Sampling from Gibbs distributions

by

Eric Joseph Vigoda

B.S. (Johns Hopkins University) 1994
M.S.E. (Johns Hopkins University) 1994

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Computer Science

in the

GRADUATE DIVISION
of the
UNIVERSITY of CALIFORNIA at BERKELEY

Committee in charge:

Professor Alistair Sinclair, Chair
Professor Yuval Peres
Professor Christos Papadimitriou

Fall 1999

The dissertation of Eric Joseph Vigoda is approved:

Chair Date

Date

Date

University of California at Berkeley

Fall 1999

Abstract

Sampling from Gibbs distributions

by

Eric Joseph Vigoda
Doctor of Philosophy in Computer Science

University of California at Berkeley

Professor Alistair Sinclair, Chair

This thesis considers computational questions concerning statistical mechanical models of idealized physical systems. The equilibrium state of the physical system is described by a probability distribution over the allowed configurations, known as a *Gibbs distribution*. By sampling at random from the Gibbs distribution one can study essentially all the thermodynamic properties of the system.

The standard approach to sampling from the Gibbs distribution is the "Markov Chain Monte Carlo" method. At the heart of this method is a Markov chain whose stationary distribution is the Gibbs distribution and which quickly converges to stationarity. For a wide class of physical systems, there is a class of very simple Markov chains known as the "Glauber dynamics," whose moves correspond to local perturbations of the current configuration. These chains are of interest because they are simple, natural and widely used in practice.

We study the properties of the Glauber dynamics in two models of particular combinatorial interest: the *Potts model*, whose configurations are the set of proper colorings of a graph; and the *hard core model*, whose configurations are the set of independent sets of a graph. For a range of parameter values we prove that the Glauber dynamics in these models quickly converges to the Gibbs distribution, while in another range of values the time to reach the stationary distribution grows exponentially with the volume of the system. Our results also address the conjectured connection between the convergence rate of the Glauber dynamics and phase transitions in the macroscopic properties of the system.

Professor Alistair Sinclair
Dissertation Committee Chair

Contents

Acknowledgements

I am particularly grateful to my advisor Alistair Sinclair who was a constant source of insight and encouragement throughout the thesis writing process. I am also very appreciative of the support and encouragement of Mike Luby who was my advisor during the beginning of my research career. I had quite a lot of fun working and hanging out with Mike during this time.

I would like to make a special thanks to a few of my friends from the department for many discussions about research and even more games of backgammon and cards, particularly: Sanjoy Dasgupta, John Byers, Edouard Servan-Schreiber, Michael Mitzenmacher, Micah Adler, Susan Lee, Ketan Mayer-Patel, Anupam Gupta, and Sean Hallgren. Many thanks to Christian Borgs, Jennifer Chayes, and Jeong-Han Kim for my many enjoyable visits to the Microsoft Research theory group. Finally, for their support and encouragement, my warmest thanks to my parents, brother, and sisters.

Chapter 1

Synopsis

1.1 Introduction

This thesis considers models from statistical mechanics used to study the macroscopic behavior of idealized physical systems in equilibrium. Typically such models are defined on the d-dimensional integer lattice $\mathbb{Z}^d$ where random variables are associated with lattice points and interact based on their proximity in the lattice. Once the forces governing the microscopic interactions are defined, it is possible in principle to study the thermodynamic properties of the system, though the necessary computations are far from trivial. This thesis is concerned with the complexity of these computations, and its relationship to the thermodynamic properties of the system.

A key property of (and justification for) these idealized models is that they exhibit *phase transitions*, similar to those that occur in nature when a small change in some parameter controlling the microscopic interactions of the system, such as temperature or pressure, causes a drastic change (in fact a discontinuity) in the macroscopic properties of the system. As an example, consider the effect on the density of increasing the temperature in a sample of water. The boiling point is a *critical* temperature at which the density suddenly decreases in a dramatic fashion; this is called a phase transition.

The laws of thermodynamics stipulate that a physical system in equilibrium is described by the *Gibbs distribution*, a probability distribution over possible configurations of the system. This thesis studies the problem of sampling at random from Gibbs distributions, a problem to which almost all thermodynamic computations can be reduced. For a discussion on the justification of the Gibbs distribution see Israel [Isr79] and for a historical perspective see Grandy [Gra87].

We shall focus on a class of models known as *spin systems*, which capture many of the most important examples in statistical mechanics. For a graph $G = (V, E)$, the space of configurations is $\Omega = S^V$, where each vertex $i \in V$ has an associated spin σ_i chosen from the discrete spin space S. Each configuration $\sigma \in \Omega$ has an energy defined by the *Hamiltonian* $H(\sigma)$. For a system in equilibrium at temperature T, the probability of being in configuration σ is proportional to

$$w(\sigma) = \exp(-\beta H(\sigma)),$$

where $\beta = \frac{1}{kT}$ and k is the Boltzmann constant. The *Gibbs* distribution (sometimes referred to as the Boltzmann distribution) is then

$$\mu(\sigma) = w(\sigma)/Z,$$

where the normalizing factor, traditionally referred to as the *partition function*, is the quantity

$$Z = \sum_{\sigma \in \Omega} w(\sigma).$$

If we could compute the partition function, we could isolate phase transitions by looking for discontinuities in its (first or higher order) derivatives with respect to different parameters of the system, and also calculate thermodynamic parameters of the macroscopic system such as free energy, entropy, and heat capacity (see, e.g., [Wal89, p. 55]). Unfortunately, the partition function is a sum over $|S|^n$ configurations, a number which grows exponentially with $n = |V|$, the volume of the system. Thus, even for a system of moderate size, computing Z by enumerating all configurations is computationally infeasible. However, it is usually possible to estimate Z within arbitrary precision by generating a sufficient number of samples from the Gibbs distribution. This is an instance of a general reduction from enumeration problems to random sampling, first formalized by Jerrum, Valiant, and Vazirani [JVV86]. Thus, by devising an efficient scheme for sampling from the Gibbs distribution, we can compute the partition function and other thermodynamic quantities of interest, and hence isolate phase transitions.

The most widely studied and perhaps simplest statistical mechanical model is the *Ising model* of a ferromagnetic solid. Each vertex has either an “up” or “down” spin, i.e., $S = \{\pm 1\}$, corresponding to the magnetic moment of an atom situated at the vertex. The Hamiltonian of a configuration $\sigma = \{\sigma(i)\}_{i\in V}$ is

$$H(\sigma) = -\sum_{\{i,j\}\in E} \sigma(i)\sigma(j). \tag{1.1}$$

Thus H assigns lower energy to configurations in which many neighboring spins are aligned.

While the definition of the Ising model holds for arbitrary graphs, most work in the statistical physics community focuses on the lattice $\mathbb{Z}^d$. For the purpose of introducing the notion of phase transitions we specialize to the lattice $\mathbb{Z}^d$ and its associated finite graph Q_L. Let $Q_L = (V, E)$ denote the graph corresponding to a finite d-dimensional cube with side length $2L + 1$ in $\mathbb{Z}^d$, i.e., $V = \{-L, \ldots, 0, \ldots, L\}^d$ and the set of edges E connects vertices that differ by 1 in exactly one coordinate. The Ising model on Q_L has two states with all spins aligned, and these have minimum energy. As a consequence, when β is very large (i.e., at very high temperatures), the Gibbs distribution is concentrated almost entirely on configurations in which most of the spins are aligned. Conversely, for small β (i.e., low temperatures), the spins are nearly independent and a typical configuration has roughly half the spins in each direction; this set of configurations has maximum entropy. In particular, let $N(\sigma) = |\sum_{i\in V}\sigma(i)|/|V|$ where $|V| = (2L+1)^d$ is the volume of Q_L. For the Ising model, it turns out that there exists a critical point β_c such that, when $\beta < \beta_c$, $E[N(\sigma)] \to 0$ as $|V| \to \infty$, while when $\beta > \beta_c$, $E[N(\sigma)] > \epsilon$ as $|V| \to \infty$ where ϵ is a positive constant independent of $|V|$. We now give a rigorous definition of the critical point β_c and the two regimes $\beta > \beta_c$ and $\beta < \beta_c$ which are referred to as *ordered* and *disordered phases*.

Let the boundary ∂Q_L of the cube Q_L denote those vertices with at least one component equal to $\pm L$. Fix the spin of every vertex on the boundary to a common value $k \in \{\pm 1\}$. Define the distribution μ_k to be the Gibbs distribution conditional on ∂Q_L having spin k. We are interested in whether the marginal distribution of the spin $\sigma(O)$ at the origin O is independent of the boundary configuration in the limit $L \to \infty$. Specifically, we say the model is in the *ordered phase* if

$$\mu_{+1}(\sigma(O) = +1) \neq \mu_{-1}(\sigma(O) = +1) \text{ as } L \to \infty;$$

otherwise, the model is said to be in the *disordered phase*. It turns out that the Ising model is monotone in β which implies that there exists a *critical point* β_c such that for $\beta < \beta_c$

the system is in the disordered phase, while for $\beta > \beta_c$ it is in the ordered phase. Note that, a priori, there may not exist a unique critical point for the transition between the ordered and disordered phases; in fact, there may not even exist one such critical point.

We now describe, for a general graph $G = (V, E)$, the two models which this thesis focuses on. The *q-state Potts model* [Pot52, Wel93, Wu82] is a natural generalization of the Ising model discussed above. The spin space, which we think of as a set of *colors*, is $S = \{1, \ldots, q\}$. For a configuration σ, the Hamiltonian for the ferromagnetic case is

$$H(\sigma) = \sum_{\{i,j\}\in E} (1 - \delta_{\sigma_i = \sigma_j});$$

where δ is the Kronecker delta function which takes value 1 when the specified condition is satisfied and 0 otherwise. Note that when $q = 2$ this Hamiltonian is equal to (1.1) up to a multiplicative factor and an additive constant. For the anti-ferromagnetic case, the Hamiltonian is

$$H(\sigma) = \sum_{\{i,j\}\in E} \delta_{\sigma_i = \sigma_j}.$$

In contrast to the ferromagnetic case, the tendency here is for adjacent spins to be different (but no distinction is made between different pairs of values).

The Gibbs distributions for the Potts model is

$$\mu_{potts}(\sigma) = w(\sigma)/Z \tag{1.2}$$

where $w(\sigma) = \exp(-\beta H(\sigma))$. In the infinite-temperature limit ($\beta \downarrow 0$) the Gibbs distribution for the Potts model (in both the ferromagnetic and anti-ferromagnetic cases) corresponds to picking a random color (spin) independently for each vertex. On the other hand, in the zero-temperature limit ($\beta \uparrow \infty$), the ferromagnetic Potts model has only q different colorings with positive weight – namely, those that assign the same color to every vertex – while the anti-ferromagnetic model is uniformly distributed over *proper colorings* (those in which all pairs of neighboring vertices have different colors). Much of this thesis focuses on this case of proper colorings since it is perhaps the richest combinatorially.

The *hard core (lattice gas) model* is a simple spatial representation of a gas as a graph $G = (V, E)$ [Dob74]. The vertices are the possible sites for particles, which are assumed to have non-negligible size. To prevent particles from overlapping, adjacent sites cannot simultaneously be occupied. Thus, valid configurations σ are *independent sets* in G, i.e., $\sigma \subseteq V$ such that no two vertices in σ are connected by an edge of E. The Hamiltonian is derived from suitable limits of the Ising model (see Dobrushin, Kolafa and Shlosman [DKS85]), leading to the following simple form for the Gibbs distribution:

$$\mu_{IS}(\sigma) = \lambda^{|\sigma|}/Z_{IS}. \tag{1.3}$$

The parameter λ is known as the *fugacity* of the system. The ordered/disordered phases of the Potts and hard core models on $\mathbb{Z}^d$ are precisely defined in Chapter 2.

1.2 Markov chain Monte Carlo

The Markov Chain Monte Carlo (MCMC) method is a simple and frequently used approach for sampling from the Gibbs distribution of a statistical mechanical system. The idea is as follows. We design a Markov chain whose state space is Ω and whose stationary distribution is the desired Gibbs distribution. Starting at an arbitrary state, we simulate the Markov chain on Ω until it is sufficiently close to its stationary distribution. We then output the final state which is a sample from (close to) the desired distribution. The required length of the simulation, in order to get close to the stationary distribution, is traditionally referred to as the *mixing time* τ; it is formally defined in chapter 2. In general, the mixing time may be as large as (or even larger than) the number of states of the chain, which as we have observed is exponentially large in the volume n. We consider the MCMC approach efficient if the mixing time is bounded by a fixed *polynomial* in n, and call such a Markov chain *rapidly mixing*. On the negative side, when proving that a certain Markov chain is not efficient, we typically show the mixing time is exponential in n^ϵ (for some fixed $\epsilon > 0$), and refer to this as *torpidly mixing*. Note that the mixing properties depend crucially on the structure of the Markov chain, not just on the model itself.

The simplest class of Markov chains for sampling from the Gibbs distribution of a spin system are usually referred to as *Glauber dynamics* [Gla63, Mar97]. This is the class of single spin-flip dynamics whose transitions consist of choosing a vertex $i \in V$ uniformly at random and changing the configuration only at i. Throughout the thesis, we will use the term "Glauber dynamics" to refer to the following specific heat-bath dynamics, but techniques for bounding the mixing time of this dynamics typically apply to other popular chains in the class of Glauber dynamics, such as the Metropolis algorithm [MRR+53].

From a configuration σ, the transitions $\sigma \mapsto \sigma'$ of the (heat-bath) Glauber dynamics are as follows:

G1 Choose a vertex i uniformly at random.

G2 Let $\sigma'(j) = \sigma(j)$ for all $j \neq i$.

G3 Choose $\sigma'(i)$ from the Gibbs distribution conditional on the fixed assignment for all vertices other than j, i.e., from $\mu(\,\cdot\,|\sigma'(j)$ for all $j \neq i)$.

As an example of Glauber dynamics, consider the zero-temperature limit of the anti-ferromagnetic Potts model, whose configurations are proper colorings. The final step [G3] corresponds to choosing a color uniformly from the set of colors which do not appear in the neighborhood of i. Meanwhile, for the hard core model the final step [G3] is the following. If no neighbors of i are in the set σ', then the vertex i is included in σ' with probability $\frac{\lambda}{1+\lambda}$, otherwise i remains out of σ'.

1.3 Results

A major motivation for the work in this thesis is a "folklore" belief that there is an intimate connection between the Glauber dynamics being rapidly/torpidly mixing and the model lying in the disordered/ordered phase. To understand the intuition for such a connection, consider the ferromagnetic Ising model. Recall that in the ordered phase we expect a typical sample from the Gibbs distribution to have most of the spins aligned. Furthermore, in this phase it is very unlikely that a configuration will have about half of the spins pointing in each direction; but the likelihood of such configurations directly affects the mixing time of the Glauber dynamics. To see this, suppose we start the dynamics in the configuration with all spins pointing up and then simulate the Markov chain. In order to get close to the stationary distribution, the chain must have a reasonable chance of reaching the set of configurations with most spins pointing down, which also has significant weight in the Gibbs distribution. En route it must pass through some configuration with exactly half the spins pointing up, since only one spin is being changed at each step. The probability of reaching a configuration with exactly half the spins pointing up is intimately related to the probability of the set of such configurations in the Gibbs distribution. We formalize this intuition for the ferromagnetic Potts model in chapter 5.1.

To be more precise, we conjecture that being in the disordered/ordered phase for the Potts model on $\mathbb{Z}^d$ is *equivalent to* rapid/torpid mixing of the Glauber dynamics on the torus $T_{L,d}$ – this is the d-dimensional cube with vertex set $\{0, \ldots, L-1\}^d$ and edges connecting vertices that differ by 1 modulo L in exactly one coordinate. Though this conjecture is far from resolved, we are able to prove torpid mixing in the entire ordered phase for sufficiently large q. For general spin systems, there are some known connections between certain properties that are stronger than rapid mixing on the one hand and than being in the disordered phase on the other. In particular, an $O(n \log n)$ mixing time of the Glauber dynamics is closely related to a stronger condition than being in the disordered phase; informally, this condition states that not only does the influence of the boundary on the origin die out in the limit $L \to \infty$, but it decays *exponentially* fast. This connection is detailed in chapter 2, after introducing the necessary definitions.

In addition to addressing the above conjecture, we also prove rapid mixing (over a significantly wider range of parameter values than previously known) of the Glauber dynamics for both the Potts and hard core models. As a corollary, we obtain improved bounds on the critical points of both models.

The first set of results addresses the zero-temperature anti-ferromagnetic q-state Potts model (proper q-colorings). Roman Kotecky (cited in [Geo88, pages 148-149,457]) showed that on $\mathbb{Z}^d$, the system is in the disordered phase when $q > 4d$. His result also applies to other lattices when the number of colors is greater than twice the degree of the lattice (note that $\mathbb{Z}^d$ has degree $2d$). For the Glauber dynamics, on the other hand, Mark Jerrum [Jer95] proved rapid mixing when $q \geq 2\Delta$, where Δ is the maximum degree of the graph. In both settings, this 2Δ barrier was broken only in specific instances by computer assisted proofs which analyzed a huge number of cases. These works are discussed in more detail in chapter 3.

In this thesis, we give a simple direct proof that breaks the 2Δ barrier for arbitrary graphs. We consider a Markov chain which we call the “flip dynamics” (which is not in fact strictly a Glauber dynamics) and is formally defined in chapter 3. The transitions of our chain consist of ‘flipping’ two-colored clusters. In particular, from a coloring σ, choose a vertex v and color c uniformly at random. Then consider the maximal cluster of vertices which contain v and are colored with c or $\sigma(v)$. With an appropriate probability, ‘flip’ this cluster by interchanging colors c and $\sigma(v)$ on it. Our main results are the following.

Theorem 1 *When $q > \frac{11}{6}\Delta$, where Δ is the maximum degree of the graph,*

(a) *the flip dynamics is rapidly mixing with mixing time $O(nq\log n)$; and*

(b) *the Glauber dynamics is rapidly mixing with mixing time $O(n^2 q\log n\log q)$.*

This is the first proof to break the 2Δ barrier that is not computer assisted, and also the first for arbitrary graphs.

Corollary 2 *The zero-temperature anti-ferromagnetic q-state Potts model on $\mathbb{Z}^d$ is in the disordered phase when $q > \frac{11}{3}d$.*

The corollary follows easily from the rapid mixing result by the general connection between an $O(n\log n)$ mixing time and exponentially fast decay of the influence of the boundary mentioned above. An analogous result also applies to other lattices of interest, such as the hexagonal and triangular lattices.

Our second set of rapid mixing results concerns the hard core model. We can prove the following optimal bounds on the mixing time of the Glauber dynamics.

Theorem 3 *The Glauber dynamics on the hard core model is rapidly mixing with mixing time $O(n\log n\Delta)$ when $\lambda < \frac{2}{\Delta-2}$, where Δ is the maximum degree of the graph. For $\lambda = \frac{2}{\Delta-2}$ the mixing time is $O(n^3\Delta^2)$.*

Corollary 4 *The hard core model on $\mathbb{Z}^d$ is in the disordered phase when $\lambda < \frac{2}{2d-2}$.*

The proofs of Theorems 1 and 3 rely on coupling arguments, which are based on the simultaneous evolution of two copies of the Markov chain. A *coupling* is a joint distribution on the transitions of the copies with the added condition that if they reach the same state then they follow the same transitions. The goal is to minimize the *coupling time*, which is the expected time till they reach the same state (maximized over all pairs of initial states). The coupling time implies a bound on the (variation) distance between their respective distributions and consequently, a bound on the mixing time. Typically, we measure the progress of our coupling with respect to a simple metric on the product state space. The proof of Theorem 3 is novel in that it relies on a non-standard metric which enables us to directly analyze the Glauber dynamics and get optimal bounds on its mixing time.

The results of Theorem 3 for the hard core model are close to optimal since Dyer, Frieze, and Jerrum [DFJ99] have recently proved that the Glauber dynamics is torpidly mixing on a random bipartite graph with maximum degree Δ when $\lambda = 1, \Delta \geq 6$, and in general for $\lambda > \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^{\Delta}} \sim \frac{e}{\Delta-2}$.

The final set of results relate directly to our conjecture about the connection between torpid mixing of the Glauber dynamics on the torus $T_{L,d}$ and being in the ordered phase for $\mathbb{Z}^d$. In particular, we prove for the ferromagnetic Potts model with sufficiently large q that throughout the ordered phase the Glauber dynamics is torpidly mixing on $T_{L,d}$. We also give a related and perhaps surprising result about the widely-used Swendsen-Wang algorithm [SW87].

In addition to the conjectured torpid mixing in the ordered phase, the Glauber dynamics is also believed to exhibit a "critical slowing down". In particular, in the disordered phase it is expected that the mixing time increases dramatically (though still only polynomially) as β approaches the critical point. The Swendsen Wang algorithm is specifically designed to overcome this slowing down, and in fact is believed to have small mixing time in the entire disordered and ordered phases, by recoloring many two-colored clusters in a single transition. The basis of the algorithm is an equivalent representation of the Potts model in the Fortuin-Kasteleyn random cluster model [FK72]. We describe the Swendsen-Wang algorithm in detail in chapter 5.1, and prove that it is in fact torpidly mixing at the critical point β_c for sufficiently large q.

Theorem 5 *For the ferromagnetic Potts model on $T_{L,d}$ with $d \geq 2$, and sufficiently large q, there exist positive constants k_1, k_2 (which depend on d) such that for the critical point $\beta_c = \beta_c(q,d)$:*

(a) *The mixing time τ_{GD} of the Glauber dynamics for $\beta \geq \beta_c$ satisfies*

$$\tau_{GD} \geq e^{k_1 L/(\log L)^2}.$$

(b) *The mixing time τ_{SW} of the Swendsen-Wang algorithm at $\beta = \beta_c$ satisfies*

$$\tau_{SW} \geq e^{k_2 L/(\log L)^2}.$$

Using analogous techniques, we also prove that for the hard core model, the Glauber dynamics is torpidly mixing for large enough values of λ.

Theorem 6 *For the hard core model on $T_{L,d}$ with $d \geq 2$, and λ sufficiently large, there exists a positive constant k_3 (which depends on d) such that the mixing time τ_{GD} of the Glauber dynamics satisfies*

$$\tau_{GD} \geq e^{k_3 L^{d-1}/(\log L)^2}.$$

In addition to the results described above, we believe that the techniques utilized in this thesis may be useful for obtaining rigorous bounds on the mixing time of other MCMC sampling schemes, isolating phase transitions in other models, and ultimately resolving the conjectured connection between the mixing time of the Glauber dynamics and phase transitions.

The remainder of the thesis is organized as follows. The next chapter provides a more comprehensive treatment of the background material sketched above. In chapter 3, we present our rapid mixing results on the zero-temperature limit of the anti-ferromagnetic Potts model (proper colorings). The rapid mixing results for the hard core model are proved in chapter 4. Finally, in chapter 5.1, we prove the torpid mixing results for the Glauber dynamics and Swendsen Wang algorithm.

Chapter 2

Background

2.1 Markov chains

Consider a stochastic process $(\sigma_t)_{t=0}^{\infty}$ on a finite state space Ω. Let P denote a non-negative matrix of size $|\Omega| \times |\Omega|$ which satisfies the constraint

$$\sum_{j \in \Omega} P_{ij} = 1 \text{ for all } i \in \Omega.$$

The process is called a *Markov chain* if for all times t and $i, j \in \Omega$

$$\Pr[\sigma_{t+1} = j | \sigma_t = i, \sigma_{t-1}, \ldots, \sigma_0] = \Pr[\sigma_1 = j | \sigma_0 = i] = P_{ij}.$$

As a consequence of the time-independence of P, the *s-step transition matrix* is P^s, i.e., $P^s_{ij} = \Pr[\sigma_{t+s} = j | \sigma_t = i]$. A distribution π is called a *stationary distribution* if it satisfies $\pi P = \pi$. A necessary and sufficient condition for a chain to have a unique stationary distribution is that the chain is

1. *irreducible:* for all $i, j \in \Omega$ there exists a time t such that $P^t_{ij} > 0$; and

2. *aperiodic:* for all $i \in \Omega$, $\gcd\{t : P^t_{ii} > 0\} = 1$.

We call a Markov chain having both of these properties ‘ergodic’.

Theorem 7 (see, e.g., [Fel68]) *An ergodic Markov chain with transition probability matrix P on a finite state space Ω has a unique limiting stationary distribution π, i.e.,*

$$\lim_{t \to \infty} P^t_{ij} = \pi_j \textit{ for all } i, j \in \Omega.$$

While the theorem guarantees the existence of the stationary distribution π, we will need to actually determine π. The chains used in MCMC simulations generally have a special property that makes this easy. Specifically, for an ergodic Markov chain, if a distribution π satisfies the detailed balance equations

$$\pi_i P_{ij} = \pi_j P_{ji} \text{ for all } i, j \in \Omega,$$

then π is the (unique) stationary distribution. Such a chain is called *(time-)reversible*. Notice in particular that if P is symmetric ($P_{ij} = P_{ji}$ for all i, j) then it is reversible with stationary distribution uniform over Ω.

2.2 Mixing time

Our goal is to bound the time until a Markov chain is sufficiently close to its stationary distribution. As a measure of the distance between two distributions μ, ν defined on a finite set Ω, we will use the *(total) variation distance*,

$$d_{TV}(\mu, \nu) = \frac{1}{2} \sum_{j \in \Omega} |\mu(j) - \nu(j)| = \max_{A \subset \Omega} |\mu(A) - \nu(A)|.$$

In our setting, the variation distance after t steps for a chain starting from state i is

$$d_i(t) = d_{TV}(P^t_{i\cdot}, \pi) = \frac{1}{2} \sum_{j \in \Omega} |P^t_{ij} - \pi_j|.$$

We are interested in bounding the distance regardless of the starting state, so we define

$$d(t) = \max_{i \in \Omega} d_i(t),$$

and

$$\tau(\epsilon) = \min\{t : d(t) \leq \epsilon\}.$$

It is in fact sufficient to consider the following simpler quantity, which we define as the *mixing time*:

$$\tau = \tau(1/2e).$$

The reason is that a bound on τ implies the following bound on $\tau(\epsilon)$ for all ϵ [Ald83]:

$$\tau(\epsilon) \leq (1 - \log \epsilon)\tau.$$

Thus we are justified in thinking of τ as the time to required to get "close to" π.

The Markov chains we consider are defined on a set Ω of combinatorial structures over some graph $G = (V, E)$, such as the set of independent sets or proper colorings of G. While the size of Ω is typically exponential in the number of vertices, $n = |V|$, in order to obtain efficient sampling algorithms we want to design Markov chains with a mixing time that is significantly smaller than $|\Omega|$. Therefore, we call a Markov chain *rapidly mixing* if τ is bounded by a polynomial in n; conversely, a *torpidly mixing* chain has mixing time at least $\exp(cn^{\epsilon})$ for constants $c, \epsilon > 0$. Note that the definition of the mixing time τ is the maximum over all initial states. Thus it is conceivable that a Markov chain can be torpidly mixing (according to our definition), but there exists a particular initial state from which the time to get close to the stationary distribution is still rapid, i.e., polynomial.

We will use two techniques for analyzing the mixing time of a Markov chain:

1. coupling to show rapid mixing; and

2. conductance to show torpid mixing.

The next two sections give an overview of these techniques.

2.3 Coupling

A *coupling* of a Markov chain on state space Ω is a stochastic process (σ_t, η_t) on $\Omega \times \Omega$ such that:

1. separately, σ_t and η_t are copies of the original Markov chain; and

2. if $\sigma_t = \eta_t$, then $\sigma_{t+1} = \eta_{t+1}$.

A trivial example of a coupling is if the copies are independent. In general, σ_t and η_t need not be independent.

In order to bound the mixing time, we will try to define a coupling so as to minimize the time until both copies of the Markov chain reach the same state,

$$T_{ij} = \min\{t : \sigma_t = \eta_t | \sigma_0 = i, \eta_0 = j\}.$$

For simplicity, we will instead bound the (expected) *coupling time*, defined as follows:

$$T = \max_{i,j \in \Omega} E[T_{ij}].$$

The coupling time implies a bound on the mixing time:

Theorem 8 ([Ald83]) *For an ergodic Markov chain,*

$$\tau \leq 2eT.$$

The proof is straightforward and relies on the following easily proven fact which is known as the coupling lemma. For $\sigma_0 = i, \eta_0 = j$, a coupling (σ_t, η_t) can be used to bound the variation distance,

$$d_{TV}(\sigma_t, \eta_t) \leq \Pr[\sigma_t \neq \eta_t] = \Pr[T_{i,j} > t].$$

To prove the theorem, it is then sufficient to show that $\Pr[T_{ij} > t] \leq T/t$, which follows easily from Markov's inequality.

We are now left with the different (though perhaps equally difficult) problem of bounding the coupling time T. In order to measure the distance between the two copies of the chain, we introduce a metric Φ on the product state space $\Omega \times \Omega$ so that $\Phi = \Phi(\sigma_t, \eta_t) = 0 \iff \sigma_t = \eta_t$. Thus, the expected time until $\Phi = 0$ is equivalent to the coupling time. As a concrete example, for spin systems a natural metric is the Hamming distance, i.e., $\Phi(\sigma, \eta) = |\{v | \sigma(v) \neq \eta(v)\}|$. The idea is to define a coupling that makes Φ decrease (in expectation) after every transition. The following toy example illustrates the approach.

Consider the infinite-temperature ferromagnetic q-state Potts model on an arbitrary graph $G = (V, E)$. Thus the model has state space $\Omega = \{1, \ldots, q\}^V$ and the Gibbs measure is uniformly distributed over Ω. For this example, a transition of the Glauber dynamics consists of choosing a color c and vertex v uniformly at random and recoloring v to color c.

We define a natural coupling (σ_t, η_t): at each step, both chains choose the same color c and vertex v and move accordingly. It is easy to verify that in fact this is a valid coupling. Notice that after a transition (at time t) that changes vertex v to color c we certainly have $\sigma_{t+1}(v) = \eta_{t+1}(v) = c$ and in fact, $\sigma_{t'}(v) = \eta_{t'}(v)$ for all times $t' > t$. Thus, if $\Phi(\sigma_t, \eta_t) = d$, then

$$\Phi(\sigma_{t+1}, \eta_{t+1}) = \begin{cases} d-1 & \text{with probability } d/n; \\ d & \text{with probability } 1 - d/n. \end{cases}$$

Since the expected number of steps before Φ decreases is n/d, we have that $E[T_{ij}] \leq \sum_{0 \leq i < n} n/(n-i) \sim n \ln n$ as $n \to \infty$. This implies that the mixing time is $O(n \ln n)$.

An important tool for helping to design couplings in complex examples is Bubley and Dyer's *path coupling* [BD97]. Using path coupling, it is sufficient to define and analyze a coupling for a suitable subset of $\Omega \times \Omega$. Both the design and analysis of the coupling automatically extend to the whole of $\Omega \times \Omega$.

We present the theorem and the necessary definitions for the specific case of spin systems, i.e., $\Omega \subseteq \{1, \ldots, k\}^V$. We consider a pair of states $\sigma, \eta \in \Omega$ *neighbors* if they differ only at a single vertex. This is denoted by $\sigma \sim \eta$. We call $\xi = (\xi^0, \ldots, \xi^k)$ a *simple path* if all ξ_i are distinct and $\xi_0 \sim \xi_1 \sim \cdots \sim \xi_k$. Define $\rho(\sigma, \eta)$ as the set of simple paths between σ and η.

Theorem 9 (Bubley and Dyer [BD97]) *Let Φ be an integer-valued metric defined on $\Omega \times \Omega$ which takes values in $\{0, \ldots, D\}$ such that, for all $\sigma, \eta \in \Omega$, there exists a path $\xi \in \rho(\sigma, \eta)$ with*

$$\Phi(\sigma, \eta) = \sum_i \Phi(\xi^i, \xi^{i+1}).$$

Suppose there exists a constant $\beta < 1$ and a coupling (σ_t, η_t) of the Markov chain such that, for all $\sigma_t \sim \eta_t$,

$$E[\Phi(\sigma_{t+1}, \eta_{t+1})] \leq \beta \Phi(\sigma_t, \eta_t). \tag{2.1}$$

Then the mixing time is bounded by

$$\tau \leq \frac{\log(2eD)}{1-\beta}.$$

Moreover, if (2.1) holds with $\beta = 1$ and in addition there exists an $\alpha > 0$ such that, for all t and arbitrary $\sigma_t, \eta_t \in \Omega$,

$$\Pr[\Phi(\sigma_{t+1}, \eta_{t+1}) \neq \Phi(\sigma_t, \eta_t)] \geq \alpha,$$

then the mixing time is bounded by

$$\tau = O\left(\frac{D^2}{\alpha}\right).$$

To see the usefulness of this theorem, reconsider the earlier example of the infinite temperature ferromagnetic q-state Potts model in the setting of the path coupling theorem. Since Φ is defined as the Hamming distance, it is easy to verify that the conditions of the theorem are satisfied and it is sufficient to consider a pair of states $\sigma_t \sim \eta_t$ that differ at just one vertex (and thus $\Phi(\sigma_t, \eta_t) = 1$). We use the same coupling (σ_t, η_t) as before, where each chain chooses the same vertex v and color c for every transition. After one transition it is clear that Φ decreases by 1 with probability $1/n$, and otherwise remains unchanged. This implies that condition (2.1) holds with $\beta = 1 - 1/n$. Since in this example $D = n$, we conclude that $\tau \leq n \log(2en)$.

For completeness, and since it is a relatively recent development, we give a proof of the path coupling theorem in the case when $\beta < 1$.

Proof of Theorem 9: We claim that it is sufficient to exhibit a coupling such that for all $\sigma_t, \eta_t \in \Omega$ (not necessarily neighbors), $E[\Phi(\sigma_{t+1}, \eta_{t+1})] \leq \beta\Phi(\sigma_t, \eta_t)$. For suppose we have such a coupling and consider two arbitrary initial states σ_0, η_0. By successively applying the coupling, the expected distance after t steps is $E[\Phi(\sigma_t, \eta_t)] \leq \Phi(\sigma_0, \eta_0)\beta^t \leq D\beta^t$. Since Φ is non-negative and integer-valued,

$$\Pr[\sigma_t \neq \eta_t] = \Pr[\Phi(\sigma_t, \eta_t) \geq 1] \leq E[\Phi(\sigma_t, \eta_t)] \leq D\beta^t.$$

This probability is at most $\frac{1}{2e}$ after $t = \frac{\log(2eD)}{\log(\beta^{-1})} < \frac{\log(2eD)}{1-\beta}$ steps.

In order to construct a coupling satisfying the above condition, consider the path $\xi_t = (\xi_t^0, \xi_t^1, \ldots, \xi_t^k)$ between σ_t and η_t such that $\Phi(\sigma_t, \eta_t) = \sum_i \Phi(\xi_t^i, \xi_t^{i+1})$. Using this path and the coupling between neighbors (ξ_t^i, ξ_t^{i+1}), we can define a coupling between σ_t and η_t as follows. Given any transition for $\sigma_t = \xi_t^0$, the coupling between neighbors (ξ_t^0, ξ_t^1) defines a move for ξ_t^1. The move for ξ_t^1 along with the coupling (ξ_t^1, ξ_t^2) defines a move for ξ_t^2 and hence a coupling (ξ_t^0, ξ_t^2). We continue along the path in this manner until we have the move for $\xi_t^k = \eta_t$. This defines a coupling for (σ_t, η_t) and in fact for all the states on the path ξ_t. After the coupled move, we have a new path ξ_{t+1} between the states σ_{t+1} and η_{t+1}. Furthermore, we know that $E[\Phi(\xi_{t+1}^i, \xi_{t+1}^{i+1})] \leq \beta\Phi(\xi_t^i, \xi_t^{i+1})$ for each i. By linearity of expectation we have the following:

$$\begin{aligned} E[\Phi(\sigma_{t+1}, \eta_{t+1})] &\leq E[\sum_i \Phi(\xi_{t+1}^i, \xi_{t+1}^{i+1})] = \sum_i E[\Phi(\xi_{t+1}^i, \xi_{t+1}^{i+1})] \\ &\leq \beta \sum_i \Phi(\xi_t^i, \xi_t^{i+1}) = \beta\Phi(\sigma_t, \eta_t). \end{aligned}$$

This completes the proof for the case when $\beta < 1$. ∎

2.4 Conductance

The notion of *conductance* was introduced by Jerrum and Sinclair [JS89] to the study of the mixing time of reversible Markov chains (see Sinclair [Sin93] for a more complete introduction). Conductance characterizes the mixing time, and thus can be used to obtain both lower and upper bounds. This thesis uses conductance solely to bound the mixing time below in order to show that particular chains are torpidly mixing. For examples that use conductance to prove rapid mixing see, e.g., [JS89, DFK91].

For a proper subset $S \subset \Omega$, let $\overline{S} = \Omega \setminus S$. We denote the *flow* out of S by

$$Q(S, \overline{S}) = \sum_{i \in S, j \in \overline{S}} \pi(i) P_{ij}.$$

The *conductance* Φ_S of a subset S is the quantity

$$\Phi_S = \frac{Q(S, \overline{S})}{\pi(S)\pi(\overline{S})}.$$

This quantity is a symmetrized version of the quantity $\frac{Q(S,\overline{S})}{\pi(S)}$, which can be viewed as the conditional probability that the Markov chain in the stationary distribution leaves the set S in one step given that it currently resides in S. The *conductance* of the Markov chain is

$$\Phi = \min_{S \neq \emptyset, \Omega} \Phi_S.$$

We use the following well-known theorem [AM85, Sin93, AF].

Theorem 10

$$\tau \geq \frac{1}{\Phi}.$$

2.5 Phase transitions

For the remainder of the chapter we specialize to the graph $\mathbb{Z}^d$ and its subgraph $Q_L = (V, E)$, which is a d-dimensional cube with side length $2L + 1$. Recall that the boundary ∂Q_L consists of those vertices in Q_L with at least one component equal to $\pm L$. For a Hamiltonian H, recall that the weight of a configuration σ is

$$w(\sigma) = \exp(-\beta H(\sigma)).$$

Also, the Gibbs distribution (at inverse temperature β) is

$$\mu(\sigma) = w(\sigma)/Z.$$

Let Ω denote the set of configurations on $\mathbb{Z}^d$ with positive weight. In particular, for the Potts model, $\Omega = \{\sigma \in \{1, \ldots, q\}^{\mathbb{Z}^d} : w(\sigma) > 0\}$; for the hard core model, the set Ω denotes the set of independent sets of $\mathbb{Z}^d$. We let $\mu_\tau = \mu_\tau^L$ denote the Gibbs distribution conditional on the fixed configuration τ on the boundary ∂Q_L, and let $\mu_\tau(\sigma(O))$ denote the corresponding marginal distribution on the spin at the origin O. We say the system is in the *disordered phase*, if for all $\tau, \tau' \in \Omega$,

$$\lim_{L\to\infty} \mu_\tau(\sigma(O)) = \lim_{L\to\infty} \mu_{\tau'}(\sigma(O)).$$

Otherwise, we say the system is in the *ordered phase*.

It is interesting to note that while the ferromagnetic Potts model, like the Ising model, is monotone in β, neither the anti-ferromagnetic Potts model or the hard core model are known to be monotone. In particular, these models may exhibit multiple critical values for the transition between the disordered and ordered phases.

To illustrate the phenomenon of a phase transition, recall from the Synopsis the definition of the ferromagnetic q-state Potts model. For very small β, the spins of the vertices are nearly independent. In fact, the influence of the boundary quickly dies out and we expect that the system is in the disordered phase for sufficiently small β. While for very large β, few configurations have non-negligible weight; those with almost all vertices having the same color. Thus, forcing all vertices on the boundary to have a particular color will tend to force the origin to have the same color, and we expect that the system is in the ordered phase.

In this thesis, we will work with the above definition of phase transition for spin systems. For completeness, we should briefly mention some other notions of phase transitions that are in wide use. In the following, we refer to the Gibbs distribution on a finite graph as the finite-volume Gibbs measure to differentiate it from an infinite-volume Gibbs measure defined on $\Omega = S^{\mathbb{Z}^d}$.

(i) *Infinite-volume Gibbs measures:* Informally, this is the set of measures μ on Ω whose conditional distribution on all finite volumes $\Lambda \subset \mathbb{Z}^d$ with all boundary configurations is the finite-volume Gibbs measure. The development of such measures began with the work of Dobrushin [Dob68] and Lanford and Ruelle [LR67], and they are often called DLR-states. In this setting, the disordered phase refers to the region with a unique infinite-volume Gibbs measure. See Georgii [Geo88] for a more comprehensive development.

(ii) *Limiting Gibbs measures:* This is a constructive way of defining measures on Ω. Roughly speaking, we fix a $\tau \in \Omega$ and consider the finite-volume Gibbs measure on Q_L with boundary configurations defined by τ. By taking the limit as $L \to \infty$ we get a measure on Ω. The existence of a unique, as opposed to multiple, limiting Gibbs measure corresponds to the system lying in the disordered phase. This approach is intimately related to infinite-volume Gibbs measures (DLR states); see Georgii, Haggstrom, and Maes [GHM99] and Liggett [Lig85, Chapter 4, section 1] for details.

(iii) *Infinite-volume Glauber dynamics:* We can define a continuous time analog of the Glauber dynamics on the infinite-volume system. The notion of the system lying in the disordered phase refers to ergodicity of the infinite-volume Glauber dynamics, which corresponds to the existence of a unique limiting invariant distribution. For details on this approach, see Liggett [Lig85].

All of the above notions (and the notion we use) are equivalent for *attractive systems*, such as the hard core model on bipartite graphs and the Ising model. Roughly speaking, a model is called 'attractive' if there exists a partial order of its state space Ω and a coupling of the Glauber dynamics that preserves the ordering. For a rigorous definition, see Liggett [Lig85, Theorem 2.4, Chapter 2]. For some equivalences and implications between the different notions, see Liggett [Lig85, Cor. 2.8 (Chpt. 2), Cor. 2.3 (Chpt. 3), Thm. 2.16 (Chpt. 4)].

2.6 Phase transitions and rapid mixing

A stronger property than being in the disordered phase is that of *strong spatial mixing*. This is a key idea in establishing a connection with rapid mixing of the Glauber dynamics. We need several preliminary definitions before formally defining the property.

We define the distance between a pair of points $x, y \in \mathbb{Z}^d$ as

$$|x-y| = \max_{i \in \{1,\dots,d\}} |x_i - y_i|.$$

Similarly, the distance between $\Delta \subset \mathbb{Z}^d$ and $y \in \mathbb{Z}^d$ is

$$d(\Delta, y) = \min_{x \in \Delta} |x - y|.$$

For a connected subset $\Lambda \subset \mathbb{Z}^d$ (not-necessarily a box Q_L), the boundary $\partial\Lambda$ is defined analogously as $\partial\Lambda = \{v \in \Lambda : w \notin \Lambda, (v, w) \in E(\mathbb{Z}^d)\}$, where $E(\mathbb{Z}^d)$ denotes the set of edges of the lattice $\mathbb{Z}^d$. Let τ and $\tau^{(y)}$ denote a pair of boundary configurations that differ only at $y \in \partial\Lambda$. We are interested in how this change at y affects the distribution of spins on a connected set $\Delta \subset \Lambda$ as a function of the distance from Δ to Λ. Let $\mu_\tau(\sigma(\Delta))$ denote the marginal distribution of μ_τ on the set Δ. We say that the system has the *strong spatial mixing property* if there exist constants c, m such that, for all $\Lambda \subset \mathbb{Z}^d$, the following condition holds, for all $\Delta \subset \Lambda$ and $y \in \partial\Lambda$,

$$\max_{\tau, \tau^{(y)}} d_{TV}(\mu^\tau(\sigma(\Delta)), \mu^{\tau^{(y)}}(\sigma(\Delta))) \leq c e^{-m d(\Delta, y)}.$$

Frigessi, Martinelli, and Stander [FMS97] prove that the strong spatial mixing property implies that the Glauber dynamics has mixing time $O(n \log n)$ for any connected subset $\Lambda \subset \mathbb{Z}^d$. Their result builds upon work of Stroock and Zegarlinski [SZ92] and holds for all of the models we have introduced. In [SZ92] there are related conditions for the reverse implication, but these involve a different measure of distance from stationarity.

For completeness, we include the following proof, based on ideas in the work of Martinelli and Olivieri [MO94], that a $O(n \log n)$ mixing time of the Glauber dynamics implies that the system is in the disordered phase. A sketch of the argument was explained to the author by J. van den Berg. For simplicity, we specialize to the case of the zero temperature anti-ferromagnetic Potts model whose configurations correspond to proper colorings. It will be clear that the proof extends to other models of interest, such as the hard core, Ising, and Potts models, and also generalizes to other lattices of interest, such as the hexagonal and triangular lattices.

Lemma 11 *For the zero-temperature anti-ferromagnetic Potts model on $\mathbb{Z}^d$ with $q \geq 2d + 1$, an $O(n \log n)$ mixing time of the Glauber dynamics, for all Q_L, implies that the system is in the disordered phase.*

Proof: For $Q_L = (V, E)$, fix a pair of colorings τ, τ' of the boundary ∂Q_L. The idea is to compare μ_τ and $\mu_{\tau'}$ by considering a pair of Markov chains $(\sigma_t), (\eta_t)$ with the Glauber dynamics having the respective fixed boundary colorings τ, τ' and thus stationary distributions μ_τ and $\mu_{\tau'}$. We run these chains until they are close to their stationary distributions; meanwhile, the chains are coupled to maintain (if possible) the same color at the origin. Observe that under the stated condition $q \geq 2d+1$ there exists a pair of colorings σ_0, η_0, with respective boundary colorings τ, τ', such that $\sigma_0(x) = \eta_0(x)$ for all $x \notin \partial Q_L$; these are the initial states of the chains.

Let $\mu_\tau(O), \mu_{\tau'}(O)$ denote the marginal distributions of the spin at the origin O in stationarity, and let

$$p_t = \Pr[\sigma_t(O) = \eta_t(O)].$$

We run the chains for T steps, a time sufficient for both to get within variation distance $1/L$ of the stationary distribution. We can then bound the variation distance between $\mu_\tau(O)$ and $\mu_{\tau'}(O)$ as follows:

$$\begin{aligned}
&d_{TV}\{\mu_\tau(O), \mu_{\tau'}(O)\} \\
&\quad\leq\; d_{TV}\{\mu_\tau(O), \sigma_T(O)\} + p_T + d_{TV}\{\eta_T(O), \mu_{\tau'}(O)\} \\
&\quad\leq\; 1/L + p_T + 1/L,
\end{aligned}$$

where the second line follows from the triangle inequality. Therefore, in order to show that the system is in the disordered phase, it is sufficient to show that $p_T \downarrow 0$ as $L \to \infty$.

Before defining the coupling that we use, recall the definition of the Glauber dynamics for proper colorings. From a coloring σ, the transitions $\sigma \mapsto \sigma'$ are the following:

[G1] Choose a vertex v uniformly at random.

[G2] Let $\sigma'(w) = \sigma(w)$ for all $w \neq v$.

[G3] Let S_σ denote the set of colors that appear in the neighborhood of v. Choose $\sigma'(v)$ uniformly at random from the set $C \setminus S_\sigma$.

The coupling (σ_t, η_t) is the following: in step [G1] both chains choose the same vertex v; while in step [G3], if $S_{\sigma_t} = S_{\eta_t}$ then both chains choose the same color for v, otherwise they choose colors independently.

Let $v \sim w$ denote a pair of adjacent vertices in Q_L. At time t, consider the vertex v chosen in step [G1] and suppose that $\sigma_{t-1}(v) = \eta_{t-1}(v)$ but $\sigma_t(v) \neq \eta_t(v)$. In order for this to occur, there must exist a vertex $w \sim v$ such that $\sigma_{t-1}(w) \neq \eta_{t-1}(w)$. Since initially the only vertices that differ are on the boundary, there must exist a "path of disagreement" from the boundary to v. More formally, let P denote a path $(w_0 \sim w_1 \sim \cdots \sim w_i = O)$ such that $w_0 \in \partial Q_L$ and similarly, let A denote a set of times $(t_1 < \cdots < t_i)$. We say the event $\mathcal{E}(P, A)$ occurs if for all $0 < j \leq i$,

$$\sigma_{t_j-1}(w_j) = \eta_{t_j-1}(w_j) \text{ and } \sigma_{t_j}(w_j) \neq \eta_{t_j}(w_j).$$

Notice that

$$p_t \leq \Pr[\mathcal{E}(P, A) \text{ occurs for some } P, A] \leq \sum_{P,A} \Pr[\mathcal{E}(P, A) \text{ occurs }].$$

In order for a specific event $\mathcal{E}(P, A)$ to occur, at each time t_j vertex w_j must be chosen by the Glauber dynamics in step [G1]. The probability of this occurring is at most $(1/2L)^d$, and thus $\Pr[\mathcal{E}(P, A)] \leq (1/2L)^{id}$. Let $\mathcal{E}(P)$ denote the event that $\mathcal{E}(P, A)$ occurs for some set of times A. Since the number of such sets A is at most $\binom{T}{i}$, we get the following bound:

$$\Pr(\mathcal{E}(P)) \leq \binom{T}{i} \left(\frac{1}{2L}\right)^{id} \leq \left(\frac{Te}{i(2L)^d}\right)^i.$$

Finally, let $\mathcal{E}$ denote the event that $\mathcal{E}(P)$ occurs for some path P. The number of such paths of length i is bounded by the number of walks (with neighbors defined by $\sim$) of length i that start at the origin, which is exactly $(2d-1)^i$. The minimum length of a path from the origin to the boundary is L, and thus

$$\Pr(\mathcal{E}) \leq \sum_{i \geq L} \left(\frac{Te(2d-1)}{i(2L)^d}\right)^i.$$

From our assumption about the mixing time of the Glauber dynamics we have $T = O((2L)^d \log L^d \log L)$ which implies the following bound:

$$\Pr(\mathcal{E}) \leq \sum_{i \geq L} \left(\frac{ed(2d-1)\log^2 L}{i}\right)^i.$$

Since this sum tends to 0 as $L \to \infty$, the proof is complete. ∎

Chapter 3

Sampling Colorings

3.1 Introduction

Recall that a *proper q-coloring* of a graph $G = (V, E)$ is a labeling σ of the vertices with colors from the set $C = \{1, \ldots, q\}$ where neighboring vertices have different colors. Mark Jerrum [Jer95] proved that the Glauber dynamics, defined in the Synopsis, is rapidly mixing when the number of colors is at least twice the maximum degree Δ of the input graph.

This 2Δ barrier also arose in related work in the statistical physics community showing that the zero-temperature anti-ferromagnetic q-state Potts model is in the disordered phase. In particular, Roman Kotecky (cited in [Geo88, pages 148-149,457]) showed that the system is in the disordered phase when the number of colors is greater than twice the degree of the lattice (i.e., $k > 2\Delta = 4d$ for $\mathbb{Z}^d$).

In both settings, this 2Δ barrier was broken in specific instances by computer-assisted proofs which analyzed a huge number of cases. Jesus Salas and Alan Sokal broke the barrier for several two-dimensional lattices [SS97]. They proved that the system is in the disordered phase for seven-colorings of the square lattice, four-colorings of the hexagonal lattice, and six-colorings of the Kagome lattice. Their proof for the square lattice, for instance, requires the computer analysis of 7^8 cases.

Bubley, Dyer, and Greenhill [BDG98] proved rapid mixing of the Glauber dynamics with five colors when Δ is at most three and seven colors on triangle-free four-regular graphs. Their proof relies on the computer solution of several hundred linear programs for the $\Delta \leq 3$ case, and over 40,000 programs for triangle-free 4-regular graphs.

In this chapter, we prove theorem 1 which breaks the 2Δ barrier for arbitrary graphs. This result appeared in [Vig99b].

Theorem 1 *When $q > \frac{11}{6}\Delta$, where Δ is the maximum degree of the graph,*

(a) *the flip dynamics is rapidly mixing with mixing time $O(nq \log n)$; and*

(b) *the Glauber dynamics is rapidly mixing with mixing time $O(n^2 q \log n \log q)$.*

The fact that the zero-temperature anti-ferromagnetic Potts model is in the disordered phase follows directly from the theorem since lemma 11 easily extends to the flip dynamics.

Corollary 2 *The zero-temperature anti-ferromagnetic q-state Potts model on $\mathbb{Z}^d$ is in the disordered phase when $q > \frac{11}{3}d$.*

This beats the previously known bound of $k > 4d$ for general d. Moreover, the result can easily be extended to other lattices that are commonly of interest, such as the hexagonal and Kagome lattice (see [SS97] for illustrations of these lattices).

3.2 Markov Chain

The state space Ω of the Markov chain for the flip dynamics is the set of all proper k-colorings. We need some notation before specifying the transitions of the chain. For a coloring σ, we will refer to a path $v = x_0, x_1, \ldots, x_l = w$ as an alternating path between vertices v and w using colors c and $\sigma(v)$ if, for all i, $(x_i, x_{i+1}) \in E$ and $\sigma(x_i) \in \{c, \sigma(v)\}, \sigma(x_i) \neq \sigma(x_{i+1})$. We let $S_\sigma(v, c)$ denote the following cluster of vertices.

$$S_\sigma(v,c) = \left\{ w \middle| \begin{array}{l} \text{there exists an alternating path between} \\ v \text{ and } w \text{ using colors } c \text{ and } \sigma(v) \end{array} \right\}$$

Also, let $S_\sigma(v, \sigma(v)) = \emptyset$. For every vertex x in the cluster $S_\sigma(v, c)$, notice that $S_\sigma(x, c) = S_\sigma(v, c)$ if $\sigma(x) = \sigma(v)$ and otherwise $S_\sigma(x, \sigma(v)) = S_\sigma(v, c)$.

For a coloring $\sigma \in \Omega$, the transitions $\sigma \mapsto \sigma'$ are defined as:

- Choose a vertex v and color c uniformly at random from the sets V, C respectively.

- Let $\alpha = |S_\sigma(v,c)|$.
  With probability $\frac{p_\alpha}{\alpha}$, 'flip' cluster $S_\sigma(v,c)$ by interchanging colors c and $\sigma(v)$ on the cluster.

The reason for dividing the flip probability by α is that, as observed above, there are exactly α ways to pick the cluster (one for each of its elements). Thus, a cluster is actually flipped with weight p_α. The parameters p_α will be defined later.

Observe that for every vertex v, the flip of cluster $S_\sigma(v,\sigma(v))$ does not change σ. Thus, the Markov chain is clearly aperiodic since $P(\sigma,\sigma) > 0$ for all $\sigma \in \Omega$.

As for irreducibility, it is sufficient to assume flips of clusters of size one have positive weight, i.e., $p_1 > 0$ and $k \geq \Delta + 2$. To go between an arbitrary pair of colorings simply consider an ordering of the vertices and attempt to recolor the vertices in that order. When recoloring a vertex, if some neighbors have the desired color then first recolor those neighbors to an arbitrary color which does not appear in its neighborhood (this requires that $k \geq \Delta + 2$). We are guaranteed that after we give a vertex its desired color, it will not interfere with the recoloring of later vertices in the ordering.

To see that the chain is symmetric and thus the stationary distribution π is uniform, let σ' denote the coloring after a flip of cluster $S_\sigma(v,c)$. Then it should be clear that a flip of cluster $S_{\sigma'}(v,\sigma(v))$ recovers σ.

To complete the description of the chain, we specify the parameters p_α. They are $p_1 = 1, p_2 = \frac{13}{42}$ and for $\alpha > 2$,

$$p_\alpha = \max(0, \frac{13}{42} - \frac{1}{7}[1 + \frac{1}{2} + \cdots + \frac{1}{\alpha - 2}])$$

Specifically, $p_3 = \frac{1}{6}, p_4 = \frac{2}{21}, p_5 = \frac{1}{21}, p_6 = \frac{1}{84}$,and $p_\alpha = 0$ for $\alpha \geq 7$.

The key properties (which will emerge in the analysis) that determined the settings for these parameters are

- $2(i-1)p_i + p_{2i+1} \leq \frac{2}{3}$, and
- $(j-1)(p_j - p_{j+1}) + i(p_i - p_{i+1}) \leq \frac{5}{6}$. This is true because $(j-1)(p_j - p_{j+1}) \leq \frac{1}{7}$, $i(p_i - p_{i+1}) \leq p_1 - p_2 = \frac{29}{42}$.

Other useful properties of these parameters that we utilize are that $ip_i \leq p_1 = 1$, $(i-1)p_i \leq 2p_3 = \frac{1}{3}$, $(i-c)p_i < \frac{1}{4}$ for $c \geq 2$.

3.3 Analysis

Recall the setting of the path coupling theorem. To use the theorem we need to define a metric Φ on $\Omega \times \Omega$ such that there exists a path between an arbitrary pair of states σ, η where the length of the path is exactly $\Phi(\sigma, \eta)$. We let Φ be the Hamming distance which is the number of vertices that are colored differently in the two states. For neighboring states σ, τ, observe that $\Phi(\sigma, \tau) = 1$. Consider a coloring σ where $\sigma(v) = 1, \sigma(w) = 2$ for adjacent vertices v and w. Let η denote the coloring which is identical to σ except $\eta(v) = 2, \eta(w) = 1$. Thus, $\Phi(\sigma, \tau) = 2$ but the shortest path in Ω between these states is of length three.

In order to apply the path coupling theorem, we redefine the state space of the Markov chain. Let the set $\Omega = C^V$, i.e., the set of all (not necessarily proper) k-colorings. Now there exists a path of length $\Phi(\sigma, \eta)$ between an arbitrary pair of states σ and η. The definition of the clusters $S_\sigma(v, c)$ and the transitions of the chain are identical for this enlarged state space.

Observe that if we start the chain at a proper coloring, we only visit proper colorings. Also, if we start at an improper coloring we eventually reach a proper coloring. (To see this simply reconsider the earlier argument for irreducibility.) Therefore, the only states with positive weight in the stationary distribution are proper colorings and the chain is still uniform over these states. Also, a bound on the mixing time of the chain on this enlarged state space will give the same bound on the mixing time of the chain restricted to just proper colorings.

To now use the path coupling theorem to get a bound on the mixing time we must first define a coupling for neighboring states σ, τ. Then we need to show that the expected change in $\Phi = \Phi(\sigma, \tau)$ under this coupling is negative. For the remainder of the analysis let σ and τ denote a pair of neighboring states such that they only differ at vertex v.

Recall that for every cluster $S_\sigma(x, c)$ there is exactly one equivalent cluster indexed by each vertex $y \in S_\sigma(v, c)$. Also, this cluster is flipped with total weight p_α where $\alpha = |S_\sigma(x, c)|$. Thus, when analyzing $E[\Delta\Phi]$ we just have to consider this cluster being flipped with weight p_α as opposed to considering the cluster being flipped with weight p_α/α for each vertex y in the cluster.

Notice that in order for a cluster $S(x, c)$ to be different in the two colorings

σ, τ it must involve v, either $v \in S_\sigma(x,c)$ and/or $v \in S_\tau(x,c)$. Recall that if $v \in S_\sigma(x,c)$ then there is an equivalent way to index the cluster with vertex v. Suppose $v \notin S_\sigma(x,c), v \in S_\tau(x,c)$. We then know that the cluster $S(x,c)$ is composed by colors $\tau(v)$ and c'. Furthermore, there exists a neighbor w of v such that: w has color c', $S_\tau(w,\tau(v)) = S_\tau(x,c) = S_\tau(v,c')$, and $S_\sigma(w,\tau(v)) = S_\sigma(x,c)$. We can conclude that the set D of clusters that might be different in the two chains are

- $S_\sigma(w,\tau(v)), S_\tau(w,\sigma(v))$ for any neighbor w of v,

- $S_\sigma(v,c), S_\tau(v,c)$ for any color c.

The moves that attempt to flip a cluster in D turn out to be the only moves that the analysis needs to consider. In particular, suppose the coupling between moves in σ and τ is simply the identity, i.e., each chain attempts the same move. The flip of a cluster $S \notin D$ does not change Φ since S is the same in both chains before and after the move. Our coupling is in fact the identity for moves that flip clusters not in D. Before stating the coupling for all moves, we partition the set D as follows. Notice that the clusters in D are composed of colors $\sigma(v)$ or $\tau(v)$ and at most one other color c. We partition D into sets D_c based on the other color c as follows, let

$$\begin{aligned} \Gamma_c &= \{w|\sigma(w) = c, w \text{ is a neighbor of } v\}, \\ D_c &= S_\sigma(v,c) \cup S_\tau(v,c) \cup \{\cup_{w\in\Gamma_c}\{S_\sigma(w,\tau(v)) \cup S_\tau(w,\sigma(v))\}\}. \end{aligned}$$

The only sets D_c that might have non-empty intersection are $D_{\sigma(v)}$ and $D_{\tau(v)}$ which both consist of clusters composed of colors $\sigma(v)$ and $\tau(v)$. We ignore this issue for now, and address this special case (*) in the analysis. Note that the sets $D_{\sigma(v)}, D_{\tau(v)}$ are simply a byproduct of redefining the state space to all (not necessarily proper) colorings.

Before defining the coupling, observe that we can think of it as a function f from a move in σ to a move in τ, i.e., we choose a move in σ and f defines the coupled move in τ. From a move in σ that flips a cluster S, the coupling f is

- For $S \notin D$, $f(S) = S$, i.e., moves that flip clusters not in the set D have the identity coupling.

- For $S \in D_c$, $f(S) \in D_c$. Moves in the set D_c for σ are coupled with moves in the same set for τ.

The specific coupling for flips of clusters in the set D_c will be defined later in the analysis. Since flips of clusters in D_c are coupled together for the chains, we can denote the effect of these moves by

$$E[\Delta_{D_c}] = E[\Delta\Phi|\sigma \text{ and } \tau \text{ flip clusters in } D_c].$$

Recall that for clusters $S \notin D$, moves that flip these clusters do not change Φ. We then have that

$$nkE[\Delta\Phi] = \sum_c E[\Delta_{D_c}\Phi]$$

The key component of the analysis is the following lemma. Let $\delta_c = |\Gamma_c|$.

Lemma 12 *For each color $c \in C$,*

(a) *If $\delta_c = 0$, then $E[\Delta_{D_c}\Phi] \leq -1$.*

(b) *If $\delta_c > 0$, then $E[\Delta_{D_c}\Phi] \leq \frac{11}{6}\delta_c - 1$.*

Based on this lemma, we get our main result.

Proof of Theorem 1(a):

Let $\delta = \delta(v)$ denote the degree of vertex v. Observe that the number of colors c with $\delta_c = 0$, i.e., that do not appear in the neighborhood of v, is exactly $k - \delta + \sum_{c':\delta_{c'}>0}(\delta_{c'} - 1)$. Together with the lemma this implies that

$$nkE[\Delta\Phi] \leq -k + \frac{11}{6}\delta.$$

Recall from the path coupling theorem that we need to bound β such that $E[\Phi(\sigma_{t+1}, \tau_{t+1})] \leq \beta\Phi(\sigma_t, \tau_t)$ for all $\sigma_t \sim \tau_t$. Letting $\sigma = \sigma_t, \tau = \tau_t$, we have a bound on $E[\Delta\Phi(\sigma_t, \tau_t)]$. Since $E[\Phi(\sigma_{t+1}, \tau_{t+1})] = \Phi(\sigma_t, \tau_t) + E[\Delta\Phi(\sigma_t, \tau_t)]$ and $\Phi(\sigma_t, \tau_t) = 1$, thus, $\beta \leq 1 - \frac{k - \frac{11}{6}\Delta}{nk}$. Applying the path coupling theorem stated earlier we get the following bound when

$k > \frac{11}{6}\Delta$,

$$\tau \leq \frac{nk}{k - \frac{11}{6}\Delta} \log(2en).$$

■

Proof of Lemma 12:

(a) Observe that $D_c = \{S_\sigma(v,c) \cup S_\tau(v,c)\}$ and furthermore, $S_\sigma(v,c) = S_\tau(v,c) = \{v\}$. Since each chain has only one cluster in D_c, the coupling for the move that flips the cluster in D_c is obviously just the identity. This move might only change v and after the move we know that $\sigma(v) = \tau(v) = c$. Thus, $E[\Delta_{D_c}\Phi] = -1$.

(b) Let $w_1, \ldots, w_{\delta_c}$ denote the set Γ_c of neighbors of v with color c. All of the clusters in the set D_c are composed of colors c and $\sigma(v)$ or c and $\tau(v)$. In fact, the clusters in the set D_c have the following relationship:

For $c \neq \sigma(v)$,

$$S_\sigma(v,c) = \{\cup_i S_\tau(w_i, \sigma(v))\} \cup \{v\}$$

For $c \neq \tau(v)$,

$$S_\tau(v,c) = \{\cup_j S_\sigma(w_j, \tau(v))\} \cup \{v\}$$

Note that in the case when $c = \sigma(v)$, we have $S_\sigma(v,c) = S_\tau(w_i, \sigma(v)) = \emptyset$. Similarly $c = \tau(v)$ implies that $S_\tau(v,c) = S_\sigma(w_j, \tau(v)) = \emptyset$. As mentioned earlier it may also occur that $D_{\sigma(v)} \cap D_{\tau(v)} \neq \emptyset$. We ignore this special case (*) until the end of the proof.

For a color c, all of the clusters in the set D_c might not be distinct. It may occur that $S_\tau(w_i, \sigma(v)) = S_\tau(w_{i'}, \sigma(v))$ or similarly for $S_\sigma(w_j, \tau(v))$. We do the following to insure that we consider the flip of each cluster exactly once. If $S_\tau(w_{i_1}, \sigma(v)) = S_\tau(w_{i_2}, \sigma(v)) = \cdots = S_\tau(w_{i_l}, \sigma(v))$, redefine $S_\tau(w_{i_{l'}}, \sigma(v)) = \emptyset$ for all $1 < l' \leq l$. Similarly for $S_\sigma(w_j, \tau(v))$.

To define our coupling, we need to distinguish the largest of the clusters $S_\tau(w_i, \sigma(v))$ and also of the clusters $S_\sigma(w_j, \tau(v))$. Let $a_i = |S_\tau(w_i, \sigma(v))|$, $A = |S_\sigma(v,c)| \leq 1 + \sum_i a_i$. In fact, $A = 1 + \sum_i a_i$ for $c \notin \{\sigma(v), tau(v)\}$. Similarly, let $b_j = |S_\sigma(w_j, \tau(v))|$, $B = |S_\tau(v,c)| \leq 1 + \sum_j b_j$. Also, let $a_{\max} = \max_i a_i$ and $i_{\max}$ is the corresponding index for $a_{\max}$ (similarly for $b_{\max}$ and $j_{\max}$). For colors $c \neq \sigma(v)$, note that $a_{\max} > 0$, while for

$c \neq \tau(v)$, $b_{\max} > 0$. In the case when $c = \sigma(v)$ we have $A = a_{\max} = 0$ and for $c = \tau(v)$, $B = b_{\max} = 0$.

We can now state the coupling for moves in M_c. The idea is to couple the big flips, $S_\sigma(v,c)$ and $S_\tau(v,c)$, with the largest of the other flips, $S_\tau(w_{i_{\max}}, \sigma(v)), S_\sigma(w_{j_{\max}}, \tau(v))$. Then for each w_i, couple together (as much as possible) the remaining weights of the flips $S_\sigma(w_i, \tau(v)), S_\tau(w_i, \sigma(v))$. More precisely, the coupling is the following:

I with weight p_A, flip $S_\sigma(v,c)$ and $S_\tau(w_{i_{\max}}, \sigma(v))$.

II with weight p_B, flip $S_\tau(v,c)$ and $S_\sigma(w_{j_{\max}}, \tau(v))$.

III For each w_l,

Let q_l (q_l') denote the remaining weight of the flip of $S_\tau(w_l, \sigma(v))$ ($S_\sigma(w_l, \tau(v))$ respectively). Specifically, let

$$q_l = \begin{cases} p_{a_l} - p_A & \text{if } l = i_{\max} \\ p_{a_l} & \text{otherwise} \end{cases}$$

$$q_l' = \begin{cases} p_{b_l} - p_B & \text{if } l = j_{\max} \\ p_{b_l} & \text{otherwise} \end{cases}$$

IIIa with weight $\min(q_l, q_l')$,

$$\text{flip } S_\tau(w_l, \sigma(v)), S_\sigma(w_l, \tau(v))$$

IIIb with weight $q_l - \min(q_l, q_l')$,

$$\text{flip } S_\tau(w_l, \sigma(v))$$

IIIc with weight $q_l' - \min(q_l, q_l')$,

$$\text{flip } S_\sigma(w_l, \tau(v))$$

Let us analyze the effect of each of these coupled moves. After coupled move (I), the colorings are still identical on the cluster which before the move was $S_\tau(w_{i_{\max}}, \sigma(v))$. Thus, their Hamming distance has increased by at most $A - a_{\max} - 1$. Similarly, coupled move (II) increases the Hamming distance by at most $B - b_{\max} - 1$.

For coupled move (IIIa), since both flips effect w_l this move increases the Hamming distance by exactly $a_l + b_l - 1$. Whereas, moves (IIIb) and (IIIc) increase the distance by a_l and b_l respectively. Let us use a function $f(w_l)$ to denote the effect of moves (IIIa), (IIIb), and (IIIc).

$$f(w_l) = a_l q_l + b_l q'_l - \min(q_l, q'_l)$$

We now have that

$$\begin{aligned} E[\Delta_{D_c}\Phi] &\leq (A - a_{\max} - 1)p_A + (B - b_{\max} - 1)p_B \\ &\quad + \sum_l f(w_l) \end{aligned} \tag{3.1}$$

We divide the remainder of the analysis into three different cases depending on the value of δ_c.

• Suppose that $\delta_c = 1$.

The situation is fairly simple: $A \leq a_1 + 1, B \leq b_1 + 1, q_1 = p_{a_1} - p_A, q'_1 = p_{b_1} - p_B$. Without loss of generality, assume that $q_1 \geq q'_1$. From (3.1), we get the following bound

$$\begin{aligned} E[\Delta_{D_c}\Phi] &\leq a_1(p_{a_1} - p_A) + (b_1 - 1)(p_{b_1} - p_B) \\ &\leq a_1(p_{a_1} - p_{a_1+1}) + (b_1 - 1)(p_{b_1} - p_{b_1+1}) \end{aligned}$$

The second key property of the parameters p_α gives us the intended bound

$$E[\Delta_{D_c}\Phi] \leq \frac{5}{6}.$$

• Suppose $\delta_c = 2$.

The following claim dramatically simplifies the situation.

Claim 13 *When $\delta_c = 2$, $E[\Delta_{D_c}\Phi]$ is maximized for $a_1 = a_2 = a \leq 3$ and $b_1 = b_2 = b = 1$.*

We can now calculate $f(w_1), f(w_2)$, and $E[\Delta_{D_c}\Phi]$.

$$\begin{aligned} f(w_1) &= (a-1)p_a + bp_b = (a-1)p_a + 1 \\ f(w_2) &= (a-1)(p_a - p_A) + b(p_b - p_B) \\ E[\Delta_{D_c}\Phi] &\leq (A - 2a)p_A + (B - 2b - 1)p_B + 2(a-1)p_a + 2bp_b \\ &= 2(a-1)p_a + p_{2a+1} + 2 \end{aligned}$$

From the first key property of the parameters p_l, we have our intended bound on $E[\Delta_{D_c}\Phi]$,

$$E[\Delta_{D_c}\Phi] \leq \frac{2}{3} + 2 = \frac{11}{6}\delta_c - 1.$$

• Suppose that $\delta_c > 2$.

Consider the following definition

$$g(w_l) = a_l p_{a_l} + b_l p_{b_l} - \min(p_{a_l}, p_{b_l})$$

Notice that $g(w_l) = f(w_l)$ for $l \neq i_{\max}, l \neq j_{\max}$. Let us look at $f(w_{i_{\max}}), f(w_{j_{\max}})$. Suppose $l = i_{\max} = j_{\max}$.

$$\begin{aligned} f(w_l) &= a_{\max}(p_{a_{\max}} - p_A) + b_{\max}(p_{b_{\max}} - p_B) - \min(p_{a_{\max}} - p_A, p_{b_{\max}} - p_B) \\ &\leq a_{\max}(p_{a_{\max}} - p_A) + b_{\max}(p_{b_{\max}} - p_B) - \min(p_{a_{\max}}, p_{b_{\max}}) + p_A + p_B \\ &= g(w_l) + p_A(-a_{\max} + 1) + p_B(-b_{\max} + 1) \end{aligned}$$

Similarly when $i_{\max} \neq j_{\max}$, we get that

$$f(w_{i_{\max}}) + f(w_{j_{\max}}) \leq g(w_{i_{\max}}) + g(w_{j_{\max}}) + p_A(-a_{\max} + 1) + p_B(-b_{\max} + 1).$$

Thus, we can bound the sum of $f(w_l)$ in terms of the sum of $g(w_l)$,

$$\sum_l f(w_l) \leq \sum_l g(w_l) + p_A(-a_{\max} + 1) + p_B(-b_{\max} + 1).$$

Plugging in this bound on the sum of $f(w_l)$ into (3.1) we get the following bound

$$E[\Delta_{D_c}\Phi] \leq (A - 2a_{\max})p_A + (B - 2b_{\max})p_B + \sum_l g(w_l). \tag{3.2}$$

We observed earlier that for our settings of p_i, $(i - c)p_i < \frac{1}{4}$ for $c \geq 2$ (or of course when $i = 0$). Thus, $(A - 2a_{\max})p_A, (B - 2b_{\max})p_B < \frac{1}{4}$. We can also easily bound $g(w_l)$. Assume $a_l \leq b_l$ and thus $p_{a_l} \geq p_{b_l}$. We then have

$$g(w_l) = a_l p_{a_l} + (b_l - 1)p_{b_l} \leq p_1 + 2p_3 = \frac{4}{3}.$$

Combining these bounds with (3.2) we can complete the case $\delta_c > 2$,

$$\begin{aligned} E[\Delta_{D_c}\Phi] &\leq \frac{1}{2} + \frac{4}{3}\delta_c \\ &\leq \frac{11}{6}\delta_c - 1 \text{ for } \delta_c > 2. \end{aligned}$$

This completes the proof except for the special case (*) when $D_{\sigma(v)} \cap D_{\tau(v)} \neq \emptyset$. Let $x_1, \ldots, x_{\delta_{\sigma(v)}}$ and $y_1, \ldots, y_{\delta_{\tau(v)}}$ denote the respective sets $\Gamma_{\sigma(v)}$ and $\Gamma_{\tau(v)}$. In particular it might occur that

$$S_\sigma(x_i, \tau(v)) = S_\sigma(v, \tau(v)), S_\tau(y_j, \sigma(v)) = S_\tau(v, \sigma(v)).$$

In order for this to happen there must exist an alternating path between x_i and y_j using colors $\sigma(v)$ and $\tau(v)$. In such a case, we let $S_\sigma(v, \tau(v)) = S_\tau(y_j, \sigma(v)) = \emptyset$. Notice that the set $D_{\sigma(v)}$ is still unchanged and in fact, it is the same as previously analyzed (with $A = a_i = a_{\max} = 0$) except that we now have $B = \sum_j b_j < 1 + \sum_j b_j$. The previous proof still holds in this case. For the set $D_{\tau(v)}$, assume $j = 1$ and we now have that $A = a_1 = 0, B = b_{\max} = 0$; while for $j > 1$ we have $a_j \geq 0$ (note that as before, if $S_\tau(y_j, \sigma(v)) = S_\tau(y_{j'}, \sigma(v))$, then we set $S_\tau(y_{j'}, \sigma(v)) = \emptyset$.) We can complete the proof as follows

$$E[\Delta_{D_{\tau(v)}} \Phi] \leq \sum_{2 \leq j \leq \delta_{\tau(v)}} a_j p_{a_j} \leq (\delta_c - 1) p_1 < \frac{11}{6} \delta_c - 1.$$

■

Proof of Claim 13:

Without loss of generality, assume that $p_{a_{\max}} - p_A \leq p_{b_{\max}} - p_B$ and $a_1 = a_{\max}$. Considering $f(w_1)$,

$$f(w_1) = \begin{cases} (a_1 - 1)(p_{a_1} - p_A) + b_1(p_{b_1} - p_B) & \text{if } b_1 = b_{\max} \\ (a_1 - 1)(p_{a_1} - p_A) + b_1 p_{b_1} & \text{otherwise} \end{cases}$$

Similarly, the other important quantities are

$$f(w_2) = \begin{cases} a_2 p_{a_2} + b_2(p_{b_2} - p_B) - \min(p_{a_2}, p_{b_2} - p_B) & \text{if } b_2 = b_{\max} \\ a_2 p_{a_2} + b_2 p_{b_2} - \min(p_{a_2}, p_{b_2}) & \text{otherwise} \end{cases}$$

$$E[\Delta_{D_c} \Phi] = (A - a_1 - 1) p_A + (B - b_{\max} - 1) p_B + f(w_1) + f(w_2)$$

Suppose that $b_1 = x, b_2 = y$ and we swap these values, i.e., let $b_1 = y$ and $b_2 = x$. Then $E[\Delta_{D_c}]$ might change only from the $\min(,)$ in $f(w_2)$. Thus, $E[\Delta_{D_c} \Phi]$ is maximized when $b_2 = \max(x, y), b_1 = \min(x, y)$. We assume from now on that $b_2 \geq b_1$ which implies the following simplified situation:

$$\begin{aligned} f(w_1) &= (a_1-1)(p_{a_1}-p_A)+b_1p_{b_1}, \\ f(w_2) &= a_2p_{a_2}+b_2(p_{b_2}-p_B)-\min(p_{a_2},p_{b_2}-p_B), \\ E[\Delta_{D_c}\Phi] &= (A-2a_1)p_A+(B-2b_2-1)p_B+(a_1-1)p_{a_1}+a_2p_{a_2} \\ &\quad +b_1p_{b_1}+b_2p_{b_2}-\min(p_{a_2},p_{b_2}-p_B). \end{aligned}$$

We can complete the proof by considering the two cases for $\min(p_{a_2}, p_{b_2}-p_B)$.

- $p_{a_2} \le p_{b_2}-p_B$: We then have

$$\begin{aligned} E[\Delta_{D_c}\Phi] &= (a_1-1)p_{a_1}+(a_2-1)p_{a_2}+(A-2a_1)p_A \\ &\quad +b_1p_{b_1}+b_2p_{b_2}+(B-2b_2-1)p_B. \end{aligned}$$

Observe that $(a_1-1)p_{a_1}$ is maximized for $a_1=3$, while $(A-2a_1)p_A > 0 \leftrightarrow a_1=a_2<3$. Thus, the terms involving a_1 and a_2 are maximized for $a_1=a_2\le 3$. Similarly, the terms $b_1p_{b_1}, b_2p_{b_2}$ are maximized for $b_1=b_2=1$, while $(B-2b_2-1)<0$ if $b_1\ne b_2$ and $(B-2b_2-1)=0$ if $b_1=b_2$. Thus, the maximum of $E[\Delta_{D_c}]$ is when $b_1=b_2=1$ and $a_1=a_2\le 3$ which completes the proof of the claim in this case.

Before considering the next case, note that when $a_1=a_2=3, b_1=b_2=1$,

$$E[\Delta_{D_c}\Phi(3,1,3,1)] = 2p_1+4p_3.$$

- $p_{a_2} > p_{b_2}-p_B$: In this case,

$$\begin{aligned} E[\Delta_{D_c}\Phi] &= (a_1-1)p_{a_1}+a_2p_{a_2}+(A-2a_1)p_A \\ &\quad +b_1p_{b_1}+(b_2-1)p_{b_2}+(B-2b_2)p_B. \end{aligned}$$

The equation is symmetric in the pair (a_1,a_2) and (b_2,b_1). Considering the terms involving a_1,a_2 we complete the proof as follows:

$$\begin{aligned} (a_1-1)p_{a_1}+a_2p_{a_2}+(A-2a_1)p_A &\le \begin{cases} 2p_3+p_1 & \text{if } a_1\ne a_2 \\ 0p_1+p_1+p_3 & \text{if } a_1=a_2 \end{cases} \\ &\le \frac{1}{2}E[\Delta_{D_c}\Phi(3,1,3,1)]. \end{aligned}$$

■

Remark

The proof showed that $E[\Delta\Phi] \leq 0$ when $k = \frac{11}{6}\Delta$. To show rapid mixing in this case, we need to bound $\alpha = \Pr[\Delta\Phi \neq 0]$. The difficulty arises when a pair of states σ, η are far apart in terms of Φ, say $\Phi(\sigma, \eta) = n$. Each vertex v may have $2\delta(v)$ colors in its neighborhood and thus no moves that decrease Φ. By some recoloring of at most $\frac{1}{6}\delta(v)$ neighbors of vertex v, we can guarantee v has some color available. Thus, $\alpha \geq \left(\frac{1}{nk}\right)^{\frac{\delta(v)}{6}+1}$ which implies the chain is rapidly mixing when the maximum degree Δ is a constant and $k = \frac{11}{6}\Delta$.

3.4 Comparison with Glauber dynamics

In this section, we prove theorem 1(b) by bounding the mixing time τ_{GD} of the Glauber dynamics in terms of the mixing time τ_{flip} of the flip dynamics. The proof relies on the comparison theorem of Diaconis and Saloff-Coste [DSC93] (see Randall and Tetali [RT98] for other examples that use this theorem).

We present the comparison theorem in our specific setting where both chains have the same state space Ω, the set of proper colorings, and uniform stationary distribution. The theorem relates the underlying graphs associated with the transition matrices P_{flip}, P_{GD} of the flip and Glauber dynamics respectively. For a reversible Markov chain with transition matrix P, the underlying graph is $G = (\Omega, E(P))$ where

$$E(P) = \{(\sigma, \tau) : P(\sigma, \tau) > 0\}.$$

Note that reversibility implies that G is undirected. For each move $(\sigma, \tau) \in E(P_{flip})$, we define an associated path of moves in $E(P_{GD})$. Instead of defining a canonical path $\gamma_{\sigma\tau}$, we define a set of fractional paths, called a *flow* (see Sinclair [Sin92] for an analogous use of flows). Let γ denote a path $(\eta_0, \eta_1, \ldots, \eta_k)$, where each $(\eta_i, \eta_{i+1}) \in E(P_{GD})$, with length $|\gamma| = k$. For $(\sigma, \tau) \in E(P_{flip})$, let $\Gamma_{\sigma\tau}$ denote the set of paths from σ to τ,

$$\Gamma_{\sigma\tau} = \{\gamma : \eta_0 = \sigma, \eta_k = \tau\}.$$

A flow is a set of functions $f = f_{\sigma\tau} : \Gamma_{\sigma\tau} \to \mathbb{R}^+$ where

$$\sum_{\gamma \in \Gamma_{\sigma\tau}} f(\gamma) = 1.$$

The idea is to define flows to minimize the (fractional) number of paths that traverse any particular edge. In particular, for $(\eta, \xi) \in E(P_{GD})$, we aim to minimize

$$A_{\eta\xi} = \frac{1}{P_{GD}(\eta,\xi)} \sum_{\substack{\gamma\in\Gamma_{\sigma\tau}:\\ (\eta,\xi)\in\gamma}} |\gamma| f(\gamma) P_{flip}(\sigma,\tau).$$

In our setting, observe that $P_{GD}(\eta,\xi) \geq \frac{1}{nk}$, while $P_{flip}(\sigma,\tau) \leq \frac{1}{nk}$. In addition, we will define flows such that if $f(\gamma) > 0$ then $|\gamma| < K_1$ for a positive constant K_1. This will follow from the fact that the flip dynamics only flips clusters of size at most 6. We can simplify the quantity $A_{\eta\xi}$ as

$$A_{\eta\xi} \leq K_1 \sum_{\substack{\gamma\in\Gamma_{\sigma\tau}:\\ (\eta,\xi)\in\gamma}} f(\gamma). \tag{3.3}$$

We are interested in the maximum over all edges,

$$A = \max_{(\eta,\xi)\in E(P_{GD})} A_{\eta\xi}.$$

We use the following theorem of Diaconis and Saloff-Coste [DSC93] (see [RT98] for the details of adapting the original theorem into the form we present below).

Theorem 14 ([DSC93])

$$\tau_{GD} \leq O(A\tau_{flip}|\Omega|)$$

Proof of Theorem 1(b):

Since $|\Omega| \leq k^n$, in order to prove the theorem it is sufficient to define a set of flows such that $A = O(1)$.

Recall that a move $\sigma \mapsto \tau$ of the flip dynamics interchanges colors $c = c_{\sigma\tau}$ and $c' = c'_{\sigma\tau}$ on a maximal two-colored cluster $S = T \cup T' = T_{\sigma\tau} \cup T'_{\sigma\tau}$, where $\sigma(v) = c$ for all $v \in T$ and $\sigma(v) = c'$ for all $v' \in T'$. A natural idea for a path $\gamma_{\sigma\tau}$ consisting of moves in the Glauber dynamics is as follows: recolor each $v \in T$ to an arbitrary color, then recolor each $v' \in T'$ to color c, and finally recolor each $v \in T$ to color c'. The problem with such paths is that by choosing an arbitrary color in the first stage, we have unnecessarily increased the 'load' through particular edges. For instance, suppose that we always try to choose color 'yellow' as the arbitrary color; meanwhile we never choose 'red', if possible. An edge e of the Glauber dynamics that recolors a vertex to color yellow will have a large 'load' (i.e., large A_e); while an edge e' that recolors a vertex to color red might have no paths that traverse it (i.e., $A_{e'} = 0$).

We instead divide the flow evenly among all such paths. In particular, denote the set of available colors for vertex v as

$$F_\sigma(v) = C \setminus \{\sigma(v) \cup \bigcup_{w \in \Gamma(v)} \sigma(w)\}.$$

Let ψ denote a set of colors for the set T where $\psi(v_i) \in F_\sigma(v_i)$ for each $v_i \in T$; the set of all such sets ψ is denoted by $\Psi_{\sigma\tau}$. Each $\psi \in \Psi$ defines a canonical path γ_ψ as follows. (Fix an arbitrary ordering on the vertices V.)

Stage i: Consider each $v_i \in T$ (in order), recolor v_i to color $\psi(v_i)$.

Stage ii: For each vertex $v' \in T'$ (in order), recolor v' to color c.

Stage iii: Finally, for each vertex $v_i \in T$ (in order), recolor v_i to color c'.

For each $\psi \in \Psi_{\sigma\tau}$, we define the flow along the path γ_ψ as

$$f(\gamma_\psi) = 1/|\Psi_{\sigma\tau}|.$$

Notice that the paths are of length $|T|+|T'|+|T|$. By the setting of the parameters for the flip dynamics, we know that $|T| + |T'| \leq 6$ and thus all paths with positive flow are of constant length.

In order to bound the flows $f()$, observe that $|F_\sigma(v)| \geq k - \Delta$, where Δ is the maximum degree of the graph. Since $k \geq \frac{11}{6}\Delta$, we have $|\Psi_{\sigma\tau}| = \Omega(k^{|T|})$ and hence

$$f(\gamma_\psi) = O(k^{-|T|}). \tag{3.4}$$

For an edge $(\eta, \xi) \in E(P_{GD})$, we can simplify the quantity $A_{\eta\xi}$ by using the upper bound on $f(\gamma)$. We partition the paths that traverse the edge based on the size of the associated set T. Let

$$R_i(\eta, \xi) = \{\gamma_\psi : (\eta, \xi) \in \gamma_\psi, \psi \in \Psi_{\sigma\tau}, |T_{\sigma\tau}| = i\}.$$

Combining (3.3) and (3.4) we get the following bound. There exists a positive constant K_2 such that

$$A_{\eta\xi} \leq K_2 \sum_i |R_i(\eta, \xi)|/k^i. \tag{3.5}$$

It remains to bound the number of paths that traverse an edge $(\eta, \xi) \in E(P_{GD})$ (i.e., $|R_i(\eta, \xi)|$). Notice that a specific path γ is defined by the sets of vertices T, T', colors c, c', set of colors ψ, as well as the colors $\sigma(x)$ for all $x \notin S$ (where $S = T \cup T'$). From the coloring η, we know $\sigma(x) = \eta(x)$ for all $x \notin S$. We need to bound the number of sets T, T', ψ and colors c, c' whose corresponding path traverses the edge (η, ξ). It turns out that many of these sets or colors are fixed. In particular, suppose the move $\eta \mapsto \xi$ recolors vertex $v \in V$. For a path γ, consider the stage during which we traverse this edge (η, ξ):

Stage ii: In this case, notice that $c = \xi(v), c' = \eta(v)$. In addition, we know that $v \in T'$. Recall that the cluster $S = T \cup T'$ is a maximal two-colored connected component with $|S| \leq 6$. The number of such clusters which contains v is at most Δ^5. Since all the vertices of T' have color c or c' in η, given a candidate set T the corresponding set T' is fixed. There are at most $O(\Delta^i)$ candidate sets T where $|T| = i$. For a specific such set T, the associated colors ψ are fixed (as well as T'). In particular, for each $w_i \in T$, $\psi(w_i) = \eta(w_i)$. Therefore, assuming that the edge (η, ξ) is traversed during stage (ii) of the path, then $|R_i(\sigma, \eta)| = O(\Delta^i)$.

Stage i: Observe that $c = \eta(v)$, $v \in T$, and $\psi(v) = \xi(v)$. There are at most k possible choices for the color c'. Let $T \setminus \{v\} = T_1 \cup T_2$ where the vertices in T_1 have already been recolored according to ψ, while the vertices in T_2 have not yet been recolored. There are at most $O(\Delta^{|T_1|})$ choices for the vertices in T_1. For each $w_i \in T_1$, we know $\psi(w_i) = \eta(w_i)$. Each of the vertices in the set T_2 (and T') still have color c (and c', respectively) in η. Thus, for a specific set T_1, we can determine the sets T_2 and T'. For the set T_2, there are $O(k^{|T_2|})$ choices for the associated colors ψ. Combining the number of choices for the color c' and sets T_1, ψ, we have $|R_i(\sigma, \eta)| = O(k^{1+|T_1|}\Delta^{|T_2|}) = O(k^i)$.

Stage iii: The situation is symmetrical with stage (i).

In general, we have $|R_i| = O(k^i)$. Combining this with (3.5) implies $A = O(1)$, which completes the proof of the theorem. ∎

Chapter 4

Sampling Independent Sets

4.1 Introduction

In this chapter we present the rapid mixing results for the Glauber dynamics on the hard core model. Recall that for a graph $G = (V, E)$, the measure μ_{IS} is defined on its set Ω of independent sets weighted by a fugacity $\lambda > 0$. Specifically, for all independent sets $\sigma \in \Omega$,

$$\mu_{IS}(\sigma) = \frac{\lambda^{|\sigma|}}{Z}$$

where $Z = \sum_{\sigma \in \Omega} \lambda^{|\sigma|}$ is the partition function.

The Glauber dynamics for this model is the following. From an independent set σ, the transitions $\sigma \mapsto \sigma'$ are

- Choose a vertex $v \in V$ uniformly at random.
- Let
$$\sigma' = \begin{cases} \sigma \cup \{v\} & \text{with probability } \frac{\lambda}{1+\lambda} \\ \sigma \setminus \{v\} & \text{with probability } \frac{1}{1+\lambda} \end{cases}$$
- If σ' is a valid independent set, move to state σ' otherwise remain at state σ.

The main result of this chapter is the following theorem.

Theorem 3 *The Glauber dynamics on the hard core model is rapidly mixing with mixing time $O(n \log n\Delta)$ when $\lambda < \frac{2}{\Delta-2}$, where Δ is the maximum degree of the graph. For $\lambda = \frac{2}{\Delta-2}$, the mixing time is $O(n^3\Delta^2)$.*

The simpler proof for triangle-free graphs, shown in section 4.2, was joint work with Michael Luby [LV99]. The case of arbitrary graphs is dealt with in section 4.3 and appeared in [Vig99a]. In an earlier work [LV97], we considered a more complicated chain whose transitions acted on edges of the graph. In that paper, we proved rapid mixing of this 'edge' chain for $\lambda \leq \frac{1}{\Delta-3}$. Independently of our improvement of this result, Dyer and Greenhill [DG97] found a modified version of the edge chain which they proved is rapid mixing when $\lambda \leq \frac{2}{\Delta-2}$. Their result also implies rapid mixing of the Glauber dynamics for the same range of λ, but gives worse time bounds.

Using lemma 11, the theorem immediately implies the following corollary.

Corollary 4 *The hard core model on $\mathbb{Z}^d$ is in the disordered phase when $\lambda < \frac{2}{2d-2}$.*

For $\mathbb{Z}^d$, it is widely believed that there is a unique critical parameter λ_c such that the system is in the disordered phase when $\lambda < \lambda_c$ and in the ordered phase when $\lambda > \lambda_c$. In two-dimensions, simulations suggest λ_c is about 3.79 [BET80], but rigorous bounds on λ_c are much worse. Most lower bounds rely on showing the Dobrushin-Shlosman [DS85] condition is satisfied. Using this approach, the best lower bounds are computer assisted and show $\lambda_c > 1.185$ [KRS89, RS87, DKS85]. A different approach has recently been used by van den Berg and Steif [BS94] who relate λ_c to the critical probability p_c for site percolation on $\mathbb{Z}^2$, proving that $\lambda_c > \frac{p_c}{1-p_c}$.

The results we present are for general graphs and imply bounds on λ_c for other lattices of interest. Work in the statistical physics community has centered on $\mathbb{Z}^2$. For this lattice their results are stronger than ours and imply fast convergence of the Glauber dynamics for larger λ than we prove.

4.2 Triangle-free graphs

We begin with the simpler proof of the theorem for triangle-free graphs.

4.2.1 Distance Function

For $\sigma \in \Omega, v \in V$ such that $v \notin \sigma$, let $\sigma_v = \sigma \cup \{v\}$ and let δ_v denote the degree of vertex v. The obvious idea for a distance function Φ between states $\sigma, \sigma_v \in \Omega$ is δ_v, the Hamming distance weighted by degree. Our distance function is a natural extension of this. Consider a vertex w which is a neighbor of v. We call w blocked if it has a neighbor which is in both independent sets. Suppose the next move of the Markov chain attempts to add w into the independent set. This might only work in one of the chains, causing an increase in the Hamming distance . Notice that this bad situation occurs if w is not blocked. Otherwise, this move is blocked from occurring in both chains. Our distance function is simply the weighted Hamming distance minus a constant $c < 1$ times the number of blocked neighbors of v.

Specifically, for $c = \frac{\Delta\lambda}{\Delta\lambda+2}$, our distance function Φ is as follows. We use $\Gamma(v)$ to denote the set of neighboring vertices of v. Denote the set of blocked neighbors of v by

$$B(\sigma, v) = \{w : w \in \Gamma(v), \Gamma(w) \cap \sigma \neq \emptyset\}.$$

The distance function for neighboring states σ, σ_v is

$$\Phi(\sigma, \sigma_v) = \delta_v - c|B(\sigma, v)|.$$

Recall that $\rho(\sigma, \eta)$ denotes the set of simple paths between σ and η. For arbitrary states σ, η, we define the distance function as

$$\Phi(\sigma, \eta) = \min_{\tau \in \rho(\sigma,\eta)} \sum_i \Phi(\tau_i, \tau_{i+1}).$$

This distance function Φ clearly satisfies the following conditions for all $\sigma, \eta \in \Omega$ and thus is a metric:

- $\Phi(\sigma, \eta) \leq \Phi(\sigma, \xi) + \Phi(\xi, \eta)$ for $\xi \in \Omega$. This is true since $\Phi(\sigma, \eta)$ is defined as a minimum over all paths including those going through ξ.
- $\Phi(\sigma, \eta) = \Phi(\eta, \sigma)$
- $\Phi(\sigma, \eta) \geq 0$ which follows from $c < 1$ and thus $\Phi(\sigma, \sigma_v) \geq \delta_v(1 - c) > 0$ for all $\sigma, \sigma_v \in \Omega$.
- $\Phi(\sigma, \eta) = 0 \iff \sigma = \eta$.

4.2.2 Analysis

For $\Phi = \Phi(\sigma, \sigma_v)$, we now analyze $E[\Delta\Phi]$. Our coupling is simply the identity, i.e., each chain attempts the same move. Notice that the only moves which might affect Φ either transition on v, a neighbor of v, or a neighbor of a neighbor of v. Let,

$$\begin{aligned}
E[\Delta^{+x}\Phi] &= E[\Delta\Phi|\text{Markov chain attempts to add } x \text{ into the independent set}] \\
E[\Delta^{-x}\Phi] &= E[\Delta\Phi|\text{Markov chain attempts to remove } x \text{ from the independent set}] \\
E[\Delta^{x}\Phi] &= \frac{\lambda}{1+\lambda}E[\Delta^{+x}\Phi] + \frac{1}{1+\lambda}E[\Delta^{-x}\Phi].
\end{aligned}$$

This gives,

$$E[\Delta\Phi] = \frac{1}{n}\left[E[\Delta^{v}\Phi] + \sum_{w\in\Gamma(v)} E[\Delta^{w}\Phi] + \sum_{x\in\Gamma(\Gamma(v))} E[\Delta^{x}\Phi]\right]$$

Consider a move which

- transitions on v:

  Since all neighbors of v are out of both independent sets, a move which transitions on v works in both chains. Afterwards, both chains are in the same state. Thus,

$$E[\Delta^{+v}\Phi] = E[\Delta^{-v}\Phi] = -\delta_v + c|B(\sigma, v)|,$$
$$E[\Delta^{v}\Phi] = -\delta_v + c|B(\sigma, v)|.$$

- transitions on w, where w is a neighbor of v:

  Since w is in neither independent set, $E[\Delta^{-w}\Phi] = 0$.

  Consider the move which attempts to add w into the set. Suppose w is not blocked. This move only works in the chain in state σ. To determine the effect of this move for such a w, observe the following:

$$\begin{aligned}
&E[\Delta^{+w}\Phi] = \Phi(\sigma_w, \sigma_v) - \Phi(\sigma, \sigma_v), \\
&\Phi(\sigma_w, \sigma_v) \le \Phi(\sigma_w, \sigma) + \Phi(\sigma, \sigma_v), \\
&\Phi(\sigma_w, \sigma) = \delta_w - c|B(\sigma, w)|.
\end{aligned}$$

  Combining these give $E[\Delta^{+w}\Phi] \le \delta_w - c|B(\sigma, w)|$.

  Note that,

$$E[\Delta^{w}\Phi] \le \begin{cases} \frac{\lambda}{1+\lambda}(\delta_w - c|B(\sigma, w)|) & \text{if } w \notin B(\sigma, v) \\ 0 & \text{otherwise.} \end{cases}$$

- transitions on x, where x is a neighbor of a neighbor of v:

  Suppose x is in both independent sets and consider the move which removes x from both sets. This move may unblock a vertex w. The set of such w are

  $$\alpha_x = \{w : w \in B(\sigma, v), \Gamma(w) \cap \sigma = \{x\}\}.$$

  We have,

  $$E[\Delta^{-x}\Phi] = \begin{cases} |\alpha_x|c & \text{if } x \in \sigma \\ 0 & \text{otherwise.} \end{cases}$$

  Consider the case when x is in neither independent set. Since the graph is triangle-free, v is not in the neighborhood of x. Thus, the move which attempts to add x into the independent set works in both or neither set. In particular, it works in both chains if no neighbor of x is in either independent set, i.e., $\Gamma(x) \cap \sigma = \emptyset$. The only possible effect of such a move is to make a vertex w blocked. The set of such w are

  $$\beta_x = \{w : w \in \Gamma(v) \cap \Gamma(x), w \notin B(\sigma, v)\}.$$

  Thus,

  $$E[\Delta^{+x}\Phi] = \begin{cases} -|\beta_x|c & \text{if } \Gamma(x) \cap \sigma = \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

  Combining these,

  $$E[\Delta^{x}\Phi] = \begin{cases} \frac{|\alpha_x|}{1+\lambda}c & \text{if } x \in \sigma \\ -\frac{\lambda|\beta_x|}{1+\lambda}c & \text{if } x \notin \sigma, \Gamma(x) \cap \sigma = \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

We now collect terms of $E[\Delta\Phi]$ in a manner that divides the contribution from x over its neighbors w. Note that for any such x, either $E[\Delta^{+x}\Phi]$ or $E[\Delta^{-x}\Phi]$, but not both are non-negative. We can amortize these over those $w \in \alpha_x$ or $w \in \beta_x$. Notice that a blocked (unblocked) vertex w can only be in α_x (β_x, respectively).

For a blocked vertex w in the neighborhood of v, let

$$\Gamma'(w) = \{x : x \in \Gamma(w) \setminus \{v\}, w \in \alpha_x\},$$

$$E[\Delta^{*w}\Phi] = E[\Delta^{w}\Phi] + \sum_{x \in \Gamma'(w)} \frac{1}{|\alpha_x|} E[\Delta^{x}\Phi].$$

Similarly, for an unblocked vertex w in the neighborhood of v, let

$\Gamma'(w) = \{x : x \in \Gamma(w) \setminus \{v\}, w \in \beta_x\}$,

$$E[\Delta^{*w}\Phi] = E[\Delta^w\Phi] + \sum_{x\in\Gamma'(w)} \frac{1}{|\beta_x|} E[\Delta^x\Phi].$$

We now have that

$$E[\Delta\Phi] = \frac{1}{n}\left[E[\Delta^v\Phi] + \sum_{w\in\Gamma(v)} E[\Delta^{*w}\Phi]\right].$$

We can bound $E[\Delta^{*w}\Phi]$ as follows.

- Suppose w is blocked.

  We know that $E[\Delta^w\Phi] = 0$. We also know that it can be in α_x for at most one x. Thus,

  $$E[\Delta^{*w}\Phi] \leq \frac{c}{1+\lambda}.$$

- Suppose w is unblocked.

  For each neighbor x of w (other than v), either $x \in B(\sigma, w)$ and thus contributes to $E[\Delta^w\Phi]$ or $x \notin B(\sigma, w)$ which implies $\Gamma(x) \cap \sigma = \emptyset$ and it contributes to $E[\Delta^x\Phi]$. From these observations we have,

  $$\begin{aligned} E[\Delta^{*w}\Phi] &= \frac{\lambda}{1+\lambda}[\delta_w - c|B(\sigma, w)|] + \sum_{x\in\Gamma(w)\setminus B(\sigma,w), x\neq v} -\frac{\lambda}{1+\lambda}c \\ &= \frac{\lambda}{1+\lambda}[\delta_w - c|B(\sigma, w)| - c(\delta_w - 1 - |B(\sigma, w)|)] \\ &= \frac{\lambda}{1+\lambda}[\delta_w - c(\delta_w - 1)]. \end{aligned}$$

  where the second equality is from noticing the summation is over a set of size exactly $\delta_w - 1 - |B(\sigma, w)|$. From algebraic manipulations and our definition of c we have that

  $$(2+\lambda)c = \lambda(\Delta - c(\Delta - 1) \geq \lambda(\delta_w - c(\delta_w - 1)).$$

  This implies that if w is unblocked then $E[\Delta^{*w}\Phi] \leq \frac{1}{1+\lambda}(2+\lambda)c$.

Using our bounds on $E[\Delta^{*w}\Phi]$ we have

$$\begin{aligned}
n(1+\lambda)E[\Delta\Phi] &= (1+\lambda)\left[E[\Delta^{v}\Phi] + \sum_{w\in\Gamma(v)} E[\Delta^{*w}\Phi]\right] \\
&\le (1+\lambda)[-\delta_v + c|B(\sigma,v)|] + \sum_{w\in B(\sigma,v)} c + \sum_{w\in\Gamma(v)\setminus B(\sigma,v)} c(2+\lambda)] \\
&= (1+\lambda)[-\delta_v] + \sum_{w\in B(\sigma,v)} c(2+\lambda) + \sum_{w\in\Gamma(v)\setminus B(\sigma,v)} c(2+\lambda) \\
&= -(1+\lambda)\delta_v + \delta_v c(2+\lambda) \\
&= \frac{\delta_v}{\Delta\lambda+2}[\lambda(\Delta-2)-2].
\end{aligned}$$

Therefore,

$$E[\Delta\Phi] \le \frac{1}{n(1+\lambda)}\frac{\delta_v}{\Delta\lambda+2}[\lambda(\Delta-2)-2].$$

Notice that $E[\Delta\Phi] < 0$ when $\lambda < \frac{2}{\Delta-2}$.

We now want to use this bound on $E[\Delta\Phi]$ with the path coupling theorem to get a bound on the mixing time. Recall that the path coupling theorem uses a bound on $\beta = \max_{\sigma,\sigma_v}\beta_{\sigma,\sigma_v}$ where

$$E[\Phi(\sigma',\sigma_v')] = \beta_{\sigma,\sigma_v}\Phi(\sigma,\sigma_v).$$

We want to determine β_{σ,σ_v} in terms of $E[\Delta\Phi(\sigma,\sigma_v)] = E[\Delta\Phi]$ as follows:

$$\begin{aligned}
\beta_{\sigma,\sigma_v}\Phi(\sigma,\sigma_v) &= E[\Phi(\sigma',\sigma_v')] \\
(\beta_{\sigma,\sigma_v}-1)\Phi(\sigma,\sigma_v) &= E[\Phi(\sigma',\sigma_v')] - \Phi(\sigma,\sigma_v) = E[\Delta\Phi] \\
\beta_{\sigma,\sigma_v} &= 1 + \frac{E[\Delta\Phi]}{\Phi}.
\end{aligned}$$

Observe that by our definition of Φ we have $\Phi \le \delta_v$. From this observation and our bound on $E[\Delta\Phi]$ we get a bound on β:

$$\beta \le 1 + \frac{1}{n(1+\lambda)}\frac{\lambda(\Delta-2)-2}{(\Delta\lambda+2)}.$$

The path coupling theorem needs a bound on β and for Φ to be integer valued on $\{0,\ldots,D\}$. At the moment, Φ can have fractional values since c is not an integer. Simply consider $\Phi' = \frac{\Phi}{c}$ which is integer-valued. Since, $\Phi(\sigma,\sigma_v) \le \delta_v \le \Delta$, we have $\Phi(\sigma,\eta) \le n\Delta$

for arbitrary $\sigma, \eta \in \Omega$. Thus, $D \leq \frac{n\Delta}{c}$. Plugging these bounds on β and D into the path coupling theorem we get that when $\lambda < \frac{2}{\Delta - 2}$,

$$\tau \leq \frac{n(1+\lambda)(\Delta\lambda + 2)}{2 - \lambda(\Delta - 2)} \log(\frac{n\Delta 2e}{c})$$

Using the fact that $\Delta \geq 3, \lambda \leq \frac{2}{\Delta-2}$, we get $\lambda \leq 2, \Delta\lambda \leq 6, c \geq \frac{1}{3}$. We can now simplify the bound on the mixing time. For $\lambda = (1-\alpha)\frac{2}{\Delta-2}$, where α is positive,

$$\tau(\epsilon) \leq \frac{48n}{\alpha} \log(n\Delta)$$

This completes the proof of the theorem for the case of triangle-free graphs.

4.3 Arbitrary Graphs

In this section, we prove the theorem for arbitrary graphs. Throughout this section, we consider an arbitrary pair of independent sets σ, η and will omit obvious references to them as parameters to functions. We begin with some notation. Let D denote the set of disagree vertices, i.e., D is the symmetric difference between σ and η. The set of agree vertices is $A = V \setminus D$. We use $D_v = \Gamma(v) \cap D$ to denote the disagree neighbors of a vertex v and d_v is the cardinality of D_v. Similarly, $A_v = \Gamma(v) \cap A$.

Let $c = \frac{\Delta\lambda}{\Delta\lambda+2}$. We now define a distance function between $\sigma, \eta \in \Omega$:

$$\alpha_v = \left\{ \begin{array}{ll} \delta_v & \text{if } v \in D \\ 0 & \text{otherwise} \end{array} \right\}$$

$$\beta_v = \left\{ \begin{array}{ll} -cd_v & \text{if there exists a neighbor } w \text{ of } v \text{ such that } w \in \sigma, w \in \eta \\ -c(d_v - 1) & \text{if there is no such } w \text{ and } d_v > 1 \\ 0 & \text{otherwise} \end{array} \right\}$$

$$\Phi = \sum_v [\alpha_v + \beta_v].$$

We use the following theorem, the proof of which is the same as theorem 9.

Theorem 15 *Let Φ be an integer-valued metric defined on $\Omega \times \Omega$ which takes values in $\{0, \ldots, D\}$. Suppose there exists a positive β such that for all $\sigma, \eta \in \Omega$ there exists a coupling (σ_t, η_t) of the Markov chain with*

$$E[\Phi(\sigma_{t+1}, \eta_{t+1})] \leq \beta\Phi(\sigma_t, \eta_t).$$

If $\beta < 1$ the mixing time is bounded by

$$\tau \leq \frac{\log(2eD)}{1-\beta}.$$

If $\beta = 1$ and there exists a positive α such that, for all t and arbitrary $\sigma_t, \eta_t \in \Omega$,

$$\Pr[\Phi(\sigma_{t+1}, \eta_{t+1}) \neq \Phi(\sigma_t, \eta_t)] \geq \alpha,$$

then the mixing time is bounded by

$$\tau = O\left(\frac{D^2}{\alpha}\right).$$

We now analyze the expected change in Φ from the next transition of the Markov chain. As before, our coupling is simply the identity, i.e., each chain attempts the same move at every step. For simplicity, we rescale everything by a factor $n(1+\lambda)$.

$$n(1+\lambda)E[\Delta\Phi] = n(1+\lambda)\sum_v E[\Delta\alpha_v] + E[\Delta\beta_v]$$

We first try to manipulate the terms in the expected change in Φ to ease the analysis. Observe that for a vertex $v \in A$ if each of its neighbors w are in A then we are guaranteed that after the next move, $v \in A$ and thus α_v is still 0. We can then amortize the expected change in α_w over its disagree neighbors as follows,

$$\sum_v E[\Delta\alpha_v] = \sum_{v\in D}\left(E[\Delta\alpha_v] + \sum_{w\in A_v}\frac{1}{d_w}E[\Delta\alpha_w]\right). \tag{4.1}$$

To simplify our accounting, we divide $E[\Delta\beta_v]$ as follows. Let,

$$E[\Delta^w\beta_v] = E[\Delta\beta_v|\text{Markov chain transitions on } w]$$

We then have the following from the definition of β_v.

$$\begin{aligned}\sum_v E[\Delta\beta_v] &= \sum_{v,w} E[\Delta^w\beta_v] \\ &= \sum_v \sum_{w\in\Gamma(v)} E[\Delta^w\beta_v]\end{aligned}$$

For an agree vertex w, we can now try to amortize the expected change in β_w over its disagree neighbors,

$$\sum_v E[\Delta\beta_v] = \sum_{v\in D}\left[\sum_{w\in\Gamma(v)} E[\Delta^w\beta_v] + \sum_{w\in A_v} E[\Delta^v\beta_w]\right] + \sum_{v\in A}\sum_{w\in A_v} E[\Delta^v\beta_w]. \qquad (4.2)$$

Observe that the following is also true,

$$\begin{aligned}
&\sum_{v\in A}\sum_{w\in A_v} E[\Delta^v\beta_w] \\
&\quad= \sum_{v\in D}\sum_{w\in A_v}\sum_{x\in A_w}\frac{1}{d_w}E[\Delta^w\beta_x] + \sum_{w\in A:d_w=0}\sum_{x\in A_w} E[\Delta^w\beta_x] \\
&\quad= \sum_{v\in D}\sum_{w\in A_v}\sum_{x\in A_w}\frac{1}{d_w}E[\Delta^w\beta_x] + \sum_{w\in A}\sum_{x\in A_w:d_x=0} E[\Delta^x\beta_w]. \qquad (4.3)
\end{aligned}$$

Notice that if a vertex $x \in A$ and $d_x = 0$ then after the move we are still guaranteed that $x \in A$ (both σ and η have the same configurations on the neighborhood of x). Consider $w \in A, x \in A_w$ such that $d_w = 0$. Currently $\beta_w = 0$ and it can only increase after one transition if d_w increases. We can conclude that $E[\Delta^x\beta_w] = 0$ and furthermore that

$$\sum_{w\in A}\sum_{x\in A_w:d_x=0} E[\Delta^x\beta_w] = \sum_{v\in D}\sum_{w\in A}\sum_{x\in A_w:d_x=0}\frac{1}{d_w}E[\Delta^x\beta_w]. \qquad (4.4)$$

Combining (4.2),(4.3), and (4.4) implies

$$\begin{aligned}
\sum_v E[\Delta\beta_v] \;=\; &\sum_{v\in D}\Bigg[\sum_{w\in\Gamma(v)} E[\Delta^w\beta_v] + \sum_{w\in A_v}\Bigg(E[\Delta^v\beta_w] + \sum_{x\in A_w}\frac{1}{d_w}E[\Delta^w\beta_x] \\
&+ \sum_{x\in A_w, d_x=0}\frac{1}{d_w}E[\Delta^x\beta_w]\Bigg)\Bigg]. \qquad (4.5)
\end{aligned}$$

Using (4.1) and (4.5), we can divide the expected change in Φ over the disagree vertices as follows.

$$E[\Delta\Phi] = \sum_{v\in D}\left[E[\Delta\alpha_v] + \sum_{w\in\Gamma(v)}\phi_v(w)\right]$$

where $\phi_v(w)$ is the following:

$$\phi_v(w) = \begin{cases} E[\Delta^w\beta_v] & \text{if } w\in D \\ \frac{1}{d_w}E[\Delta\alpha_w] + E[\Delta^w\beta_v] + E[\Delta^v\beta_w] & \\ \quad +\frac{1}{d_w}\left(\sum_{x\in A_w} E[\Delta^w\beta_x] + \sum_{x\in A_w, d_x=0} E[\Delta^x\beta_w]\right) & \text{if } w\in A \end{cases}$$

The analysis will show that for each disagree vertex v, $E[\Delta\alpha_v]+\sum_{w\in\Gamma(v)}\phi_v(w)\le 0$ when $\lambda\le\frac{2}{\Delta-2}$.

For the remainder of the proof, consider a disagree vertex v and without loss of generality we assume that $v\in\sigma, v\notin\eta$. Recall that $\alpha_v=\delta_v$. Notice that the move which attempts to remove v from both independent sets definitely causes v to agree and has weight 1. We also know that the move which attempts to add v to both sets works in both sets if $d_v=0$. This move occurs with weight λ. Thus,

$$E[\Delta\alpha_v]=\begin{cases}-(1+\lambda)\delta_v & \text{if } d_v=0\\ -\delta_v & \text{otherwise}\end{cases}\tag{4.6}$$

Let w denote a neighbor of $v\in D$, we analyze $\phi_v(w)$ based on the following cases:

- $w\in D$:

  We know that $w\notin\sigma, w\in\eta$ and need to consider $E[\Delta^w\beta_v]$. Notice that the move that attempts to add w into both independent sets does not change σ or η; while the move that attempts to remove w works in both sets. Since no neighbor of v is in both independent sets, the move that removes w only changes β_v if $d_v>1$. In this case, we then have that

  $$\phi_v(w)=\begin{cases}c & \text{if } d_v>1\\ 0 & \text{otherwise}\end{cases}\tag{4.7}$$

- $w\in A$ and w has a neighbor z such that $z\in\eta$:

  In this scenario we have $w\notin\sigma, w\notin\eta$. Also, there are no moves that change the configuration at w (since $v\in\sigma, z\in\eta$). Thus for all neighbors x of w,

  $$E[\Delta\alpha_w]=E[\Delta^w\beta_v]=E[\Delta^w\beta_x]=0$$

  Now let us consider $E[\Delta^v\beta_w]$. Notice that if $d_v=0$ then v agrees after the moves which attempt to add or remove v from both independent sets. Whereas if $d_v>0$ then v only agrees after the move which attempts to remove it from both sets. In the worst case, these moves cause β_w to increase by c,

  $$E[\Delta^v\beta_w]\le\begin{cases}c(1+\lambda) & \text{if } d_v=0\\ c & \text{otherwise.}\end{cases}$$

We still need to consider $E[\Delta^x \beta_w]$ where x is an agree neighbor of w and all the neighbors of x agree. Attempting to add x to both independent sets can only decrease β_w. Also, attempting to remove x from both sets can only have an effect if x is already in both sets. In fact, if there is one such x that is in both sets, then removing it from both sets might increase β_w by c,

$$\sum_{x \in A_v, d_x = 0} E[\Delta^x \beta_w] \leq c.$$

Since $\phi_v(w)$ is maximized for $d_w = 1$, we have that

$$\phi_v(w) \leq \begin{cases} c(2+\lambda) & \text{if } d_v = 0 \\ 2c & \text{otherwise.} \end{cases}$$

- $w \in A$, and no neighbors of w are in η:

We know that $\alpha_w = 0$, $w \notin \sigma$, $w \notin \eta$, but attempting to add w into the independent set only works in η. Thus, this move increases α_w by δ_w. Similarly, if $d_v > 0$ this move will decrease β_v by c.

$$E[\Delta \alpha_w] = \lambda \delta_w$$

$$E[\Delta^w \beta_v] = \begin{cases} -c\lambda & \text{if } d_v > 0 \\ 0 & \text{otherwise} \end{cases}$$

Now consider $E[\Delta^v \beta_w]$. If $d_w > 1$ then the move which attempts to remove v from both independent sets will increase β_w by c. Whereas attempting to add v to both sets does not effect β_w.

$$E[\Delta^v \beta_w] = \begin{cases} c & \text{if } d_w > 1 \\ 0 & \text{otherwise} \end{cases}$$

Consider a vertex $x \in A_w$. Furthermore suppose that x either has $d_x > 0$ or it has a neighbor which is in both independent sets. Then the move which attempts to add w to both independent sets and causes w to disagree will decrease β_x by c. In the other case where each neighbor of x is out of both independent sets, then the move that adds x to both sets decreases β_w by c and occurs with weight λ. Thus, for all agree neighbors x of w, either: (i) $E[\Delta^w \beta_x] = -c\lambda$ or (ii) $E[\Delta^x \beta_w] = -c\lambda$ and $d_x = 0$.

$$\sum_{x \in A_v} E[\Delta^w \beta_x] + \sum_{x \in A_v, d_x = 0} E[\Delta^x \beta_w] = -c\lambda(\delta_w - d_w)$$

For this case we have that

$$\phi_v(w) = \begin{cases} \lambda\delta_w - c\lambda(\delta_w - 1) & \text{if } d_w = 1, d_v = 0 \\ \lambda\delta_w - c\lambda(\delta_w - 1) - c\lambda & \text{if } d_w = 1, d_v > 0 \\ \frac{1}{d_w}(\lambda\delta_w - c\lambda(\delta_w - d_w)) + c & \text{if } d_w > 1, d_v = 0 \\ \frac{1}{d_w}(\lambda\delta_w - c\lambda(\delta_w - d_w)) + c - c\lambda & \text{if } d_w > 1, d_v > 0 \end{cases}$$

Recall our setting of $c = \frac{\Delta\lambda}{\Delta\lambda + 2}$.

We leave it to the reader to verify that once again the worst case is when $d_w = 1$.

$$\phi_v(w) \leq \begin{cases} \lambda\delta_w - c\lambda(\delta_w - 1) & \text{if } d_v = 0 \\ \lambda\delta_w - c\lambda(\delta_w - 1) - c\lambda & \text{otherwise} \end{cases}$$

Notice that for our setting of c, the following are true:

$$\begin{aligned} (2+\lambda)c &\geq \lambda\delta_w - c\lambda(\delta_w - 1) \\ 2c &\geq \lambda\delta_w - c\lambda(\delta_w - 1) - c\lambda \end{aligned}$$

We then have that for $w \in A$,

$$\phi_v(w) \leq \begin{cases} (2+\lambda)c & \text{if } d_v = 0 \\ 2c & \text{otherwise} \end{cases} \tag{4.8}$$

Using (4.6), (4.7), and (4.8), we have the following:

$$E[\Delta\alpha_v] + \sum_{w \in \Gamma(v)} \phi_v(w) \leq \begin{cases} -(1+\lambda)\delta_v + (2+\lambda)c\delta_v & \text{if } d_v = 0 \\ -\delta_v + 2c(\delta_v - 1) & \text{if } d_v = 1 \\ -\delta_v + 2c(\delta_v - d_v) + d_v c & \text{otherwise} \end{cases}$$

Using the facts that for $d_v > 1$,

$$-(1+\lambda)\delta_v + (2+\lambda)c\delta_v = -\delta_v + 2c(\delta_v - 1) \geq -\delta_v + 2c(\delta_v - d_v) + d_v c$$

$$\frac{\delta_v}{\Delta\lambda + 2}[\lambda(\Delta - 2) - 2] \geq -(1+\lambda)\delta_v + (2+\lambda)c\delta_v$$

Therefore,

$$E[\Delta\Phi] \leq \frac{1}{n(1+\lambda)} \frac{\sum_{v \in D} \delta_v}{\Delta\lambda + 2}[\lambda(\Delta - 2) - 2]$$

Notice that $E[\Delta\Phi] < 0$ when $\lambda < \frac{2}{\Delta - 2}$. Using theorem 15, the remainder of the proof proceeds as in the triangle-free case. ∎

Remark

When $\lambda = \frac{2}{\Delta-2}$ we have that $\beta = 1$. To prove rapid mixing we need to bound $\alpha = \Pr[\Delta\Phi \neq 0]$. For any disagree vertex v, consider the move that attempts to remove v from both sets. This move reduces α_v by δ_v; meanwhile, it might increase β_w for each $w \in \Gamma(v)$. Since $c < 1$, we know that this move changes Φ by at least $\delta_v(1-c) > 0$. Therefore, $\alpha \geq \frac{1}{n(1+\lambda)}$ and from theorem 15 we conclude that when $\lambda = \frac{2}{\Delta-2}$,

$$\tau = O\left(\frac{n^3\lambda^2(1+\lambda)}{c^2}\right) = O(n^3\Delta^2).$$

Chapter 5

Torpid Mixing

5.1 Introduction

In this chapter, we prove the torpid mixing results on the Glauber dynamics and the Swendsen-Wang algorithm quoted in the Synopsis . The work described in this chapter is joint work with Borgs, Chayes, Frieze, Kim, Tetali, and Vu [BCF+99]. The Swendsen-Wang algorithm for the ferromagnetic Potts model is described in detail in the next section.

Gore and Jerrum [GJ97] proved that the Swendsen-Wang algorithm is torpidly mixing on the complete graph K_n at the critical point β_c when $q \geq 3$. Their arguments were extended to random graphs by Cooper and Frieze [CF99]. For the most interesting case of lattice graphs, however, only much weaker lower bounds on the mixing time were known: specifically, Li and Sokal [LS89] proved that the mixing time is at least linear in the number of vertices for finite boxes Q_L in $\mathbb{Z}^d$. We prove the following theorem concerning the torus $T_{L,d}$, which is defined in the Synopsis.

Theorem 5 *For the ferromagnetic Potts model on $T_{L,d}$ with $d \geq 2$ and sufficiently large q, there exist positive constants k_1, k_2 (which depend on d) such that if $\beta_c = \beta_c(q,d)$ is the critical point:*

(a) *The mixing time τ_{GD} of the Glauber dynamics for $\beta \geq \beta_c$ satisfies*

$$\tau_{GD} \geq e^{k_1 L/(\log L)^2}.$$

(b) *The mixing time τ_{SW} of the Swendsen-Wang algorithm at $\beta = \beta_c$ satisfies*

$$\tau_{SW} \geq e^{k_2 L/(\log L)^2}.$$

We also prove the following related theorem for the Glauber dynamics on the hard core model. Prior to our work, Dyer, Frieze, and Jerrum [DFJ99] proved the existence of a graph for which the Glauber dynamics is torpidly mixing for sufficiently large λ.

Theorem 6 *For the hard core model on $T_{L,d}$ with $d \geq 2$ and λ sufficiently large, there exists a positive constant k_3 (which depends on d) such that the mixing time τ_{GD} of the Glauber dynamics satisfies*

$$\tau_{GD} \geq e^{k_3 L^{d-1}/(\log L)^2}.$$

5.2 The Swendsen-Wang Algorithm

Consider a graph $G = (V, E)$ and the corresponding state space $\Omega = \{1, \ldots, q\}^V$ for the ferromagnetic q-state Potts model. The Swendsen-Wang (SW) algorithm uses the equivalence of the ferromagnetic Potts model and the random-cluster (FK) model of Fortuin and Kasteleyn [FK72]. The FK-model is a probability distribution defined on subsets of edges, i.e., $\Omega = 2^E$. It corresponds to independent bond percolation where each edge is present independently with probability p, but weighted by the number of clusters in the resulting subgraph. More precisely, in the FK model with parameters p and q, the probability of a subgraph $S \subset E$ is

$$\mu_{FK}(S) = \frac{p^{|S|}(1-p)^{|E \setminus S|} q^{\mathcal{C}(S)}}{Z_{FK}} \tag{5.1}$$

where $\mathcal{C}(S)$ denotes the number of connected components in the graph (V, S) and Z_{FK} is the appropriate normalizing factor (partition fuction).

Setting $p = 1 - \exp(-\beta)$ and keeping the same q in both models, we get the equivalence $Z_{potts} = Z_{FK}$. Edwards and Sokal [ES88] showed an interesting way to view the Swendsen-Wang algorithm by defining a joint distribution $\mu_{FKSW}(\sigma, S)$ on colorings and subgraphs. We will see this joint distribution as one of the intermediate steps in the following proof that $Z_{potts} = Z_{FK}$. Let δ denote the Kronecker delta function, which takes value 1 if the specified condition is true and 0 otherwise. Moreover, let $i \sim j$ denote a pair of adjacent vertices and A_σ denote the set of monochromatic edges in the coloring σ, i.e., $A_\sigma = \{i \sim j | \sigma(i) = \sigma(j)\}$. Then, from the definition of the Potts model 1.2 we have

$$\begin{aligned}
Z_{potts} &= \sum_{\sigma\in\Omega} \exp(-\beta \sum_{i\sim j} \delta_{\sigma(i)\neq\sigma(j)}) \\
&= \sum_{\sigma\in\Omega} \prod_{\substack{i\sim j:\\ \sigma(i)\neq\sigma(j)}} (1-p) \\
&= \sum_{\sigma\in\Omega} \prod_{\substack{i\sim j:\\ \sigma(i)\neq\sigma(j)}} (1-p) \left[\sum_{S\subset A_\sigma} p^{|S|}(1-p)^{|A_\sigma\setminus S|}\right] \\
&= \sum_{\sigma\in\Omega} \sum_{S\subset A_\sigma} p^{|S|}(1-p)^{|E\setminus S|} \qquad (5.2) \\
&= \sum_{S\subset E} p^{|S|}(1-p)^{|E\setminus S|} |\{\sigma\in\Omega : S\subset A_\sigma\}| \\
&= \sum_{S\subset E} p^{|S|}(1-p)^{|E\setminus S|} q^{\mathcal{C}(S)} \\
&= Z_{FK}
\end{aligned}$$

The basis of the Swendsen-Wang algorithm is line (5.2) above. Define

$$Z_{FKSW} = \sum_{\sigma\in\Omega} \sum_{S\subset A_\sigma} p^{|S|}(1-p)^{|E\setminus S|}. \qquad (5.3)$$

This suggests a joint distribution on colorings and subgraphs:

$$\mu_{FKSW}(\sigma, S) = \frac{p^{|S|}(1-p)^{|E\setminus S|}\delta_{S\subset A_\sigma}}{Z_{FKSW}}.$$

The Markov chain for the Swendsen-Wang algorithm alternately generates a sample S from the distribution μ_{FKSW} conditional on the current coloring σ, and then generates a sample σ' from μ_{FKSW} conditional on S. More precisely, from a coloring σ, the transitions $\sigma \mapsto \sigma'$ are:

- Create a subgraph $S \subset A_\sigma$ in the following manner: Independently for each $\{i,j\} \in A_\sigma$, include $\{i,j\}$ in the set S with probability $p = 1 - e^{-\beta}$.
- Independently for each component C of the subgraph S, choose a color c uniformly at random from $S = \{1, \dots, q\}$ and recolor every vertex in C with color c. This defines the new coloring σ'.

As the reader may easily verify, this Markov chain is ergodic and reversible with stationary distribution μ_{FKSW}.

5.3 Minimal Cutsets

In this section, we define minimal cutsets and present some technical lemmas about such sets. For a connected graph $G = (V, E)$, we let G_W denote the graph induced on $W \subset V$. Moreover, we say that $C \subset W$ is a component of W if C is the vertex set of a component of G_W. As usual, we define a subset $\gamma \subset E$ to be a cutset if $(V, E \setminus \gamma)$ is disconnected. We define γ to be a minimal cutset if all cutsets contained in γ are identical to γ. If γ is minimal, $(V, E \setminus \gamma)$ has exactly two connected components. For $W \subset V$, we let $\overline{W}$ denote the complement of W, i.e., $\overline{W} = V \setminus W$. We denote the set of edges between two disjoint sets of vertices W and W' by $(W : W')$. Finally, we use $\mathcal{C}(W)$ to denote the set of connected components of W.

We consider the cutset $\partial W = (W : \overline{W})$ and decompose it as $\partial W = \cup_{C \in \mathcal{C}(W)} \partial C$. We will further decompose ∂C into minimal cutsets, see Lemma 16 below. In order to state the lemma, we introduce the sets

$$\Gamma_C = \{(C : D) | D \in \mathcal{C}(\overline{C})\} = \{(\overline{D} : D) | D \in \mathcal{C}(\overline{C})\}$$

and

$$\Gamma(W) = \bigcup_{C \in \mathcal{C}(W)} \Gamma_C .$$

Lemma 16 *Consider $W \subset V$.*

(a) *Let C, C' be different connected components of W. There exist unique $D \in \mathcal{C}(\overline{C})$ and $D' \in \mathcal{C}(\overline{C'})$ such that $\overline{D} \subsetneq D'$ or equivalently $\overline{D'} \subsetneq D$.*

(b) *For $C \in \mathcal{C}(W)$, ∂C has a unique decomposition into minimal cutsets as $\partial C = \cup_{\gamma \in \Gamma_C} \gamma$.*

(c) *If γ, $\gamma' \in \Gamma(W)$ are distinct then they are disjoint.*

(d) *Let C and C' be two (not necessarily distinct) connected components of $W \subset V$. If X or $\overline{X}$ is a component of $\overline{C}$ and Y or $\overline{Y}$ is a component of $\overline{C'}$ then*

$$X \cap Y = \emptyset, X \cap \overline{Y} = \emptyset, \overline{X} \cap Y = \emptyset, \text{ or } \overline{X} \cap \overline{Y} = \emptyset.$$

Proof:

(a) We will first prove uniqueness. Since $C \cap C' = \emptyset$, $C' \subset \overline{C}$ and $C \subset \overline{C'}$. Furthermore, C is connected. Hence, there exists a unique $D' \in \mathcal{C}(\overline{C'})$ such that $C \subset D'$. For all $D \in \mathcal{C}(\overline{C})$, $C \subset \overline{D}$. Therefore, if there exists a $D' \in \mathcal{C}(\overline{C'})$ with $\overline{D} \subset D'$, D' must be the unique component containing C. The uniqueness of D is proved similarly. Next, we prove existence. Let D' be as above, so that $\overline{D'} \subset \overline{C}$. Since $C \cup D''$ is connected for all $D'' \in \mathcal{C}(\overline{C'})$, the set $\overline{D'} = C' \cup \bigcup\{D'' \in \mathcal{C}(\overline{C'}) : D'' \neq D'\}$ is connected. As a consequence, $\overline{D'} \subset \overline{C}$ must lie in one of the connected components D of $\overline{C}$.

(b) Obviously, $\partial C = \cup_{\gamma \in \Gamma_C}$ is a decomposition of ∂C into minimal cutsets of G. To prove uniqueness, assume that $\gamma \subset \partial C$ is a minimal cutset of G. Then there exists a $D \in \mathcal{C}(\overline{C})$ such that $(D : \overline{D}) \subset \gamma$. Otherwise, $C \cup D$ is connected in $G \setminus \gamma$ for every $D \in \mathcal{C}(\overline{C})$, which would imply that $G \setminus \gamma$ is connected. Since γ is minimal, $(D : \overline{D}) \subset \gamma$ implies $(D : \overline{D}) = \gamma$.

(c) For cutsets γ and γ' corresponding to the same component C, disjointness follows from the explicit form given in (b). Assume that $\gamma \cap \gamma' \neq \emptyset$ for two different connected components C and C'. This would imply that $\partial C \cap \partial C' \neq \emptyset$, which in turn implies that C and C' are connected in G, and hence in G_W. But this contradicts the assumption that C and C' are different connected components of G_W.

(d) Without loss of generality $X \in \mathcal{C}(\overline{C})$ and $Y \in \mathcal{C}(\overline{C'})$. We consider several cases:

- If $X = Y$ then $X \cap \overline{Y} = \emptyset$.
- If $C = C'$ and $X \neq Y$, then X and Y are different connected components of $\overline{C}$ which implies that $X \cap Y = \emptyset$.
- If $C \neq C'$ then we use part (a) of this lemma. We condition on whether X and/or Y are the unique $D \in \mathcal{C}(\overline{C})$ and $D' \in \mathcal{C}(\overline{C}')$ such that $\overline{D} \subseteq D'$.
  - $X \neq D, Y \neq D'$: Since $Y \subset \overline{D'}$ and part (a) implies that $X \subset \overline{D} \subset D'$, so we have that $X \cap Y = \emptyset$.
  - $X \neq D, Y = D'$: We saw in the previous case $X \subset D'$ and thus $X \cap \overline{Y} = \emptyset$. The case when $X = D, Y \neq D'$ is symmetric.
  - $X = D, Y = D'$: Since $\overline{X} \subset Y$ by part (a), $\overline{X} \cap \overline{Y} = \emptyset$.

■

Let $\gamma = (D : \overline{D})$ be a minimal cutset of G, in particular D and $\overline{D}$ are connected. We then define $\operatorname{Int}\gamma$ as the smaller (in terms of cardinality) of D and $\overline{D}$. If D and $\overline{D}$ have the same size, we can define $\operatorname{Int}\gamma$ as either D or $\overline{D}$. For definiteness, we define $\operatorname{Int}\gamma$ as the one containing a fixed point $x_o \in V$. For a cutset γ we define $\operatorname{Ext}\gamma = V \setminus \operatorname{Int}\gamma$, and for a collection Γ of minimal cutsets, we define the *interior of* Γ and the *common exterior of* Γ as

$$\operatorname{Int}\Gamma = \bigcup_{\gamma\in\Gamma} \operatorname{Int}\Gamma \qquad \text{and} \qquad \operatorname{Ext}\Gamma = \bigcap_{\gamma\in\Gamma} \operatorname{Ext}\gamma.$$

Note that $\operatorname{Int}\Gamma \cup \operatorname{Ext}\Gamma = V$ for all sets Γ of minimal cutsets. The following is a technical lemma about the interior of cutsets.

Lemma 17 *Let $W \subset V$.*

(a) *Let $\gamma, \gamma' \in \Gamma(W)$. If $\operatorname{Int}\gamma \cap \operatorname{Int}\gamma' \neq \emptyset$, then either $\operatorname{Int}\gamma \subset \operatorname{Int}\gamma'$ or $\operatorname{Int}\gamma' \subset \operatorname{Int}\gamma$.*

(b) *Either W or $\overline{W}$ is a subset of $\operatorname{Int}\Gamma(W)$.*

Proof:

(a) Let $X = \operatorname{Int}\gamma$ and $Y = \operatorname{Int}\gamma'$, and assume without loss of generality that $X \cap Y \neq \emptyset$. Applying the previous lemma, we have three cases:

(i) $X \cap \overline{Y} = \emptyset$, which is equivalent to $X \subset Y$,

(ii) $\overline{X} \cap Y = \emptyset$, which is equivalent to $Y \subset X$, and

(iii) $\overline{X} \cap \overline{Y} = \emptyset$ which is equivalent to $\overline{X} \subset Y$. Notice that $|\overline{X}| \geq |V|/2$ which implies that $|Y| \geq |V|/2$ and $|\overline{Y}| \leq |V|/2$. This contradicts the fact that $|Y| = |\operatorname{Int}\gamma'| \leq |\overline{Y}|$ unless equality holds, i.e. unless $|Y| = |\overline{Y}| = |X| = |\overline{X}| = |V|/2$. Together with $\overline{X} \subset Y$, this implies $\overline{X} = Y$ in contradiction to our assumption $X \cap Y \neq \emptyset$.

(b) We consider two cases. Suppose that for every $C \in \mathcal{C}(W)$ there is a cutset $\gamma \in \Gamma(W)$ with $C \subset \operatorname{Int}\gamma$. Then, clearly

$$W = \bigcup_{C\in\mathcal{C}(W)} C \subset \bigcup_{\gamma\in\Gamma(W)} \operatorname{Int}\gamma.$$

Suppose instead that there is $C \in \mathcal{C}(W)$ such that $C \not\subset \operatorname{Int}\gamma$ for all γ. Then since C is a subset of $\overline{D}$ for every component D of $\overline{C}$, the interior of the corresponding cutset $\gamma_D = (D : \overline{D})$ must be D. Thus $\overline{C} = \cup_{D\in\mathcal{C}(\overline{C})} \operatorname{Int}\gamma_D$. In particular, since C is a component of W,

$$\overline{W} \subset \overline{C} \subset \bigcup_{\gamma\in\Gamma(W)} \operatorname{Int}\gamma.$$

∎

Next we specialize to the torus $T_{L,d} = (V_{L,d}, E_{L,d})$. Consider a set $W \subset V_{L,d}$ and a fixed minimal cutset γ corresponding to W. For $e \in \gamma$ we define a *dual* $(d-1)$-dimensional cube e^* which is (i) orthogonal to e and (ii) bisects e, when $T_{L,d}$ is considered as immersed in the continuum torus $(\boldsymbol{R}/\boldsymbol{Z})^d$. (In dimension $d = 3$, the two-dimensional dual cells are referred to as *plaquettes*). We define a graph $\Gamma^* = (\gamma^*, E^*)$ where $\gamma^* = \{e^* : e \in \gamma\}$ and $(e_1^*, e_2^*) \in E^*$ iff $e_1^* \cap e_2^*$ is a cube of dimension $d-2$. The connected components of Γ^* are called the *co-components* of γ. These co-components are connected hypersurfaces of dual $(d-1)$-dimensional cells.

In the following, we will call cutsets with one co-component topologically trivial, and cutsets with more than one co-component topologically non-trivial. Small connected components which can be embedded in Z^d give rise to cutsets with only a single co-component, which are therefore topologically trivial. Topologically non-trivial cutsets arise from certain connected components which are large enough to "feel" the non-trivial topology of the torus. For example, the component $C = \{x \in V_{L,d} \mid 1 \le x_1 \le L/2\}$ gives rise to a cutset whose two co-connected components are two parallel interfaces, each of which has size L^{d-1}. The following lemma will be useful in order to bound the number of cutsets.

Lemma 18 **(a)** *Given a fixed edge $e \in E_{L,d}$ there are at most ν^k, $\nu = \min\{3, d^{64/d}\}$, distinct co-components γ of size k with $e \in \gamma$.*
(b) *If a cutset is non-trivial, each of its co-components contains at least L^{d-1} edges.*

Proof:

(a) This follows from the observation that the proofs in [Rue69] and [LM98] may be applied without changes to the torus.

(b) We need some notation. Consider a set of edges X and its dual X^*. Define the boundary ∂X^* of X^* as the set of $(d-2)$-dimensional hypercubes which belong to an odd number of $(d-1)$-dimensional cells in X^*. If $\partial X^* = \emptyset$, define the Z_2 winding vector of X^* as the vector $\mathbf{N}(X^*) = (N_1, \ldots, N_d)$, where N_i is the number of times X^* intersects an elementary loop in the i^{th} lattice direction mod 2.

Let X be a cutset, $X = (W : \overline{W})$, where $W \subset V$. Let $\mathcal{W} \subset (\boldsymbol{R}/\boldsymbol{Z})^d$ be the union of all closed unit cubes with center $w \in W$. Then X^* is the boundary of the set $\mathcal{W}$, and hence $\partial X^* = \emptyset$. Obviously, each elementary loop must leave and enter the set $\mathcal{W}$ the same number of times, implying that the winding vector of X^* is 0.

On the other hand, it is not difficult to prove that each set of edges X with $\partial X^* = \emptyset$ and $\mathbf{N}(X^*) = 0$ is a cutset for some set of points $W \subset V$, $X = (W : \overline{W})$. Indeed, the assumptions $\partial X^* = \emptyset$ and $\mathbf{N}(X^*) = 0$ imply that every closed loop in $T_{L,d}$ intersects X^* an even number of times. Considering an arbitrary vertex $w_0 \in V$ and the set of all "walks" of the form $(w_0, w_1, \ldots, w_k)$, $\{w_i, w_{i+1}\} \in \mathbf{E}_{L,d}$, we then define W as the set of points which can be reached from w_0 by a walk which intersects X^* an odd number of times.

Consider now a non-trivial minimal cutset γ and one of its co-components $\tilde{\gamma}$. Since γ is a cutset, $\partial\gamma^* = \emptyset$. This property is inherited by all its co-components, implying that $\partial\tilde{\gamma}^* = \emptyset$. Obviously, $\mathbf{N}(\tilde{\gamma}^*)$ is different from zero, since otherwise $\tilde{\gamma}$ would be a cutset itself, in contradiction to the assumption that γ is minimal. Let j be a direction for which $N_j(\tilde{\gamma}^*) \neq 0$. Then $\tilde{\gamma}^*$ intersects any fundamental loop in the j-direction an odd number of times, giving that $\tilde{\gamma}^*$ contains at least L^{d-1} dual $(d-1)$-dimensional cells. ■

5.4 The hard core model

In this section, we give a proof of Theorem 6. We start with some notation. For a bipartite graph $G = (V, E)$ we arbitrarily call the vertices in one part of the partition even, and those in the other part odd. We write V_{even} for the set of even vertices in V, and V_{odd} for the set of odd vertices in V. We denote the collection of independent sets of G by Ω. For an independent set $I \in \Omega$, define $W_{odd}(I)$ as the set of vertices in or adjacent to a vertex in the set $I \cap V_{odd}$ with an analogous definition for $W_{even}(I)$. We define the set $\Gamma_{odd}(I)$ as the set of minimal cutsets corresponding to $W_{odd}(I)$, $\Gamma_{odd}(I) = \Gamma(W_{odd}(I))$, and similarly for the set $\Gamma_{even}(I)$. Finally, for a cutset γ, we define $V(\gamma) = \bigcup_{\{x,y\}\in\gamma}\{x, y\}$. We present the following technical lemma for cutsets in this setting of independent sets.

Lemma 19 **(a)** *If $\gamma \in \Gamma_{odd}(I)$, then $V(\gamma) \cap I = \emptyset$.*
(b) *For $\gamma \in \Gamma_{odd}(I)$, the vertices in the set $V(\gamma) \cap \operatorname{Int}\gamma$ are either all even or all odd.*
(c) *For $\gamma \in \Gamma_{odd}(I)$, there exists an independent set I_γ such that $\Gamma_{odd}(I_\gamma) = \{\gamma\}$.*
(d) *Either $I \cap V_{odd}$ or $I \cap V_{even}$ is a subset of* $\operatorname{Int}\Gamma_{odd}(I)$.

Proof:

(a) We have to prove that $\{x, y\} \cap I = \emptyset$ whenever $\{x, y\} \in \gamma \subset \partial W_{odd}(I)$. First notice that for an odd vertex v, $v \in W_{odd}(I) \Leftrightarrow v \in I$, whereas if v is even then $v \in W_{odd}(I) \Leftrightarrow v$ has a neighbor $w \in I$. Suppose that $x \in I, y \notin I$. If x is odd then

$x, y \in W_{odd}(I)$. If x is even, then $x, y \notin W_{odd}(I)$. In either case, we have the contradiction that $\{x, y\} \notin \partial W_{odd}(I)$.

(b) If $\gamma \in \Gamma_{odd}(I)$, then $\gamma = (\overline{D} : D) = (C : D)$ for some component C of $W_{odd}(I)$ and some component D of $\overline{C}$. As a consequence, either $(V(\gamma) \cap \operatorname{Int} \gamma) \subseteq W_{odd}(I)$, or $(V(\gamma) \cap \operatorname{Int} \gamma) \subseteq \overline{W_{odd}(I)}$. If an odd vertex v is in the set $W_{odd}(I)$ then $v \in I$ and $w \in W_{odd}(I)$ for all neighbors w of v. Thus an odd vertex $v \in W_{odd}(I)$ cannot be incident to an edge in $\partial W_{odd}(I)$. As a consequence, the vertices of $V(\gamma) \cap \operatorname{Int} \gamma$ are even if $(V(\gamma) \cap \operatorname{Int} \gamma) \subseteq W_{odd}(I)$ and odd otherwise.

(c) If the vertices of the set $V(\gamma) \cap \operatorname{Int} \gamma$ are even then let $I_\gamma = (V_{odd} \cap \operatorname{Int} \gamma) \cup (V_{even} \cap \overline{\operatorname{Int} \gamma})$. Otherwise, exchange the sets V_{odd} and V_{even} in the definition of I_γ.

(d) Lemma 17 implies that either

$$W_{odd}(I) \subset \operatorname{Int} \Gamma_{odd}(I) \text{ or } \quad \overline{W_{odd}(I)} \subset \operatorname{Int} \Gamma_{odd}(I).$$

Since $I \cap V_{odd} \subset W_{odd}(I)$ and $I \cap V_{even} \subset \overline{W_{odd}(I)}$, the result follows. ∎

From now on, we specialize to the graph $T_{L,d}$. For a vertex $v = (v_1, \ldots, v_d) \in V$ and a "direction" $\alpha \in \{\pm 1, \ldots, \pm d\}$, we define the *shift* $\sigma_\alpha(v)$ as the vertex with coordinates v_i for $i \neq |\alpha|$ and $v_i + \operatorname{sign}(\alpha) \pmod{L}$ for $i = |\alpha|$, where $\operatorname{sign}(\alpha) = \alpha/|\alpha|$. For a cutset $\gamma \in \Gamma_{odd}(I)$, we define $\gamma_\alpha = \{(v, w) | (v, w) \in \gamma, v \in \operatorname{Int} \gamma, w = \sigma_\alpha(v)\}$. The following technical lemma is used in the proof of lemma 21.

Lemma 20 *For any cutset $\gamma \in \Gamma_{odd}(I)$ and any direction α, $|\gamma_\alpha| = |\gamma|/2d$.*

Proof: We first prove the lemma for $d = 2$. Let γ^* be the set of edges dual to the edges in γ. The set γ^* is a union of cycles, and each edge in the $+1$ or -1 direction in any of these loops is followed by an edge in the $+2$ or -2 direction by Lemma 19 (b). We therefore have that $|\gamma_i| + |\gamma_{-i}|$ is independent of the direction i. Since γ is a cutset, $|\gamma_i|$ must be equal to $|\gamma_{-i}|$, which implies the claim. For $d > 2$, we consider the intersection of $\operatorname{Int} \gamma$ with a two-dimensional plane $S(\{k_i\}) = \{x \in T \mid x_i = k_i, i \notin \{1, 2\}\}$. Since also the points in $(V(\gamma) \cap \operatorname{Int} \gamma) \cap S(\{k_i\})$ are all even or all odd, the above arguments can be applied to the intersection of γ and $S(\{k_i\})$, implying that $|\gamma_1| = |\gamma_{-1}| = |\gamma_2| = |\gamma_{-2}|$ since it is true for the intersection of these sets with any of the hyperplanes $S(\{k_i\})$. Applying this argument for an arbitrary pair of directions, we get the lemma. ∎

The next lemma is a generalization of a lemma first proved by Dobrushin in [Dob74].

Lemma 21 *Let Γ be a set of minimal cutsets, and let $\Omega_\Gamma = \{I : \Gamma \subset \Gamma_{odd}(I)\}$. Then*

$$\mu(\Omega_\Gamma) \leq \lambda^{-\sum_{\gamma\in\Gamma}(|\gamma|/2d)}$$

Proof: We first note that it is enough to prove there exists an injective map $\phi_\Gamma : \Omega_\Gamma \to \Omega$ such that

$$\mu(I) = \lambda^{-\sum_i |\gamma_i|/2d}\mu(\phi_\Gamma(I)).$$

Indeed, given such a map, we have

$$\mu(\Omega_\Gamma) = \lambda^{-\sum_i |\gamma_i|/2d}\mu(\phi_\Gamma(\Omega_\Gamma)) \leq \lambda^{-\sum_i |\gamma_i|/2d}.$$

In order to construct such a map ϕ_Γ, we introduce the partial order $\gamma \leq \gamma' \Leftrightarrow \operatorname{Int}\gamma \subset \operatorname{Int}\gamma'$. We then observe that, by induction, it is enough to prove that for any Γ and any $\gamma \in \Gamma$ such that γ is minimal in Γ with respect to the partial order, we have an injective map $\phi_\gamma : \Omega_\Gamma \to \Omega_{\Gamma\setminus\{\gamma\}}$ such that $\mu(I) = \lambda^{-|\gamma|/2d}\mu(\phi_\gamma(I))$.

We will now construct such a map. Consider $I \in \Omega_\Gamma$. Let $\sigma = \sigma_\alpha$. The proof holds for any choice of α. Defining

$$\phi_\gamma(I) = (I \cap \overline{\operatorname{Int}\gamma}) \cup \sigma(I \cap \operatorname{Int}\gamma) \cup (\operatorname{Int}\gamma \setminus \sigma(\operatorname{Int}\gamma)),$$

we will have to show that ϕ_γ is an injection, that $I' = \phi_\gamma(I)$ is an independent set with $\mu(I') = \mu(I)\lambda^{|\gamma|/2d}$ and that $I' \in \Omega_{\Gamma\setminus\{\gamma\}}$.

The first statement is obvious from the fact that the three sets $I_1 = I \cap \overline{\operatorname{Int}\gamma}$, $I_2 = \sigma(I \cap \operatorname{Int}\gamma)$ and $I_3 = \operatorname{Int}\gamma \setminus \sigma(\operatorname{Int}\gamma)$ are pairwise disjoint (use Lemma 19 (a) to see that I_1 and I_2 are disjoint).

I_1, I_2 are obviously independent and the independence of I_3 follows from $I_3 \subseteq V(\gamma)$ and Lemma 19(b). To then prove that I' is an independent set, we use that, again by Lemma 19 (a), the sets $I_1 \cup I_2$ and $I_1 \cup I_3$ are independent sets. It remains to show that $I_2 \cup I_3$ is also an independent set. Consider $v \in \operatorname{Int}\gamma \setminus \sigma(\operatorname{Int}\gamma)$ and $w \in \sigma(I \cap \operatorname{Int}\gamma)$. Then $v \notin \sigma(\operatorname{Int}\gamma)$ and hence $\sigma_{-\alpha}(v) \notin \operatorname{Int}\gamma$. On the other hand, $\sigma_{-\alpha}(w) \in I \cap \operatorname{Int}\gamma$. Therefore, $\sigma_{-\alpha}(v)$ and $\sigma_{-\alpha}(w)$ cannot be adjacent by Lemma 19 (a), which implies that v and w cannot be adjacent.

To prove $\mu(I') = \mu(I)\lambda^{|\gamma|/2d}$, we notice that $|(I \cap \overline{\operatorname{Int}\gamma}) \cup \sigma(I \cap \operatorname{Int}\gamma)| = |I|$. Thus ϕ_γ has increased the size of the independent set by exactly $|\operatorname{Int}\gamma \setminus \sigma(\operatorname{Int}\gamma)|$ which is $|\gamma_{-\alpha}| = |\gamma|/2d$ by Lemma 20.

To see that $I' \in \Omega_{\Gamma\setminus\{\gamma\}}$ note that $W_{odd}(I') = W_{odd}(I) \setminus \operatorname{Int}\gamma$. There are two possibilities for $\gamma' \in \Gamma \setminus \{\gamma\}$: $\operatorname{Int}\gamma \cap \operatorname{Int}\gamma' = \emptyset$ implying that $\operatorname{dist}(\operatorname{Int}\gamma, \operatorname{Int}\gamma') \geq 2$ and $I \cap \operatorname{Int}\gamma' = I' \cap \operatorname{Int}\gamma'$. Otherwise, $\operatorname{Int}\gamma \subseteq \operatorname{Int}\gamma'$ implying $\operatorname{dist}(\operatorname{Int}\gamma, \operatorname{Ext}\gamma') \geq 2$ and $I \cap \operatorname{Ext}\gamma' = I' \cap \operatorname{Ext}\gamma'$. ∎

Lemma 22 *Let $\Omega(k_1,\ldots,k_t)$ be the set of independent sets $I \in \Omega$ which contain a set of odd trivial cutsets of sizes $k_1,\ldots,k_t$. Then for $a = \min_{1\leq i\leq t} k_i$, $b = \max_{1\leq i\leq t} k_i$ and $k = \sum_{i=1}^t k_i$, we have*

$$\mu(\Omega(k_1,\ldots,k_t)) \leq (e(b-a+1)L^d/t)^t(\nu\lambda^{-1/(2d)})^k.$$

Let Ω_{nt} be the set of $I \in \Omega$ such that $\Gamma_{odd}(I)$ contains at least one non-trivial cutset. Then

$$\mu(\Omega_{nt}) \leq \left(L^d \frac{(\nu\lambda^{-1/2d})^{L^{d-1}}}{1-\nu\lambda^{-1/2d}}\right)^2 \exp\left(L^d \frac{(\nu\lambda^{-1/2d})^{L^{d-1}}}{1-\nu\lambda^{-1/2d}}\right).$$

Proof: To generate $\{\gamma_1,...,\gamma_t\}$ with $|\gamma_i| = k_i$, we first choose edges e_i in a certain fixed direction, e.g. direction 1, and then cutsets $\gamma_i \ni e_i$. (Every cut set contains an edge in direction 1 – see Lemma 20). In this way, each $\{\gamma_1,...,\gamma_t\}$ is counted $\prod_{j=a}^b t_j!$ times, where t_j is the number of k_i with $k_i = j$. The previous lemma and Lemma 18 (a) yield

$$\mu(\Omega(k_1,\ldots,k_t)) \leq \frac{L^{dt}}{\prod_{j=a}^b (t_j!)}(\nu\lambda^{-1/(2d)})^k.$$

(Note that it is safe to use the bound from Lemma 18(a) to bound the number of *trivial* cutsets, since for each trivial cutset, the dual is a single co-component.) Since $\sum_{j=a}^b t_j = t$,

$$\prod_{j=a}^b (t_j!) \geq \prod_{j=a}^b \left(\frac{t_j}{e}\right)^{t_j} \geq \left(\frac{t}{e(b-a+1)}\right)^t$$

and hence the result follows.

To prove the second statement, we use the previous lemma and the fact that each non-trivial cutset has at least two co-connected components to bound

$$\mu(\Omega_{nt}) \leq \sum_{k=2}^{\infty} \sum_{\gamma_{nt}^{(k)}} \lambda^{-|\gamma_{nt}^{(k)}|/2d}.$$

Here the sum $\sum_{\gamma_{nt}^{(k)}}$ goes over minimal cutsets with k co-components. Using Lemma 18, and the fact that there are at most L^{kd} possibilities for the k starting edges for the k

co-components of $\gamma_{nt}^{(k)}$, we conclude that

$$\begin{aligned}\mu(\Omega_{nt}) &\leq \sum_{k=2}^{\infty}\frac{1}{k!}\Big(\sum_{\ell=L^{d-1}}^{\infty}L^d(\nu\lambda^{-2d})^{\ell}\Big)^k\\ &\leq \sum_{k=0}^{\infty}\frac{1}{k!}\Big(\sum_{\ell=L^{d-1}}^{\infty}L^d(\nu\lambda^{-2d})^{\ell}\Big)^{k+2},\end{aligned}$$

which concludes the proof of the second statement. ∎

Lemma 23 *Let $0 < \alpha < 1$, and let*

$$\begin{aligned}\Omega_\alpha = \{I \in \Omega : \ & \Gamma_{odd}(I) \ \ \textit{contains only trivial cutsets,}\\ & \textit{and } \big|\mathrm{Int}\,\Gamma_{odd}(I)\big| \geq \alpha L^d\}.\end{aligned}$$

If λ is sufficiently large, say $\lambda^{1/(2d)} \geq 200\nu/\alpha$, then

$$\mu(\Omega_\alpha) \leq 2^{-c_\alpha L^{d-1}/(\log L)^2}$$

for some constant c_α depending on α and d.

Proof: For $I \in \Omega_\alpha$, the isoperimetric inequality of Bollobás and Leader [BL91] implies that $|\gamma| \geq |\mathrm{Int}\,\gamma|^{(d-1)/d}$ and hence

$$\begin{aligned}\sum_{\gamma\in\Gamma_{odd}(I)}|\gamma|^{d/(d-1)} &\geq \sum_{\gamma\in\Gamma_{odd}(I)}|\mathrm{Int}\,\gamma|\\ &\geq |\bigcup_{\gamma\in\Gamma_{odd}(I)}\mathrm{Int}\,\gamma|\\ &\geq \alpha L^d.\end{aligned}$$

If there is a cutset in $\Gamma_{odd}(I)$ of size at least L^{d-1}, then Lemma 22 directly gives the desired bound. Assume all cutsets are of size at most L^{d-1}. Let $\Gamma_i(I) = \{\gamma \in \Gamma_{odd}(I) : 2^{i-1} \leq |\gamma| < 2^i\}$, $i = 1, 2, ..., r = \lceil \log_2 L^{d-1} + 1\rceil$. Then since $\sum_{i=1}^{\infty}\frac{1}{i^2} = \pi^2/6$, there exists i such that

$$\sum_{\gamma\in\Gamma_i(I)}|\gamma|^{d/(d-1)} \geq c_\alpha^* L^d/i^2$$

where $c_\alpha^* = 6\alpha/\pi^2$. Thus I is in $\Omega(k_1, ..., k_t)$ for some t and $k_1, ..., k_t$ with $2^{i-1} \leq k_j \leq 2^i$ and $\sum_{j=1}^{t} k_j^{d/(d-1)} \geq c_\alpha^* L^d/i^2$. Let $s_i = c_\alpha^* L^d/(i^2 2^{id/(d-1)})$. The fact that $k_j \leq 2^i$ implies that $t \geq s_i$. This together with Lemma 22 gives

$$\begin{aligned}
\mu(\Omega_\alpha) &\leq \sum_{i=1}^{r} \sum_{t \geq s_i} \mu(\Omega(k_1, ..., k_t)) \\
&\leq \sum_{i=1}^{r} \sum_{t \geq s_i} (e2^i L^d/t)^t (\nu\lambda^{-1/(2d)})^{\sum k_j}.
\end{aligned}$$

Since $\sum k_j \geq 2^{i-1}t$ and there are at most 2^{it} choices for $k_1, k_2, \ldots, k_t$,

$$\begin{aligned}
\mu(\Omega_\alpha) &\leq \sum_{i=1}^{r} \sum_{t \geq s_i} (e2^{2i} L^d/t)^t (\nu\lambda^{-1/(2d)})^{2^{i-1}t} \\
&\leq \sum_{i=1}^{r} \sum_{t \geq s_i} \left(ei^2 2^{2i} 2^{id/(d-1)} (c_\alpha^*)^{-1} (\nu\lambda^{-1/(2d)})^{2^{i-1}}\right)^t \\
&\leq \sum_{i=1}^{r} \sum_{t \geq s_i} \left(\left[ei^2 2^{2i} 2^{id/(d-1)}\right]^{2^{1-i}} \nu\lambda^{-1/(2d)}/c_\alpha^*\right)^{2^{i-1}t},
\end{aligned}$$

where we have used the fact that $(c_\alpha^*)^{-1} \geq 1$ in the last step. Bounding $\left[ei^2 2^{2i} 2^{id/(d-1)}\right]^{2^{1-i}} \leq \left[ei^2 (16)^i\right]^{2^{1-i}} \leq [e(16e)^i]^{2^{1-i}} \leq 16e^2$ we see that for λ large enough (e.g. for $\lambda^{1/2d} \geq 32e^2\nu/c_\alpha^*$) , one gets

$$\begin{aligned}
\mu(\Omega_\alpha) &\leq \sum_{i=1}^{r} \sum_{t \geq s_i} 2^{-2^{i-1}t} \\
&\leq r2^{1-c_\alpha^* 2^{r-1} L^d/(r^2 2^{rd/(d-1)})} \\
&\leq 2^{-c_\alpha L^{d-1}/(\log L)^2}.
\end{aligned}$$

∎

We show next that if I is chosen from the probability distribution (1.3), then $|I|$ is unlikely to be small.

Lemma 24 *Let $0 < \delta < 1$. Then*

$$\mu\left(|I| \leq (1-\delta)\frac{L^d}{2}\right) \leq (2\lambda^{-\delta/2})^{L^d}.$$

Proof: There are at most 2^{L^d} independent subsets in $T_{L,d}$ and so the weight of those of size at most $(1-\delta)L^d/2$ is at most $2^{L^d}\lambda^{(1-\delta)L^d/2}$. On the other hand, the set of all even points has weight $\lambda^{L^d/2}$. The lemma follows immediately. ∎

Lemma 25 *For any $\rho < 1$, there is a constant c^*_ρ such that for λ sufficiently large,*

$$\begin{aligned} &\mu(||I \cap V_{odd}| - |I \cap V_{even}|| \leq \rho L^d/2) \\ \leq\ & \exp(-c^*_\rho L^{d-1}/(\log L)^2). \end{aligned}$$

Proof: Let $\delta = (1-\rho)/2$. Lemma 22 and Lemma 24 imply that $\mu(\Omega_{nt})$ and $\mu(|I| \leq (1-\delta)L^d/2)$ are small enough. Moreover, Lemma 23 for $\alpha = (1-\rho)/8$ implies that $\mu(\Omega_\alpha)$ is also small enough. If none of the three events whose probabilities we discuss above occurs, then $|I| > (1-\delta)L^d/2$ and $|\mathrm{Int}\,\Gamma_{odd}(I)| < L^d(1-\rho)/8$. The latter and Lemma 19(d) imply that either $|I \cap V_{odd}| < L^d(1-\rho)/8$ or $|I \cap V_{even}| < L^d(1-\rho)/8$. This together with the former yields that either

$$\begin{aligned} |I \cap V_{odd}| - |I \cap V_{even}| &= |I| - 2|I \cap V_{even}| \\ &> (1-\delta)L^d/2 - L^d(1-\rho)/4 \end{aligned}$$

or

$$\begin{aligned} |I \cap V_{even}| - |I \cap V_{odd}| &= |I| - 2|I \cap V_{odd}| \\ &> (1-\delta)L^d/2 - L^d(1-\rho)/4. \end{aligned}$$

Since $(1-\delta)L^d/2 - L^d(1-\rho)/4 = \rho L^d/2$, this concludes the proof. ∎

Proof of Theorem 6: We now partition $\Omega = \Omega^{(\rho)}_{\text{odd}} \cup \Omega^{(\rho)}_{\text{even}} \cup \Omega^{(\rho)}_{\text{rest}}$ where

$$\begin{aligned} \Omega^{(\rho)}_{\text{odd}} &= \{I \in \Omega : \ |I \cap V_{odd}| - |I \cap V_{even}| > \rho L^d/2\} \\ \Omega^{(\rho)}_{\text{even}} &= \{I \in \Omega : \ |I \cap V_{even}| - |I \cap V_{odd}| > \rho L^d/2\} \\ \Omega^{(\rho)}_{\text{rest}} &= \Omega \setminus (\Omega^{(\rho)}_{\text{odd}} \cup \Omega^{(\rho)}_{\text{even}}) \end{aligned}$$

By the last lemma $\mu(\Omega^{(\rho)}_{\text{rest}}) \leq \exp(-c^*_\rho L^{d-1}/(\log L)^2)$, and by symmetry $\mu(\Omega^{(\rho)}_{\text{odd}}) = \mu(\Omega^{(\rho)}_{\text{even}})$. Now consider the Glauber dynamics. Clearly, if $I \in \Omega^{(\rho)}_{\text{odd}}$ and I' is obtained by a single transition then $I' \in \Omega^{(\rho)}_{\text{odd}} \cup \Omega^{(\rho)}_{\text{rest}}$. To complete our proof by estimating Φ_S (see Theorem 10 in Chapter 2) for $S = \Omega_{\text{odd}}$, first notice that $\mu(S)\mu(\overline{S}) \geq 1/5$. Furthermore,

$$\begin{aligned} Q(S, \overline{S}) &= \sum_{\substack{I \in \Omega_{\text{odd}} \\ J \in \Omega_{\text{rest}}}} \mu(I)P(I,J) \\ &= \sum_{\substack{I \in \Omega_{\text{odd}} \\ J \in \Omega_{\text{rest}}}} \mu(J)P(J,I) \\ &\leq \mu(\Omega_{\text{rest}}). \end{aligned}$$

The theorem now follows. ∎

5.5 The Potts model

In this section we combine the methods and results of [BKMS91] and [BK90] with those of the last section to prove Theorem 5. Recall the definition of the ferromagnetic Potts model from chapter 1, as well as the definition of the critical point β_c for $\mathbb{Z}^d$.

As a first step towards proving Theorem 5, we define the contours corresponding to a configuration $A \in \Omega = 2^E$. To this end, we embed the vertex set V of the torus $T = (V, E)$ into the set $\mathbf{V} = (\mathbf{R}/(L\mathbf{Z}))^d$. For a set $X \subset \mathbf{V}$, we define its diameter $\operatorname{diam}(X) = \inf_{y \in \mathbf{V}} \sup_{x \in X} \operatorname{dist}(x, y)$, where $\operatorname{dist}(x, y)$ is the ℓ_∞-distance between the two points x and y in the torus $\mathbf{V}$. For an edge $e = \{x, y\} \in E$, let $\mathbf{e}$ be the set of points in $\mathbf{V}$ that lie on the line between x and y. Given A, we call a closed k-dimensional unit hypercube $c \subset \mathbf{V}$ with vertices in V *occupied* if all edges e with $\mathbf{e} \subset c$ are in A. We then define the set $\mathbf{V}(A) \subset \mathbf{V}$ as the 1/3-neighborhood of the union of all occupied k-dimensional hypercubes, $k = 1, \ldots, d$, i.e., $\mathbf{V}(A) = \{x \in \mathbf{V} : \exists c \text{ occupied, such that} \operatorname{dist}(x, c) < 1/3\}$, and the set $V(A)$ as the intersection of $\mathbf{V}(A)$ with the vertex set V of the discrete torus T. Note that $V(A) = \bigcup_{\{x,y\} \in A} \{x, y\}$. The set $\Gamma(A)$ of contours corresponding to a configuration $A \in \Omega$ are then the connected components of the boundary of $\mathbf{V}(A)$.

Following [BKMS91], we decompose the set of configurations Ω into three sets Ω_{ord}, Ω_{dis} and Ω_{Big}. To this end, we define a contour γ to be small if $\operatorname{diam}(\gamma) \leq L/3$. The set Ω_{Big} is then just the set of configurations $A \in \Omega$ for which $\Gamma(A)$ contains at least one contour that is not small. Next, restricting ourselves to small contours γ, we define the set $\mathbf{Ext}\,\gamma$ as the larger of the two connected components of $\mathbf{V} \setminus \gamma$, the set $\operatorname{Ext}\gamma$ as the intersection of $\mathbf{Ext}\,\gamma$ with V, and the set $\operatorname{Int}\gamma$ as $V \setminus \operatorname{Ext}\gamma$. For $A \in \Omega \setminus \Omega_{\text{Big}}$, let $\operatorname{Int} A = \bigcup_{\gamma \in \Gamma(A)} \operatorname{Int}\gamma$ and $\operatorname{Ext} A = V \setminus \operatorname{Int} A$. The sets Ω_{ord}, Ω_{dis} and Ω_{Big} are then defined as

$$
\begin{aligned}
\Omega_{\text{Big}} &= \{A \subset E : \ \exists \gamma \in \Gamma(A) \text{ such that} \\
&\qquad \operatorname{diam}(\gamma) > L/3\} \\
\Omega_{\text{ord}} &= \{A \subset E : \ \operatorname{diam}(\gamma) \leq L/3 \ \forall \gamma \in \Gamma(A) \\
&\qquad \text{and } V(A) \cap \operatorname{Ext} A \neq \emptyset\} \\
\Omega_{\text{dis}} &= \{A \subset E : \ \operatorname{diam}(\gamma) \leq L/3 \ \forall \gamma \in \Gamma(A) \\
&\qquad \text{and } V(A) \cap \operatorname{Ext} A = \emptyset\}.
\end{aligned}
$$

Lemma 26 *Let $A \in \Omega_{\text{ord}}$, and let $A_{\operatorname{Ext} A} = \{b \in E : b \subset \operatorname{Ext} A\}$. Then*
(a) $\operatorname{Ext} A = V(A) \cap \operatorname{Ext} A \neq \emptyset$*, and*

(b) $(\text{Ext}\, A, A_{\text{Ext}\, A})$ *is connected.*

Proof:

(a) Proceeding as in the proof of Lemma 17 (b), we obtain that either $V(A) \subset \text{Int}\, A$ or $\overline{V(A)} \subset \text{Int}\, A$. Since $A \in \Omega_{\text{ord}}$, we conclude that the latter is the case, which is equivalent to the statement that $\text{Ext}\, A = V(A) \cap \text{Ext}\, A$.

(b) The proof of this statement, which is implicit in [BKMS91], is straightforward but tedious. We leave it to the reader. ∎

In the next lemma we summarize some of the results of [BKMS91] used in this work. We begin with some notation. Let $A \in \Omega \setminus \Omega_{\text{Big}}$, and let $\gamma \in \Gamma(A)$. We say that γ is an exterior contour in $\Gamma(A)$ if $\gamma \subset \text{Ext}\, \gamma'$ for all $\gamma' \in \Gamma(A) \setminus \{\gamma\}$, and denote by $\Gamma_{\text{ext}}(A)$ the set of exterior contours in $\Gamma(A)$. Also, we define the size $\|\gamma\|$ of a contour γ as the number of times γ intersects the set $\bigcup_{e \in E} \mathbf{e}$. In order to motivate this definition, assume for a moment that the definition of the set $\mathbf{V}(A)$ had involved an ϵ-neighborhood, instead of the 1/3-neighborhood used above. With such a definition, the $(d-1)$-dimensional area of a contour γ would actually converge to $\|\gamma\|$ as $\epsilon \to 1/2$.

Lemma 27 *For all $d \geq 2$ there are constants $c > 0$ and $q_0 < \infty$ such that the following statements hold for $q \geq q_0$.*
(a) $\beta_c = \log q/d + O(q^{-c})$.
(b) *For all $\beta > 0$,*

$$\mu(\Omega_{\text{Big}}) \leq q^{-cL}.$$

(c) *If $\beta = \beta_c$, then*

$$\begin{aligned} \mu(\Omega_{\text{ord}}) &= \frac{q}{q+1} + O(q^{-cL}), \quad \textit{and} \\ \mu(\Omega_{\text{dis}}) &= \frac{1}{q+1} + O(q^{-cL}). \end{aligned}$$

(d) *If $\beta \geq \beta_c$, then*

$$\mu(\Omega_{\text{ord}}) \geq \frac{q}{q+1} + O(q^{-cL}).$$

(e) *If $\beta \geq \beta_c$ and Γ is a set of contours, then*

$$\mu\Big(A \in \Omega \setminus \Omega_{\text{big}} \textit{ and } \Gamma \subset \Gamma_{\text{ext}}(A)\Big) \leq q^{-c \sum_{\gamma \in \Gamma} \|\gamma\|}.$$

Observing that for $A \in \Omega \setminus \Omega_{\text{Big}}$, the set $\operatorname{Ext} A$ can be written as $\bigcup_{\gamma \in \Gamma_{\text{ext}}} \operatorname{Ext} \gamma$, which in turn implies that $\operatorname{Int} A = \bigcap_{\gamma \in \Gamma_{\text{ext}}} \operatorname{Int} \gamma$, we can now continue as in Section 5.4 to prove an analog of Lemma 23. Defining

$$\begin{aligned}
\Omega_{\text{ord}}^{(\alpha)} &= \{A \in \Omega_{\text{ord}} : |\{b \in A : b \subset \operatorname{Ext} A\}| \geq (1-\alpha) d L^d\}, \\
\Omega_{\text{dis}}^{(\alpha)} &= \{A \in \Omega_{\text{dis}} : |\operatorname{Int} A| \leq \alpha L^d\},
\end{aligned}$$

and $\Omega_{\text{Big}}^{(\alpha)} = \Omega \setminus (\Omega_{\text{ord}}^{(\alpha)} \cup \Omega_{\text{dis}}^{(\alpha)})$, we therefore get the following lemma.

Lemma 28 *Let $d \geq 2$ and $0 < \alpha < 1$. Then there are constants $c > 0$ and $c_\alpha > 0$ such that for q large enough the following statements hold.*
(a) *If $\beta \geq \beta_c$, then*

$$\mu(\Omega_{\text{Big}}^{(\alpha)}) = O(q^{-cL}) + O(q^{-c_\alpha L^{d-1}/(\log L)^2})$$

and

$$\mu(\Omega_{\text{ord}}^{(\alpha)}) \geq \frac{q}{q+1} + O(q^{-cL}) + O(q^{-c_\alpha L^{d-1}/(\log L)^2}).$$

(b) *If $\beta = \beta_c$, then*

$$\mu(\Omega_{\text{ord}}^{(\alpha)}) = \frac{q}{q+1} + O(q^{-cL}) + O(q^{-c_\alpha L^{d-1}/(\log L)^2}).$$

Proof of Theorem 5(a): Let $S = \Omega_{\text{ord}}^{(\alpha)}$. The conductance Φ_{SW} of the Swendsen-Wang chain can then be estimated as follows:

$$\Phi_{SW} \leq \Phi_S = \frac{1}{\mu(S)} \Pr(A' \notin \Omega_{\text{ord}}^{(\alpha)} \mid A \in \Omega_{\text{ord}}^{(\alpha)}). \tag{5.4}$$

Here A is chosen according to the measure μ defined in (5.1) and A' is constructed from A by one step of the Swendsen-Wang algorithm. We have

$$\begin{aligned}
\Pr(A' \notin \Omega_{\text{ord}}^{(\alpha)} \mid A \in \Omega_{\text{ord}}^{(\alpha)}) = \\
\Pr(A' \in \Omega_{\text{dis}}^{(\alpha)} \mid A \in \Omega_{\text{ord}}^{(\alpha)}) \\
+ \Pr(A' \in \Omega_{\text{Big}}^{(\alpha)} \mid A \in \Omega_{\text{ord}}^{(\alpha)}).
\end{aligned}$$

Observing that $A \in \Omega_{\text{ord}}^{(\alpha)}$ implies $|A| \geq (1-\alpha) d L^d$ while $A' \in \Omega_{\text{dis}}^{(\alpha)}$ implies $|A'| \leq d|V(A')| \leq d|\operatorname{Int} A'| \leq d\alpha L^d$, we see that A' can only be in $\Omega_{\text{dis}}^{(\alpha)}$ if at least $(1-2\alpha) d L^d$ edges are deleted in Step (SW1) of Swendsen-Wang. But the number of edges deleted is dominated by the binomial $B(dL^d, 1-p_c)$ and so

$$\begin{aligned}\Pr(A' \in \Omega_{\text{dis}}^{(\alpha)} \mid A \in \Omega_{\text{ord}}^{(\alpha)}) &\le \binom{dL^d}{(1-2\alpha)dL^d}(1-p_c)^{(1-2\alpha)dL^d} \\ &\le \left(\frac{e(1-p_c)}{1-2\alpha}\right)^{(1-2\alpha)dL^d} \\ &= e^{-\Omega((\log q)L^d)},\end{aligned}$$

where we have used Lemma 27(a) to bound $1-p_c = e^{-\beta_c} = e^{-\Omega(\log q)}$. Also

$$\begin{aligned}&\Pr(A' \in \Omega_{\text{Big}}^{(\alpha)} \mid A \in \Omega_{\text{ord}}^{(\alpha)}) \\ &\le \frac{\Pr(A' \in \Omega_{\text{Big}}^{(\alpha)})}{\Pr(A \in \Omega_{\text{ord}}^{(\alpha)})} \\ &= O(q^{-cL}) + O(q^{-c_\alpha L^{d-1}/(\log L)^2}),\end{aligned}$$

by Lemma 28(a). Using Lemma 28(b) to bound $\mu(\overline{S}) = 1-\mu(\Omega_{\text{ord}}^{(\alpha)})$ from below, we obtain that

$$\Phi_{SW} = O(q^{-cL}) + O(q^{-c_\alpha L^{d-1}/(\log L)^2}).$$

■

Proof of Theorem 5(b):

Let $\widehat{\Omega} = [q]^V$ be the set of colorings, and let $V_k(\boldsymbol{\sigma}) = \{x \in V : \sigma_x = k\}$ be the set of vertices that have color k in the coloring $\boldsymbol{\sigma} \in \widehat{\Omega}$. We then define the sets

$$\widehat{\Omega}_k^{(\alpha)} = \{\boldsymbol{\sigma} \in \widehat{\Omega} : |V_k| \ge (1-\alpha)|V|\}, \quad , k \in [q],$$

$$\widehat{\Omega}_{\text{ord}}^{(\alpha)} = \bigcup_{k\in[q]} \widehat{\Omega}_k^{(\alpha)},$$

$$\widehat{\Omega}_{\text{dis}}^{(\alpha)} = \{\boldsymbol{\sigma} \in \widehat{\Omega} : |V_k| \ge \frac{(1-\alpha)^2}{q}|V| \text{ for all } k \in [q]\},$$

and

$$\widehat{\Omega}_{\text{Rest}}^{(\alpha)} = \widehat{\Omega} \setminus (\widehat{\Omega}_{\text{ord}}^{(\alpha)} \cup \widehat{\Omega}_{\text{dis}}^{(\alpha)}).$$

To estimate the probability of $\widehat{\Omega}_{\text{ord}}^{(\alpha)}$ in the measure (1.2), we use the fact that both the measure (1.2) (denoted $\widehat{\mu}$ in this section) and the measure (5.1) (denoted μ in this section) are marginals of the Edwards-Sokal measure (5.3). Thus

$$\widehat{\mu}(\widehat{\Omega}_{\text{ord}}^{(\alpha)}) = \sum_{A\in\Omega} \pi(\widehat{\Omega}_{\text{ord}}^{(\alpha)} \mid A)\,\mu(A), \tag{5.5}$$

where $\pi(\widehat{\Omega}^{(\alpha)}_{\rm ord} \mid A)$ is the conditional measure of $\widehat{\Omega}^{(\alpha)}_{\rm ord}$, given $A \in \Omega$. Observing that $A \in \Omega^{(\alpha)}_{\rm ord}$ implies that all vertices in $\mathrm{Ext}\, A$ have the same color by Lemma 26 and the definition (5.3) of π, we have that

$$\pi(\widehat{\Omega}^{(\alpha)}_{\rm ord} \mid A) = 1 \qquad \text{if} \qquad A \in \Omega^{(\alpha)}_{\rm ord}. \tag{5.6}$$

For $A \in \Omega^{(\alpha)}_{\rm dis}$, on the other hand, all $n_A = |\mathrm{Ext}\, A| \geq (1-\alpha)|V|$ vertices in $\mathrm{Ext}\, A$ are colored independently of each other, so that

$$\begin{aligned}
&\pi\Big(|V_k(\boldsymbol{\sigma}) \cap V| \leq (1-\alpha)^2 V) \Bigm| A\Big) \\
&\leq \ \pi\Big(|V_k(\boldsymbol{\sigma}) \cap \mathrm{Ext}\, A| \leq (1-\alpha) n_A) \Bigm| A\Big) \\
&= \ \sum_{k \leq (1-\alpha)|n_A} \binom{n_A}{k} \Big(\frac{1}{q}\Big)^k \Big(1 - \frac{1}{q}\Big)^{n_A - k} \\
&\leq \ e^{-c^* L^d}
\end{aligned}$$

for some constant c^* depending on q and α. As a consequence,

$$\pi(\widehat{\Omega}^{(\alpha)}_{\rm dis} \mid A) \geq 1 - O(e^{-c^* L^d}) \text{ if } A \in \Omega^{(\alpha)}_{\rm dis}. \tag{5.7}$$

Combining (5.5) – (5.7) with Lemma 28 and the fact that $\widehat{\Omega}^{(\alpha)}_{\rm dis} \cap \widehat{\Omega}^{(\alpha)}_{\rm ord} = \emptyset$ if α is chosen small enough, we then get

$$\begin{aligned}
\widehat{\mu}(\widehat{\Omega}^{(\alpha)}_k) &= \frac{1}{q}\widehat{\mu}(\widehat{\Omega}^{(\alpha)}_{\rm ord}) \\
&= \frac{1}{q}\mu(\Omega^{(\alpha)}_{\rm ord}) + O(e^{-c^* L^d}) + O(q^{-cL}) \\
&\quad + O(q^{-c_\alpha L^{d-1}/(\log L)^2}), \\
\widehat{\mu}(\widehat{\Omega}^{(\alpha)}_{\rm dis}) &= \mu(\Omega^{(\alpha)}_{\rm dis}) + O(e^{-c^* L^d}) + O(q^{-cL}) \\
&\quad + O(q^{-c_\alpha L^{d-1}/(\log L)^2}), \\
\widehat{\mu}(\widehat{\Omega}^{(\alpha)}_{\rm Rest}) &= O(e^{-c^* L^d}) + O(q^{-cL}) \\
&\quad + O(q^{-c_\alpha L^{d-1}/(\log L)^2}).
\end{aligned}$$

We complete our proof by estimating Φ_S (see Theorem 10) for $S = \widehat{\Omega}^{(\alpha)}_1$. First notice $\widehat{\mu}(S)\widehat{\mu}(\overline{S}) \geq (1-1/q)/2q$. Since the heat bath algorithm can only change one vertex at a time, it does not make transitions between the different sets $\widehat{\Omega}^{(\alpha)}_k$, nor does it make transitions between $\widehat{\Omega}_1^{(\alpha)}$ and $\Omega_{\text{dis}}^{(\alpha)}$. Thus

$$\begin{aligned} Q(S,\overline{S}) &= \sum_{I\in\widehat{\Omega}_1^{(\alpha)},J\in\Omega_{\text{Rest}}^{(\alpha)}} \widehat{\mu}(I)P(I,J) \\ &= \sum_{I\in\widehat{\Omega}_1^{(\alpha)},J\in\Omega_{\text{Rest}}^{(\alpha)}} \widehat{\mu}(J)P(J,I) \\ &\le \widehat{\mu}(\Omega_{\text{Rest}}^{(\alpha)}). \end{aligned}$$

The theorem now follows. ∎

Bibliography

[AF] D. J. Aldous and J. Fill. Reversible Markov chains and random walks on graphs. Some chapters available at http://stat-www.berkeley.edu/pub/users/aldous/book.html. In preparation.

[Ald83] D. J. Aldous. Random walks on finite groups and rapidly mixing Markov chains. In *Séminaire de Probabilites XVII*, pages 243–297. Springer-Verlag, 1983. Lecture Notes in Mathematics 986.

[AM85] N. Alon and V.D. Milman. λ_1, isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory Series B*, 38:73–88, 1985.

[BCF$^+$99] C. Borgs, J. T. Chayes, A. M. Frieze, J. H. Kim, P. Tetali, E. Vigoda, and V. Vu. Torpid mixing of some Monte-Carlo Markov chain algorithms from statistical physics. In *40th Annual Symposium on Foundations of Computer Science*, 1999.

[BD97] R. Bubley and M.E. Dyer. Path coupling, Dobrushin uniqueness, and approximate counting. In *38th Annual Symposium on Foundations of Computer Science*, Miami Beach, FL, october 1997. IEEE.

[BDG98] R. Bubley, M. E. Dyer, and C. Greenhill. Beating the 2Δ bound for approximately counting colourings: A computer-assisted proof of rapid mixing. In *Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, San Francisco, CA, January 1998.

[BET80] R.J. Baxter, I.G. Entig, and S.K. Tsang. Hard-square lattice gas. *Journal of Statistical Physics*, 22(4):465–489, 1980.

[BK90] C. Borgs and R. Kotecký. A rigorous theory of finite-size scaling at first-order transitions. *Journal of Statistical Physics*, 61:79–110, 1990.

[BKMS91] C. Borgs, R. Kotecký, and S. Miracle-Solé. Finite-size scaling for potts models. *Journal of Statistical Physics*, 62:529–551, 1991.

[BL91] B. Bollobás and I. Leader. Edge-isoperimetric inequalities in the grid. *Combinatorica*, 11:299–314, 1991.

[BS94] J. van den Berg and J.E. Steif. Percolation and the hard-core lattice gas model. *Stochastic Processes and their Applications*, 49:179–197, 1994.

[CF99] C. Cooper and A. M. Frieze. Mixing properties of the Swendsen-Wang process on classes of graphs. *Random Structures and Algorithms*, 15, 1999.

[DFJ99] M. E. Dyer, A. M. Frieze, and M. R. Jerrum. On counting independent sets in sparse graphs. In *40th Annual Symposium on Foundations of Computer Science*, 1999.

[DFK91] M. E. Dyer, A. M. Frieze, and R. Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. *Journal of the ACM*, 38:1–17, 1991.

[DG97] M. E. Dyer and C. Greenhill. On Markov chains for independent sets. preprint, October 1997.

[DKS85] R.L. Dobrushin, J. Kolafa, and S.B. Shlosman. Phase diagram of the two-dimensional Ising antiferromagnet, computer-assisted proof. *Communications in Mathematical Physics*, 102(1):89–103, 1985.

[Dob68] R. L. Dobrushin. The problem of uniqueness of a Gibbsian random field and the problem of phase transitions. *Functional Analysis and its Applications*, 2(4):302–312, 1968.

[Dob74] R. L. Dobrushin. An investigation of Gibbs states for three-dimensional lattice systems (russian with english summary). *Teor. Verojatnost. i Primenen. (Theor. Probability Appl.)*, 18:261–279 (253–271), 1973 (1974).

[DS85] R.L. Dobrushin and S.B. Shlosman. Constructive criterion for the uniqueness of Gibbs field. In *Statistical physics and dynamical systems*, volume 10 of *Progress in Physics*, pages 347–370. Birkhäuser, Boston, 1985.

[DSC93] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible Markov chains. *The Annals of Applied Probability*, 3(3):696–730, 1993.

[ES88] R.G. Edwards and A.D. Sokal. Generalizations of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithm. *Phys. Rev. D*, 38:2009–2012, 1988.

[Fel68] W. Feller. *An introduction to probability theory and its applications Volume I (3rd ed.)*. John Wiley and Sons, Inc., 1968.

[FK72] C. Fortuin and P. Kasteleyn. On the random cluster model I: Introduction and relation to other models. *Physica*, 57:536–564, 1972.

[FMS97] A. Frigessi, F. Martinelli, and J. Stander. Computational complexity of Markov chain Monte Carlo methods for finite Markov random fields. *Biometrika*, 84(1):1–18, 1997.

[Geo88] H. O. Georgii. *Gibbs measures and phase transitions*, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter and Co., Berlin-New York, 1988.

[GHM99] H. O. Georgii, O. Häggström, and C. Maes. *Phase Transitions and Critical Phenomena*, chapter The random geometry of equilibrium phases. Academic Press, 1999.

[GJ97] V. Gore and M. R. Jerrum. The Swendsen-Wang process does not always mix rapidly. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 674–681, 1997.

[Gla63] R. Glauber. Time-dependent statistics of the Ising model. *J. Math. Phys.*, 4:294–307, 1963.

[Gra87] W. T. Grandy. *Foundations of Statistical Mechanics*, volume I: Equilibrium Theory. D. Reidel Publishing Company, 1987.

[Isr79] R. B. Israel. *Convexity in the theory of lattice gases.* Princeton University Press, Princeton, N.J., 1979. Princeton Series in Physics, With an introduction by Arthur S. Wightman.

[Jer95] M. R. Jerrum. A very simple algorithm for estimating the number of k-colorings of a low-degree graph. *Random Structures and Algorithms*, 7(2), 1995.

[JS89] M. R. Jerrum and A. J. Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18:1149–1178, 1989.

[JVV86] M. R. Jerrum, L. G. Valiant, and V. V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43(2–3):169–188, 1986.

[KRS89] A. B. Kirillov, D. C. Radulescu, and D. F. Styer. Vassertein distances in two-state systems. *Journal of Statistical Physics*, 56:931–937, 1989.

[Lig85] T. M. Liggett. *Interacting Particle Systems.* Springer-Verlag, 1985.

[LM98] J. L. Lebowitz and A. E. Mazel. Improved peierls argument for higher dimensional ising models. *Journal of Statistical Physics*, 90:1051–1059, 1998.

[LR67] O. Lanford and D. Ruelle. Integral representations of invariant states on b* algebras. *J. Math. Phys.*, 8:1460–1463, 1967.

[LS89] X-J. Li and A. D. Sokal. Rigorous lower bound on the dynamic critical exponents of the swendsen-wang algorithm. *Phys. Rev. Lett.*, 63:827–830, 1989.

[LV97] M. Luby and E. Vigoda. Approximately counting up to four (extended abstract). In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 682–687, El Paso, Texas, 1997.

[LV99] M. Luby and E. Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets. *Random Structures and Algorithms*, 15, 1999.

[Mar97] F. Martinelli. Lectures on Glauber dynamics for discrete spin models, 1997. Available from http://www.ma.utexas.edu/mp_arc.

[MO94] F. Martinelli and E. Olivieri. Approach to equilibrium of Glauber dynamics in the one phase region. I. The attractive case. *Communications in Mathematical Physics*, 161(3):447–486, 1994.

[MRR+53] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21:1087–1092, 1953.

[Pot52] R. Potts. Some generalized order-disorder transformations. *Proceedings of the Cambridge Philosophical Society*, 48:106–109, 1952.

[RS87] D.C. Radulescu and D.F. Styer. The Dobrushin-Shlosman phase uniqueness criterion and applications to hard squares. *Journal of Statistical Physics*, 49:281–295, 1987.

[RT98] D. Randall and P. Tetali. Analyzing Glauber dynamics by comparison of Markov chains. In *Third Latin American Symposium on Theoretical Informatics*, Campinas, Brazil, April 1998. UNICAMP.

[Rue69] D. Ruelle. *Statistical mechanics: Rigorous results.* W. A. Benjamin, Inc., New York-Amsterdam, 1969.

[Sin92] A. J. Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability, and Computing*, 1:351–370, 1992.

[Sin93] A. J. Sinclair. *Algorithms for Random Generation and Counting: a Markov chain approach.* Birkhäuser, Boston, 1993.

[SS97] J. Salas and A. D. Sokal. Absence of phase transition for antiferromagnetic Potts models via the Dobrushin uniqueness theorem. *Journal of Statistical Physics*, 86(3-4):551–579, February 1997.

[SW87] R. Swendsen and J.S. Wang. Non-universal critical dynamics in Monte-Carlo simulation. *Phys. Rev. Lett.*, 58:86–88, 1987.

[SZ92] D. W. Stroock and B. Zegarlinski. The logarithmic Sobolev inequality for discrete spin systems on a lattice. *Communications in Mathematical Physics*, 149:175–194, 1992.

[Vig99a] E. Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets: Part II. Technical Report TR-99-003, International Computer Science Institute, January 1999.

[Vig99b] E. Vigoda. Improved bounds for sampling colorings. In *40th Annual Symposium on Foundations of Computer Science*, 1999.

[Wal89] J. D. Walecka. *Fundaments of Statistical Mechanics: Manuscript and Notes of Felix Bloch*. Stanford University, 1989.

[Wel93] D. Welsh. *Complexity:Knots, Colourings, and Counting*. London Mathematical Society Lecture Note Series 186, 1993.

[Wu82] F. Y. Wu. The Potts model. *Reviews of Modern Physics*, 54(1), 1982.