilibrium. This leads to a complete understanding of the stationary distributions of the voter model.

Finally, Chapter 6 summarizes the main conclusions and offers a discussion of possible extensions. These include perturbations such as biased or threshold voter models, as well as scaling limits that connect the voter model to measure-valued diffusions and super processes.

2

Markov Processes and Interacting Particle Systems

This chapter lays the probabilistic groundwork for the analysis of the voter model. We begin by introducing the general theory of continuous-time Markov processes, including their generators, semigroups, and stationary distributions. These concepts are essential for describing the time evolution of stochastic systems and form the basis for interacting particle systems (IPS), which model the collective behavior of locally interacting agents on a lattice.

We then focus on spin-flip systems, a class of IPS where each site updates its state based on the configuration of its neighbors. The voter model is a prominent example of such a system. To rigorously construct and analyze these models, we introduce the Harris graphical construction [Harris, 1972], which provides a visual and probabilistic representation of the system's evolution. This construction is particularly useful for coupling arguments and for establishing properties like attractiveness and monotonicity.

The chapter concludes with the concept of duality, a powerful analytical tool that allows us to relate the dynamics of the voter model to coalescing random walks. This duality simplifies the analysis of long-term behavior and is central to many of the results in later chapters.

The material in this chapter draws primarily from Liggett's texts on interacting particle systems [Liggett, 1985; Liggett, 1999], as well as Durrett's treatment of Markov processes and stochastic modeling [Durrett, 2019], and the original construction by Harris [1972].

2.1. Probability spaces and configuration space

To describe the voter model rigorously, we begin by defining the underlying *configuration space* and associated *probability space*. The voter model is an example of a *spin system* on the lattice $\mathbb{Z}^d$, where each site carries a *spin* taking values in the set $S = \{0, 1\}$. A *configuration* is a function $\eta : \mathbb{Z}^d \to S$, assigning a spin to each site. The full configuration space is then

$$\Omega = S^{\mathbb{Z}^d},$$

equipped with the product topology and the corresponding *product σ-algebra*. This σ-algebra is generated by *cylinder sets*, which are events that depend only on the values of finitely many coordinates. For example, a cylinder set might specify the spin values at a finite subset of sites, while leaving the rest unspecified.

To turn this into a probability space, we often equip Ω with a product measure, such as the Bernoulli product measure ν_ρ, where each site independently takes value 1 with probability $\rho \in [0, 1]$. This defines the full probability space $(\Omega, \mathscr{F}, \nu_\rho)$, where $\mathscr{F}$ is the product σ-algebra.

A function $f : \Omega \to \mathbb{R}$ is called a *local observable* if it depends only on the configuration at finitely many sites. These functions are important because they serve as natural *test functions* for the generator of the Markov process: they allow us to probe the infinitesimal behavior of the system in a mathematically controlled way. Since the dynamics of spin systems are local it is natural to study their effect on local observables.

2.2. Continuous-time Markov processes

The voter model is a special case of a *continuous-time Markov process* on a large configuration space. To study it rigorously, we first introduce the general theory of Markov processes, which provides the probabilistic framework for interacting particle systems.

A *Markov process* models the evolution of a system in which the future depends only on the present state, not on the past history. This *memoryless property* makes Markov processes both tractable and natural for modeling local interactions. In the context of the voter model, this means that once we know the current configuration of spins, the probabilities of future updates depend only on that configuration, not on how it was reached.

Before we define the continuous-time case, we briefly illustrate the idea with a simple example.Imagine a simple model of the weather. Each day, the weather can be either *sunny* or *rainy*. Suppose that:

- If today is sunny, there is a 90% chance that tomorrow will also be sunny, and a 10% chance it will be rainy.
- If today is rainy, there is a 50% chance that tomorrow will be sunny, and a 50% chance it will remain rainy.

This is a *Markov process*, because the probability of tomorrow's weather depends only on today's weather and not on how long it has been sunny or rainy before.

2.2.1. Basic definitions

We now formalize the notion of a *Markov process* in continuous time. These processes are central to the study of interacting particle systems, as they model systems whose future evolution depends only on the current state, not on the past history.

Definition 2.2.1 (Markov Process). *Let $(X_t)_{t\geq 0}$ be a stochastic process taking values in a measurable state space $(E,\mathscr{E})$. We say that (X_t) is a* Markov process *if for all $t,s \geq 0$, $x \in E$, and $A \in \mathscr{E}$,*

$$P_x(X_{t+s} \in A \mid \mathscr{F}_t) = P_x(X_{t+s} \in A \mid X_t),$$

where P_x denotes the probability measure under which $X_0 = x$, and $\mathscr{F}_t = \sigma(X_u : 0 \leq u \leq t)$ is the natural filtration *of the process.*

The notation P_x refers to the law of the process starting from $X_0 = x$, and E_x denotes the corresponding expectation. The *natural filtration* $(\mathscr{F}_t)_{t\geq 0}$ captures all information generated by the process up to time t.

Definition 2.2.2 (Transition Semigroup). *Let (X_t) be a Markov process with state space E. The family of operators $(P_t)_{t\geq 0}$ defined by*

$$P_t f(x) = E_x[f(X_t)] = \int_E f(y) P_t(x, dy),$$

for all bounded measurable functions $f : E \to \mathbb{R}$, is called the transition semigroup *of the process.*

This semigroup satisfies the *Chapman–Kolmogorov equations*:

$$P_{t+s} = P_t P_s.$$

This identity expresses the consistency of the process over time: the probability of transitioning from state x to a set A in time $t+s$ is the same as first transitioning from x to an intermediate state y in time t, and then from y to A in time s, averaged over all possible intermediate states y. This equation must hold because it reflects the *Markov property* itself. The Chapman–Kolmogorov equations ensure that the transition probabilities are consistent across different time intervals and that the process can be built incrementally over time.

Definition 2.2.3 (Strong Markov Property). *Let $(X_t)_{t\geq 0}$ be a Markov process with natural filtration $(\mathscr{F}_t)_{t\geq 0}$, and let τ be a* stopping time, *meaning that for every $t \geq 0$, the event $\{\tau \leq t\}$ is measurable with respect to $\mathscr{F}_t$. Then (X_t) satisfies the* strong Markov property *if for all bounded measurable functions $f : E \to \mathbb{R}$ and all $s \geq 0$,*

$$E_x[f(X_{\tau+s}) \mid \mathscr{F}_\tau] = E_{X_\tau}[f(X_s)].$$

This property strengthens the usual Markov property by allowing the process to "restart" not just at fixed times, but also at random times τ that are determined by the evolution of the process itself.

The strong Markov property is essential when analyzing events such as *hitting times* (the first time the process enters a given set) or *return times*, which are random by nature. We will use this property later when studying recurrence and transience of random walks in Chapter 3, and when analyzing dual processes in the voter model.

2.2.2. Feller Processes

In the context of continuous-time Markov processes, it is useful to distinguish a class of processes that behave well analytically. These are known as Feller processes.

Definition 2.2.4 (Feller Process). *A Markov process $(X_t)_{t\geq 0}$ with state space E is called a* Feller process *if its associated semigroup $(P_t)_{t\geq 0}$ satisfies:*

- *For each $t \geq 0$, P_t maps $C_b(E)$ (bounded continuous functions on E) into itself.*
- *For each $f \in C_b(E)$, $P_t f \to f$ uniformly as $t \to 0$.*

Feller processes preserve continuity and allow the generator to be defined via pointwise limits. They form an important analytical foundation for the study of interacting particle systems such as the voter model.

2.2.3. Generators

In continuous-time Markov processes, the evolution of the system is governed by infinitesimal changes. These changes are captured by the generator of the process, which plays a role analogous to a derivative in calculus.

Definition 2.2.5 (Generators). *Let $(X_t)_{t\geq 0}$ be a Markov process with transition semigroup $(P_t)_{t\geq 0}$. The* generator *L is defined by:*

$$Lf(x) = \lim_{t\to 0} \frac{P_t f(x) - f(x)}{t}$$

for all functions f for which this limit exists. The set of such functions is called the domain of L.

Intuitively, $Lf(x)$ describes the instantaneous rate of change of the expected value of $f(X_t)$ starting from state x.

This leads us to a fundamental result in stochastic analysis:

Definition 2.2.6 (Martingale). *A stochastic process $(M_t)_{t\geq 0}$ is called a* martingale *with respect to a filtration $(\mathscr{F}_t)$ if for all $s \leq t$:*

- *M_s is $\mathscr{F}_s$-measurable,*
- *$\mathbb{E}[|M_t|] < \infty$,*
- *$\mathbb{E}[M_t \mid \mathscr{F}_s] = M_s$.*

Martingales model their expected future value, given current information, is equal to their present value.

Theorem 2.2.1 (Dynkin's Formula). *Let f be a function in the domain of L. Then the process*

$$M_t^f := f(X_t) - f(X_0) - \int_0^t Lf(X_s)\,ds$$

is a martingale with respect to the natural filtration of (X_t).

Proof. We define the transition semigroup $(P_t)_{t\geq 0}$ by $P_t f(x) = \mathbb{E}_x[f(X_t)]$. The generator L is given by:

$$Lf(x) = \lim_{h\to 0} \frac{P_h f(x) - f(x)}{h}.$$

Let f be in the domain of L. For small $h > 0$, we approximate the change in $f(X_t)$ over time by:

$$f(X_{t+h}) - f(X_t) \approx h \cdot Lf(X_t).$$

Summing over small increments and taking expectations, we obtain:

$$\mathbb{E}_x[f(X_t)] - f(x) = \mathbb{E}_x\left(\sum_{k=0}^{n-1}\left(f(X_{t_{k+1}}) - f(X_{t_k})\right)\right) \approx \mathbb{E}_x\left(\sum_{k=0}^{n-1} h\cdot Lf(X_{t_k})\right).$$

In the limit as $h \to 0$, this becomes the identity:

$$\mathbb{E}_x[f(X_t)] - f(x) = \mathbb{E}_x\left(\int_0^t Lf(X_s)\,ds\right),$$

which expresses the expected change in $f(X_t)$ in terms of the generator L.
Now, define the process:

$$M_t^f := f(X_t) - f(X_0) - \int_0^t Lf(X_s)\,ds.$$

To show that M_t^f is a martingale, we compute its conditional expectation given $\mathscr{F}_s$:

$$\mathbb{E}[M_t^f \mid \mathscr{F}_s] = \mathbb{E}\left[f(X_t) - f(X_0) - \int_0^t Lf(X_r)\,dr \mid \mathscr{F}_s\right].$$

We split this into:

$$M_s^f + \mathbb{E}\left[f(X_t) - f(X_s) - \int_s^t Lf(X_r)\,dr \mid \mathscr{F}_s\right].$$

By the Markov property, the future evolution of the process depends only on the current state X_s, so:

$$\mathbb{E}[f(X_t) \mid \mathscr{F}_s] = P_{t-s} f(X_s),$$

and similarly, Using the Markov property and the fact that we can interchange expectation and integration, since $Lf(X_r)$ is integrable:

$$\mathbb{E}\left[\int_s^t Lf(X_r)\,dr \mid \mathscr{F}_s\right] = \int_s^t \mathbb{E}[Lf(X_r) \mid \mathscr{F}_s]\,dr = \int_0^{t-s} P_u Lf(X_s)\,du.$$

Therefore:

$$\mathbb{E}[M_t^f \mid \mathscr{F}_s] = M_s^f + P_{t-s} f(X_s) - f(X_s) - \int_0^{t-s} P_u Lf(X_s)\,du.$$

The last two terms cancel due to the definition of the generator:

$$P_{t-s} f(X_s) - f(X_s) = \int_0^{t-s} P_u Lf(X_s)\,du,$$

which follows from integrating the generator over time.
Hence:

$$\mathbb{E}[M_t^f \mid \mathscr{F}_s] = M_s^f,$$

which confirms that M_t^f is a martingale. □

Dynkin's formula is a powerful tool: it allows us to translate generator calculations into martingale properties. In the context of the voter model, applying this formula to local observables shows that the density of 1's evolves as a martingale — a key fact used in Chapter 5 to analyze invariant measures and long-term behavior.

2.2.4. Stationary distributions

Definition 2.2.7 (Stationary distribution). *A probability measure μ on E is called stationary if*

$$\mu P_t = \mu, \qquad \forall t \geq 0.$$

That is, if $X_0 \sim \mu$ then $X_t \sim \mu$ for all $t \geq 0$.

A stationary distribution describes a state of balance: if the system starts in this distribution, it will look statistically the same at all future times. In other words, the probabilities of being in different states do not change over time. For the voter model, this means that the overall pattern of opinions remains stable, even though individual sites may continue to update. In Chapters 4 and 5, the invariant measures of the voter model are precisely its stationary distributions.

2.3. Interacting particle systems

Interacting particle systems (IPS) are a class of stochastic processes that model the collective behavior of locally interacting agents on a lattice. Each site carries a state (often called a spin), and the system evolves over time according to local update rules. These models describe phenomena such as the spread of diseases, opinion dynamics, and phase transitions.

A central subclass of IPS is the *spin-flip systems*, where each site updates its state by flipping its spin based on the configuration of its neighbors. The voter model is a prominent example of such a system, but other important models include the *contact process*, which models infection and recovery, and the *Ising model*, which describes ferromagnetic interactions.

Definition 2.3.1 (Spin-flip system). *A spin-flip system is a continuous-time Markov process on the configuration space $\Omega = \{0,1\}^{\mathbb{Z}^d}$ with generator*

$$Lf(\eta) = \sum_{x \in \mathbb{Z}^d} c(x,\eta) \left[f(\eta^x) - f(\eta) \right],$$

where η^x denotes the configuration η with the spin at site x flipped, and $c(x,\eta) \geq 0$ is the flip rate at site x given configuration η.

The flip rate $c(x,\eta)$ is assumed to be:

- *local*: it depends only on the values of η in a finite neighborhood of x,
- *uniformly bounded*: there exists $C < \infty$ such that $c(x,\eta) \leq C$ for all x and η.

A spin-flip system models a collection of agents (or spins) that update their state based on local interactions. Each site on the lattice decides whether to flip its state depending on the configuration of its neighbors. The generator encodes these dynamics by specifying the rate at which each site flips. The term $f(\eta^x) - f(\eta)$ measures the effect of flipping the spin at x on the observable f.

The generator L acts on functions $f : \Omega \to \mathbb{R}$, typically assumed to be *local observables*, meaning they depend only on finitely many coordinates. These functions serve as test functions to probe the infinitesimal behavior of the system.

A common choice for the flip rate is based on a symmetric, finite-range probability kernel $p(x, y)$, which describes the interaction structure. For example, in the voter model, the flip rate is:

$$c(x,\eta) = \sum_y p(x, y) \cdot \mathbf{1}\{\eta(x) \neq \eta(y)\},$$

meaning that site x copies the state of a randomly chosen neighbor y at rate $p(x, y)$.

Example 2.3.1 (Contact Process). *In the contact process, each site can be either infected (1) or healthy (0). The flip rates are:*

- *Infection:* $c(x,\eta) = \lambda \sum_y p(x, y) \cdot \eta(y)$ *if* $\eta(x) = 0$,

- *Recovery:* $c(x,\eta) = 1$ *if* $\eta(x) = 1$.

Here, λ *is the infection rate, and* $p(x, y)$ *defines the neighborhood. This model captures the spread of infection through local contact.*

Example 2.3.2 (Ising Model Dynamics)**.** *In Glauber dynamics for the Ising model, spins flip based on energy considerations. The flip rate depends on the sum of neighboring spins and a temperature parameter, modeling thermal fluctuations. This model is used to study phase transitions and magnetization.*

These examples illustrate the diversity of IPS: while the voter model is neutral and symmetric, the contact process introduces asymmetry (infection vs. recovery), and the Ising model incorporates energy-based interactions

2.3.1. Kolmogorov's Theorems

To rigorously define interacting particle systems such as the voter model, we need to construct stochastic processes on infinite product spaces. Kolmogorov's theorems provide the formal foundation for this.

Theorem 2.3.1 (Kolmogorov Consistency Theorem)**.** *Let* $\{P_{t_1,\dots,t_n}\}$ *be a family of probability measures on* E^n *for all finite sequences* $t_1 < \dots < t_n$ *in* $[0,\infty)$*, where* E *is a Polish space. Suppose this family is consistent in the sense that:*

- ***Permutation invariance****:* $P_{t_1,\dots,t_n}$ *is invariant under permutations of indices.*
- ***Marginalization****: For any* $t_1 < \dots < t_n$ *and* $k < n$*, the marginal of* $P_{t_1,\dots,t_n}$ *on the first* k *coordinates equals* $P_{t_1,\dots,t_k}$.

Then there exists a probability measure P *on the product space* $E^{[0,\infty)}$ *such that for all finite sets of times, the finite-dimensional distributions of the coordinate projections agree with* $\{P_{t_1,\dots,t_n}\}$.

Theorem 2.3.2 (Kolmogorov Extension Theorem)**.** *Let* $\{P_{t_1,\dots,t_n}\}$ *be a consistent family of finite-dimensional distributions on a Polish space* E*. Then there exists a stochastic process* $(X_t)_{t\geq 0}$ *with state space* E *such that for all finite sets of times* $t_1 < \dots < t_n$*, the joint distribution of* $(X_{t_1},\dots,X_{t_n})$ *is* $P_{t_1,\dots,t_n}$.

Proof. Let E be a Polish space (i.e., a complete separable metric space), and let $\{P_{t_1,\dots,t_n}\}$ be a family of finite-dimensional distributions indexed by finite increasing time tuples $t_1 < \dots < t_n$.

Assume that this family is *consistent*, meaning:

- (Symmetry) $P_{t_1,\dots,t_n}$ is invariant under permutations of indices.
- (Marginalization) For any $t_1 < \dots < t_n$ and any $k < n$, the marginal of $P_{t_1,\dots,t_n}$ on the first k coordinates equals $P_{t_1,\dots,t_k}$.

By Kolmogorov's consistency theorem, there exists a unique probability measure $\mathbb{P}$ on the product space $E^{[0,\infty)}$ (equipped with the product σ-algebra) such that for every finite set of times $t_1 < \dots < t_n$, the marginal distribution of $(X_{t_1},\dots,X_{t_n})$ under $\mathbb{P}$ is $P_{t_1,\dots,t_n}$.

This measure $\mathbb{P}$ defines a stochastic process $(X_t)_{t\geq 0}$ with state space E and the desired finite-dimensional distributions. The process is measurable and satisfies the required consistency conditions. □

These two theorems are closely related but serve distinct purposes:

- The **Consistency Theorem** ensures that a family of finite-dimensional distributions is compatible. That is, they can be stitched together without contradiction.
- The **Extension Theorem** guarantees that such a consistent family actually defines a stochastic process on the infinite product space.

Together, they provide the rigorous foundation for constructing processes on infinite lattices, such as the voter model and its dual. The consistency condition is a prerequisite for the extension result, and both are essential when working with interacting particle systems defined via local rules but evolving globally.

2.4. The Harris graphical construction

To analyze interacting particle systems such as the voter model and the contact process, it is often useful to have a concrete, visual representation of the system's evolution. One powerful method for this is the *Harris graphical construction*, introduced by T.E. Harris in the 1970s. This construction provides a probabilistic coupling of the process across all initial configurations and allows us to visualize the ancestry and influence structure of the system.

Definition 2.4.1 (Graphical representation of the voter model)**.** *For each ordered pair (x, y) with $p(x, y) > 0$, let $N^{x,y}$ be a Poisson process of rate $p(x, y)$. At each event time of $N^{x,y}$, draw an arrow $y \to x$. The evolution is defined by: whenever an arrow $y \to x$ occurs, set $\eta(x) \leftarrow \eta(y)$. The resulting process is the voter model with kernel p.*

Example 2.4.1 (Nearest-neighbor voter model)**.** *Let $p(x, y) = \frac{1}{2d}$ if x and y are nearest neighbors in $\mathbb{Z}^d$, and* 0 *otherwise. Then $N_{x,y}$ is a Poisson process with rate $\frac{1}{2d}$, and the graphical construction consists of arrows between neighboring sites. At each arrow time, site x copies the state of its neighbor y.*

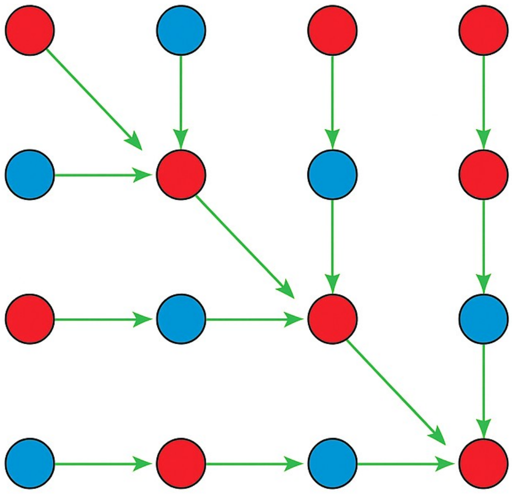

Figure 2.1: Graphical representation of the nearest-neighbor voter model. Nodes are colored red or blue to indicate their state, and green arrows represent update events where a site copies the state of a neighbor.

This construction is particularly well-suited for the voter model, where each site updates by imitating a randomly chosen neighbor. The graphical representation makes the dynamics explicit and provides a natural way to couple multiple processes starting from different initial configurations.

The arrows in the graphical construction encode the ancestry of each site: by tracing arrows backward in time, one can determine which sites influenced the current state. This ancestry structure is crucial for understanding duality and long-term behavior.

Proposition 2.4.1 (Consistency)**.** *The Harris construction produces a well-defined Markov process with generator*

$$Lf(\eta) = \sum_{x \in \mathbb{Z}^d} \sum_{y \in \mathbb{Z}^d} p(x, y) \cdot \mathbf{1}\{\eta(x) \neq \eta(y)\} \cdot \left[f(\eta^x) - f(\eta) \right],$$

simultaneously for all initial configurations η. The arrows are locally finite in time, ensuring that the evolution is well-defined.

The Harris construction enables powerful coupling arguments. It allows us to compare processes with different initial states, prove monotonicity and attractiveness, and construct dual processes. These properties are essential for analyzing convergence, invariant measures, and phase transitions.

2.5. Duality

Duality is a fundamental concept in the theory of Markov processes and interacting particle systems. It provides a powerful method for analyzing complex stochastic systems by relating them to simpler or more tractable dual processes. In many cases, duality allows us to reduce infinite-dimensional problems to finite ones, or to relate long-time behavior to short-time computations.

Suppose we are interested in the behavior of a Markov process $(\eta_t)_{t\geq 0}$ on a state space E. If we can find another process $(\xi_t)_{t\geq 0}$ on a possibly different state space F, and a function $H : E \times F \to \mathbb{R}$ such that the expected value of H under one process equals the expected value under the other (with time reversed), then we can often transfer results between the two processes.

Definition 2.5.1 (Duality). *Let (η_t) and (ξ_t) be Markov processes with state spaces E and F, and generators L and $L^\dagger$, respectively. We say that (η_t) and (ξ_t) are* dual *with respect to a function $H : E \times F \to \mathbb{R}$ if for all $\eta \in E$, $\xi \in F$, and $t \geq 0$,*

$$\mathbb{E}_\eta[H(\eta_t, \xi)] = \mathbb{E}_\xi[H(\eta, \xi_t)].$$

The function H is called the *duality function.* It often has a product structure or indicator form, depending on the processes involved.

3

Random Walks and Potential Theory

This chapter introduces the analytic tools needed to study spatial stochastic processes, focusing on random walks and potential theory. These concepts are central to understanding recurrence, transience, and long-term behavior in lattice-based systems.

We begin by defining discrete and continuous-time random walks and their generators. We then explore recurrence and transience, Green's functions, harmonic functions, and ergodic theory on $\mathbb{Z}^d$. These results will be essential in Chapter 4, where we apply them to analyze the voter model.

3.1. Continuous-time random walks

Random walks are fundamental stochastic processes that model the movement of a particle through space. In this section, we define both discrete-time and continuous-time versions on the lattice $\mathbb{Z}^d$, and describe their generators. These walks will later serve as the dual objects to the voter model, but for now we focus on their intrinsic properties.

Definition 3.1.1 (Discrete-time random walk)**.** *Let $p : \mathbb{Z}^d \times \mathbb{Z}^d \to [0,1]$ be a transition kernel with $\sum_y p(x, y) = 1$. A discrete-time random walk $(S_n)_{n\geq 0}$ on $\mathbb{Z}^d$ is a Markov chain with transition probabilities*

$$\mathbb{P}(S_{n+1} = y \mid S_n = x) = p(x, y).$$

In continuous time, the particle waits a random amount of time before jumping. Specifically, the waiting times are exponentially distributed with mean 1.

Definition 3.1.2 (Continuous-time random walk)**.** *Given a kernel $p(x, y)$ as above, a continuous-time random walk (CTRW) $(X_t)_{t\geq 0}$ is a process that waits an exponential(1) time between jumps and then moves according to p. Equivalently, $X_t = S_{N_t}$ where (N_t) is a rate-1 Poisson process and (S_n) is the discrete-time walk.*

We use X_t to denote the position of the particle at time t, and S_n to denote its position after n jumps. The Poisson process N_t counts the number of jumps up to time t.

The infinitesimal behavior of a continuous-time Markov process is described by its generator.

Proposition 3.1.1 (Generator of CTRW)**.** *The generator of (X_t) acts on bounded functions $f : \mathbb{Z}^d \to \mathbb{R}$ as*

$$(\mathscr{L}_{\mathrm{RW}} f)(x) = \sum_y p(x, y)(f(y) - f(x)).$$

To make this concrete, consider the following example:

Example 3.1.1 (Nearest-Neighbor CTRW)**.** *Let $p(x, y) = \frac{1}{2d}$ if $\|x - y\|_1 = 1$ and $p(x, y) = 0$ otherwise. Here $\|\cdot\|_1$ denotes the ℓ^1-norm, i.e., the sum of absolute coordinate differences. Then (X_t) jumps at rate 1 to each of its $2d$ nearest neighbors with equal probability. This is the simple symmetric nearest-neighbor walk on $\mathbb{Z}^d$.*

3.2. Recurrence and transience

In this section, we study the long-term behavior of random walks on the lattice $\mathbb{Z}^d$. A central question is whether a random walk tends to return to its starting point repeatedly (recurrence), or whether it eventually escapes and never comes back (transience).

Definition 3.2.1 (Recurrence and transience)**.** *A discrete-time random walk* (S_n) *is recurrent if*

$$\mathbb{P}_0(S_n = 0 \text{ i.o.}) = 1,$$

i.e. it returns to the origin infinitely often with probability 1. Otherwise it is transient. The same definitions apply to CTRWs via their jump chains.

To understand recurrence and transience, we analyze the probability $p_n(0)$ that the walk is at the origin after n steps. For the simple symmetric random walk, this is given by the Fourier transform of the step distribution.

Theorem 3.2.1 (Classical dichotomy)**.** *Let* $(S_n)_{n\geq 0}$ *be a simple symmetric random walk on* $\mathbb{Z}^d$*. Then:*

- *The walk is **recurrent** for* $d = 1, 2$*,*
- *The walk is **transient** for* $d \geq 3$.

Proof. We analyze this using the Fourier transform of the transition kernel.
The *characteristic function* of a probability distribution is the Fourier transform of its probability measure. For a random variable X taking values in $\mathbb{Z}^d$, the characteristic function $\phi(\theta)$ is defined as

$$\phi(\theta) = \mathbb{E}[e^{i\theta\cdot X}], \quad \theta \in [-\pi, \pi]^d.$$

In the case of the simple symmetric random walk, this becomes

$$\phi(\theta) = \frac{1}{d}\sum_{j=1}^{d} \cos(\theta_j).$$

Then, the probability that the walk returns to the origin after n steps is

$$p_n(0) = \frac{1}{(2\pi)^d}\int_{[-\pi,\pi]^d} \left[\phi(\theta)\right]^n d\theta.$$

Summing over n, we get the expected number of returns:

$$\sum_{n=0}^{\infty} p_n(0) = \frac{1}{(2\pi)^d}\int_{[-\pi,\pi]^d} \sum_{n=0}^{\infty} \left[\phi(\theta)\right]^n d\theta = \frac{1}{(2\pi)^d}\int_{[-\pi,\pi]^d} \frac{1}{1-\phi(\theta)} d\theta.$$

Now we analyze the convergence of this integral:

- For $d = 1, 2$, the integrand $\frac{1}{1-\phi(\theta)}$ has a non-integrable singularity at $\theta = 0$, and the integral diverges.
- For $d \geq 3$, the singularity is integrable, and the integral converges.

□

The divergence of the integral for $d = 1, 2$ reflects the fact that the random walk tends to revisit its starting point frequently. In low dimensions, the walker "gets trapped" in its local neighborhood due to the limited number of directions it can escape to. In contrast, in higher dimensions ($d \geq 3$), the walker has more "room" to move, and the probability of returning to the origin decreases rapidly enough for the total expected number of returns to remain finite.

3.3. Green's functions and potential kernel

A useful tool to analyze recurrence and correlations is the Green's function.
Let $(X_t)_{t\geq 0}$ be a continuous-time random walk (CTRW) on $\mathbb{Z}^d$ with transition kernel $p(x,y)$. The Green's function is defined as:

$$G(x,y) = \int_0^\infty P_x(X_t = y)\,dt.$$

For the discrete-time random walk (S_n), the Green's function is:

$$G(x,y) = \sum_{n=0}^{\infty} P_x(S_n = y).$$

Proposition 3.3.1 (Characterization of Recurrence)**.** *The walk is recurrent if and only if* $G(0,0) = \infty$, *and transient if and only if* $G(0,0) < \infty$.

Proof. The quantity $G(0,0)$ represents the expected number of visits to the origin:

$$G(0,0) = \sum_{n=0}^{\infty} P_0(S_n = 0).$$

If this sum diverges, the walk returns to the origin infinitely often with probability 1, hence it is recurrent. If the sum converges, the expected number of returns is finite, and the walk is transient. □

Theorem 3.3.1. *For the simple symmetric random walk on* $\mathbb{Z}^d$ *with* $d \geq 3$, *there exists a constant* $c_d > 0$ *such that:*

$$G(0,x) \sim c_d |x|^{2-d}, \quad \textit{as } |x| \to \infty.$$

This asymptotic behavior shows that the Green's function decays algebraically in high dimensions. In the context of the voter model, this implies that correlations between spins at sites 0 and x decay like $|x|^{2-d}$, reflecting the critical nature of the model.

3.4. Harmonic functions and the Liouville theorem

In this section, we introduce harmonic functions, which play a central role in the classification of invariant measures for the voter model.

Definition 3.4.1 (Harmonic function)**.** *Let* p *be a transition kernel on* $\mathbb{Z}^d$. *A function* $h : \mathbb{Z}^d \to \mathbb{R}$ *is called* harmonic *(with respect to* p*) if*

$$h(x) = \sum_{y\in\mathbb{Z}^d} p(x,y)h(y), \quad \textit{for all } x \in \mathbb{Z}^d.$$

This means that $h(x)$ is the expected value of h after one step of the random walk starting at x. Harmonic functions are precisely the fixed points of the averaging operator induced by p.

Proposition 3.4.1 (Maximum principle)**.** *Let* h *be a bounded harmonic function on* $\mathbb{Z}^d$. *Then for every finite subset* $\Lambda \subset \mathbb{Z}^d$,

$$\min_{x\in\partial\Lambda} h(x) \leq h(y) \leq \max_{x\in\partial\Lambda} h(x), \quad \textit{for all } y \in \Lambda.$$

Proof. Since $h(y)$ is the average of its neighbors, it cannot exceed the maximum or be less than the minimum of them. □

We now turn to a fundamental result known as the Liouville theorem, which characterizes harmonic functions in low dimensions.

Theorem 3.4.1 (Liouville property)**.** *Let* p *be a symmetric finite-range kernel on* $\mathbb{Z}^d$. *If* $d \leq 2$, *then every bounded harmonic function is constant.*

Proof. Fix $y \in \mathbb{Z}^d$ and consider the hitting time τ_y of y for a simple symmetric random walk (X_n). Since the walk is recurrent for $d \leq 2$, we have

$$\mathbb{P}_x(\tau_y < \infty) = 1 \quad \text{for all } x \in \mathbb{Z}^d.$$

Now consider the process $(h(X_n))$ and apply the *optional stopping theorem* to the stopping time $\tau_y \wedge n$, which is the minimum of τ_y and n. This theorem states that if (M_n) is a martingale and τ is a bounded stopping time, then

$$\mathbb{E}_x[M_\tau] = M_0.$$

In our case, $(h(X_n))$ is a bounded martingale because h is harmonic and bounded. Therefore,

$$h(x) = \mathbb{E}_x[h(X_{\tau_y \wedge n})].$$

Taking the limit $n \to \infty$ and using the recurrence of the walk, we get

$$h(x) = \mathbb{E}_x[h(X_{\tau_y})] = h(y),$$

since $X_{\tau_y} = y$ almost surely. As x and y were arbitrary, it follows that h is constant. □

This result has direct consequences for the voter model. In dimensions $d \leq 2$, the recurrence of random walks forces all bounded harmonic functions to be constant.
In contrast, for $d \geq 3$, the random walk is transient, and nontrivial bounded harmonic functions exist. For example, the function

$$h(x) = \mathbb{P}_x(\text{the walk never returns to } 0)$$

is bounded between 0 and 1, nonconstant, and harmonic. This reflects the richer structure of invariant measures in high dimensions, where coexistence is possible.

3.5. Ergodic theory on $\mathbb{Z}^d$

We conclude this chapter with some results of ergodic theory, which allows us to understand the long-term behavior of translation-invariant measures in spin systems.

Definition 3.5.1 (Translation invariance and ergodicity)**.** *A probability measure μ on Ω is called* translation-invariant *if*

$$\mu \circ \tau_x^{-1} = \mu \quad \textit{for all } x \in \mathbb{Z}^d,$$

where the shift operator τ_x acts on configurations by $(\tau_x\eta)(y) = \eta(y+x)$.
The measure μ is called ergodic *if every translation-invariant event has probability* 0 *or* 1 *under μ.*

Ergodicity ensures that spatial averages converge to expectations under μ, which allows us to interpret $\mu[\eta(0)]$ as the density of 1's in the system.

Theorem 3.5.1 (Birkhoff's ergodic theorem for $\mathbb{Z}^d$)**.** *Let μ be a translation-invariant probability measure on Ω, and let $f : \Omega \to \mathbb{R}$ be a local function with $\int |f|\, d\mu < \infty$. Then*

$$\frac{1}{|B_n|} \sum_{x \in B_n} f(\tau_x \eta) \xrightarrow[n\to\infty]{\mu\text{-a.s.}} \int f \, d\mu,$$

where $B_n = [-n, n]^d$ is the box of side length $2n+1$ centered at the origin.

Corollary 3.5.1 (Spin averages)**.** *Let μ be translation-invariant. Then*

$$\frac{1}{|B_n|} \sum_{x \in B_n} \eta(x) \xrightarrow[n\to\infty]{\mu\text{-a.s.}} \mu[\eta(0)].$$

This result justifies interpreting $\mu[\eta(0)]$ as the average density of 1's in the system. In Chapter 5, we need this for the extremal invariant measures ν_ρ, which are parameterized by this density.

4

The Voter Model

In this chapter we introduce the linear voter model carefully. We start by defining the model with general interaction kernels, with emphasis on the k-nearest-neighbour (k-NN) version as an illustrative example. We then establish basic properties such as attractiveness and conservation of density. Next, we develop the graphical construction and the duality with coalescing random walks. Equipped with this, we can prove the central dichotomy: *clustering* in dimensions $d \leq 2$ and *coexistence* in $d \geq 3$. Along the way, examples and intuition will clarify why these results matter. We also discuss finite systems and preview nonlinear variants for the next chapter.

This chapter is based primarily on the foundational works of Liggett [1985, Ch. V] and Durrett [2008, Ch. 1], which provide a comprehensive treatment of the voter model and its probabilistic structure.

4.1. Definition

A voter model is nothing more than a Markov process η_t on $\{0,1\}^{Z^d}$ whose generator has the form of

$$\Omega f(\eta) = \sum_x c(x,\eta)\left[f(\eta_x) - f(\eta)\right],$$

where the rate function $c(x,\eta)$ has the following properties:

(a) $c(x,\eta) = 0$ for every $x \in \mathbb{Z}^d$ if $\eta \equiv 0$ or if $\eta \equiv 1$,

(b) $c(x,\eta) = c(x,\zeta)$ for every $x \in \mathbb{Z}^d$ if $\eta(y) + \zeta(y) = 1$ for all $y \in \mathbb{Z}^d$,

(c) $c(x,\eta) \leq c(x,\zeta)$ if $\eta \leq \zeta$ and $\eta(x) = \zeta(x) = 0$, and

(d) $c(x,\eta)$ is invariant under shifts in $\mathbb{Z}^d$.

Property (a) just implies that if we are in a constant configuration, $\eta \equiv 0$ or if $\eta \equiv 1$, than the flip rate of an individual voter is 0. The second property states that if we interchange the roles of 0 and 1 the system will evolve the same way. Property (c) makes the process attractive. The last property states that if we shift in space the process is invariant.

Definition 4.1.1 (Attractiveness)**.** *A process is attractive or monotone if it satisfies the following equivalent conditions.*

$$f \text{ increasing implies } S(t)f \text{ increasing for all } t \geq 0$$

and

$$\mu_1 \leq \mu_2 \text{ implies } \mu_1 S(t) \leq \mu_2 S(t) \text{ for all } t \geq 0.$$

This is useful because it allows for coupling arguments and comparison between configurations. In attractive systems, larger initial configurations tend to remain larger over time.

Definition 4.1.2. *In the linear voter model are the flip rates given by*

$$c(x,\eta) = \begin{cases} \sum_y p(x,y)\,\eta(y), & \text{if } \eta(x) = 0, \\ \sum_y p(x,y)\,[1-\eta(y)], & \text{if } \eta(x) = 1, \end{cases} \tag{4.1}$$

where $\eta \in \{0,1\}^{\mathbb{Z}^d}$ *is a configuration and* $p(x,y)$ *is a probability kernel satisfying*

$$p(x,y) \geq 0, \qquad \sum_y p(x,y) = 1 \quad \textit{for all}$$

This definition indeed satisfies the structural properties listed at the beginning of this subsection. These properties are directly implied by the form of $c(x,\eta)$ above, and they make the model analytically tractable. This formulation defines the specific version of the voter model we will analyze throughout the chapter.

4.2. Basic Properties of the Voter Model

Having defined the voter model via its flip rates $c(x,\eta)$, we now turn to its most fundamental structural properties. These properties explain why the model is central to the theory of interacting particle systems and why it serves as a prototype for more general spin systems.

4.2.1. Duality

The voter model exhibits a powerful duality with coalescing random walks. This duality allows us to analyze infinite systems by studying finite ones, and is closely tied to the monotonicity of the process.

Theorem 4.2.1 (Duality and monotonicity, adapted from Liggett)**.** *Let* $(\eta_t)_{t\geq 0}$ *be a Feller process on* $\{0,1\}^{\mathbb{Z}^d}$ *such that* $\{1\}$ *is an absorbing state and* $P_t[\eta_t = 1] = 0$ *for all* $t > 0$. *Then there exists a Feller process* $(\xi_t)_{t\geq 0}$ *which is dual to* (η_t) *with respect to the function* $H(\eta,\xi) = \eta^\xi$ *if and only if* (η_t) *is monotone.*

Proof. ($\Rightarrow$) Suppose a dual process (ξ_t) exists with respect to $H(\eta,\xi) = \eta^\xi$. Then for all $t \geq 0$,

$$\mathbb{E}_\eta[\eta^{\xi_t}] = \mathbb{E}_\xi[\eta_t^\xi].$$

Fix ξ and consider the function $\eta \mapsto \mathbb{E}_\eta[\eta^{\xi_t}]$. Since $\eta \mapsto \eta^\xi$ is increasing in η for fixed ξ, and the expectation preserves monotonicity, it follows that $\eta \mapsto \mathbb{E}_\eta[\eta^{\xi_t}]$ is increasing. Hence, (η_t) must be monotone.
($\Leftarrow$) Suppose (η_t) is monotone. We want to construct a dual process (ξ_t) such that

$$\mathbb{E}_\eta[\eta^{\xi_t}] = \mathbb{E}_\xi[\eta_t^\xi]$$

for all η, ξ, and $t \geq 0$. For each fixed ξ, define the function

$$F_\eta(t) := \mathbb{E}_\eta[\eta^{\xi_t}].$$

Because $\eta \mapsto \eta^\xi$ is increasing and (η_t) is monotone, the function $\eta \mapsto F_\eta(t)$ is increasing for all t. This family of functions defines a consistent set of finite-dimensional distributions for (ξ_t) via the relation

$$\mathbb{E}_\eta[\eta^{\xi_t}] := \mathbb{E}_\xi[\eta_t^\xi].$$

By Kolmogorov's extension theorem, this defines a stochastic process (ξ_t) with the desired duality property. The Feller property follows from the continuity of the semigroup and the monotonicity of the process. □

4.2.2. Clustering and Coexistence

A central question in the analysis of the voter model is whether the system converges to consensus or allows for long-term coexistence of different opinions. The answer depends critically on the dimension d and the properties of the associated random walk.

Theorem 4.2.2 (Clustering vs. Coexistence)**.** *Let the voter model be defined with a symmetric, finite-range interaction kernel* $p(x,y)$.

- *If* $d \leq 2$, *the model* ***clusters****: with probability 1, all sites eventually agree, and the system converges to one of the absorbing states* δ_0 *or* δ_1.

- *If $d \geq 3$, the model **coexists**: there exists a nontrivial stationary distribution in which both opinions persist indefinitely.*

This result applies specifically to the *nearest-neighbor voter model*, where $p(x,y)$ assigns equal probability to adjacent sites. For more general kernels, the outcome depends on whether the associated random walk is *recurrent* or *transient*. In low dimensions, recurrence ensures that ancestral lineages meet almost surely, forcing consensus. In higher dimensions, transience allows lineages to escape each other, enabling coexistence.

It is important to note that this dichotomy relies on the symmetry of the kernel. If $p(x,y)$ is not symmetric, the associated random walk may behave differently, and the clustering/coexistence behavior may change. Thus, the long-term behavior of the voter model is determined by the probabilistic properties of the dual random walk, which in turn depend on the structure of p.

4.2.3. Examples: Nearest-Neighbor and k-th Neighbor Voter Models

So far, we have considered a general kernel $p(x,y)$. Let us now study concrete examples.

Let $S = \mathbb{Z}^d$ and
$$p(x,y) = \frac{1}{2d}\mathbf{1}_{\{|x-y|=1\}},$$
i.e. each site updates by imitating one of its $2d$ nearest neighbors with equal probability. This is the standard voter model considered in most of the literature. By Theorem 4.2.2, it clusters in dimensions $d \leq 2$ and coexists in dimensions $d \geq 3$.

Fix $k \geq 1$. Define
$$p(x,y) = \frac{1}{N_k}\mathbf{1}_{\{|x-y|=k\}},$$
where N_k is the number of sites at distance k from x. Here each site copies uniformly from its k-th nearest neighbors.

This model illustrates how the range of interactions affects clustering and coexistence:

- In one dimension, using only second- or third-nearest neighbors does not change recurrence of random walks, the model still clusters.
- In higher dimensions, the larger interaction range makes the coalescing random walks even more transient. Intuitively, when each site can copy from more distant neighbors, the paths of the dual random walks spread out faster and are less likely to intersect. This reduces the chance that different lineages coalesce, allowing distinct opinions to survive longer and increasing the likelihood of coexistence.

The duality theorem gives us the mathematical tool to understand long-term behavior, while the clustering vs. coexistence dichotomy shows the dramatic dependence on dimension. The examples above help us visualize how the choice of kernel $p(x,y)$ influences outcomes. Together, these results form the conceptual foundation of voter model theory.

4.3. Absorbing States and Invariant Measures

Having established the duality and the clustering–coexistence dichotomy, we now turn to the possible long-term states of the voter model. These are described by the *invariant measures* of the process.

Definition 4.3.1. *The configurations* $\mathbf{0}$ *(all sites in state* 0*) and* $\mathbf{1}$ *(all sites in state* 1*) are called the absorbing states of the voter model.*

Once the system reaches $\mathbf{0}$ or $\mathbf{1}$ it never leaves, since no updates can change any site. In finite systems, these are the only possible limiting states. In infinite systems, the situation is richer.

Definition 4.3.2 (Invariant measure)**.** *A probability measure μ on $\{0,1\}^{\mathbb{Z}^d}$ is called an invariant measure for the voter model if $\mu P_t = \mu$ for all $t \geq 0$, where P_t is the transition semigroup of the process.*

Theorem 4.3.1 (Extremal invariant measures)**.**

- *For $d \leq 2$, the only extremal invariant measures are $\delta_{\mathbf{0}}$ and $\delta_{\mathbf{1}}$.*

- *For $d \geq 3$, there exists a one-parameter family $\{\nu_\rho : \rho \in [0,1]\}$ of translation-invariant ergodic invariant measures, where $\nu_\rho(\eta(0) = 1) = \rho$.*

In low dimensions, the only "equilibria" are the absorbing states, reflecting the inevitability of consensus. In high dimensions, a whole family of nontrivial equilibria exists, parameterized by the density ρ of 1's. These measures represent coexistence of opinions at all densities.

4.3.1. Connection with duality

These invariant measures can be constructed directly from the dual random walks. For $d \geq 3$, duality implies that for initial product measure ν_ρ, the distribution at time t is given by

$$\mathbb{E}_{\nu_\rho}\left[\prod_{x\in A}\eta_t(x)\right] = \mathbb{E}\left[\rho^{|A_t|}\right],$$

where A_t is the set of coalescing random walks started from A at time 0.

The set A_t evolves as a stochastic process: each particle in A performs an independent random walk and coalesces with others upon meeting. Thus, A_t tracks the positions of the remaining (non-coalesced) particles at time t, and its distribution determines the structure of the invariant measure ν_ρ via the duality relation.

4.4. Convergence to Equilibrium

Having identified the invariant measures of the voter model, we now ask: does the system converge to one of these measures as $t \to \infty$? The answer again depends on the dimension d.

Theorem 4.4.1 (Convergence to equilibrium).

- *If $d \leq 2$, the process (η_t) converges weakly to a* mixture *of the absorbing states δ_0 and δ_1. That is, the limiting distribution is a convex combination $\alpha\delta_0 + (1-\alpha)\delta_1$, where α depends on the initial density of 1's.*
- *If $d \geq 3$ and $\eta_0 \sim \nu_\rho$, then η_t converges weakly to ν_ρ as $t \to \infty$.*

Proof. We use duality with coalescing random walks. Fix a finite set $A \subset \mathbb{Z}^d$. By duality, we have

$$\mathbb{E}_{\eta_0}\left[\prod_{x\in A}\eta_t(x)\right] = \mathbb{E}\left[\prod_{x\in A_t}\eta_0(x)\right],$$

where A_t is the set of coalescing random walks started from A at time 0.

In dimensions $d \leq 2$, the random walk is recurrent, so all walks in A_t eventually coalesce into a single site. Thus, the product $\prod_{x\in A_t}\eta_0(x)$ becomes either 0 or 1, depending on the value at the coalesced site. This implies that the system converges to either δ_0 or δ_1, and the limiting distribution is a convex combination $\alpha\delta_0 + (1-\alpha)\delta_1$, where $\alpha = \mathbb{E}[\eta_0(x)]$ is the initial density of 1's.

In dimensions $d \geq 3$, the random walk is transient, so the number of distinct particles in A_t remains positive with high probability. If $\eta_0 \sim \nu_\rho$, then each $\eta_0(x)$ is an independent Bernoulli variable with mean ρ, and

$$\mathbb{E}\left[\prod_{x\in A_t}\eta_0(x)\right] = \mathbb{E}\left[\rho^{|A_t|}\right].$$

Since the law of A_t does not depend on η_0, and ν_ρ is invariant, it follows that $\eta_t \sim \nu_\rho$ for all t, and hence η_t converges weakly to ν_ρ. □

In low dimensions, clustering forces the system into one of the absorbing states. The mixture $\alpha\delta_0 + (1-\alpha)\delta_1$ reflects the probability that the system ends up in consensus state 0 or 1, depending on the initial configuration. This mixture is not an evolving distribution but a fixed convex combination determined by the initial density.

In high dimensions, the invariant measures ν_ρ are translation-invariant and ergodic. Since they are stationary, starting the process in ν_ρ means the distribution does not change over time. The point of the second bullet is that ρ represents the density of 1's, and since ρ can be strictly between 0 and 1, the system exhibits genuine coexistence.

4.4.1. Consensus times in finite systems

On finite graphs, the system almost surely reaches consensus. Duality again gives estimates for the consensus time.

Proposition 4.4.1 (Consensus time on the torus)**.** *On the d-dimensional torus with n^d sites, the expected time to consensus is of order*

$$\begin{cases} n^2, & d=1, \\ n^2 \log n, & d=2, \\ n^d, & d \geq 3. \end{cases}$$

This result follows from the time it takes for n^d coalescing random walks to merge into one. In one dimension, this is diffusive (n^2). In two dimensions, the recurrence of random walks produces a logarithmic correction. In higher dimensions, walks are less likely to meet, so the time is of order the system volume n^d.

4.5. Extensions of the Voter Model

The linear voter model is neutral: it treats both opinions symmetrically. Many variations modify the update rules to model different types of interactions.

4.5.1. Biased voter model

In the biased voter model, one opinion has an advantage. Formally, the flip rates are

$$c_\alpha(x,\eta) = \begin{cases} (1-\alpha)\sum_y p(x,y)\eta(y), & \eta(x)=0, \\ \alpha \sum_y p(x,y)(1-\eta(y)), & \eta(x)=1, \end{cases}$$

where $\alpha \in [0,1]$ is a bias parameter. For $\alpha < 1/2$, opinion 0 is favored; for $\alpha > 1/2$, opinion 1 is favored. This destroys coexistence: eventually the favored opinion wins out. You can think about this bias as a stubborn friend. Who will not change his mind no matter the arguments.

4.5.2. Threshold voter model

Another variation is the threshold voter model, where a site changes state only if sufficiently many neighbors disagree. For example, in the majority rule version, a site flips if more than half of its neighbors are in the opposite state. These models show very different behavior: coexistence can occur even in one and two dimensions.

For example, in the majority rule version, a site flips if more than half of its neighbors are in the opposite state. This corresponds to a flip rate

$$c(x,\eta) = 1\left\{ \sum_{y\sim x} 1\{\eta(y) \neq \eta(x)\} > \frac{1}{2}\deg(x) \right\},$$

where the sum is over neighbors y of x, and $\deg(x)$ is the number of neighbors. Unlike the linear voter model, this rate is not linear in η and introduces a nonlinearity that can lead to very different behavior, including phase transitions and coexistence even in low dimensions.

5

Invariant Measures and Long-Term Behavior

In Chapter 4 we described the voter model, its duality with coalescing random walks, and the clustering–coexistence dichotomy. We saw that the long-term behavior of the system is captured by its invariant measures: in low dimensions the absorbing states are the only invariants, while in higher dimensions a continuum of nontrivial invariant measures arises.

In this chapter we deepen the analysis. We first introduce the connection between invariant measures and harmonic functions. We then state and prove theorems about extremal invariant measures. Finally, we use ergodic theorems to describe convergence of the system towards equilibrium, both in infinite and finite systems.

The results presented here are based on foundational work by Liggett and Durrett, whose treatments of interacting particle systems and stochastic processes provided the theoretical backbone for our approach.

5.1. Harmonic Functions and Martingales

We begin by asking a fundamental question: if the voter model evolves for a long time, what kind of equilibria can it reach? These equilibria are precisely the invariant measures of the process. But how can we characterize them?

One of the most powerful methods is to relate invariant measures of spin systems to *harmonic functions* of the underlying random walk. This connection emerges naturally from the generator of the voter model, and it can be sharpened by martingale arguments. In this section we develop this link in detail, starting from definitions, then proving properties, and finally building up to classification results.

5.1.1. Harmonic Functions

Definition 5.1.1 (Harmonic function)**.** *Let (X_t) be a random walk on $\mathbb{Z}^d$ with transition kernel $p(x,y)$. A function $h:\mathbb{Z}^d \to \mathbb{R}$ is called* harmonic *if*

$$h(x) = \sum_y p(x,y)h(y), \qquad \forall x \in \mathbb{Z}^d.$$

The meaning of this definition is that $h(x)$ is the expected value of h after one step of the random walk. In other words, harmonic functions are precisely the fixed points of the averaging operator induced by p.

Harmonic functions are not just abstract objects: they encode equilibrium behavior of Markov processes. To see why they matter here, we connect them to invariant measures of the voter model.

5.1.2. Invariant Measures and Harmonic Functions

Proposition 5.1.1 (Harmonic functions from invariant measures). *If μ is a translation-invariant invariant measure for the voter model. Define $\eta \sim \mu$ as a random configuration drawn from μ. Then the function*

$$h(x) = \mu[\eta(x)]$$

is harmonic for p.

Proof. Let $\eta \sim \mu$ be a random configuration drawn from the invariant measure μ. Since μ is invariant, the distribution of η_t remains equal to μ for all $t \geq 0$. This means that for any site x,

$$\frac{d}{dt}\mathbb{E}_\mu[\eta_t(x)] = 0.$$

On the other hand, the generator L of the voter model acts on the spin at site x as

$$L\eta(x) = \sum_y p(x,y)\left(\eta(y) - \eta(x)\right).$$

Taking expectations under μ and using linearity, we get

$$\frac{d}{dt}\mathbb{E}_\mu[\eta_t(x)] = \sum_y p(x,y)\left(\mathbb{E}_\mu[\eta_t(y)] - \mathbb{E}_\mu[\eta_t(x)]\right).$$

Since the left-hand side is zero (by invariance), this implies

$$\sum_y p(x,y)\left(h(y) - h(x)\right) = 0,$$

where $h(x) := \mathbb{E}_\mu[\eta(x)]$. Rearranging gives

$$h(x) = \sum_y p(x,y)\, h(y),$$

which is precisely the definition of harmonicity. So the invariance of μ ensures that the expected spin value at each site satisfies the harmonic condition. □

Invariant measures correspond to harmonic functions. Thus the classification of invariant measures reduces to the classification of bounded harmonic functions of the kernel p.

5.1.3. Martingale Properties

A second way to see this connection is through martingale methods. Martingales naturally arise in interacting particle systems when we look at averages of conserved quantities. For the voter model, the relevant quantity is the density of 1's.

Proposition 5.1.2 (Density martingale). *Let η_t be the voter model. Define the empirical density of 1's in a finite box $B_n = [-n,n]^d$ by*

$$M_t^{(n)} = \frac{1}{|B_n|}\sum_{x \in B_n} \eta_t(x).$$

Then $(M_t^{(n)})_{t\geq 0}$ is a martingale with respect to the natural filtration, and in particular

$$\mathbb{E}[M_t^{(n)}] = \mathbb{E}[M_0^{(n)}], \qquad \forall\, t \geq 0.$$

In the infinite-volume limit, the global density $\rho_t = \mathbb{E}[\eta_t(0)]$ satisfies $\rho_t = \rho_0$ for all t.

Proof. Fix n and consider $M_n(t)$. The generator of the voter model acts on $\eta(x)$ as

$$L\eta(x) = \sum_y p(x,y)\left(\eta(y) - \eta(x)\right).$$

Taking expectations and differentiating in time gives

$$\frac{d}{dt}\mathbb{E}[M_n(t)] = \frac{1}{|B_n|}\sum_{x \in B_n}\sum_y p(x,y)\left(\mathbb{E}[\eta_t(y)] - \mathbb{E}[\eta_t(x)]\right).$$

Interchanging the sums and using symmetry of $p(x, y)$, the double sum cancels:

$$\sum_{x,y} p(x, y)\left(\mathbb{E}[\eta_t(y)] - \mathbb{E}[\eta_t(x)]\right) = 0.$$

Hence, $\mathbb{E}[M_n(t)]$ is constant in t. To conclude that $M_n(t)$ is a martingale, we note that it is adapted to the natural filtration, has bounded increments, and satisfies

$$\mathbb{E}[M_n(t) \mid \mathscr{F}_s] = M_n(s) \quad \text{for } s \le t,$$

by the Markov property and linearity of expectation.
For the global density, translation invariance of the measure implies that the expected value of $\eta_t(x)$ is the same for all x. Therefore,

$$\rho_t := \mathbb{E}[\eta_t(x)] = \mathbb{E}[\eta_t(0)] \quad \text{for all } x,$$

and since the expectation is constant in time, we have $\rho_t = \rho_0$. □

This proposition formalizes the neutrality of the voter model: the expected density of 1's never changes. In fact, the martingale convergence theorem implies that densities converge almost surely, a fact that will be crucial in later sections.

5.1.4. The Liouville Property

We now return to harmonic functions. The key question is: what harmonic functions exist? The answer depends dramatically on the dimension of the underlying lattice.

Theorem 5.1.1 (Liouville property)**.** *If $d \le 2$, the only bounded harmonic functions for simple random walk are constants.*

Proof. Let h be a bounded harmonic function, i.e.,

$$h(x) = \sum_y p(x, y)h(y) \quad \text{for all } x \in \mathbb{Z}^d.$$

This condition implies that the process $(h(X_n))_{n\ge 0}$, where X_n is a simple symmetric random walk, is a bounded martingale. Indeed, the harmonicity ensures that

$$\mathbb{E}[h(X_{n+1}) \mid X_n] = h(X_n),$$

which is the defining property of a martingale.
Now fix $y \in \mathbb{Z}^d$ and let τ_y be the hitting time of y. Since the walk is recurrent for $d \le 2$, we have

$$\mathbb{P}_x(\tau_y < \infty) = 1 \quad \text{for all } x.$$

Apply the optional stopping theorem to the bounded martingale $h(X_n)$ and the stopping time $\tau_y \wedge n$:

$$h(x) = \mathbb{E}_x[h(X_{\tau_y \wedge n})].$$

Taking the limit as $n \to \infty$ and using recurrence, we get

$$h(x) = \mathbb{E}_x[h(X_{\tau_y})] = h(y),$$

since $X_{\tau_y} = y$ almost surely. As x and y were arbitrary, it follows that h is constant. □

Corollary 5.1.1 (Trivial invariant measures in low dimensions)**.** *For $d \le 2$, the only translation-invariant invariant measures of the voter model are $\delta_{\mathbf{0}}$ and $\delta_{\mathbf{1}}$, or their convex mixtures.*

This result reveals the first major dichotomy: in low dimensions, recurrence of random walks forces all harmonic functions to be trivial, and therefore the only possible equilibria of the voter model are consensus states.

6
Conclusion

In this thesis, we have developed a rigorous and self-contained analysis of the linear voter model, a fundamental example of an interacting particle system. Our primary goal was to understand the long-term behavior of the model, particularly the dichotomy between clustering and coexistence, and to explore the role of duality in deriving these results.

Chapter 2 laid the probabilistic foundation by introducing continuous-time Markov processes and the framework of interacting particle systems. We defined key concepts such as generators, semigroups, and invariant measures, and discussed the Harris graphical construction as a tool for visualizing the evolution of spin systems.

In Chapter 3, we turned to random walks and potential theory. We studied recurrence and transience, Green's functions, and harmonic functions, culminating in the Liouville property. These results provided the analytic tools necessary to understand the dual behavior of the voter model in different dimensions.

Chapter 4 introduced the voter model itself. We defined the model formally, discussed its basic properties, and established its duality with coalescing random walks. Using this duality, we proved the central result: the model clusters in dimensions $d \leq 2$ and coexists in dimensions $d \geq 3$. We also examined absorbing states, invariant measures, and convergence to equilibrium, and briefly discussed extensions such as the biased and threshold voter models.

Finally, Chapter 5 connected invariant measures to harmonic functions and martingale properties. We showed how the classification of extremal invariant measures reduces to the classification of bounded harmonic functions, and how the Liouville property explains the absence of nontrivial invariant measures in low dimensions.

Overall, this thesis demonstrates how probabilistic and analytic techniques can be combined to yield a comprehensive understanding of the voter model. The interplay between duality, dimension, and long-term behavior highlights the richness of interacting particle systems and their connections to classical probability theory.

7

Discussion and Further Research

While this thesis provides a comprehensive analysis of the linear voter model, several aspects remain open for further exploration and refinement.

Chapter 2 introduced the probabilistic framework underlying interacting particle systems. Although the general theory was presented clearly, some definitions could benefit from additional motivation or illustrative examples.

In Chapter 3, we developed the analytic tools necessary to study the voter model, including recurrence, transience, and harmonic functions. While the connection between recurrence and clustering was established, the derivation of Green's functions and their asymptotic behavior could be expanded on.

Chapter 4 presented the core results on the voter model, including its duality with coalescing random walks. There could be more elaboration on the connection to coalescing random walks. A more explicit construction of the dual process would clarify this relationship.

The discussion of invariant measures in Chapter 5 highlighted the connection between harmonic functions and long-term behavior. While the classification of extremal invariant measures was outlined, the construction of the family $\{\nu_\rho\}_{\rho\in[0,1]}$ in high dimensions could be made more concrete.

Future research directions include:

- Extending the analysis to nonlinear variants of the voter model, such as the biased or threshold models, and studying their phase transitions.
- Investigating the scaling limits of the voter model and its connection to measure-valued diffusions or superprocesses.
- Exploring numerical simulations to visualize clustering, coexistence, and consensus times in finite systems.
- Developing a more detailed treatment of duality, including explicit constructions and applications to other interacting particle systems.

These directions offer promising avenues for deepening our understanding of the voter model and its role in the broader theory.

Bibliography

[1] T. M. Liggett. *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*. Grundlehren der mathematischen Wissenschaften, Vol. 324. Springer, Berlin; New York, 1999. ISBN: 978-3-540-65995-2. DOI: 10.1007/978-3-662-03990-8.

[2] T. M. Liggett. *Interacting Particle Systems*. Classics in Mathematics. Springer, Berlin; New York, 2005. Originally published in 1985. ISBN: 978-3-540-26962-5.

[3] R. Durrett. *Essentials of Stochastic Processes*. 2nd edition. Springer Texts in Statistics. Springer, New York, 2012. ISBN: 978-1-4614-3615-7. DOI: 10.1007/978-1-4614-3615-7.

[4] T. E. Harris. Nearest-neighbor Markov interaction processes on multidimensional lattices. *Advances in Mathematics*, 9(1):66–89, 1972. ISSN: 0001-8708. DOI: 10.1016/0001-8708(72)90030-8.

[5] R. Durrett. *Ten Lectures on Particle Systems*. Lecture Notes from the 1993 Saint-Flour Probability Summer School. Springer, Berlin, 1995.