

Tight Bounds for Mixing of the Swendsen-Wang Algorithm at the Potts Transition Point

Christian Borgs* Jennifer T. Chayes* Prasad Tetali†

August 22, 2008

Abstract

We study two widely used algorithms for the Potts model on rectangular subsets of the hypercubic lattice \mathbb{Z}^d – heat bath dynamics and the Swendsen-Wang algorithm – and prove that, under certain circumstances, the mixing in these algorithms is *torpid* or slow. In particular, we show that for heat bath dynamics throughout the region of phase coexistence, and for the Swendsen-Wang algorithm at the transition point, the mixing time in a box of side length L with periodic boundary conditions has upper and lower bounds which are exponential in L^{d-1} . This work provides the first upper bound of this form for the Swendsen-Wang algorithm, and gives lower bounds for both algorithms which significantly improve the previous lower bounds that were exponential in $L/(\log L)^2$.

1 Introduction

Convergence to equilibrium of heat bath dynamics and other dynamics for several lattice spin models of statistical mechanics has been of significant interest, for well over a decade, in probability theory, statistical physics, combinatorics and theoretical computer science. While the excellent monograph [33] provides a testament to this, many exciting new results and new techniques have since been developed. Fine examples of this development include, on the fast mixing front, results for Glauber dynamics on trees [5], for Swendsen-Wang algorithm on various classes of graphs [17], for a simple random walk on the super critical percolation cluster [4, 23]. On the slow mixing side, results on the Swendsen-Wang for the Potts model on the complete graph [24], the heat bath algorithm for the Ising model at low temperature [39], and on quasi-local algorithms for the hardcore lattice gas model at low temperature [8] form similarly interesting and technically challenging examples.

In this paper, we study two Monte Carlo Markov chains (MCMC), heat bath dynamics and the empirically more rapid Swendsen-Wang algorithm, for the q -state Potts model. Our

*Microsoft Research, 1 Memorial Drive, Cambridge, MA 02124, {borgs, jchayes}@microsoft.com

†School of Mathematics and School of Computer Science, Georgia Tech, Atlanta, GA 30332-0160, tetali@math.gatech.edu; research supported in part by the NSF Grants DMS-9800351, DMS-0401239, DMS-0701043.

work is a continuation of the work begun some time ago in collaboration with several other authors [8], where we obtained weaker bounds than those we establish here.

The point of our previous work was to relate the mixing times of MCMC in several models, including the Potts model, to the phase structure of the underlying equilibrium models. The Potts model is known to undergo a phase transition from a so-called disordered phase with a unique equilibrium state to an ordered phase with the coexistence of multiple equilibrium states. In our previous work we showed that, for the q -state Potts model on rectangular subsets of the hypercubic lattice \mathbb{Z}^d with periodic boundary conditions, heat bath dynamics is slow or *torpid* throughout the region of phase coexistence, while the Swendsen-Wang algorithm is torpid at the transition point, provided that q is large enough. There the lower bounds on the mixing time in a box of side length L with periodic boundary conditions were exponential in $L/(\log L)^2$. In this paper, we show that the mixing is even slower, obtaining essentially optimal results: both lower and upper bounds on the mixing time which are exponential in L^{d-1} .

Slowness of the Swendsen-Wang algorithm for the Potts model at the transition point was proved first on the complete graph [24]. This result initially came as a surprise to many physicists who had tacitly assumed that the algorithm was fast at all temperatures. Our previous work [8] was the first to establish such a result on subsets of the hypercubic lattice, a case which is both more physically relevant and technically much more challenging than the complete graph. To overcome these difficulties, we used some deep results from mathematical physics, which we now extend. In particular, our work brings to bear and extends, statistical physics expansion techniques for the problem of controlling the number of cutsets in graphical expansions of these models. Specifically, we use the so-called Pirogov-Sinai theory [35] from the statistical physics literature, in the form adapted to the Potts model by Borgs, Kotecký and Miracle-Sole ([10], [11]). We also use the isoperimetric inequalities of Bollobás and Leader [6], as well as a large deviations technique borrowed from [7].

For Markov chains that change the value of only a bounded number of spins – such as the heat bath algorithm, it is easy to obtain upper bounds exponential in L^{d-1} using either refined canonical path arguments as in [5] or recursive bounds on the Dirichlet form as in [14], see [33] for a review. However, for the Swendsen-Wang algorithm, which is highly non-local in the spin representation, such an upper bound is not obvious. However, it turns out that a refinement of the bounds on the Dirichlet form can be used to obtain the desired upper bound. More generally, introducing a new graph parameter which we call the “decomposition width,” we derive upper bounds for Swendsen-Wang on arbitrary graphs, which as a special case proves that Swendsen-Wang on trees is polynomial in the number of vertices for all temperatures.

The real challenge is to obtain a *lower* bound which is exponential in L^{d-1} , significantly improving the $e^{-cL/\log^2 L}$ lower bound of [8]. While the previous bound required that we consider only contours which can be embedded into \mathbb{Z}^d , this optimal lower bound requires that we deal explicitly with the topology of the torus. In particular, we must distinguish between surfaces with vanishing and non-vanishing winding numbers, which we call contours and interfaces respectively. Moreover, in order to deal optimally with the contours, we need to define an appropriate notion of exterior and interior which allows us to establish a partial order on the set of contours. This in turn is used to develop the appropriate Pirogov-Sinai

theory on the torus.

In order to state our results precisely, we need a few definitions. Let $G = (V, E)$ be a finite graph and let $\beta > 0$. For a positive integer q , let $[q] = \{1, 2, \dots, q\}$. The Gibbs measure of the (ferromagnetic) q -state Potts model on G at inverse temperature β is a measure on $[q]^V$ with density

$$\mu_G(\boldsymbol{\sigma}) = \frac{e^{-\beta H_G(\boldsymbol{\sigma})}}{Z_G}, \quad (1.1)$$

where $\boldsymbol{\sigma} \in [q]^V$ is a spin configuration. Here

$$H_G(\boldsymbol{\sigma}) = \sum_{xy \in E} (1 - \delta(\sigma_x, \sigma_y)) \quad (1.2)$$

is the Hamiltonian, and the normalization factor

$$Z_G = \sum_{\boldsymbol{\sigma} \in [q]^V} e^{-\beta H_G(\boldsymbol{\sigma})} \quad (1.3)$$

is the partition function. (In the above, δ denotes the Kronecker delta function.)

For a finite $\Lambda \subset \mathbb{Z}^d$ we define the measure $\mu_{\Lambda,1}$ with the “1-boundary condition” by setting all the spins at the external boundary of Λ to 1. Explicitly, for $\boldsymbol{\sigma} \in \{1, 2, \dots, q\}^\Lambda$, let

$$H_{\Lambda,1}(\boldsymbol{\sigma}) = H_{G[\Lambda]}(\boldsymbol{\sigma}) + \sum_{\substack{x \in \Lambda \\ y \in \Lambda^c \\ |x-y|=1}} (1 - \delta(\sigma_x, 1)),$$

where $G[\Lambda]$ is the induced subgraph of Λ and $|\cdot|$ is the l_1 distance.

Then $\mu_{\Lambda,1}$ is the probability measure with density

$$\mu_{\Lambda,1}(\boldsymbol{\sigma}) = \frac{e^{-\beta H_{\Lambda,1}(\boldsymbol{\sigma})}}{Z_{\Lambda,1}}, \quad (1.4)$$

where $Z_{\Lambda,1}$ is the appropriate normalization factor.

The infinite volume magnetization is defined as $M(\beta) = \lim_{L \rightarrow \infty} M_{\Lambda_L}(\beta)$ where $\Lambda_L = \{1, \dots, L\}^d$ and

$$M_{\Lambda}(\beta) = \frac{1}{|\Lambda|} \sum_{x \in \Lambda} \left(\mu_{\Lambda,1}(\{\sigma_x = 1\}) - \frac{1}{q} \right).$$

By standard correlation inequalities, the limit $\lim_{L \rightarrow \infty} M_{\Lambda_L}(\beta)$ is known to exist and to be monotone nondecreasing in β . The transition point is defined by

$$\beta_0 = \beta_0(\mathbb{Z}^d) = \inf\{\beta : M(\beta) > 0\}.$$

The model is known to have a unique (infinite-volume) Gibbs state for $\beta < \beta_0$ and q extremal translation-invariant Gibbs states for $\beta > \beta_0$. For q small enough (depending on d), the model is believed to have a unique Gibbs state at β_0 , whereas for q large enough, it is known to have $q + 1$ extremal translation-invariant Gibbs states at $\beta_0 = \frac{1}{d} \log q + O(q^{-1/d})$.

Next we define the mixing time of a finite Markov chain with state space Ω . Let P denote the transition probability matrix of an irreducible Markov chain with the (unique) stationary measure μ . The (variational) mixing time of such a chain is defined as

$$\tau = \min \left\{ t : d(t) \leq \frac{1}{2e} \right\}, \quad (1.5)$$

where

$$d(t) = \max_{\sigma \in \Omega} \max_{A \subset \Omega} \left| \mu(A) - \sum_{\sigma' \in A} P^t(\sigma, \sigma') \right|.$$

In this paper, we will consider several Markov chains for the Potts model on the torus

$$T_{L,d} = (\mathbb{Z}/L\mathbb{Z})^d.$$

The chains we consider are the heat bath (or Glauber) dynamics, and the presumably much faster SW algorithm; see Section 2.1 of the definition of these chains. We denote the mixing times of these algorithms for the q -state Potts model on the torus $T_{L,d}$ by $\tau^{\text{HB}} = \tau^{\text{HB}}(T_{L,d})$ and $\tau^{\text{SW}} = \tau^{\text{SW}}(T_{L,d})$.

Theorem 1.1 *There are universal constants $k_1, k_2 < \infty$ such that, for $\beta > 0$, $d \geq 2$ and any positive integer L , the following bounds hold*

$$\tau^{\text{HB}}(T_{L,d}) \leq e^{(k_1+k_2\beta)L^{d-1}} \quad (1.6)$$

$$\tau^{\text{SW}}(T_{L,d}) \leq e^{(k_1+k_2\beta)L^{d-1}}. \quad (1.7)$$

In order to prove this, in Section 3, we introduce the “partition width” of a graph, a notion that may be of independent interest. As a corollary of our proof, we also obtain that the mixing time of SW on a tree with n vertices, maximum degree d_{\max} , and depth $O(\log n)$ is bounded by $n^{1+O(\beta d_{\max})}$, see Corollary 3.2 for the precise statement. This generalizes the result of [5], which gives polynomial mixing for the HB algorithm on trees, to the SW algorithm.

Theorem 1.2 *Let $d \geq 2$. Then there exists a constant $k_3 = k_3(d) > 0$ such that, for q and L sufficiently large, the following bounds hold:*

$$\tau^{\text{HB}}(T_{L,d}) \geq e^{k_3\beta L^{d-1}} \quad \text{for all } \beta \geq \beta_0(\mathbb{Z}^d) \quad (1.8)$$

$$\tau^{\text{SW}}(T_{L,d}) \geq e^{k_3\beta L^{d-1}} \quad \text{for all } \beta = \beta_0(\mathbb{Z}^d). \quad (1.9)$$

Very roughly speaking, the reason for the heat bath lower bound is that this algorithm cannot move quickly among (the finite analogs of) the q translation-invariant extremal states present for $\beta > \beta_0$. On the other hand, the Swendsen-Wang algorithm can move quickly among these q states, but cannot move quickly between these q states and the one additional translation-invariant extremal state present at $\beta = \beta_0$. Whenever the algorithm cannot move quickly between states, in order for the system to mix, it must pass through a configuration with a “separating surface” of size at least L^{d-1} between the relevant states.

The organization of this paper is as follows. In Section 2, we define the algorithms and the necessary notions from the theory of MCMC. In Section 3, we introduce the notion of partition width and establish the upper bound on the mixing time. In order to obtain the corresponding lower bound, we need some preparation: in Section 4 we construct the contour representation of the model, while Section 5 establishes the required geometric properties of contours and interfaces. This section can perhaps be skipped on first reading. In the next section, we state the necessary bounds from Pirogov-Sinai (to be proved in the appendix), and use these to establish two key estimates needed for the main proof: a bound on the probability of interfaces, and a suitable large deviation bound. Using these bounds, we then prove our main result in Section 7.

The reader only interested in the main flavor of our proofs should perhaps start with Section 2.3, where we explain the main proof strategy, and then immediately jump to Section 7, glancing back at Section 4 and Section 6.1 as necessary.

2 MCMC Preliminaries

2.1 Algorithms for the Potts Model

There are several MCMC algorithms that are used to generate a random sample from the distribution corresponding to the ferromagnetic Potts model. The heat bath is perhaps the simplest such Markov chain. Its transitions are as follows: Choose a vertex at random, and modify the spin of that vertex by choosing from the distribution conditional on the spins of the other vertices remaining the same. In contrast, the Swendsen-Wang algorithm can alter the spins on many vertices in each iteration.

Throughout this section $G = (V, E)$ is a fixed finite graph. For a subgraph $\tilde{G} = (\tilde{V}, \tilde{E})$ of G , we denote the set of (connected) components of \tilde{G} by $\mathcal{C}(\tilde{G}) = \mathcal{C}(\tilde{V}, \tilde{E})$, and its cardinality by $c(\tilde{G}) = c(\tilde{V}, \tilde{E})$. Finally, for a spin configuration $\sigma \in [q]^V$, let $E(\sigma)$ be the set of “monochromatic edges” $xy \in E$ with $\sigma_x = \sigma_y$.

Heat Bath: From a spin configuration $\sigma \in [q]^V$, we construct a new configuration σ' as follows:

HB1 Choose v uniformly at random from V .

HB2 Take $\sigma'_w = \sigma_w$, for all $w \in V \setminus \{v\}$, and change σ_v to σ'_v with probability

$$\mu_G(\sigma'_v | \sigma_{V \setminus \{v\}}) = \frac{\exp\left\{\beta \sum_{\substack{w \in V \\ vw \in E}} \delta(\sigma'_v, \sigma_w)\right\}}{q \sum_{k=1}^q \exp\left\{\beta \sum_{\substack{w \in V \\ vw \in E}} \delta(k, \sigma_w)\right\}}.$$

For future reference we denote the transition matrix of this chain by

$$P_G^{\text{HB}}(\sigma, \sigma') = \frac{1}{|V|} \sum_{v \in V} \left(\mu_G(\sigma'_v | \sigma_{V \setminus \{v\}}) \prod_{w \neq v} \delta(\sigma'_w, \sigma_w) \right).$$

In practice, an alternative method, the Swendsen-Wang algorithm [38], is often preferred.

Swendsen-Wang Algorithm: For $\sigma \in [q]^V$:

SW1 Let $E(\sigma) \subset E$ be the set of monochromatic edges. Delete each edge of $E(\sigma)$ independently with probability $1 - p$, where $p = 1 - e^{-\beta}$. This gives a random subset $A \subset E(\sigma)$.

SW2 The graph (V, A) consists of connected components. For each component, choose a color (spin) k uniformly at random from $[q]$, and for all vertices v within that component, set $\sigma'_v = k$.

Again for future reference, we denote the transition matrix of this chain by

$$P_G^{\text{SW}}(\sigma, \sigma') = \sum_{A \subset E(\sigma)} p^{|A|} (1-p)^{|E(\sigma) \setminus A|} \prod_{C \in \mathcal{C}(V, A)} \left(\frac{1}{q} \sum_{k=1}^q \prod_{v \in V(C)} \delta(\sigma'_v, k) \right).$$

The Swendsen-Wang algorithm was motivated by the equivalence of the ferromagnetic q -state Potts model and the *random cluster* model of Fortuin and Kasteleyn [22], which we now describe. Fortuin and Kasteleyn realized that the Potts model partition function Z_G and expectations with respect to the measure μ_G can be rewritten in terms of a weighted graph model on spanning subgraphs $(V, A) \subset G$ with weights

$$\nu_G(A) = \frac{1}{Z_G} p^{|A|} (1-p)^{|E \setminus A|} q^{c(V, A)}. \quad (2.1)$$

The relationship between the two models is elucidated in a paper by Edwards and Sokal [21]. The Potts and random cluster models are defined on a joint probability space $[q]^V \times 2^E$. The joint probability $\pi(\sigma, A)$ is defined by

$$\pi_G(\sigma, A) = \frac{1}{Z_G} p^{|A|} (1-p)^{|E \setminus A|} \prod_{xy \in A} \delta(\sigma_x, \sigma_y). \quad (2.2)$$

By summing over σ or A we see that the marginal distributions are ν_G or μ_G respectively.

A step $\sigma \rightarrow \sigma'$ of the Swendsen-Wang algorithm can be seen as (i) choose a random A' according to $\pi_G(\sigma, \cdot) / \mu_G(\sigma)$ and then (ii) choose a random σ' according to $\pi(\cdot, A') / \nu_G(A')$. After Step SW1, we say that we are in the random cluster representation of the chain.

2.2 Mixing Time and Related Quantities

Throughout this section, let P be the transition matrix of an irreducible Markov chain on a finite state space Ω , let μ be the stationary distribution of P , i.e., $\mu(\sigma') = \sum_{\sigma \in \Omega} \mu(\sigma) P(\sigma, \sigma')$, and let $\mu_{\min} = \min_{\sigma \in \Omega} \mu(\sigma)$. We denote the mixing time, defined in (1.5), by $\tau(P)$, and define the inverse gap (or eigentime) $\tilde{\tau}(P)$ as

$$\tilde{\tau}(P) = \sup_g \frac{\text{Var } g}{\mathcal{E}(g, g)}, \quad (2.3)$$

where the supremum is over all real valued functions g on Ω with $\text{Var } g > 0$. Here as usual,

$$\text{Var } g = \text{Var}_\mu g = \frac{1}{2} \sum_{\sigma, \sigma'} (g(\sigma) - g(\sigma'))^2 \mu(\sigma) \mu(\sigma'),$$

and

$$\mathcal{E}(g, g) = \mathcal{E}_P(g, g) = \frac{1}{2} \sum_{\sigma, \sigma'} (g(\sigma) - g(\sigma'))^2 \mu(\sigma) P(\sigma, \sigma').$$

If P is reversible, $\tilde{\tau}(P)$ is just $(1 - \beta_2(P))^{-1}$, where $\beta_2(P)$ is the second largest eigenvalue of P .

It is well known that the inverse gap can be bounded above in terms of the mixing times; if the chain is lazy, i.e., if the minimal self-loop probability $\min_\sigma P(\sigma, \sigma)$ is uniformly bounded from below, a bound in the opposite direction is also not very hard to prove, see, e.g., [2]. However, the SW chain is not lazy. Instead of the standard upper bound on $\tau(P)$ in terms of $\tilde{\tau}(P)$, we therefore use a bound from [34]. For reversible chains, this bound gives

$$\tau(P) \leq \tilde{\tau}(P^2) \log\left(\frac{e^2}{\mu_{\min}}\right), \quad (2.4)$$

where, as usual, $P^2(\sigma, \sigma') = \sum_{\sigma'' \in \Omega} P(\sigma, \sigma'') P(\sigma'', \sigma')$, denotes the transition matrix of the two-step chain.

We will also need an identity for the mixing time of a product chain. Let Ω_1, Ω_2 be finite sets, and let P_1, P_2 be the transition matrices of two irreducible Markov chains on Ω_1 and Ω_2 with stationary distributions μ_1 and μ_2 respectively. Let $\Omega_1 \times \Omega_2$ be the set of all pairs $\sigma = (\sigma^{(1)}, \sigma^{(2)})$ with $\sigma^{(1)} \in \Omega_1$ and $\sigma^{(2)} \in \Omega_2$. Then the product chain is defined as the chain with the transition matrix

$$(P_1 \times P_2)(\sigma, \tilde{\sigma}) = P_1(\sigma^{(1)}, \tilde{\sigma}^{(1)}) P_2(\sigma^{(2)}, \tilde{\sigma}^{(2)}). \quad (2.5)$$

Let

$$(\mu_1 \times \mu_2)(\sigma) = \mu_1(\sigma^{(1)}) \mu_2(\sigma^{(2)}).$$

If both P_1 and P_2 have non-negative eigenvalues, then $P_1 \times P_2$ has non-negative eigenvalues and $\beta_2(P_1 \times P_2) = \max\{\beta_2(P_1), \beta_2(P_2)\}$. Using this fact, one immediately shows that $P_1 \times P_2$ is irreducible with stationary distribution $\mu_1 \times \mu_2$ and obeys the bound

$$\tilde{\tau}(P_1 \times P_2) = \max\{\tilde{\tau}(P_1), \tilde{\tau}(P_2)\}. \quad (2.6)$$

For our lower bounds, we use the notion of conductance and its relation to the mixing time. Setting

$$Q(S, S') = \sum_{\sigma \in S} \sum_{\sigma' \in S'} \mu(\sigma) P(\sigma, \sigma'),$$

the conductance of a set $S \subset \Omega$ can be defined as

$$\Phi_S = \frac{Q(S, S^c)}{\mu(S) \mu(S^c)}. \quad (2.7)$$

Finally the conductance of a Markov chain with the transition matrix P is

$$\Phi(P) = \min_{S:0<\mu(S)<1} \Phi_S. \quad (2.8)$$

The mixing time can be easily bounded from below in terms of the conductance, see [7],[15], [19]:

$$\tau(P) \geq \frac{e-1}{e} \frac{1}{\Phi(P)}. \quad (2.9)$$

2.3 Proof Strategy

In order to prove Theorem 1.1, we will want to give an upper bound on the inverse gap $\tilde{\tau}$ defined in (2.3). To this end, it will be convenient to consider the SW algorithm on a general graph G . We then iteratively partition the set V of vertices of G into two sets V_1 and V_2 , and show that the inverse gap of $(P_G^{\text{SW}})^2$ is bounded by the inverse gap of the product chain for SW on the two induced graphs $G[V_1]$ and $G[V_2]$, times a factor which is exponential in the number of edges with one endpoint in V_1 and one endpoint in V_2 . With the help of (2.6) and (2.4), this allows us to bound the mixing time of SW by a number which is exponential in a quantity we call the partition width of the graph G . Applied to the torus $T_{L,d}$, this gives a bound which is exponential in L^{d-1} , and applied to a tree, this will give a bound which is polynomial in the number of vertices. The bound for the HB algorithm follows a similar strategy.

To prove Theorem 1.2, we will use the lower bound (2.9). To this end, we will want to find a set of spin configurations S such that Φ_S is exponentially small in L^{d-1} . Recalling that the Potts model at the transition point exhibits the coexistence of q ordered phases and one translation invariant phase, we will want to exploit the fact that the SW algorithm cannot transition easily between (the finite volume analogue of) the ordered phases and the disordered phase. To make this precise, we define S to be the set $S = \{\sigma : |E(\sigma)| \geq (1-\alpha)dL^d\}$, where $\alpha > 0$ is a small constant, say $\alpha = 1/3$. Thus S consists of the configurations whose set of monochromatic edges form almost all of E .

For large q the inverse transition temperature β_0 is large as well, implying that at $\beta = \beta_0$, the probability of deleting an edge in the first step of the SW algorithms is small; starting from a configuration $\sigma \in S$ it is therefore unlikely that after one step of the Markov chain, the new configuration σ' is such that the number of edges in $E(\sigma')$ is much smaller than $(1-\alpha)dL^d$. (The probability that it is smaller than, say αdL^d , is actually exponentially small in L^d). But it is also unlikely that a configuration $\sigma' \in S^c$ has a number of monochromatic edges which is larger than αdL^d , since both requirements together imply that the number of edges lies between αdL^d and $(1-\alpha)dL^d$. This corresponds neither to an ordered phase (which would have more than $(1-\alpha)dL^d$ monochromatic edges), nor to a disordered phase (which would have less than αdL^d monochromatic edges), and thus to a configuration which has low probability. Thus with high probability a configuration $\sigma \in S$ leads to a new configuration σ' which is in S as well, showing that ϕ_S is small.

To make this quantitative, we will have to show that the weight of the configurations in $S_0 = \{\sigma' : \alpha dL^d \leq |E(\sigma')| \leq (1-\alpha)dL^d\}$ is exponentially small in L^{d-1} . To this end, we will first switch to the FK representation (2.1), and then describe an edge configuration $A \subset E$

in terms of a hypersurface separating regions with edges in A from regions with edges in $E \setminus A$. We then decompose this hypersurface into connected components, some of which to be called contours, and others to be called interfaces. While contours can have small or large size, interfaces will always have size at least L^{d-1} .

To prove our desired bounds on $\mu(S_0)$, we will show that the probability of a configuration with an interface is exponentially small in the size of the interface, leaving us with the analysis of configurations without interfaces. These will come in two classes: configurations describing perturbations of $A = E$ (we denote the set of these configurations by Ω_{ord}), and configurations describing perturbations of $A = \emptyset$ (to be denoted by Ω_{dis}). Our last step then consists of a large deviations bound showing that with high probability, for configurations $A \in \Omega_{\text{ord}}$, the common exterior of a set of contours has size at least $(1 - \frac{1}{2}\alpha)L^d$, and similarly for configurations $A \in \Omega_{\text{dis}}$. This will imply that with high probability, the set of monochromatic edges of an ordered configuration σ has at least $(1 - \alpha)dL^d$ edges, while that of a disordered configuration has at most αdL^d edges. Put together, these estimates give the desired bound for $\mu(S_0)$.

Our approach differs in several aspects from the approach taken in [8], which led to a conductance bound that was exponentially small in $L/(\log^2 L)$.

First, the bounds in [8] relied on a combination of Pirogov-Sinai theory with the finite-size scaling theory developed in [10] and [11], where contours were by definition objects that could be embedded in \mathbb{R}^d . This allowed for an immediate application of standard Pirogov-Sinai results, but produced error terms that were only exponential in L , which is not strong enough for our current purpose. To avoid this problem, we define contours in a purely topological manner, by requiring that their \mathbb{Z}_2 winding number with respect to the torus is equal to zero. This has the advantage that the objects which cannot be classified as contours (we call them interfaces) must have size at least L^{d-1} , since all surfaces of smaller size have winding number zero. The price we have to pay is that the set of contours now contains objects like long tubes winding around the torus which cannot be embedded into \mathbb{R}^d , preventing us from applying the standard Pirogov-Sinai machinery as more or less a black box.

Instead, we will show that Pirogov-Sinai theory does not really rely on the topology of \mathbb{R}^d , but rather on the implied structure of partial order on contours. More precisely, it relies on the fact that for any set of pairwise non-overlapping contours and interfaces, this partial order leads to a Hasse diagram that is a forest – this is expressed in Lemma 5.5, see also Lemma 4.3 and Definition 4.4.

Second, we will use a large deviation bound obtained by adding an artificial magnetic field to the contour model, see Section 6.3 for details. This turns out to be much more efficient than the iterative, combinatorial arguments from [8], allowing us to improve a bound that is exponentially small in $L/(\log L)^2$ to one which is exponentially small in L^{d-1} .

3 Upper Bound on Mixing Time

In order to prove our upper bound on the mixing time, it will be convenient to prove a more general theorem, involving a new notion, called the “partition width” of a graph, which we expect may be of independent interest. We need some notation.

Given a graph $G = (V, E)$, we define a hierarchical partition \mathcal{P} of V by first dividing V into two non-empty subsets V_1, V_2 , and then successively subdividing each set with more than one element into two further subsets until all sets contain only one element. Each such partition can be described by a rooted binary tree as follows: the vertices are subsets of V , with the root being V , the leaves being the singletons $\{x\}$, $x \in V$. In addition, we have the constraint that for any vertex V_i with children V_{i1}, V_{i2} , we have $V_i = V_{i1} \cup V_{i2}$ and $V_{i1} \cap V_{i2} = \emptyset$. If V_i has children V_{i1} and V_{i2} , we define the weight $w_{\mathcal{P}}(V_i)$ of V_i as the number of edges between V_{i1} and V_{i2} in G ; if V_i is a leaf, we set its weight to zero.

For $x \in V$, we now define the separation cost $\text{sep}_{\mathcal{P}}(x)$ of x as the sum of all vertex weights along the path from the root to x . The cost of a partition \mathcal{P} is then defined as $\text{sep}(\mathcal{P}) = \max_{x \in V} \text{sep}_{\mathcal{P}}(x)$, and the *partition width* of G is defined as $\omega(G) = \min_{\mathcal{P}} \text{sep}(\mathcal{P})$.

Theorem 3.1 *For any finite graph $G = (V, E)$, the mixing time of the SW algorithm obeys the upper bound*

$$\tau(P_G^{\text{SW}}) \leq e^{5\beta\omega(G)} (2 + |V| \log 2 + \beta|E|).$$

Before proving the theorem, we state (and prove) the following corollary, which illustrates the usefulness of our notion of partition width.

Corollary 3.2 *Let Λ be a rectilinear subset of \mathbb{Z}^d , let $T_{L,d}$ be the d -dimensional torus of side length L , and let T be a tree on n vertices with maximum degree d_{\max} and depth $c \log n$. Then*

$$\tau(P_{\Lambda}^{\text{SW}}) \leq e^{45\beta|A(\Lambda)|} (2 + (\log 2 + d\beta)|\Lambda|), \quad (3.1)$$

$$\tau(P_{T_{L,d}}^{\text{SW}}) \leq e^{75\beta L^{d-1}} (2 + (\log 2 + d\beta)L^d), \quad (3.2)$$

$$\tau(P_T^{\text{SW}}) \leq n^{5c\beta d_{\max}} (2 + (\log 2 + \beta)n). \quad (3.3)$$

Here $A(\Lambda)$ is the volume of Λ divided by the minimal side length.

Note that the second bound of corollary implies the bound (1.7) in Theorem 1.1. The bound (1.6) of this theorem can either be proved by generalizing Theorem 3.1 to the heat bath algorithm (the proof is actually easier for this case), or by using the canonical path techniques of [5].

Proof of Corollary 3.2. Let $G = (V, E)$ be an arbitrary finite graph, and let $V = V_1 \cup V_2$ be a decomposition of V into two disjoint subsets. Using the definition of the partition width, one easily verifies that

$$\omega(G) \leq |E_{12}| + \max\{\omega(G[V_1]), \omega(G[V_2])\}, \quad (3.4)$$

where E_{12} is the set of edges between V_1 and V_2 . Using this bound, it is easy to verify by induction, that for a tree T of maximal degree d_{\max} and depth D , one has

$$\omega(T) \leq d_{\max} D,$$

which in turn gives the bound in the corollary for trees.

We are thus left with proving upper bounds for the partition width for cubic subsets of \mathbb{Z}^d and the torus $T_{L,d}$. We start with a rectilinear subset of side-lengths $L_1 \geq L_2 \geq \dots \geq L_d$, which we denote by $[L_1, \dots, L_d]$. To this end, we first note that

$$\omega(G') \leq \omega(G),$$

whenever G' is a spanning subgraph of G , implying that $\omega([L_1, \dots, L_d])$ is non-decreasing in the side-lengths L_i . Consider a set of sidelength $L_1 \geq \dots \geq L_d$ with $L_1 \geq 2$. Using the bound (3.4), the monotonicity of $\omega([L_1, \dots, L_d])$ and the fact that $\lceil L_i/2 \rceil \leq 2L_i/3$, whenever $L_i \geq 2$, we then bound

$$\begin{aligned} \omega([L_1, \dots, L_d]) &\leq L_2 \cdots L_d + \omega\left(\left[\lceil \frac{L_1}{2} \rceil, L_2, \dots, L_d\right]\right) \\ &\leq L_2 \cdots L_d + \lceil \frac{L_1}{2} \rceil L_3 \cdots L_d + \omega\left(\left[\lceil \frac{L_1}{2} \rceil, \lceil \frac{L_2}{2} \rceil, \dots, L_d\right]\right) \\ &\leq \dots \\ &\leq \left(1 + \frac{2}{3} + \left(\frac{2}{3}\right)^2 + \dots + \left(\frac{2}{3}\right)^{k-1}\right) L_1 \cdots L_{d-1} \\ &\quad + \omega\left(\left[\lceil \frac{L_1}{2} \rceil, \lceil \frac{L_2}{2} \rceil, \dots, \lceil \frac{L_d}{2} \rceil\right]\right) \\ &\leq 3L_1 \cdots L_{d-1} + \omega\left(\left[\lceil \frac{L_1}{2} \rceil, \lceil \frac{L_2}{2} \rceil, \dots, \lceil \frac{L_d}{2} \rceil\right]\right). \end{aligned} \tag{3.5}$$

where k is the smallest i such that $L_i \geq 2$. Using this bound, it is now easy to prove by induction that

$$\omega([L_1, \dots, L_d]) \leq 9A([L_1, \dots, L_d]),$$

implying the desired bound for the inverse gap on rectilinear sets.

Next, we would like to bound the partition width of the torus $T_{L,d}$ using the just established bound for rectilinear sets. To this end we successively cut the torus in the d different coordinate directions, proceeding as in the proof above. Here, however, since we have a torus rather than a rectilinear set, we need two cuts rather than one cut in each direction to obtain two components. Keeping this in mind, we get

$$\begin{aligned} \omega(T_{L,d}) &\leq 2 \left(1 + \frac{2}{3} + \left(\frac{2}{3}\right)^2 + \dots + \left(\frac{2}{3}\right)^{d-1}\right) L_1 \cdots L_{d-1} \\ &\quad + \omega\left(\left[\lceil \frac{L}{2} \rceil, \lceil \frac{L}{2} \rceil, \dots, \lceil \frac{L}{2} \rceil\right]\right) \\ &\leq 6L^{d-1} + 9L^{d-1} = 15L^{d-1}, \end{aligned} \tag{3.6}$$

which implies the desired bound on the mixing time. \square

The proof of Theorem 3.1 is based on the following lemma.

Lemma 3.3 *Let $G = (V, E)$, let $P_G = P_G^{\text{SW}}$ and $G_0 = (V, E_0)$, where E_0 is an arbitrary subset of E . Then*

$$\tilde{\tau}(P_G^2) \leq \tilde{\tau}(P_{G_0}^2) e^{5\beta|E \setminus E_0|}.$$

Proof. Recall that a single transition of the SW dynamics consists of two steps. Given a current Potts configuration σ , the first step identifies $e^{-\beta}$ connected components of color classes and performs random edge deletion with probability $e^{-\beta}$ independently for each monochromatic edge. We denote the resulting configuration by E' . The second step assigns colors independently for each new connected component (cluster) in E' , resulting in a new Potts configuration σ' .

Let $E(\sigma) = \{xy \in E : \sigma_x = \sigma_y\}$. Let G and G_0 be as in the statement of the lemma, and let $E_1 = E \setminus E_0$. Then

$$\begin{aligned} P_G(\sigma, \sigma') &= \sum_{E' \subseteq E(\sigma)} (1 - e^{-\beta})^{|E'|} e^{\beta|E(\sigma) \setminus E'|} \prod_{C \in \mathcal{C}(V, E')} \frac{1}{q} \prod_{x, y \in C} \delta(\sigma'_x, \sigma'_y) \\ &\geq e^{-\beta|E_1|} \sum_{E' \subseteq E(\sigma) \setminus E_1} (1 - e^{-\beta})^{|E'|} e^{\beta|\{E(\sigma) \setminus E_1\} \setminus E'|} \prod_{C \in \mathcal{C}(V, E')} \frac{1}{q} \prod_{x, y \in C} \delta(\sigma'_x, \sigma'_y) \\ &\geq e^{-\beta|E_1|} P_{G_0}(\sigma, \sigma'), \end{aligned}$$

implying that

$$P_G^2(\sigma, \sigma') \geq e^{-2\beta|E_1|} P_{G_0}^2(\sigma, \sigma'). \quad (3.7)$$

Next we observe that, by the definition of H_G , we have

$$e^{-\beta H_{G_0}(\sigma)} e^{-\beta|E_1|} \leq e^{-\beta H_G(\sigma)} \leq e^{-\beta H_{G_0}(\sigma)}.$$

implying that

$$\mu_{G_0}(\sigma) e^{-\beta|E_1|} \leq \mu_G(\sigma) \leq \mu_{G_0}(\sigma) e^{\beta|E_1|}.$$

Combined with (3.7) and the definition of variance and the Dirichlet form, this proves that

$$\frac{\text{Var}_{\mu_G} g}{\mathcal{E}_{P_G^2}(g, g)} \leq \frac{e^{2\beta|E_1|} \text{Var}_{\mu_{G_0}} g}{e^{-3\beta|E_1|} \mathcal{E}_{P_{G_0}^2}(g, g)} \leq e^{5\beta|E_1|} \frac{\text{Var}_{\mu_{G_0}} g}{\mathcal{E}_{P_{G_0}^2}(g, g)},$$

and hence

$$\tilde{\tau}(P_G^2) = \sup_g \frac{\text{Var}_{\mu_G} g}{\mathcal{E}_{P_G^2}(g, g)} \leq e^{5\beta|E_1|} \tilde{\tau}(P_{G_0}^2).$$

□

Having established Lemma 3.3, we are now ready to prove Theorem 3.1.

Proof of Theorem 3.1. Given a graph $G = (V, E)$ and a decomposition of V into two disjoint subsets V_1 and V_2 , let E_{12} be the set of edges in E that join V_1 and V_2 , and let $E_0 = E \setminus E_{12}$. Let $G_1 = G[V_1]$, $G_2 = G[V_2]$ and $G_0 = (V, E_0)$. Observing that no edge in G_0 joins V_1 and V_2 , we clearly have $P_{G_0} = P_{G_1} \times P_{G_2}$ and thus $P_{G_0}^2 = P_{G_1}^2 \times P_{G_2}^2$. Combining Lemma 3.3 with the identity (2.6) for $P_{G_1}^2$ and $P_{G_2}^2$, we thus have

$$\tilde{\tau}(P_G^2) \leq \max\{\tilde{\tau}(P_{G_1}^2), \tilde{\tau}(P_{G_2}^2)\} e^{5\beta|E_{12}|}.$$

Applying this bound recursively for the decompositions in a hierarchical partition \mathcal{P} of G , we obtain that

$$\tilde{\tau}(P_G^2) \leq e^{5\beta\omega(G)} \prod_{x \in V} \tilde{\tau}(P_{G[\{x\}]}^2).$$

Since the inverse gap for the SW algorithms on a single site is one, we get $\tilde{\tau}(P_G^2) \leq e^{5\beta\omega(G)}$. Combined with the bound (2.4), this proves the theorem. □

4 Contour Representation

In this section, we derive a representation for the Potts model in terms of contours and interfaces. This representation is a modified version of the representation of [11]. We use $T = T_{L,d}$ to denote the d -dimensional discrete torus of sidelength L , with vertex set $V = V_{L,d} = (\mathbb{Z}/L\mathbb{Z})^d$ and edge set $E = E_{L,d}$, and Ω to denote the SW configuration space 2^E , i.e., the set of all edge configurations $A \subset E$.

We start from the random cluster representation (2.1). Given $A \subset E$, let $V(A) = \cup_{\{x,y\} \in A} \{x, y\}$, and let $\tilde{G}(A) = (V(A), A)$. Recalling the definition of $c(A)$ as the number of connected components of the graph $G(A) = (V, A)$, we introduce $\tilde{c}(A)$ as the number of connected components of $\tilde{G}(A)$. Observe that $c(A) = \tilde{c}(A) + |V \setminus V(A)|$ and

$$2|A| = 2d|V(A)| - |\delta_1 A| - 2|\delta_2 A|,$$

where

$$\delta_k A = \{\{x, y\} \in E \setminus A; |\{x, y\} \cap V(A)| = k\} \quad k = 1, 2.$$

Using the notation $\delta A = \{\{x, y\} \in E \setminus A; |\{x, y\} \cap V(A)| > 0\}$ and $\|\delta A\| = |\delta_1 A| + 2|\delta_2 A|$, we rewrite the weight $w(A) = Z_T \nu_T(A)$ of a configuration A in (2.1) as

$$\begin{aligned} w(A) &= \left((1-p)^d q \right)^{|V \setminus V(A)|} p^{d|V(A)|} \left(\frac{1-p}{p} \right)^{\|\delta A\|/2} q^{\tilde{c}(A)} \\ &= q^{\tilde{c}(A)} e^{-e_{\text{dis}}|V \setminus V(A)|} e^{-e_{\text{ord}}|V(A)|} e^{-\kappa \|\delta A\|} \end{aligned} \quad (4.1)$$

where

$$e_{\text{dis}} = -\log\left((1-p)^d q\right) = d\beta - \log q, \quad (4.2)$$

$$e_{\text{ord}} = -d \log p = -d \log(1 - e^{-\beta}) = O(e^{-\beta}), \quad (4.3)$$

$$\kappa = -\frac{1}{2} \log\left(\frac{1-p}{p}\right) = \frac{1}{2} \log(e^\beta - 1) = \frac{\beta}{2} + O(e^{-\beta}). \quad (4.4)$$

This already shows that there are three regions of interest: for $\beta \ll \frac{1}{d} \log q$, the dominant configurations are those with most vertices belonging to $V \setminus V(A)$, i.e., most vertices are isolated, corresponding to a “disordered high-temperature phase;” for $\beta \gg \frac{1}{d} \log q$, the dominant configurations have most vertices in the “ordered region” $V(A)$ corresponding to an “ordered low-temperature phase;” and for $\beta \approx \frac{1}{d} \log q$ and q (and hence κ) large, the dominant configurations fall into two classes, one with mostly isolated vertices and the other with most vertices in $V(A)$. As we will see, even the SW algorithm has difficulties transitioning between these two classes, which leads to slow mixing at $\beta_0 = \frac{1}{d} \log q + O(q^{-1/d})$.

We will decompose δA into several pieces called “interfaces” and “contours”. More precisely, we will first “fatten” the set A into a subset $\mathbf{V}(A)$ of the continuum torus $\mathbf{V} = (\mathbb{R}/(L\mathbb{Z}))^d$ and then decompose the boundary of $\mathbf{V}(A)$ into components, some of which will be called interfaces, and some of which will be called contours.

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 $\mathbf{c} \subset \mathbf{V}$ with corners in

V occupied if all edges e with $\mathbf{e} \subset \mathbf{c}$ are in A . We then define the set $\mathbf{V}(A) \subset \mathbf{V}$ as the $1/4$ -neighborhood of the union of all occupied k -dimensional hypercubes, $k = 1, \dots, d$, i.e., $\mathbf{V}(A) = \{x \in \mathbf{V} : \exists \mathbf{c} \text{ occupied, such that } \text{dist}(x, \mathbf{c}) \leq 1/4\}$, where $\text{dist}(x, y)$ is the ℓ_∞ -distance between the two points x and y in the torus \mathbf{V} and $\text{dist}(x, \mathbf{c}) = \inf_{y \in \mathbf{c}} \text{dist}(x, y)$. Note that the set $\mathbf{V}(A)$ is a union of cubes of side-length $1/2$ with centers in $V_{1/2} = (\frac{1}{2}\mathbb{Z}/L\mathbb{Z})^d$, and that the set $V(A)$ defined at the beginning of this section is just the intersection of $\mathbf{V}(A)$ with the vertex set V of the discrete torus T .

Alternatively, one can define the set $\mathbf{V}(A)$ by constructing its complement, the “disordered region” $\mathbf{V} \setminus \mathbf{V}(A)$ as follows: Let $(E \setminus A)^*$ be the set of $d - 1$ dimensional unit hypercubes dual to the edges in $E \setminus A$. It is then easy to see that the union of the open $3/4$ -neighborhood of $V \setminus V(A)$ and the open $1/4$ -neighborhood of $(E \setminus A)^*$ is just the disordered region $\mathbf{V} \setminus \mathbf{V}(A)$ (see Lemma 5.1 in Section 5).

For $i = 1, 2, \dots, d$, let L_i be the fundamental loop $L_i = \{y \in \mathbf{V} | y_i = 1\}$. If \mathbf{A} is a union of cubes with diameter $1/2$ and centers in $V_{1/2}$ and γ is a component of $\partial \mathbf{A}$, then the winding vector $\mathbf{N}(\gamma) \in \{0, 1\}^d$, with its i th component being equal to the number of intersections (mod 2) of γ with L_i .

Definition 4.1 *Let A be a configuration in Ω . The contours corresponding to the configuration A are defined as those connected components γ of the boundary of $\mathbf{V}(A)$ which have winding number zero, $\mathbf{N}(\gamma) = \mathbf{0}$; the remaining connected components of the boundary of $\mathbf{V}(A)$ are called the interfaces corresponding to A ; the set of these interfaces is called the interface network corresponding to A . We denote the set of contours corresponding to A by $\Gamma(A)$, and the interface network corresponding to A by $\mathcal{S}(A)$.*

Without reference to a configuration, we say that γ is a contour if there exists a configuration A such that $\gamma \in \Gamma(A)$, and similarly for an interface and an interface network. We define two contours (or two interfaces, or one interface and one contour) γ, γ' to be compatible, if $\text{dist}(\gamma, \gamma') \geq 1/2$. We define a contour γ and an interface network \mathcal{S} to be compatible if γ is compatible with all interfaces in \mathcal{S} .

Note that the contours and interfaces corresponding to a configuration $A \in \Omega$ are orientable in the standard topological sense (see e.g. [12] or [3]); in fact, they are oriented, with an “ordered side” facing $\mathbf{V}(A)$, and a “disordered side”, facing the complement of $\mathbf{V}(A)$. Thus our contours are “labeled contours”, with labels $\ell \in \{\text{ord}, \text{dis}\}$ indicating which side is ordered, and which is disordered.

Note also that the contours and interfaces corresponding to a configuration A are pairwise compatible. It is not true, however, that any set of pairwise compatible contours and interfaces correspond to a configuration $A \in \Omega$. In order to get a one-to-one correspondence, we define the notion of “matching labels.”

Definition 4.2 *Let \mathcal{S} be an interface network, and let Γ be a set of contours. We say that $\mathcal{S} \cup \Gamma$ is a set of matching contours and interfaces if the following conditions hold:*

- (i) *The contours and interfaces in $\Gamma \cup \mathcal{S}$ are pairwise compatible.*
- (ii) *The labels are matching in the sense that, for each component \mathbf{C} of $\mathbf{V} \setminus \bigcup_{\gamma \in \Gamma \cup \mathcal{S}} \gamma$, there exists a label $\ell(\mathbf{C}) \in \{\text{ord}, \text{dis}\}$ such that the ordered side of each contour (respectively,*

interface) faces an ordered component (i.e., a component with label $\ell(C) = \text{ord}$), and similarly for the disordered sides.

For a set of matching contours and interfaces, we denote the union of the ordered components by \mathbf{V}_{ord} , and the union of the disordered components by \mathbf{V}_{dis} .

With this definition, the set of contours and interfaces corresponding to a configuration $A \in \Omega$ are clearly matching. It turns out (see Corollary 5.11 in Section 5) that the converse is also true, namely that any set of matching contours and interfaces corresponds to exactly one configuration $A \in \Omega$. Using this fact, we rewrite the partition function Z as a sum over sets of matching contours and interfaces.

To this end, we first note that the number of components $\tilde{c}(A)$ is clearly equal to the number of components $c(\mathbf{V}(A))$ of the continuum set $\mathbf{V}(A)$, and hence also to the number of components of the set \mathbf{V}_{ord} . Note also that both \mathbf{V}_{ord} and \mathbf{V}_{dis} are functions of the matching contours and interfaces (Γ, \mathcal{S}) . As a consequence, the weight (4.1) can be rewritten as a function of (Γ, \mathcal{S}) according to:

$$\begin{aligned} w(A) &= q^{c(\mathbf{V}(A))} e^{-e_{\text{dis}}|V \setminus \mathbf{V}(A)|} e^{-e_{\text{ord}}|V \cap \mathbf{V}(A)|} e^{-\kappa \|\partial \mathbf{V}(A)\|} \\ &= q^{c(\mathbf{V}_{\text{ord}})} e^{-e_{\text{dis}}|\mathbf{V}_{\text{dis}} \cap V|} e^{-e_{\text{ord}}|\mathbf{V}_{\text{ord}} \cap V|} \prod_{S \in \mathcal{S}} e^{-\kappa \|S\|} \prod_{\gamma \in \Gamma} e^{-\kappa \|\gamma\|}, \end{aligned} \quad (4.5)$$

where $\|\partial \mathbf{V}(A)\|$ is the number of intersections of $\partial \mathbf{V}(A)$ with the continuum set $\mathbf{E} = \bigcup_{e \in E} \mathbf{e}$, and similarly for $\|S\|$ and $\|\gamma\|$. Together with the already mentioned Corollary 5.11, about the one-to-one correspondence between configurations and sets of matching contours and interfaces, this leads to the desired representation of the partition function $Z = Z_T = \sum_A w(A)$:

$$Z = \sum_{\mathcal{S}, \Gamma} q^{c(\mathbf{V}_{\text{ord}})} e^{-e_{\text{dis}}|\mathbf{V}_{\text{dis}} \cap V|} e^{-e_{\text{ord}}|\mathbf{V}_{\text{ord}} \cap V|} \prod_{S \in \mathcal{S}} e^{-\kappa \|S\|} \prod_{\gamma \in \Gamma} e^{-\kappa \|\gamma\|}, \quad (4.6)$$

where the sum runs over sets of matching contours and interfaces.

Next we define the interior and exterior of a contour. To this end, we need the following geometric lemma. Its proof is deferred to Section 5.

Lemma 4.3 *Let A be a configuration in Ω , and let γ be a contour of A .*

- i) The set $\mathbf{V} \setminus \gamma$ has exactly two components.*
- ii) Let \mathbf{C} and \mathbf{D} be the two components of $\mathbf{V} \setminus \gamma$, and let S_1, S_2 be two (not necessarily compatible) interfaces that are both compatible with γ . Then both S_1 and S_2 lie either in \mathbf{C} or \mathbf{D} .*

For the purpose of the next definition, it is convenient to define the size of a set $\mathbf{W} \subset \mathbf{V}$ as the cardinality of $\mathbf{W} \cap V$.

Definition 4.4 *Let γ be a contour. If there exists an interface S (not necessarily corresponding to the same configuration) that is compatible with γ , we define the exterior, $\mathbf{Ext} \gamma$, of γ as the component of $\mathbf{V} \setminus \gamma$ that contains S ; otherwise we choose the larger of the two components; finally, if both of these components have the same size, we choose that containing a distinguished point, $x_0 \in \mathbf{V}$. The interior is defined as the set $\mathbf{Int} \gamma = \mathbf{V} \setminus (\gamma \cup \mathbf{Ext} \gamma)$.*

Given a set of pairwise compatible contours Γ , we define a contour $\gamma \in \Gamma$ to be an external contour in Γ if there exists no contour $\gamma' \in \Gamma \setminus \{\gamma\}$ such that $\mathbf{Int} \gamma \subset \mathbf{Int} \gamma'$. We also define the exterior of Γ as the set

$$\mathbf{Ext} \Gamma = \bigcap_{\gamma \in \Gamma} \mathbf{Ext} \gamma. \quad (4.7)$$

Finally, we say that $\gamma_1, \dots, \gamma_n$ are mutually external if they are pairwise compatible and $\mathbf{Int} \gamma_i \cap \mathbf{Int} \gamma_j = \emptyset$ for all $i \neq j$.

The next lemma states several properties of the set $\mathbf{Ext} \Gamma(A)$ of a set of contours corresponding to a configuration A without interfaces. Its proof is again deferred to Section 5.

Lemma 4.5 *Let A be a configuration with $\mathcal{S}(A) = \emptyset$, and let A_{ext} be the set of edges with both endpoints in $\mathbf{Ext} \Gamma(A)$. Then the following statements hold*

- (i) $\mathbf{Ext} \Gamma(A)$ is a connected subset of the continuum torus \mathbf{V} .
- (ii) $\mathbf{Ext} \Gamma(A)$ is either a subset of the ordered region $\mathbf{V}(A)$ or the disordered region $\mathbf{V} \setminus \mathbf{V}(A)$.
- (iii) If $\mathbf{Ext} \Gamma(A) \subset \mathbf{V}(A)$, then $(V(A_{\text{ext}}), A_{\text{ext}})$ is a connected subgraph of $T_{L,d}$.

5 The Geometry of Contours and Interfaces

5.1 Elementary Topological Notions

We start by reviewing some standard notions from algebraic topology. Let $T_{1/2} = ((\frac{1}{2}\mathbb{Z})/L\mathbb{Z})^d$ be the torus with two points connected by an edge if they differ by $1/2 \pmod{L}$ in one component. Its vertex and edge sets will be denoted by $V_{1/2}$ and $E_{1/2}$, respectively.

We define k -cells in $V_{1/2}$ as the k -dimensional elementary cubes in $V_{1/2}$, so that the points in $V_{1/2}$ are 0-cells, the edges are 1-cells, etc. We also consider the dual $V_{1/2}^*$ of $V_{1/2}$, consisting of the barycenters of the d -cells in $V_{1/2}$, when considered as subsets of the continuum torus $(\mathbb{R}/L\mathbb{Z})^d$. As usual, given a k -cell c in $V_{1/2}$, we define its dual as the $d - k$ cell c^* in $V_{1/2}^*$ that has the same barycenter as c , and similarly for the dual of a k -cell in $V_{1/2}^*$. Note that $(c^*)^* = c$.

Given a k -cell c , we define its boundary as the set of all $(k - 1)$ -cells that are subcubes of c (note that there are $2k$ such subcubes). More generally, for a set K of k -cells, we define the (\mathbb{Z}_2) -boundary of K as the set of all $(k - 1)$ -cells which are in the boundaries of an odd number of cells in K . We denote this boundary by ∂K . The co-boundary of a k -cell c is defined to be the set of $(k + 1)$ -cells which have c in their boundaries (there are $2(d - k)$ such $(k + 1)$ -cells).

We often identify a k -cell with the closed k -dimensional continuum cube with corners being the vertices of the discrete cell. In this context, ∂K is identified with the corresponding continuum boundary.

The $(d - 1)$ -cells in $V_{1/2}^*$ are called facets; the set of all such facets is denoted by $F_{1/2}^*$. We say two facets are connected (or adjacent) if they share a $(d - 2)$ -dimensional cell in their boundaries.

A sequence of points $x_0, \dots, x_k \in V_{1/2}$ is called a loop in $T_{1/2}$ if $x_0 = x_k$ and $\{x_i, x_{i+1}\} \in E_{1/2}$ for all $i \in \{0, k - 1\}$. Such a loop is called a fundamental loop in the i^{th} lattice direction

if $k = 2L$ and all edges point in the i^{th} lattice direction, and it is called an elementary loop if $k = 4$ and neither $x_0 = x_2$ nor $x_1 = x_3$. Note that there are exactly $2d(2d - 2)(2L)^d$ elementary loops in $T_{1/2}$.

Consider now a set of edges $X \subset E_{1/2}$ and its dual $X^* = \{e^* : e \in X\}$, and assume that X is orientable. If $\partial X^* = \emptyset$, then we say that X^* is an orientable closed surface, and we define the \mathbb{Z}_2 -winding vector of X^* as the vector $\mathbf{N}(X^*) = (N_1, \dots, N_d) \in \{0, 1\}^d$ with N_i equal to the number of times X^* intersects a fundamental loop in the i^{th} lattice direction mod 2.

5.2 Preliminaries

Our first lemma summarizes several simple properties of the construction used in the definition of contours. It involves both facets in $V_{1/2}^*$, the objects dual to the 1-cells in $V_{1/2}$, and $(d-1)$ -dimensional unit hypercubes dual to the edges in E . While the first will be considered to be abstract objects in the sense of algebraic topology, the second will be considered to be closed hypercubes in \mathbf{V} . We trust that this does not cause any confusion to the reader.

Lemma 5.1 *i) For $A \in \Omega$, the boundary $\partial\mathbf{V}(A)$ of $\mathbf{V}(A)$ is regular in the sense that each $(d-2)$ -cell with corners in $V_{1/2}^*$ is shared by either zero or two facets in $\partial\mathbf{V}(A)$.*

ii) Let $A \in \Omega$, and let \mathbf{C} be a component of $\mathbf{V} \setminus \partial\mathbf{V}(A)$. Then \mathbf{C} is either a component of $\mathbf{V}(A)$ or of $\mathbf{V} \setminus \mathbf{V}(A)$.

iii) Let $\mathbf{C}_1, \dots, \mathbf{C}_k$ be the connected components of $\mathbf{V}(A)$, let $V_i = V \cap \mathbf{C}_i$, and let A_i be the set of edges whose endpoints lie in C_i . Then $(V_1, A_1), \dots, (V_k, A_k)$ are the connected components of (V, A) , and $\mathbf{C}_i = \mathbf{V} \setminus \mathbf{V}(A_i)$.

iv) Let $(E \setminus A)^$ be the union of all $d-1$ dimensional unit cubes that are dual to the edges in $E \setminus A$, and let $\mathbf{V}_{\text{dis}}(E \setminus A)$ be the union of the open 3/4-neighborhood of $V \setminus V(A)$ and the open 1/4-neighborhood of $(E \setminus A)^*$. Then $\mathbf{V} \setminus \mathbf{V}(A) = \mathbf{V}_{\text{dis}}(E \setminus A)$.*

v) Let \mathbf{C} be a component of $\mathbf{V}_{\text{dis}}(E \setminus A)$, and let $E_{\mathbf{C}}$ be the set of edges whose midpoint lies in \mathbf{C} . Then $\mathbf{C} = \mathbf{V}_{\text{dis}}(E_{\mathbf{C}} \setminus A)$.

Proof.

(i) Given a configuration $A \in \Omega$, let $V_{1/2}(A)$ be the intersection of $\mathbf{V}(A)$ with the vertex set $V_{1/2}$ of the discrete torus $T_{1/2}$. The boundary of $\mathbf{V}(A)$ can then be rewritten as the union of all facets that are dual to an edge $e \in E_{1/2}$ joining $V_{1/2}(A)$ to its complement in $V_{1/2}$. Using this fact, we easily prove the first statement of the lemma.

Indeed, let e be a $(d-2)$ -cell with corners in $V_{1/2}^*$, and let f_1, f_2, f_3, f_4 be the four facets in the co-boundary of e . Then f_1, f_2, f_3, f_4 are dual to edges $x_1x_2, x_2x_3, x_3x_4, x_4x_1$, where the points x_1, x_2, x_3, x_4 comprise a closed path of length four in the torus $T_{1/2}$. Let $c(x_1), \dots, c(x_4)$ be the four d -cells with centers x_1, \dots, x_4 . Exactly one of these four cubes, say the cube $c(x_1)$, will be a cube whose center lies in the original vertex set V (recall that vertices in V have twice the spacing of those in $V_{1/2}$). Our convention of filling in only those hypercubes in the original torus $T_{L,d}$ whose edges are all in A implies that either none or all or exactly one of the four cubes $c(x_1), \dots, c(x_4)$ lies in $\mathbf{V}(A)$. In the first two cases, none of the facets f_1, f_2, f_3, f_4 are in the boundary of $\mathbf{V}(A)$, and in the third case exactly two are in

the boundary of $\mathbf{V}(A)$, which proves that each $(d - 2)$ -cell with corners in $V_{1/2}^*$ is shared by either zero or two facets, as claimed.

(ii) This is obvious.

(iii) This now follows immediately from our fattening procedure for the ordered region, which respects the component structure of (V, A) .

(iv) We first prove that $\mathbf{V}_{\text{dis}}(E \setminus A) \subset \mathbf{V} \setminus \mathbf{V}(A)$. Consider first an edge $e \in E \setminus A$ and its dual e^* . We claim that all points with distance less than $1/4$ from e^* lie in $\mathbf{V} \setminus \mathbf{V}(A)$. Since the set $\mathbf{V}(A)$ increases if A increases, it is clearly enough to prove this statement for $A = E \setminus \{e\}$, in which case it follows immediately from the way we set up our fattening procedure for $\mathbf{V}(A)$. In a similar way, one proves that all points with distance less than $3/4$ from the vertices in $V \setminus V(A)$ lie in $\mathbf{V} \setminus \mathbf{V}(A)$.

To prove that $\mathbf{V} \setminus \mathbf{V}(A) \subset \mathbf{V}_{\text{dis}}(E \setminus A)$, let $x \in \mathbf{V} \setminus \mathbf{V}(A)$, and let $\mathbf{c} \subset \mathbf{V}$ be a d -dimensional unit cube with corners in V such that $x \in \mathbf{c}$. If x has distance less than $1/4$ from the center of \mathbf{c} , then at least one edge $e \subset \mathbf{c} \cap V$ must lie in $E \setminus A$, since otherwise \mathbf{c} would have been filled in our fattening procedure, contradicting $x \in \mathbf{V} \setminus \mathbf{V}(A)$. But the $1/4$ neighborhood $\mathcal{N}_{1/4}(e^*)$ of e^* contains all points with distance less than $1/4$ from the center of \mathbf{c} , implying that $x \in \mathcal{N}_{1/4}(e^*) \subset \mathbf{V}_{\text{dis}}(E \setminus A)$. If x has distance $1/4$ or more from the center of \mathbf{c} , it must have distance less than $1/4$ from a $d - 1$ dimensional unit cube \mathbf{c}_1 in the boundary of \mathbf{c} . If the projection, x_1 , of x onto \mathbf{c}_1 has distance less than $1/4$ from the center of \mathbf{c}_1 , then one of the edges in $\mathbf{c}_1 \cap V$ must lie in $E \setminus A$, which again implies $x \in \mathcal{N}_{1/4}(e^*) \subset \mathbf{V}_{\text{dis}}(E \setminus A)$. Continuing inductively, we are left with the case that x has distance at most $1/4$ from one of the corners, y , of \mathbf{c} . But this means that none of the edges containing y can lie in A , thus $y \in V \setminus V(A)$, and hence $x \in \mathcal{N}_{3/4}(V \setminus V(A)) \subset \mathbf{V}_{\text{dis}}(E \setminus A)$.

(v) Noting that the midpoint of an edge lies in $\mathbf{V}_{\text{dis}}(E \setminus A)$ if and only if its dual lies in $\mathbf{V}_{\text{dis}}(E \setminus A)$, this is an immediate consequence of our fattening procedure for the disordered region. \square

Lemma 5.2 *Let S_1 and S_2 be two interfaces with $S_1 \cap S_2 = \emptyset$. Then $\mathbf{N}(S_1) = \mathbf{N}(S_2)$.*

Proof. The interfaces S_1 and S_2 are closed orientable submanifolds of the torus $(\mathbb{R}/L\mathbb{Z})^d$. In the language of algebraic topology, the winding numbers $\mathbf{N}(S_1)$ and $\mathbf{N}(S_2)$ are the Poincare duals of the submanifolds S_1 and S_2 . The Poincare dual of the transverse intersection of two such submanifolds is then given by the wedge product of the Poincare duals of the submanifolds, see [12], Section 6. Since empty intersection is a special case of transverse intersection, we conclude that the wedge product of $\mathbf{N}(S_1)$ and $\mathbf{N}(S_2)$ must be zero. Let \vec{e}_i be the unit vector whose j^{th} coordinate is $\delta_{i,j}$. Recalling that $\vec{e}_i \wedge \vec{e}_i = 0$ and $\vec{e}_i \wedge \vec{e}_j = -\vec{e}_j \wedge \vec{e}_i$ for all i, j , the condition $\mathbf{N}(S_1) \wedge \mathbf{N}(S_2) = 0$ is equivalent to the $\binom{d}{2}$ conditions $N_i(S_1)N_j(S_2) - N_j(S_1)N_i(S_2) = 0$, which implies that $\mathbf{N}(S_1)$ and $\mathbf{N}(S_2)$ are multiples of each other. Since both are different from 0, we conclude that $\mathbf{N}(S_1) = \mathbf{N}(S_2)$. \square

We close this section with an elementary lemma about ‘‘cutsets’’. It is best formulated in the context of a general connected graph $G = (V, E)$. As usual, a subset $E' \subset E$ is a *cutset* if $(V, E \setminus E')$ is disconnected. It is called a *minimal* cutset if no proper subset of E' is a cutset. The following lemma is elementary; its proof is left to the reader.

Lemma 5.3 *Let $G = (V, E)$ be a connected graph, and let $W \subset V$. If the edge-boundary of W is a minimal cutset then both the induced graph on W and $V \setminus W$ are connected.*

5.3 Proofs of Lemma 4.3 and 4.5

We start with the proof of Lemma 4.3.

Proof of Lemma 4.3. i) γ has no boundary, is orientable, and has winding number zero. Therefore any closed path intersects γ an even number of times, implying that γ is the boundary of some open set \mathbf{C} . Let \mathbf{D} be the set $\mathbf{D} = \mathbf{V} \setminus (\mathbf{C} \cup \gamma)$. Then both \mathbf{C} and \mathbf{D} must be connected. Indeed, considering γ as dual to a minimal cutset on the half-integer lattice, the connectedness of \mathbf{C} and \mathbf{D} follows immediately from the corresponding statement (Lemma 5.3) for minimal cutsets.

ii) The interfaces S_1, S_2 are connected subsets of \mathbf{V} . Since they do not intersect γ , each of them must lie in one of the connected components of $\mathbf{V} \setminus \gamma$. We will have to prove that they both lie in the same component of $\mathbf{V} \setminus \gamma$.

Assume, by contradiction, that $S_1 \subset \mathbf{C}$ and $S_2 \subset \mathbf{D}$. Since $\mathbf{C} \cap \mathbf{D} = \emptyset$, this implies $S_1 \cap S_2 = \emptyset$, which in turn, by Lemma 5.2 implies that $\mathbf{N}(S_1) = \mathbf{N}(S_2)$. Together with the fact that $\mathbf{N}(\gamma) = 0$ we conclude that the winding number of $\gamma \cup S_1 \cup S_2$ is zero, implying that any loop in \mathbf{V} must intersect $\gamma \cup S_1 \cup S_2$ an even number of times.

Consider now a component $N_i(S_1)$ of the winding vector $\mathbf{N}(S_1)$ that is not equal to 0, and let ω be a fundamental loop in the i^{th} direction, oriented in an arbitrary but fixed fashion. To reach our contradiction, we will modify ω in such a way that it does not intersect γ or S_2 , while intersecting S_1 an odd number of times. First, we note that the original loop ω intersects S_1 an odd number of times. If it does not intersect γ , then it does not intersect S_2 either, since S_2 lies in \mathbf{D} , while S_1 lies in \mathbf{C} . If ω intersects γ , let x be one of the intersection points, and let y be the next intersection point (since $\mathbf{N}(\gamma) = 0$, the number of these intersection points must be even, so there must be such a y). Recalling that γ is connected, we now replace the segment of ω that joins x to y by a path in γ , and then deform this segment in such a way that it lies completely in \mathbf{C} without intersecting S_1 (this is possible since $\text{dist}(S_1, \gamma) \geq 1/2$). Given that the original segment from x to y did not intersect S_1 since $S_1 \subset \mathbf{C}$ while the segment was a path in \mathbf{D} , we have not changed the number of intersections with S_1 , so this number is still odd. Repeating this step until we have no intersections with γ , we remove all parts of ω that lie outside of \mathbf{C} , ending up with a path inside \mathbf{C} that intersects S_1 an odd number of times, as desired. Note that the final ω may not be a lattice path, but nevertheless it gives the desired contradiction, since winding numbers are topological invariants in the continuum as well. \square

Definition 5.4 *Let γ and γ' be two contours. We say that γ and γ' are mutually external (and write $\gamma \perp \gamma'$) if γ is compatible with γ' and $\mathbf{Int} \gamma \cap \mathbf{Int} \gamma' = \emptyset$, and we say that γ lies inside of γ' (and write $\gamma < \gamma'$) if γ is compatible with γ' and $\mathbf{Int} \gamma \subset \mathbf{Int} \gamma'$.*

Lemma 5.5 *Let γ, γ' and γ'' be contours.*

(i) *If γ and γ' are compatible, then exactly one of the following three holds: $\gamma < \gamma'$, $\gamma' < \gamma$, or $\gamma \perp \gamma'$.*

(ii) If $\gamma \perp \gamma'$, then $\text{dist}(\mathbf{Int} \gamma, \mathbf{Int} \gamma') \geq 1/2$, and if $\gamma < \gamma'$, then $\text{dist}(\mathbf{Int} \gamma, \mathbf{Ext} \gamma') \geq 1/2$.

(iii) If $\gamma < \gamma'$ and $\gamma' < \gamma''$ then $\gamma < \gamma''$.

(iv) If $\gamma < \gamma'$ and $\gamma' \perp \gamma''$ then $\gamma \perp \gamma''$.

Proof. (i) Let $\mathbf{Int} \gamma$ and $\mathbf{Ext} \gamma$ be the two components of $\mathbf{V} \setminus \gamma$. Since γ' is connected and $\gamma \cap \gamma' = \emptyset$, γ' must lie in one of the two components of $\mathbf{V} \setminus \gamma$; therefore we have that either $\gamma' \subset \mathbf{Int} \gamma$ or $\gamma' \subset \mathbf{Ext} \gamma$. Since the same statement holds with the roles of γ and γ' exchanged, we get that exactly one of the following four cases must hold:

$$\gamma \subset \mathbf{Ext} \gamma' \quad \text{and} \quad \gamma' \subset \mathbf{Ext} \gamma \quad (5.1)$$

$$\gamma \subset \mathbf{Int} \gamma' \quad \text{and} \quad \gamma' \subset \mathbf{Ext} \gamma \quad (5.2)$$

$$\gamma \subset \mathbf{Ext} \gamma' \quad \text{and} \quad \gamma' \subset \mathbf{Int} \gamma \quad (5.3)$$

$$\gamma \subset \mathbf{Int} \gamma' \quad \text{and} \quad \gamma' \subset \mathbf{Int} \gamma \quad (5.4)$$

We claim that the last case is impossible.

To prove this, we first show that (5.4) implies that $\mathbf{Ext} \gamma \cap \mathbf{Ext} \gamma' \neq \emptyset$. Recalling Definition 4.4 of $\mathbf{Ext} \gamma$, note that $\mathbf{Ext} \gamma$ is defined differently in several distinct cases. Let us first consider the case that there exists an interface S with $\text{dist}(S, \gamma) \geq 1/2$ and $S \subset \mathbf{Ext} \gamma$. If $\text{dist}(S, \gamma') < 1/2$, then also $\text{dist}(S, \mathbf{Int} \gamma) < 1/2$ by our assumption that $\gamma' \subset \mathbf{Int} \gamma$. But this is not compatible with $\text{dist}(S, \gamma) \geq 1/2$ and $S \subset \mathbf{Ext} \gamma$. Thus we have $\text{dist}(S, \gamma') \geq 1/2$. But if $\text{dist}(S, \gamma') \geq 1/2$, then $S \subset \mathbf{Ext} \gamma'$ as well, implying in particular that $\mathbf{Ext} \gamma \cap \mathbf{Ext} \gamma' \neq \emptyset$.

Let us now consider the cases in the definitions of $\mathbf{Ext} \gamma$ and $\mathbf{Ext} \gamma'$ such that there is no interface S compatible with either γ or γ' . Consider the subcase that $\mathbf{Ext} \gamma$ is defined by size, so that $|\mathbf{Ext} \gamma \cap V| > |\mathbf{Int} \gamma \cap V|$ and $|\mathbf{Ext} \gamma' \cap V| \geq |\mathbf{Int} \gamma' \cap V|$. This implies that $|\mathbf{Ext} \gamma \cap V| > |V|/2$ and $|\mathbf{Ext} \gamma' \cap V| \geq |V|/2$ which in turn implies that $\mathbf{Ext} \gamma \cap \mathbf{Ext} \gamma' \neq \emptyset$. The case where $\mathbf{Ext} \gamma'$ is defined by size is strictly analogous, so we are left with the subcase where both $\mathbf{Ext} \gamma$ and $\mathbf{Ext} \gamma'$ are defined by containing the distinguished point x_0 . Again, this implies that $\mathbf{Ext} \gamma \cap \mathbf{Ext} \gamma' \neq \emptyset$.

Now we show that the condition $\mathbf{Ext} \gamma \cap \mathbf{Ext} \gamma' \neq \emptyset$ rules out the case (5.4). Let $u \in \mathbf{Ext} \gamma \cap \mathbf{Ext} \gamma'$, and let $x \in \mathbf{Int} \gamma$. Consider a path ω from x to u , and let y be the last exist point from $\mathbf{Int} \gamma$ along ω . This implies $y \in \gamma$, and by our assumption (5.4), we therefore have $y \in \mathbf{Int} \gamma'$. But this implies there exists a point $z \in \gamma' \subset \mathbf{Int} \gamma$ after y , which is a contradiction. Thus case (5.4) is impossible.

Next we prove that the three remaining cases imply that $\mathbf{Int} \gamma' \cap \mathbf{Int} \gamma = \emptyset$, $\mathbf{Ext} \gamma' \cap \mathbf{Int} \gamma = \emptyset$ and $\mathbf{Int} \gamma' \cap \mathbf{Ext} \gamma = \emptyset$, respectively, showing that either $\gamma \perp \gamma'$, $\gamma < \gamma'$ or $\gamma' < \gamma$, respectively.

We prove all three statements in one sweep, by setting $A = \mathbf{Ext} \gamma$ and $A' = \mathbf{Ext} \gamma'$ in the first case, $A = \mathbf{Ext} \gamma$ and $A' = \mathbf{Int} \gamma'$ in the second case, and $A = \mathbf{Int} \gamma$ and $A' = \mathbf{Ext} \gamma'$ in the third case. Our assumption then reads $\gamma \subset A'$ and $\gamma' \subset A$, and our claim is $B \cap B' = \emptyset$, where $B = A^c \setminus \gamma$ and $B' = (A')^c \setminus \gamma'$. In a preliminary step, we prove that $\gamma \subset A'$ implies $B \setminus B' \neq \emptyset$. Indeed, assume the contrary, i.e., $B \subset B'$. Taking the closure on both sides, this gives $\gamma \cup B \subset \gamma' \cup B'$, and hence $\gamma \subset \gamma' \cup B' = (A')^c$, a contradiction. Now we prove the main claim $B \cap B' = \emptyset$. Assume the contrary, that there exists an $x \in B \cap B'$. From our preliminary claim, we also know that there exists a $y \in B \setminus B'$. Since B is connected,

we conclude that there must be a path $\omega \subset B$ from x to y . Let z be the first time this path exits B' . Then $z \in B \cap \partial B' = B \cap \gamma'$. Thus $B \cap \gamma' \neq \emptyset$, a contradiction.

(ii) Both statements follow from the observation that if $A \cap A' = \emptyset$, then $\text{dist}(\partial A, \partial A') \leq \text{dist}(A, A')$.

(iii) Assume that $\gamma < \gamma'$ and $\gamma' < \gamma''$. Then $\mathbf{Int} \gamma \subset \mathbf{Int} \gamma' \subset \mathbf{Int} \gamma''$, so we need only prove that γ and γ'' are compatible, i.e. that $\text{dist}(\gamma, \gamma'') \geq 1/2$. On the other hand, $\mathbf{Int} \gamma \subset \mathbf{Int} \gamma'$. Taking the closure of both sides and using a trivial inclusion, this implies that $\gamma \subset \gamma \cup \mathbf{Int} \gamma \subset \gamma' \cup \mathbf{Int} \gamma'$. Thus

$$\text{dist}(\gamma, \gamma'') \geq \text{dist}(\gamma' \cup \mathbf{Int} \gamma', \gamma'' \cup \mathbf{Ext} \gamma'') = \text{dist}(\mathbf{Int} \gamma', \mathbf{Ext} \gamma'') \geq 1/2$$

where we used (ii) in the last step.

(iv) This is proved strictly analogously to the proof of (iii). \square

The next lemma is an easy corollary of Lemma 5.5.

Lemma 5.6 *Let Γ be a set of pairwise compatible contours. Then $\mathbf{Ext} \Gamma$ is a connected subset of \mathbf{V} .*

Proof. Let Γ_{ext} be the set of external contours in Γ . It follows immediately from the definition of $\mathbf{Ext} \Gamma$ and the last lemma that $\mathbf{Ext} \Gamma = \mathbf{Ext} \Gamma_{\text{ext}}$. It is therefore enough to consider a set Γ of mutually external contours. We prove the statement by induction on the number of contours in Γ . The statement is trivial if $\Gamma = \emptyset$. Assume the statement is proved for $\Gamma = \{\gamma_1, \dots, \gamma_{n-1}\}$. Adding an additional mutually external contour γ_n will not change the connectivity. Indeed, let $x, y \in \mathbf{Ext}(\Gamma \cup \{\gamma_n\})$, and let ω be a path in $\mathbf{Ext} \Gamma$ that joins x to y . If ω does not intersect $\gamma_n \cup \mathbf{Int} \gamma_n$ there is nothing to prove. Otherwise, let $x' \in \gamma_n$ be the first entry point into $\gamma_n \cup \mathbf{Int} \gamma_n$, and let $y' \in \gamma_n$ be the last exit point from $\gamma_n \cup \mathbf{Int} \gamma_n$. Since γ_n is connected, we can replace the path ω from x' to y' by a path in γ_n , leading to a path ω' joining x and y in $\gamma_n \cup \mathbf{Ext}(\Gamma \cup \{\gamma_n\})$. By deforming ω' this immediately leads to a path in $\mathbf{Ext}(\Gamma \cup \{\gamma_n\})$, proving the lemma. \square

We are now ready to prove Lemma 4.5.

Proof of Lemma 4.5.

(i) Since the contours corresponding to a configuration A are pairwise compatible, this statement follows immediately from the previous lemma.

(ii) Using the fact that the boundary of $\mathbf{V}(A)$ is equal to the union over all contours in $\Gamma(A)$, we conclude that $\mathbf{Ext} \Gamma(A)$ is a connected subset of $\mathbf{V} \setminus \partial \mathbf{V}(A)$. But this implies that $\mathbf{Ext} \Gamma(A) \subset \mathbf{V}(A)$ or $\mathbf{Ext} \Gamma(A) \subset \mathbf{V} \setminus \mathbf{V}(A)$, as claimed.

(iii) This follows immediately from (i) and the third statement of Lemma 5.1. \square

5.4 Isoperimetric Estimates

We need the notion of the diameter, $\text{diam} \gamma$, of a contour γ . It is defined as the side-length of the minimum cube $C \subset V$ such that $\mathbf{Int} \gamma \subset \{x \in \mathbf{V} : \text{dist}(x, C) \leq 3/4\}$, where, as usual, a cube $C \subset V$ is considered a cube of sidelength k if C contains k^d points.

Lemma 5.7 *For every contour γ , we have*

$$|\mathbf{Int} \gamma \cap V| \leq \frac{1}{2} \|\gamma\| \text{diam} \gamma \quad (5.5)$$

and

$$\|\gamma\| \geq \text{diam} \gamma. \quad (5.6)$$

Proof. We start with the proof of the first inequality. Let $W = \mathbf{Int} \gamma \cap V$, and let $\partial_{\text{edge}} W$ be the edge-boundary of W . If $|W| \leq L^d/2$, then

$$|\partial_{\text{edge}} W| \geq \min_{i=1, \dots, d} 2i |W|^{1-1/i} L^{d/i-1} \geq 2 \min_i |W|^{1-1/i} L^{d/i-1} = 2|W|^{1-1/d},$$

by the isoperimetric inequality of Bollobás-Leader [6]. Since $|W| \leq (\text{diam} \gamma)^d$, this implies that

$$|W| \leq \frac{1}{2} |\partial_{\text{edge}} W| |W|^{1/d} \leq \frac{1}{2} \|\gamma\| \text{diam} \gamma.$$

If $|W| > L^d/2$, then, by the definition of $\mathbf{Ext} \gamma$, there exists an interface S such that $S \subset \mathbf{Ext} \gamma$. Without loss of generality, let us assume that the first component of the winding vector $\mathbf{N}(S)$ is 1. This implies that every line in the 1-direction intersects S at least once. Hence, for all $x \in W = \mathbf{Int} \gamma \cap V$, a line through x in the 1-direction intersects γ at least twice, implying that

$$|\mathbf{Int} \gamma \cap V| \leq \frac{1}{2} \|\gamma\| \text{diam} \gamma.$$

Now we prove the second inequality of the lemma. Let $C \subset V$ be a minimal cube of sidelength $k = \text{diam} \gamma$ such that $\mathbf{Int} \gamma \subset \bar{C} := \{x : \text{dist}(x, C) \leq 3/4\}$, and let

$$D_i^+ = \{x \in \bar{C} : \exists y \in V \setminus \bar{C}, x_i > y_i - 1\},$$

and

$$D_i^- = \{x \in \bar{C} : \exists y \in V \setminus \bar{C}, x_i < y_i + 1\}.$$

Since C is minimal, there exists an i such that $\mathbf{Int} \gamma \cap D_i^+ \neq \emptyset$ as well as $\mathbf{Int} \gamma \cap D_i^- \neq \emptyset$.

Next, let $B_{i,j}$, $j = 1, \dots, k$ be the j th layer of C , where in particular, $B_{i,1} = D_i^+ \cap V$, and $B_{i,k} = D_i^- \cap V$. If γ does not intersect an edge in $B_{i,j}$, then

$$\bar{B}_{i,j} = \{x \in \bar{C} : \exists y \in B_{i,l}, \text{dist}(x, y) \leq 1/4\}$$

has the same label as the exterior of γ . Together with the connectedness of γ , this contradicts the fact that both $\mathbf{Int} \gamma \cap D_i^+ \neq \emptyset$, and $\mathbf{Int} \gamma \cap D_i^- \neq \emptyset$.

5.5 Matching Contours and Interfaces

In this section we will show that the partition function Z can be written as a sum over sets of matching contours and interfaces. To this end, we establish a sequence of lemmas. We start with the following lemma.

Lemma 5.8 *Let γ be a contour. Then there exists a configuration A such that $\mathcal{S}(A) = \emptyset$ and $\Gamma(A) = \{\gamma\}$.*

Proof. By definition, there exists a configuration A_1 such that γ is one of the contours corresponding to A_1 . Thus γ is a connected component of $\partial\mathbf{V}(A_1)$, and hence a connected component of the boundary of one of the components, \mathbf{C} , of $\mathbf{V}(A_1)$. Let A_2 be the set of edges in A_1 whose endpoints both lie in \mathbf{C} . Then $(V(A_2), A_2)$ is a component of (V, A_1) , and, by Lemma 5.1 (iii), the set $\mathbf{V}(A_2)$ is nothing but the component \mathbf{C} of $\mathbf{V}(A_1)$.

Thus γ is a component of $\partial\mathbf{V}(A_2)$. Consider now the complement $\mathbf{D} = \mathbf{V} \setminus \mathbf{V}(A_2)$, and its components $\mathbf{D}_0, \dots, \mathbf{D}_k$. Then γ is the boundary of one of these components, say \mathbf{D}_0 . Let A_0 be the set of edges whose midpoint lies in $\mathbf{V} \setminus \mathbf{D}_0$. By Lemma 5.1 (iv) and (v), we have that $\mathbf{D}_0 = \mathbf{V} \setminus \mathbf{V}(A_0)$, which in turn implies that $\gamma = \partial\mathbf{V}(A_0)$. \square

Lemma 5.9 *Let \mathcal{S} be an interface network. Then there exists a configuration A such that $\mathcal{S}(A) = \mathcal{S}$ and $\Gamma(A) = \emptyset$.*

Proof. The proof is identical to that of the previous lemma.

Recall that each contour has an ordered and a disordered side. We call γ a contour with external label *ord* if the side facing $\mathbf{Ext} \gamma$ is ordered. Otherwise it is called a contour with external label *dis*, for disordered.

Lemma 5.10 *Let $A \subset \Omega$. If γ is a contour with $\text{dist}(\mathbf{V}(A), \mathbf{Int} \gamma) \geq 1/2$ and external label *dis*, or a contour with $\text{dist}(\mathbf{V} \setminus \mathbf{V}(A), \mathbf{Int} \gamma) \geq 1/2$ and external label *ord*, then there exists a configuration $A' \subset \Omega$ such that $\Gamma(A') = \Gamma(A) \cup \{\gamma\}$ and $S(A) = S(A')$.*

Proof. Consider first the case that $\ell = \text{dis}$. Define A_1 to be the set of edges with both endpoints in $\mathbf{Int} \gamma$. By Lemma 5.1 (iii), we have that $\mathbf{V}(A \cup A_1) = \mathbf{V}(A) \cup \mathbf{V}(A_1) = \mathbf{V}(A) \cup \mathbf{Int} \gamma$, which shows that $A \cup A_1$ is the desired configuration.

For $\ell = \text{ord}$, we define E_1 to be the set of edges whose midpoint lies in $\mathbf{Int} \gamma$. Using Lemma 5.1 (iv) and (v), we now conclude that $A \setminus E_1$ is the desired configuration. \square

Corollary 5.11 *Let $A \subset \Omega$. Then the contours and interfaces corresponding to a configuration $A \in \Omega$ are matching. Conversely, any set of matching contours and interfaces corresponds to exactly one configuration $A \in \Omega$.*

Proof. The first statement is obvious. The second follows from Lemmas 5.9 and the previous lemma by induction on the number of contours. Indeed, by Lemma 5.5, the partial order on contours leads to a forest on any set $\Gamma = \{\gamma_1, \dots, \gamma_n\}$ of pairwise compatible contours, in such a way that $\gamma' < \gamma$ whenever γ' is a child of γ . Adding the interface network as a common root, we may proceed by induction from this root to add contours in such a way that the new contour added always obeys the condition from the previous lemma. \square

Note that this corollary, together with the representation (4.5) for the weights of a configuration A , established the representation (4.6) for the partition function Z .

6 Key Ingredients for the Lower Bound

As explained in Section 2.3, we will prove our lower bound on the mixing time by proving an upper bound on the conductance. This in turn will require both a bound on the probability of the set of configurations with at least one interface, and a large deviation bound on configurations for which the joint exterior of all contours contains less than $(1 - \alpha)L^d$ points in V , see Lemmas 6.1 and 6.2 below. We will start with the decomposition of the partition function into terms *with* and *without* interfaces.

6.1 Decomposition of the Partition Function

Let

$$\Omega_{\text{tunnel}} = \{A \in \Omega \mid \mathcal{S}(A) \neq \emptyset\}, \quad (6.1)$$

$$\Omega_{\text{ord}} = \{A \in \Omega \setminus \Omega_{\text{tunnel}} \mid \mathbf{Ext} \Gamma(A) \subset \mathbf{V}(A)\} \quad (6.2)$$

and

$$\Omega_{\text{dis}} = \{A \in \Omega \setminus \Omega_{\text{tunnel}} \mid \mathbf{Ext} \Gamma(A) \subset \mathbf{V} \setminus \mathbf{V}(A)\}. \quad (6.3)$$

By Lemma 4.5 (ii), $\Omega = \Omega_{\text{ord}} \cup \Omega_{\text{dis}} \cup \Omega_{\text{tunnel}}$. As a consequence

$$Z = Z_{\text{dis}} + qZ_{\text{ord}} + Z_{\text{tunnel}}, \quad (6.4)$$

where

$$Z_{\text{tunnel}} = \sum_{A \in \Omega_{\text{tunnel}}} w(A), \quad Z_{\text{dis}} = \sum_{A \in \Omega_{\text{dis}}} w(A), \quad (6.5)$$

and

$$Z_{\text{ord}} = \frac{1}{q} \sum_{A \in \Omega_{\text{ord}}} w(A). \quad (6.6)$$

Note the extra factor of q in (6.4), which accounts for the fact that there are q different ordered phases.

The results of this section are summarized in the next two lemmas. The first is a finite-size scaling bound analogous to those proved in [9, 10, 11]. The second is a large deviations bound.

Lemma 6.1 *For all $d \geq 2$, there are constants $c > 0$, $q_0 < \infty$ and $L_0 < \infty$ such that the following statements hold for $q \geq q_0$ and $L \geq L_0$:*

(a) *If $\beta \geq \beta_0$, then*

$$\nu(\Omega_{\text{tunnel}}) \leq e^{-c\beta L^{d-1}} \quad (6.7)$$

and

$$\nu(\Omega_{\text{ord}}) \geq \frac{q}{q+1} - e^{-c\beta L}. \quad (6.8)$$

(b) *If $\beta = \beta_0$, then*

$$\left| \nu(\Omega_{\text{ord}}) - \frac{q}{q+1} \right| \leq e^{-c\beta L}. \quad (6.9)$$

To state the next lemma, we define

$$\begin{aligned}\Omega_{\text{ord}}^{(\alpha)} &= \{A \in \Omega_{\text{ord}} : |\mathbf{Ext} \Gamma(A) \cap V| \geq (1 - \alpha)L^d\}, \\ \Omega_{\text{dis}}^{(\alpha)} &= \{A \in \Omega_{\text{dis}} : |\mathbf{Ext} \Gamma(A) \cap V| \geq (1 - \alpha)L^d\}.\end{aligned}$$

Lemma 6.2 *Let $d \geq 2$ and $0 < \alpha < 1$. Then there are constants $c = c(\alpha) > 0$ and $q_0 = q_0(\alpha)$ such that for $q \geq q_0$ and $\beta \geq \beta_0$ we have*

$$\nu(\Omega_{\text{ord}} \setminus \Omega_{\text{ord}}^{(\alpha)}) \leq e^{-c\beta L^{d-1}} \quad (6.10)$$

and

$$\nu(\Omega_{\text{dis}} \setminus \Omega_{\text{dis}}^{(\alpha)}) \leq e^{-c\beta L^{d-1}}. \quad (6.11)$$

In order to prove Lemma 6.1, we will need upper bounds on Z_{tunnel} and Z_{dis} , as well as upper and lower bounds on Z_{ord} . Since contours and interfaces are suppressed if β (and hence κ) is large, the leading configurations to Z_{ord} and Z_{dis} are those without contours, giving a contribution of $e^{-e_{\text{ord}}L^d}$ and $e^{-e_{\text{dis}}L^d}$, respectively. For Z_{tunnel} , the leading configurations have a single pair of parallel interfaces of area L^{d-1} each, and no contours, giving a contribution of at most $e^{-2\kappa L^{d-1}} \max\{qe^{-e_{\text{ord}}L^d}, e^{-e_{\text{dis}}L^d}\}$. If we took only the leading configurations into account, we therefore would get that Z_{tunnel}/Z is exponentially suppressed like $e^{-2\kappa L^{d-1}}$, as required for the first bound in Lemma 6.1. But of course, this is too naive, since subleading contributions have to be taken into account. A systematic way to do this is provided by the powerful theory of Pirogov and Sinai.

6.2 Ingredients from Pirogov-Sinai Theory

In this section we will prove Lemma 6.1. To this end, we will express Z_{tunnel} in terms of partition functions that are analogs of Z_{ord} and Z_{dis} for a subset $\mathbf{\Lambda} \subset \mathbf{V}$ such that

$$\mathbf{\Lambda} \text{ is a connected component of } \mathbf{V} \setminus \partial\mathbf{V}(A_0) \text{ for some } A_0 \in \Omega. \quad (6.12)$$

We say that a contour γ is a contour in $\mathbf{\Lambda}$ if $\text{dist}(V(\gamma), \mathbf{V} \setminus \mathbf{\Lambda}) \geq 1/2$, and we say that a set of contours Γ with matching labels has external label $\ell \in \{\text{ord}, \text{dis}\}$, if the external contours in Γ have external label ℓ . We then set

$$Z_{\text{dis}}(\mathbf{\Lambda}) = \sum_{\Gamma} q^{c(\mathbf{V}_{\text{ord}})} e^{-e_{\text{dis}}|\mathbf{V}_{\text{dis}} \cap V \cap \mathbf{\Lambda}|} e^{-e_{\text{ord}}|\mathbf{V}_{\text{ord}} \cap V \cap \mathbf{\Lambda}|} \prod_{\gamma \in \Gamma} e^{-\kappa \|\gamma\|}, \quad (6.13)$$

where the sum goes over sets of contours Γ with matching labels such that the external label of Γ is dis, and similarly for $Z_{\text{ord}}(\mathbf{\Lambda})$:

$$qZ_{\text{ord}}(\mathbf{\Lambda}) = \sum_{\Gamma} q^{c(\mathbf{V}_{\text{ord}})} e^{-e_{\text{dis}}|\mathbf{V}_{\text{dis}} \cap V \cap \mathbf{\Lambda}|} e^{-e_{\text{ord}}|\mathbf{V}_{\text{ord}} \cap V \cap \mathbf{\Lambda}|} \prod_{\gamma \in \Gamma} e^{-\kappa \|\gamma\|}. \quad (6.14)$$

Note that in this notation, the partition functions Z_{dis} and Z_{ord} introduced in (6.4) are just the partition functions $Z_{\text{dis}}(\mathbf{V})$ and $Z_{\text{ord}}(\mathbf{V})$.

With these definitions, the partition function Z_{tunnel} can be rewritten as

$$Z_{\text{tunnel}} = \sum_{\mathcal{S}} \prod_{S \in \mathcal{S}} e^{-\kappa \|S\|} \prod_{\Lambda \in \mathcal{C}_{\text{dis}}(\mathcal{S})} Z_{\text{dis}}(\Lambda) \prod_{\Lambda \in \mathcal{C}_{\text{ord}}(\mathcal{S})} (q Z_{\text{ord}}(\Lambda)), \quad (6.15)$$

where the first sum goes over interface networks, while $\mathcal{C}_{\text{ord}}(\mathcal{S})$ is the set of components of $\mathbf{V} \setminus \bigcup_{S \in \mathcal{S}} S$ with $\ell(C) = \text{ord}$, and $\mathcal{C}_{\text{dis}}(\mathcal{S})$ is the set of components of $\mathbf{V} \setminus \bigcup_{S \in \mathcal{S}} S$ with $\ell(C) = \text{dis}$. The formal proof of this, by now almost obvious, identity uses again the forest structure of sets of pairwise compatible contours established in Lemma 5.5, and is similar to that of Corollary 5.11.

6.2.1 An alternative representation

We will need a representation for the partition functions $Z_{\text{dis}}(\Lambda)$ and $Z_{\text{ord}}(\Lambda)$ which does not involve the restriction of matching labels. To this end, first sum all terms in (6.13) and (6.14) which lead to the same set, Γ_{ext} , of external contours. Taking into account the forest structure of sets of pairwise compatible contours established in Lemma 5.5, this leads to the identities

$$Z_{\text{ord}}(\Lambda) = \sum_{\Gamma_{\text{ext}}} e^{-e_{\text{ord}} |\Lambda \cap \mathbf{Ext} \Gamma_{\text{ext}} \cap V|} \prod_{\gamma \in \Gamma_{\text{ext}}} e^{-\kappa \|\gamma\|} Z_{\text{dis}}(\mathbf{Int} \gamma) \quad (6.16)$$

and

$$Z_{\text{dis}}(\Lambda) = \sum_{\Gamma_{\text{ext}}} e^{-e_{\text{dis}} |\Lambda \cap \mathbf{Ext} \Gamma_{\text{ext}} \cap V|} \prod_{\gamma \in \Gamma_{\text{ext}}} e^{-\kappa \|\gamma\|} q Z_{\text{ord}}(\mathbf{Int} \gamma), \quad (6.17)$$

where the sums run over sets of mutually external contours in Λ which all have external label ord and dis , respectively.

Defining

$$K_{\text{ord}}(\gamma) = e^{-\kappa \|\gamma\|} \frac{Z_{\text{dis}}(\mathbf{Int} \gamma)}{Z_{\text{ord}}(\mathbf{Int} \gamma)}, \quad (6.18)$$

we rewrite (6.16) as

$$Z_{\text{ord}}(\Lambda) = \sum_{\Gamma_{\text{ext}}} e^{-e_{\text{ord}} |\Lambda \cap \mathbf{Ext} \Gamma_{\text{ext}} \cap V|} \prod_{\gamma \in \Gamma_{\text{ext}}} K_{\text{ord}}(\gamma) Z_{\text{ord}}(\mathbf{Int} \gamma). \quad (6.19)$$

Inserting (6.19) inductively into itself, and using the forest structure of sets of pairwise compatible contours one last time, we finally arrive at the representation

$$Z_{\text{ord}}(\Lambda) = e^{-e_{\text{ord}} |\Lambda \cap V|} \sum_{\Gamma} \prod_{\gamma \in \Gamma} K_{\text{ord}}(\gamma) \quad (6.20)$$

where the sum now runs over sets Γ of pairwise compatible contours in Λ which all have external label ord . In a similar way, one shows that

$$Z_{\text{dis}}(\Lambda) = e^{-e_{\text{dis}} |\Lambda \cap V|} \sum_{\Gamma} \prod_{\gamma \in \Gamma} K_{\text{dis}}(\gamma), \quad (6.21)$$

where the sum runs over sets Γ of pairwise compatible contours in Λ which all have external label dis , and

$$K_{\text{dis}}(\gamma) = e^{-\kappa\|\gamma\|} \frac{qZ_{\text{ord}}(\mathbf{Int} \gamma)}{Z_{\text{dis}}(\mathbf{Int} \gamma)}. \quad (6.22)$$

The representations (6.20) and (6.21) give $Z_{\text{ord}}(\Lambda)$ and $Z_{\text{dis}}(\Lambda)$ as partition functions of the so-called abstract polymer systems, see, e.g., [13] or [7], for a review. As a consequence, the logarithms of $Z_{\text{ord}}(\Lambda)$ and $Z_{\text{dis}}(\Lambda)$ can be analyzed by absolutely convergent expansions (so-called Mayer-expansions), provided the weights (6.18) and (6.22) are sufficiently small.

The following lemma gives the bounds needed to apply these expansions. Given the geometric preparations of the last section, its proof follows by a careful extension of the methods of [10, 11]. For the convenience of the reader, we give it in the appendix.

Lemma 6.3 *Let $d \geq 2$. Then there are constants $q_0 > 0$ and $c > 0$, as well as two real-valued functions $f_{\text{ord}} = f_{\text{ord}}(q, \beta)$ and $f_{\text{dis}} = f_{\text{dis}}(q, \beta)$ such that the following statements hold for $\ell \in \{\text{ord}, \text{dis}\}$, $q \geq q_0$, and $\beta \geq \beta_0$:*

(i) *Let $f = \min\{f_{\text{ord}}, f_{\text{dis}}\}$, and let $a_\ell = f_\ell - f$. If γ is a contour with external label ℓ and $a_\ell \text{diam} \gamma \leq c\beta$, then*

$$K_\ell(\gamma) \leq e^{-c\beta\|\gamma\|}. \quad (6.23)$$

(ii) *If $\Lambda \subset \mathbf{V}$ is of the form (6.12), then*

$$Z_\ell(\Lambda) \geq e^{-(f_\ell + \varepsilon_L)|\Lambda \cap V|} e^{-\|\partial\Lambda\|} \quad (6.24)$$

and

$$Z_\ell(\Lambda) \leq e^{(-f + \varepsilon_L)|\Lambda \cap V|} e^{2\|\partial\Lambda\|} \max_{\Gamma_{\text{ext}}} e^{-\frac{a_\ell}{2}|\Lambda \cap \mathbf{Ext} \Gamma_{\text{ext}} \cap V|} \prod_{\gamma \in \Gamma_{\text{ext}}} e^{-\frac{\varepsilon}{2}\beta\|\gamma\|}, \quad (6.25)$$

where the maximum goes over sets of mutually external contours in Λ which all have external label ℓ , and $\varepsilon_L = 2e^{-c\beta L}$.

(iii) $f_{\text{ord}} \leq f_{\text{dis}}$ if $\beta \geq \beta_0$, with equality if $\beta = \beta_0$.

6.2.2 Proof of Lemma 6.1

We start by noting that $\beta \geq \beta_0$ implies that $\beta \geq \frac{1}{d} \log q - 1$ and $\kappa \geq \frac{1}{2d} \log q - 1/2$ provided q is large enough (depending on d).

To prove (6.7), we combine (6.15) with (6.25), to conclude that

$$\begin{aligned} Z_{\text{tunnel}} &\leq \sum_S \prod_{S \in S} e^{-\kappa\|S\|} \prod_{\Lambda \in \mathcal{C}_{\text{dis}}(S)} e^{-f|\Lambda \cap V|} e^{2\|\partial\Lambda\| + L^d \varepsilon_L} \prod_{\Lambda \in \mathcal{C}_{\text{ord}}(S)} q e^{-f|\Lambda \cap V|} e^{2\|\partial\Lambda\| + L^d \varepsilon_L} \\ &= e^{-fL^d} \sum_S \prod_{S \in S} q e^{2L^d \varepsilon_L - (\kappa - 4)\|S\|}, \end{aligned} \quad (6.26)$$

where the sum goes over interface networks. In the last step we used that each interface bounds precisely one ordered and one disordered component. Using the facts that $q \leq e^{2d\kappa} e^d$, that there are at most $2dL(Cd)^k$ interfaces of size k in $V_{L,d}$ (for an appropriate universal

constant $C < \infty$), and that the sum over interface networks contains at least two interfaces, the bound (6.26) implies that for q and L large enough (depending on d), we have

$$\begin{aligned} Z_{\text{tunnel}} &\leq e^{-fL^d} \sum_{n \geq 2} \left(\sum_{k \geq L^{d-1}} e^{2d\kappa+d} 2dL(Cd)^k e^{2L^d \varepsilon_L - (\kappa-4)k} \right)^n \\ &\leq e^{-fL^d} e^{-\frac{\beta}{2}L^{d-1}}. \end{aligned} \quad (6.27)$$

Applying the bound (6.25) to $Z_{\text{dis}} = Z_{\text{dis}}(V_{L,d})$, we get

$$Z_{\text{dis}} \leq e^{-fL^d + L^d \varepsilon_L} \leq e^{-fL^d} (1 + e^{-c\beta L/2}),$$

while the bound (6.24), together with the fact that $f_{\text{ord}} = f$ if $\beta \geq \beta_0$, gives

$$Z_{\text{ord}} \geq e^{-fL^d - L^d \varepsilon_L} \geq e^{-fL^d} (1 - e^{-c\beta L/2}).$$

Together with (6.4), this gives the first statement of the lemma.

To prove the second statement, we use that $f = f_{\text{dis}} = f_{\text{ord}}$ if $\beta = \beta_0$ which, together with (6.24) and (6.25) gives

$$|\log Z_{\text{dis}} + fL^d| \leq L^d \varepsilon_L \leq e^{-c\beta L/2} \quad \text{and} \quad |\log Z_{\text{ord}} + fL^d| \leq L^d \varepsilon_L \leq e^{-c\beta L/2}.$$

6.3 Large Deviation Bounds

We now prove Lemma 6.2, by starting with the proof of (6.10). Let

$$Z_{\text{ord}}^{<\alpha} = \frac{1}{q} \sum_{\substack{A \in \Omega_{\text{ord}}: \\ |\mathbf{Ext} \Gamma(A) \cap V| < (1-\alpha)L^d}} w(A),$$

so that

$$\nu(\Omega_{\text{ord}} \setminus \Omega_{\text{ord}}^{(\alpha)}) = \frac{q Z_{\text{ord}}^{<\alpha}}{Z} \leq \frac{Z_{\text{ord}}^{<\alpha}}{Z_{\text{ord}}}. \quad (6.28)$$

Proceeding as in the derivation of (6.20), we rewrite

$$Z_{\text{ord}}^{<\alpha} = e^{-e_{\text{ord}} L^d} \sum_{\Gamma: |\mathbf{Ext} \Gamma \cap V| < (1-\alpha)L^d} \prod_{\gamma \in \Gamma} K_{\text{ord}}(\gamma),$$

where the sum runs over sets Γ of pairwise compatible contours in $V_{L,d}$, all of which have external label ord . Using the fact that $|\mathbf{Ext} \Gamma \cap V| < (1-\alpha)L^d$ implies that $\sum_{\gamma \in \Gamma} |\mathbf{Int} \gamma \cap V| \geq \alpha L^d$, we then bound

$$Z_{\text{ord}}^{<\alpha} \leq e^{-e_{\text{ord}} L^d} \sum_{\Gamma} e^{h(\sum_{\gamma} |\mathbf{Int} \gamma \cap V| - \alpha L^d)} \prod_{\gamma \in \Gamma} K(\gamma) = e^{-\alpha h L^d} Z_{\text{ord}}^{(h)}, \quad (6.29)$$

where

$$Z_{\text{ord}}^{(h)} = e^{-e_{\text{ord}} L^d} \sum_{\Gamma} \prod_{\gamma \in \Gamma} e^{h|\mathbf{Int} \gamma \cap V|} K_{\text{ord}}(\gamma).$$

Next we estimate the dependence of $Z_{\text{ord}}^{(h)}$ on h . To this end, we set $K_h(\gamma) = e^{h|\mathbf{Int} \gamma \cap V|} K_{\text{ord}}(\gamma)$ and bound

$$\begin{aligned} \frac{d}{dh} Z_{\text{ord}}^{(h)} &= e^{-e_{\text{ord}} L^d} \sum_{\Gamma} \left(\sum_{\gamma \in \Gamma} |\mathbf{Int} \gamma \cap V| \right) \prod_{\gamma \in \Gamma} K_h(\gamma) \\ &= e^{-e_{\text{ord}} L^d} \sum_{\gamma \subset V_{L,d}} |\mathbf{Int} \gamma \cap V| \sum_{\Gamma \ni \gamma} \prod_{\gamma \in \Gamma} K_h(\gamma) \\ &\leq \left(\sum_{\gamma \subset V_{L,d}} |\mathbf{Int} \gamma \cap V| K_h(\gamma) \right) Z_{\text{ord}}^{(h)}, \end{aligned}$$

where we used a standard Peierl's argument in the last step, and the sums over $\gamma \subset V_{L,d}$ denote sums over contours in $V_{L,d}$ with external label ord . Hence,

$$\begin{aligned} \frac{d}{dh} \log Z_{\text{ord}}^{(h)} &\leq \sum_{\gamma \subset V_{L,d}} |\mathbf{Int} \gamma \cap V| K_h(\gamma) \\ &\leq \sum_{\gamma \subset V_{L,d}} \frac{1}{2} \|\gamma\|_L e^{h\|\gamma\|_L} e^{-c\beta\|\gamma\|}, \end{aligned}$$

where we used the bounds (5.5) and (6.23), together with the fact that $f_{\text{ord}} = f$ for $\beta \geq \beta_0$, in the last step.

Now assume that $h \leq c\beta/(2L)$. Since there are most $dL^d(Cd)^k$ contours of size k in $V_{L,d}$, we conclude that

$$\frac{d}{dh} \log Z_{\text{ord}}^{(h)} \leq \frac{1}{2} \sum_{k \geq 2} dL^d(Cd)^k k L e^{-c\beta k/2} \leq \frac{\alpha}{4} L^d,$$

and thus

$$Z_{\text{ord}}^{(h)} \leq Z_{\text{ord}} e^{\alpha h L^d/4},$$

provided β is large enough (depending on d and α .) Inserted into (6.29) and (6.28), this gives

$$\nu(\Omega_{\text{ord}} \setminus \Omega_{\text{ord}}^\alpha) \leq e^{-3\alpha h L^d/4} = e^{-3\alpha\beta L^{d-1}/8},$$

where we have set h equal to $c\beta/(2L)$ in the last step.

This proves (6.10). For $a_{\text{dis}}L \leq c\beta$, the bound (6.11) is proved in exactly the same way, but for $a_{\text{dis}}L > c\beta$, this strategy does not work, since (6.23) is not at our disposal anymore. In stead, we use (6.24), (6.25) and (5.5) to bound

$$\begin{aligned} \nu(\Omega_{\text{dis}} \setminus \Omega_{\text{dis}}^{(\alpha)}) &\leq \nu(\Omega_{\text{dis}}) \leq \frac{Z_{\text{dis}}}{Z_{\text{ord}}} \leq e^{2L^d \varepsilon_L} \max_{\Gamma_{\text{ext}}} e^{-\frac{\alpha_{\text{dis}}}{2} |\mathbf{Ext} \Gamma_{\text{ext}} \cap V|} \prod_{\gamma \in \Gamma_{\text{ext}}} e^{-\frac{\alpha}{2} \beta \|\gamma\|} \\ &\leq e^{2L^d \varepsilon_L} \max_{\Gamma_{\text{ext}}} e^{-\frac{c\beta}{2L} |\mathbf{Ext} \Gamma_{\text{ext}} \cap V|} \prod_{\gamma \in \Gamma_{\text{ext}}} e^{-\frac{c\beta}{L} |\mathbf{Int} \gamma \cap V|} \\ &\leq e^{2L^d \varepsilon_L} e^{-\frac{c\beta}{2L} L^d}, \end{aligned}$$

where we used that $|\mathbf{Ext} \Gamma_{\text{ext}} \cap V| = L^d - \sum_{\gamma \in \Gamma_{\text{ext}}} |\mathbf{Int} \gamma \cap V|$ in the last step. This concludes the proof of the lemma.

7 Lower bounds on the mixing time

Lemma 6.1 and Lemma 6.2 give us the necessary ingredients to prove our main result. We start by proving a key theorem which expresses the statements of these lemmas in terms of the probability measure μ on spin configurations, rather than the probability measure ν on SW configurations.

7.1 An Important Probabilistic Estimate

Recall that for a spin configuration $\sigma \in [q]^V$, we defined $E(\sigma)$ to be the set of all edges whose two endpoints have the same color. We also introduce the set, $\mathcal{C}(\sigma)$, of connected components of the graph $(V, E(\sigma))$, that is, the set of monochromatic components.

Theorem 7.1 *Let $d \geq 2$ and $0 < \alpha < 1/2$. Then there are constants $c > 0$, $q_0 < \infty$ and $L_0 < \infty$ such that for $q \geq q_0$ and $L \geq L_0$ the following statements hold:*

(a) *If $\beta = \beta_0$, then*

$$\mu\left(\alpha dL^d \leq |E(\sigma)| \leq (1 - \alpha)dL^d\right) \leq e^{-c\beta L^{d-1}} \quad (7.1)$$

and

$$\left|\mu\left(|E(\sigma)| \geq (1 - \alpha)dL^d\right) - \frac{q}{q+1}\right| \leq e^{-c\beta L}. \quad (7.2)$$

(b) *If $\beta \geq \beta_0$, then*

$$\mu\left(\alpha L^d < \max_{C \in \mathcal{C}(\sigma)} |V(C)| < (1 - \alpha)L^d\right) \leq e^{-c\beta L^{d-1}} \quad (7.3)$$

and

$$\mu\left(\max_{C \in \mathcal{C}(\sigma)} |V(C)| \geq (1 - \alpha)L^d\right) \geq \frac{q}{q+1} - e^{-c\beta L}. \quad (7.4)$$

Proof. To relate the statements of Lemmas 6.1 and 6.2 to the theorem, we use that both the spin measure μ and the FK-measure ν are marginals of the Edwards-Sokal measure π . Consider thus a configuration (σ, A) with positive measure $\pi((\sigma, A))$. Under this condition, all spins in a component of (V, A) must have the same color, implying in particular that

$$A \subset E(\sigma) \quad \text{and} \quad \max_{C \in \mathcal{C}(\sigma)} |V(C)| \geq \max_{C \in \mathcal{C}(V, A)} |V(C)|. \quad (7.5)$$

It turns out that with high probability, $|A|$ is not much smaller than $|E(\sigma)|$ either. More precisely, we will prove that for $0 < \alpha < 1$,

$$\pi\left(|E(\sigma)| \geq |A| + \alpha dL^d\right) \leq e^{-\alpha\beta L^{d-1}}. \quad (7.6)$$

To prove this bound, we rewrite the left hand side as

$$\pi\left(|E(\sigma)| \geq |A| + \alpha dL^d\right) = \sum_{\sigma} \pi\left(|E(\sigma)| \geq |A| + \alpha dL^d \mid \sigma\right) \mu(\sigma).$$

But given σ , a configuration A according to the conditional measure $\pi(\cdot \mid \sigma)$ is obtained by deleting the edges in $E(\sigma)$ independently with probability $e^{-\beta}$. The number of deleted edges is therefore equal to a binomial random variable with parameters m and $e^{-\beta}$, where $m = |E(\sigma)| \leq dL^d$. We now bound the probability that the number of deleted edges X is larger than αdL^d as follows:

$$\Pr(X \geq \alpha dL^d) = \sum_{k \geq \alpha dL^d} \binom{m}{k} e^{-\beta k} (1 - e^{-\beta})^{m-k} \leq e^{-\beta \alpha dL^d} 2^{dL^d}. \quad (7.7)$$

Recalling that $\beta \geq \beta_0 = \frac{1}{d} \log q + O(q^{-1/d})$, we see that for q (and hence β) large enough, the right hand side is bounded by $e^{-\beta \alpha dL^d/2} \leq e^{-\beta \alpha L^{d-1}}$. This implies the bound (7.6).

Next we observe that if $A \in \Omega_{\text{ord}}^{(\alpha)}$ then $\mathbf{V}(A) \supset \mathbf{Ext} \Gamma(A)$ (by the definition of Ω_{ord}) and thus $|V(A)| = |\mathbf{V}(A) \cap V| \geq |\mathbf{Ext} \Gamma(A) \cap V| \geq (1 - \alpha)L^d$. This in turn implies $|V \setminus V(A)| \leq \alpha L^d$, and thus

$$|A| = dL^d - |\{xy \in E \mid \{x, y\} \cap V \setminus V(A) \neq \emptyset\}| \geq dL^d - 2d\alpha L^d = (1 - 2\alpha)dL^d.$$

In summary

$$A \in \Omega_{\text{ord}}^{(\alpha)} \implies |V(A)| \geq (1 - \alpha)L^d \implies |A| \geq (1 - 2\alpha)dL^d. \quad (7.8)$$

In a similar way

$$A \in \Omega_{\text{dis}}^{(\alpha)} \implies |V \setminus V(A)| \geq (1 - \alpha)L^d \implies |A| \leq \alpha dL^d. \quad (7.9)$$

We now turn to the first bound of the theorem. Taking into account the bounds (7.5) and (7.6), we have that the left hand side of (7.1) is bounded by

$$\nu\left(\frac{1}{2}\alpha dL^d < |A| < (1 - \alpha)dL^d\right) + e^{-\frac{1}{2}\alpha\beta L^{d-1}}.$$

But by (7.8) and (7.9), the probability of the event $\frac{1}{2}\alpha dL^d < |A| < (1 - \alpha)dL^d$ is bounded $\mu(\Omega \setminus (\Omega_{\text{ord}}^{(\alpha/2)} \cup \Omega_{\text{dis}}^{(\alpha/2)}))$, which in turn can be bounded with the help of the bounds (6.7), (6.10) and (6.11). Putting everything together, we therefore obtain that there exists a constant $c > 0$ depending on α and d such that

$$\mu\left(\alpha dL^d < |E(\sigma)| < (1 - \alpha)dL^d\right) \leq 4e^{-c\beta L^{d-1}} \leq e^{-\frac{c}{2}\beta L^{d-1}},$$

provided q (and hence β) is large enough. This proves the bound (7.1).

The proof of the bound (7.2) is similar. Indeed, using again that $E(\sigma) \supset A$ if $\pi((\sigma, A)) > 0$, we have

$$\begin{aligned} \mu\left(|E(\sigma)| \geq (1 - \alpha)dL^d\right) &\geq \nu\left(|A| \geq (1 - \alpha)dL^d\right) \\ &\geq \nu(\Omega_{\text{ord}}^{(\alpha/2)}) \geq \frac{q}{q+1} - e^{-c\beta L^{d-1}} - e^{-c\beta L}. \end{aligned}$$

where we used the implication (7.8) in the second to last step, and the bounds (6.8) and (6.10) in the last step. To prove the upper bound in (7.2), we start with the bound (7.6), and that $(1 - \alpha)^2 = 1 - \alpha - \alpha(1 - \alpha)$, in saying that

$$\mu\left(|E(\boldsymbol{\sigma})| \geq (1 - \alpha)dL^d\right) \leq \nu\left(|A| \geq (1 - \alpha)^2dL^d\right) + e^{-c\beta L^{d-1}},$$

where $c = \alpha(1 - \alpha)$. Let $\alpha' = 1 - (1 - \alpha)^2$ and $\tilde{\alpha} = \min\{\alpha', 1 - \alpha'\}$. If $|A| \geq (1 - \alpha')dL^d$, then $A \in \Omega \setminus \Omega_{\text{dis}}^{(\tilde{\alpha})}$ by the implication (7.9). Combining the bounds (6.7), (6.10) and (6.11), we may now conclude that for $\beta = \beta_0$, we have

$$\begin{aligned} \mu\left(|E(\boldsymbol{\sigma})| \geq (1 - \alpha)dL^d\right) &\leq \nu\left(\Omega_{\text{ord}}^{(\alpha')}\right) + 4e^{-\beta c L^{d-1}} \\ &\leq \frac{q}{q+1} + 4e^{-\beta c L^{d-1}} + e^{-c\beta L} \leq e^{-c\beta L/2}, \end{aligned}$$

where we also used the bound (6.9) in the second to last step. This completes the proof of the bound (7.2).

The bounds of part (b) are proved in a similar way. Indeed, let $A \in \Omega_{\text{ord}}^{(\alpha)}$ and let A_{ext} be the set of edges with both endpoints in $\mathbf{Ext} \Gamma(A)$. By Lemma 4.5 (iii), the graph $(V(A_{\text{ext}}), A_{\text{ext}})$ is connected, implying that $\max_{C \in \mathcal{C}(V, A)} |V(C)| \geq (1 - \alpha)L^d$ whenever $A \in \Omega_{\text{ord}}^{(\alpha)}$. Taking into account the second bound in (7.5) we get that

$$A \in \Omega_{\text{ord}}^{(\alpha)} \text{ and } \pi((\boldsymbol{\sigma}, A)) > 0 \implies \max_{C \in \mathcal{C}(\boldsymbol{\sigma})} |V(C)| \geq (1 - \alpha)L^d. \quad (7.10)$$

Together with (6.8) and (6.10), this immediately gives the bound (7.4).

We are thus left with the bound (7.3). To prove this bound, we note that $|V(A)| \leq \frac{1}{2}\alpha L^d$ and $|E(\boldsymbol{\sigma})| \leq |A| + \frac{1}{2}\alpha L^d$ implies the largest component of $(V, E(\boldsymbol{\sigma}))$ is at most of size $|V(A)| + \frac{1}{2}\alpha L^d \leq \alpha L^d$. Combined with the bound (7.6), we conclude that there exists a constant $c > 0$ (depending on d and α) such that

$$\pi\left(A \in \Omega_{\text{dis}}^{(\alpha/2)} \text{ and } \max_{C \in \mathcal{C}(\boldsymbol{\sigma})} |V(C)| \geq \alpha L^d\right) \leq e^{-c\beta L^{d-1}}.$$

Combined with (7.10), this implies that the left hand side of (7.3) can be bounded by

$$\mu\left(\Omega \setminus (\Omega_{\text{ord}}^{(\alpha)} \cup \Omega_{\text{dis}}^{(\alpha/2)})\right) + e^{-c\beta L^{d-1}}.$$

Together with the bounds (6.7), (6.10) and (6.11), this gives the desired bound (7.3). \square

7.2 Proof of the SW bound in Theorem 1.2

Let $\beta = \beta_0$. Let $S = \{\boldsymbol{\sigma} : |E(\boldsymbol{\sigma})| \geq (1 - \alpha)dL^d\}$. We will show that Φ_S is exponentially small in βL^{d-1} , which will establish the theorem. For q (and hence β_0) large enough, $\pi(S) \geq 1/2$, using (7.2). Also, $\pi(S^c) \geq 1/q - e^{-c\beta L} \geq 1/2q$, if L is large enough. Thus

$$\Phi_S = \frac{Q(S, S^c)}{\mu(S)\mu(S^c)} \leq 4qQ(S, S^c). \quad (7.11)$$

Let $S_0 = \{\boldsymbol{\sigma} : \alpha dL^d \leq |E(\boldsymbol{\sigma})| \leq (1 - \alpha)L^d\}$. Then

$$Q(S, S^c) = Q(S, S^c \setminus S_0) + Q(S, S_0). \quad (7.12)$$

Now

$$Q(S, S_0) = Q(S_0, S) \leq \pi(S_0) \leq e^{-c\beta L^{d-1}},$$

using (7.1), while

$$Q(S, S^c \setminus S_0) = \pi(S) \Pr(E(\boldsymbol{\sigma}') < \alpha dL^d \mid E(\boldsymbol{\sigma}) \geq (1 - \alpha)dL^d),$$

where $\boldsymbol{\sigma}$ is a μ -random spin configuration, and $\boldsymbol{\sigma}'$ is constructed from one step of the SW algorithm. The above probability is in turn the probability that at least $(1 - 2\alpha)dL^d$ edges are deleted in one step of the SW algorithm, which is at most $(2e^{-\beta(1-2\alpha)})^{dL^d}$, using (7.7). Choosing $\alpha = 1/3$, and q large (so that β is large) and L sufficiently large, yields:

$$Q(S, S^c \setminus S_0) \leq \pi(S) (2e^{-\beta(1-2\alpha)})^{dL^d} \leq e^{-c\beta L^{d-1}}.$$

This implies the desired bound on conductance,

$$\Phi(P^{\text{SW}}) \leq e^{-c\beta L^{d-1}},$$

which together with (2.9) concludes the proof of the lower bound (1.9) on the mixing time of the SW algorithm. \square

7.3 Proof of the HB bound in Theorem 1.2

Let $\boldsymbol{\sigma} \in [q]^V$ and let C be a component of $(V, E(\boldsymbol{\sigma}))$. We say that C has color k , if $\sigma_x = k$ for all $x \in V(C)$, and denote the set of components of color k by $\mathcal{C}_k(\boldsymbol{\sigma})$. For $k = 1, \dots, q$ and $0 < \alpha < 1/2$, we then define

$$\widehat{\Omega}_k^{(\alpha)} = \left\{ \boldsymbol{\sigma} \in [q]^V : \exists C \in \mathcal{C}_k(\boldsymbol{\sigma}) \text{ s.t. } |V(C)| \geq (1 - \alpha)|V| \right\}.$$

Note that the sets $\widehat{\Omega}_k^{(\alpha)}$ are mutually disjoint, so by symmetry and the bound (7.4), we have

$$\mu(\widehat{\Omega}_k^{(\alpha)}) \geq \frac{1}{q+1} - \frac{1}{q} e^{-c\beta L}. \quad (7.13)$$

Finally, let

$$\widehat{\Omega}_{\text{dis}}^{(\alpha)} = \left\{ \boldsymbol{\sigma} \in [q]^V : \max_{C \in \mathcal{C}(\boldsymbol{\sigma})} |V(C)| \leq \alpha L^d \right\}.$$

We complete our proof by estimating Φ_S (see (2.7)) for $S = \widehat{\Omega}_1^{(\alpha)}$. First notice that for q (and hence also β) sufficiently large $\mu(S)\mu(S^c) \geq 1/4q$ so that

$$\Phi_S \leq 4qQ(S, S^c).$$

Since the heat bath algorithm can only change one vertex at a time, it does not make transitions between the different sets $\widehat{\Omega}_k^{(\alpha)}$. For α small enough, it cannot make transitions

between $\widehat{\Omega}_1^{(\alpha)}$ and $\Omega_{\text{dis}}^{(\alpha)}$ either. Indeed, changing the color of a single vertex can not break a component $C \in \mathcal{C}(\sigma)$ into more than $2d + 1$ new components: in the worst case, C gets broken into a single component of size one and $2d$ components of size $((1 - \alpha)L^d - 1)/(2d)$. For α sufficiently small (say $\alpha = 1/(4d)$), the heat bath algorithm therefore cannot make transitions between $\widehat{\Omega}_1^{(\alpha)}$ and $\Omega_{\text{dis}}^{(\alpha)}$. Defining S_0 as the set of configurations which are neither in $\Omega_{\text{dis}}^{(\alpha)}$ nor in one of the sets $\widehat{\Omega}_k^{(\alpha)}$, we thus have

$$Q(S, S^c) = Q(S, S_0) = Q(S_0, S) \leq \mu(S_0) \leq e^{-c\beta L^{d-1}},$$

where we have used the bound (7.3) in the last step. Recalling that $\beta \geq \beta_0 = \frac{1}{d} \log q + O(q^{-1/d})$, we see that for L sufficiently large, we have $\Phi_S \leq 4qQ(S, S^c) \leq e^{-c\beta L^{d-1}/2}$, as required. \square

Acknowledgements: The authors wish to thank Marek Biskup, Michael Freedman, Roman Kotecký, Fabio Martinelli and Jacob Lurie for numerous helpful discussions. The authors also thank their collaborators Alan Frieze, Jeong Han Kim, Eric Vigoda, and Van Vu for their efforts on earlier research (reported in [8]) related to the present contribution.

References

- [1] M. Aizenman, J.T. Chayes, L. Chayes, and C.M. Newman, Discontinuity of the magnetization in the one-dimensional $1/|x - y|^2$ Ising and Potts models. *J. Stat. Phys.* **50** 1–40, 1988.
- [2] D. Aldous and J. Fill, *Reversible Markov Chains and Random Walks on Graphs*, in preparation. Some chapters available at <http://stat-www.berkeley.edu/pub/users/aldous/book.html>.
- [3] P.S. Alexandrov, *Combinatorial Topology*, Dover Publications, New York, 1998.
- [4] I. Benjamini and E. Mossel, On the mixing time of a simple random walk on the supercritical percolation cluster, *Prob. Theor. Rel. Fields*, 125:408-420, 2003.
- [5] N. Berger, C. Kenyon, E. Mossel and Y. Peres, Glauber dynamics on trees and hyperbolic graphs, *Prob. Theor. Rel. Fields*, 131:311-340, 2005.
- [6] B. Bollobás and I. Leader, Edge-isoperimetric inequalities in the grid, *Combinatorica* 11:299-314, 1991.
- [7] C. Borgs, Statistical Physics Expansion Methods in Combinatorics and Computer Science, CBMS Lectures, Memphis 2003.
- [8] C. Borgs, J. T. Chayes, A. Frieze, J. H. Kim, P. Tetali, E. Vigoda, V. Vu, Torpid Mixing of Some Monte Carlo Markov Chain Algorithms in Statistical Physics, *Proceedings of the 40th IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE Computer Society Press, 1999, 218-229.

- [9] C. Borgs and J. Imbrie, A unified approach to phase diagrams in field theory and statistical mechanics, *Commun. Math. Phys.* 123:305-328, 1989.
- [10] C. Borgs and R. Kotecký, A rigorous theory of finite-size scaling at first-order transitions, *Jour. Statist. Phys.* 61:79-110, 1990.
- [11] C. Borgs, R. Kotecký and S. Miracle-Solé, Finite-size scaling for Potts models, *Jour. Statist. Phys.*, 62:529-551, 1991.
- [12] R. Bott, L.W. Tu, *Differential Forms in Algebraic Topology*, Graduate Texts in Mathematics **82**, Springer Verlag, New York, Heidelberg, Berlin, 1986.
- [13] D.C. Brydges: A short course on cluster expansions, *Phénomènes critiques, systèmes aléatoires, théories de jauge, Part I, II (Les Houches, 1984)*, 129–183, North-Holland, Amsterdam, 1986.
- [14] F. Cesi, C. Maes, and F. Martinelli: Relaxation to equilibrium for two-dimensional disordered Ising systems in the Griffiths phase, *Comm. Math. Phys.*, **189** (1997) 323–335
- [15] F. Chen, L. Lovász, I. Pak, Unpublished appendix for *Lifting Markov chains to speed up mixing*. Proc. 31st ACM Symp. on Theory of Comp. (1999) 275–281.
- [16] C. Cooper, M. Dyer, A. M. Frieze and R. Rue, Mixing properties of the Swendsen-Wang process on the complete graph and narrow grids, *J. Math. Phys.* 41(3):1499–1527, 2000.
- [17] C. Cooper and A. M. Frieze, Mixing properties of the Swendsen-Wang process on classes of graphs, *Random Structures and Algorithms*, 15 (3-4): 242-261, 1999.
- [18] R. L. Dobrushin, An investigation of Gibbs states for three-dimensional lattice systems (Russian with English summary), *Teor. Veroyatnost. i Primenen.* 18: 261–279, 1973.
English translation: *Theor. Probability Appl.* 18: 253–271, 1974.
- [19] M. E. Dyer, A. M. Frieze and M. R. Jerrum, On counting independent sets in sparse graphs, *Siam J. Comput.* 31 (2002) 1527-1541.
- [20] M. E. Dyer and C. Greenhill, On Markov chains for independent sets, *J. Algorithms*, 35(1):17-49, 2000.
- [21] 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.
- [22] C. Fortuin and P. Kasteleyn, On the random cluster model I: Introduction and relation to other models, *Physica* 57: 536-564, 1972.
- [23] N. Fountoulakis and B.A. Reed, Faster mixing and small bottlenecks, *Prob. Theor. Rel. Fields*, 137:475-486, 2007.

- [24] V. Gore and M. Jerrum, The Swendsen-Wang process does not always mix rapidly, *J. Statist. Phys.* 97(1-2):67-86, 1999.
- [25] M. Huber, Efficient exact sampling from the Ising model using Swendsen-Wang, *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1999, 921-922.
- [26] M.R. Jerrum, Talk given at Workshop on Randomised Approximation and Stochastic Simulation in Warwick, England, 1998.
- [27] G.F. Lawler and A.D. Sokal, Bounds on the L_2 spectrum for Markov chains and Markov processes: a generalization of Cheeger's inequality, *Trans. Amer. Math. Soc.* **309**, (1988), 557-580.
- [28] L. Laanait, A. Messenger, S. Miracle-Solé, J. Ruiz and S. Shlossman, Interfaces in the Potts models I: Pirogov-Sinai theory of the Fortuin-Kasteleyn representation, *Commun. Math. Phys.* 140(1):81-91, 1991.
- [29] J.L. Lebowitz and A.E. Mazel, Improved Peierls argument for higher dimensional Ising models, *Jour. Statist. Phys.* 90:1051-1059, 1998.
- [30] 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.
- [31] M. Luby and E. Vigoda, Fast convergence of the Glauber dynamics for sampling independent sets, *Random Structures and Algorithms*, 15 (3-4): 229-241, 1999.
- [32] F. Martinelli, Dynamical Analysis of the low temperature cluster algorithms, *Jour. Statist. Phys.*, 66:1245-1276, 1992.
- [33] F. Martinelli: Lectures on Glauber dynamics for discrete spin models, *Lectures on probability theory and statistics (Saint-Flour, 1997)*, Lecture Notes in Math., **1717**, 93-191, Springer, Berlin, 1999.
- [34] R. Montenegro and P. Tetali, *Mathematical Aspects of Mixing Times in Markov Chains*, Foundations & Trends in Theoretical Computer Science, 1 (3):237-354, 2006.
- [35] S. A. Pirogov and Ya.G. Sinai, Phase diagrams of classical lattice systems, *Theor. Math. Phys.* , 25:1185-1192, 1975. Phase diagrams of classical lattice systems. (Continuation), *Theor. Math. Phys.* 26:39-49, 1976.
- [36] D. Ruelle, *Statistical Mechanics: Rigorous Results*, W. A. Benjamin, 1969.
- [37] A. J. Sinclair and M. R. Jerrum, Approximate counting, uniform generation and rapidly mixing Markov chains, *Information and Computation*, 82:93-133, 1989.
- [38] R. Swendsen and J-S. Wang, Non-universal critical dynamics in Monte-Carlo simulation, *Phys. Rev. Lett.*, 58:86-88, 1987.

- [39] L. Thomas, Bound on the mass gap for finite volume stochastic Ising models at low temperature, *Commun. Math. Phys.* 126:1-11, 1989.
- [40] D. J. A. Welsh, *Complexity: Knots, Colourings and Counting*, London Mathematical Society Lecture Note Series 186, 1993.
- [41] F. Y. Wu, The Potts model, *Rev. Mod. Phys.* 54:235-268, 1982.

A Technical Estimates using Truncation

Throughout this section, we will assume that

$$\kappa \geq \max\left\{C_1, C_2 \log d, \frac{1}{2d} \log q - \frac{1}{2}\right\}, \quad (\text{A.1})$$

where C_1, C_2 are universal constants to be determined in the course of the proof below. In fact, we will prove statements (i) and (ii) of Lemma 6.3 for all β such that (A.1) holds.

Also, for the purpose of this appendix, we will use the symbol $|\mathbf{\Lambda}|$ for the cardinality of the set $V \cap \mathbf{\Lambda}$, so that expressions of the form $|V \cap \mathbf{Ext} \Gamma \cap \mathbf{\Lambda}|$ can be simplified to $|\mathbf{\Lambda} \cap \mathbf{Ext} \Gamma|$.

A.1 Truncated Contour Models

We need some preparation. We start by bounding the factor $qe^{-\kappa\|\gamma\|}$ appearing in the weight (6.22) of a contour with external label dis . To this end, we observe that the smallest contour with disordered external label has size $\|\gamma\| \geq 4d - 2$. Combined with the assumption (A.1) and the assumption $d \geq 2$, this gives

$$\begin{aligned} qe^{-\kappa\|\gamma\|} &\leq e^{2d\kappa+d} e^{-\kappa\|\gamma\|} \leq \exp\left(\left(\frac{2d\kappa+d}{4d-2} - \kappa\right)\|\gamma\|\right) \\ &\leq \exp\left(\left(\frac{4\kappa+2}{6} - \kappa\right)\|\gamma\|\right) \leq e^{-\frac{\kappa}{4}\|\gamma\|} \end{aligned} \quad (\text{A.2})$$

provided $C_1 \geq 8$.

As usual in Pirogov-Sinai theory, we next introduce a truncated model. It is given in terms of the truncated activities

$$K'_\ell(\gamma) = \min\{K_\ell(\gamma), e^{-(\kappa/4-4-c\beta)\|\gamma\|}\} \quad (\text{A.3})$$

and the corresponding partition functions

$$Z'_\ell(\mathbf{\Lambda}) = e^{-e_\ell|\mathbf{\Lambda}|} \sum_{\Gamma} \prod_{\gamma \in \Gamma} K'_\ell(\gamma), \quad (\text{A.4})$$

where $\ell \in \{\text{ord}, \text{dis}\}$, the sum in (A.4) runs over sets Γ of pairwise compatible contours in $\mathbf{\Lambda}$ which all have external label ℓ , and c is a suitable constant (we will choose $c = 1/20$).

Recall that there are no more than $(Cd)^k$ contours of size k which intersect a given edge \mathbf{e} . As a consequence, there exists a universal constant $C_2 < \infty$ such that the following holds whenever κ obeys the bound (A.1) and $\ell \in \{\text{ord}, \text{dis}\}$. If $x \in V$, then

$$\sum_{\gamma: \mathbf{Int} \gamma \ni x} K'_\ell(\gamma) e^{(c\beta+1)\|\gamma\|} \leq 1, \quad (\text{A.5})$$

and if γ_0 is a contour or an interface, then

$$\sum_{\gamma: \text{dist}(\gamma, \gamma_0) < 1/2} K'_\ell(\gamma) e^{(c\beta+1)\|\gamma\|} \leq \|\gamma_0\|. \quad (\text{A.6})$$

The above bounds imply absolute convergence of the cluster expansions for abstract polymer systems, which in turn gives the existence of the limits

$$f_\ell = \lim_{L \rightarrow \infty} f_\ell^{(L)} \quad \text{with} \quad f_\ell^{(L)} = -\frac{1}{L^d} \log Z'_\ell(V_{L,d}), \quad (\text{A.7})$$

where $\ell \in \{\text{ord}, \text{dis}\}$ and $V_{L,d}$ denotes the d -dimensional torus of sidelength L , see, e.g., [13, 7] for a review of cluster expansions for abstract polymer systems. These methods also imply that, for $\ell \in \{\text{ord}, \text{dis}\}$, and $\mathbf{\Lambda} \subset \mathbf{V}$ of the form (6.12), we have $|f_\ell - f_\ell^{(L)}| \leq \varepsilon_L$ and

$$\left| \log Z'_\ell(\mathbf{\Lambda}) + f_\ell |\mathbf{\Lambda}| \right| \leq \|\partial \mathbf{\Lambda}\| + \varepsilon_L |\mathbf{\Lambda}| \quad (\text{A.8})$$

where $\varepsilon_L = 2e^{-c\beta L}$. We will assume that C_1 in (A.1) has been chosen in such a way that $L\varepsilon_L \leq 1$.

A.2 Proof of Lemma 6.3 (i) and (ii)

We first note that $Z_\ell(\mathbf{\Lambda}) \geq Z'_\ell(\mathbf{\Lambda})$, so in view of (A.8), we have

$$Z_\ell(\mathbf{\Lambda}) \geq e^{-(f_\ell + \varepsilon_L)|\mathbf{\Lambda}|} e^{-\|\partial \mathbf{\Lambda}\|}. \quad (\text{A.9})$$

Furthermore, for κ large enough, $K'_\ell(\gamma) \leq e^{-c\beta\|\gamma\|}$, implying that $K_\ell(\gamma) \leq e^{-c\beta\|\gamma\|}$ whenever $K_\ell(\gamma) = K'_\ell(\gamma)$. To prove Lemma 6.3 (i) and (ii), it is therefore enough to establish the following lemma:

Lemma A.1 *Under the condition (A.1), we have that*

- (i) $K_\ell(\gamma) = K'_\ell(\gamma)$ whenever γ is a contour with external label ℓ and $a_\ell \text{diam} \gamma \leq c\beta$.
- (ii) For all $\mathbf{\Lambda}$ of the form (6.12),

$$Z_\ell(\mathbf{\Lambda}) \leq e^{(\varepsilon_L - f)|\mathbf{\Lambda}|} e^{2\|\partial \mathbf{\Lambda}\|} \max_{\Gamma_{\text{ext}}} e^{-\frac{a_\ell}{2}|\mathbf{\Lambda} \cap \text{Ext} \Gamma_{\text{ext}}|} \prod_{\gamma \in \Gamma_{\text{ext}}} e^{-\frac{c}{2}\beta\|\gamma\|} \quad (\text{A.10})$$

where the maximum goes over sets of mutually external contours in $\mathbf{\Lambda}$ which all have external label ℓ .

Proof. We prove the lemma by induction on the level of Λ and γ . Here the level of a set Λ is defined to be zero if there are no contours γ such that γ is a contour in Λ . The level of a set Λ is defined to be k if the highest level of a contour γ in Λ is $k - 1$, with the level of a contour inductively defined to be 1 plus the level of its interior. Note that with this definition, the levels of two contours γ, γ' with $\gamma < \gamma'$ differ by at least two.

Assume that Λ has level 0. Then $Z_\ell(\Lambda) = e^{-e_\ell|\Lambda|} \leq e^{-f_\ell|\Lambda|} = e^{-f|\Lambda|}e^{-a_\ell|\Lambda|}$, which proves (A.10) for sets Λ of level 0.

Next assume that the bound (A.10) in (ii) has been proven for all sets Λ of level k or less. If γ has level $k + 1$ or less, then

$$\begin{aligned} K_{\text{dis}}(\gamma) &\leq e^{-\frac{1}{4}\kappa\|\gamma\|} \frac{Z_{\text{ord}}(\mathbf{Int} \gamma)}{Z_{\text{dis}}(\mathbf{Int} \gamma)} \\ &\leq e^{(a_\ell + 2\varepsilon_L)|\mathbf{Int} \gamma|} e^{-(\kappa/4 - 3)\|\gamma\|} \leq e^{-(\kappa/4 - 3 - \frac{1}{2}(a_\ell + 2\varepsilon_L)\text{diam}\gamma)\|\gamma\|} \end{aligned}$$

where we used (A.2) in the first inequality, the inductive assumption (A.10) and the bound (A.9) in the second, and the bound (5.5) in the last. Bounding $\varepsilon_L \text{diam}\gamma$ by $L\varepsilon_L \leq 1$ and using the assumption $a_\ell \text{diam}\gamma \leq c\beta$, this gives $K_{\text{dis}}(\gamma) \leq e^{-(\kappa/4 - 4 - c\beta)\|\gamma\|}$ and hence $K_{\text{dis}}(\gamma) = K'_{\text{dis}}(\gamma)$. The bound for contours with ordered external label is exactly the same.

Finally, assume that Λ has level $k + 2$, that (i) has been proven for all contours of level at most $k + 1$, and that (ii) has been proven for all sets of level at most k . Define a contour γ with external label ℓ to be *small* if $a_\ell \text{diam}\gamma \leq c\beta$, and *large* otherwise. Consider the representation (6.13) for $Z_{\text{dis}}(\Lambda)$, and fix, for a moment, the set Γ_{large} of all large external contours contributing to the right hand side. Summing over the remaining contours, we get a factor of $Z_{\text{ord}}(\mathbf{Int} \gamma)$ for the interior of each contour $\gamma \in \Gamma_{\text{large}}$, as well as a factor $Z_{\text{dis}}^{(\text{small})}(\mathbf{Ext} \Gamma_{\text{large}})$ for the exterior of Γ_{large} , where $Z_{\text{dis}}^{(\text{small})}(\Lambda')$ is obtained from $Z_{\text{dis}}(\Lambda')$ by dropping all configurations with large external contours. Thus

$$Z_{\text{dis}}(\Lambda) = \sum_{\Gamma_{\text{large}}} Z_{\text{dis}}^{(\text{small})}(\mathbf{Ext} \Gamma_{\text{large}}) \prod_{\gamma \in \Gamma_{\text{large}}} q e^{-\kappa\|\gamma\|} Z_{\text{ord}}(\mathbf{Int} \gamma), \quad (\text{A.11})$$

where the sum goes over sets of mutually external, large contours with disordered external label.

Since γ is small whenever $\gamma < \gamma'$ and γ' is a contour with the same external label as γ , the representation (6.21) for $Z_{\text{dis}}^{(\text{small})}(\mathbf{Ext} \Gamma_{\text{large}})$ contains only small contours, implying that for all these contours $K_{\text{dis}}(\gamma) = K'_{\text{dis}}(\gamma)$. As a consequence,

$$Z_{\text{dis}}^{(\text{small})}(\mathbf{Ext} \Gamma_{\text{large}}) \leq Z'_{\text{dis}}(\mathbf{Ext} \Gamma_{\text{large}}),$$

which allows us to use the estimate (A.8) to estimate the factor $Z_{\text{dis}}^{(\text{small})}(\mathbf{Ext} \Gamma_{\text{large}})$ in (A.11). Using the inductive assumption (ii) to bound the factors $Z_{\text{ord}}(\mathbf{Int} \gamma)$ and the bound (A.2) to estimate the factors $q e^{-\kappa\|\gamma\|}$, this gives

$$\begin{aligned} Z_{\text{dis}}(\Lambda) &\leq \sum_{\Gamma_{\text{large}}} e^{(\varepsilon_L - f_{\text{dis}})|\Lambda \cap \mathbf{Ext} \Gamma_{\text{large}}|} e^{\|\partial \mathbf{Ext} \Gamma_{\text{large}}\|} \prod_{\gamma \in \Gamma_{\text{large}}} e^{-(\kappa/4 - 2)\|\gamma\|} e^{(\varepsilon_L - f)|\mathbf{Int} \gamma|} \\ &= e^{(\varepsilon_L - f)|\Lambda| + \|\partial \Lambda\|} \sum_{\Gamma_{\text{large}}} e^{-a_{\text{dis}}|\Lambda \cap \mathbf{Ext} \Gamma_{\text{large}}|} \prod_{\gamma \in \Gamma_{\text{large}}} e^{-(\kappa/4 - 3)\|\gamma\|}. \end{aligned}$$

In order to prove statement (ii), we will have to show that

$$\sum_{\Gamma_{\text{large}}} e^{-\frac{1}{2}a_{\text{dis}}|\Lambda \cap \mathbf{Ext} \Gamma_{\text{large}}|} \prod_{\gamma \in \Gamma_{\text{large}}} e^{-(\kappa/4-3-c\beta/2)\|\gamma\|} \leq e^{\|\partial\Lambda\|} \quad (\text{A.12})$$

To this end, we define

$$\tilde{K}(\gamma) = \begin{cases} e^{-(\kappa/4-4-c\beta/2)\|\gamma\|} & \text{if } \gamma \text{ is a large contour with external label dis} \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\tilde{Z}(\Lambda') = \sum_{\Gamma} \prod_{\gamma \in \Gamma} \tilde{K}(\gamma)$$

where the sum runs over sets Γ of pairwise compatible contours in Λ' which all have external label dis. We also define

$$\tilde{f} = -\frac{1}{L^d} \log \tilde{Z}(V_{L,d}).$$

Note that, as before, one can use the forest structure of sets of pairwise compatible contours to rewrite $\tilde{Z}(\Lambda')$ as a sum over sets of mutually external contours in Λ' :

$$\tilde{Z}(\Lambda') = \sum_{\Gamma_{\text{ext}}} \prod_{\gamma \in \Gamma_{\text{ext}}} \tilde{K}(\gamma) \tilde{Z}(\mathbf{Int} \gamma).$$

For C_1 and C_2 sufficiently large, the partition function $\tilde{Z}(\Lambda')$ can again be analyzed by convergent Mayer expansion, leading to the bound

$$\left| \log \tilde{Z}(\Lambda') + \tilde{f}|\Lambda'| \right| \leq \|\partial\Lambda'\| \quad (\text{A.13})$$

(the term proportional to ε_L is absent since we defined \tilde{f} without taking the limit $L \rightarrow \infty$). Furthermore, since all contours γ contributing to \tilde{f} are large, which by (5.6) implies that $\|\gamma\| \geq k_0 = c\beta/a_{\text{dis}}$, we have that

$$-\tilde{f} = |\tilde{f}| \leq e^{-c\beta k_0} \leq \frac{a_{\text{dis}}}{(c\beta)^2} \leq \frac{a_{\text{dis}}}{2}.$$

As a consequence, the left hand side of (A.12) can be bounded by

$$\begin{aligned} \sum_{\Gamma_{\text{ext}}} e^{\tilde{f}|\Lambda \cap \mathbf{Ext} \Gamma_{\text{ext}}|} \prod_{\gamma \in \Gamma_{\text{ext}}} e^{-\|\gamma\|} \tilde{K}(\gamma) &\leq \sum_{\Gamma_{\text{ext}}} e^{\tilde{f}|\Lambda \cap \mathbf{Ext} \Gamma_{\text{ext}}|} \prod_{\gamma \in \Gamma_{\text{ext}}} \tilde{K}(\gamma) e^{\tilde{f}|\mathbf{Int} \gamma|} \tilde{Z}(\mathbf{Int} \gamma) \\ &= e^{\tilde{f}|\Lambda|} \sum_{\Gamma_{\text{ext}}} \prod_{\gamma \in \Gamma_{\text{ext}}} \tilde{K}(\gamma) \tilde{Z}(\mathbf{Int} \gamma) = e^{\tilde{f}|\Lambda|} \tilde{Z}(\Lambda) \leq e^{\|\partial\Lambda\|}, \end{aligned}$$

proving (A.12). Here the sums run over sets of mutually external contours, all of which have external label dis, the two inequalities follow from the bound (A.13).

This concludes the proof of (ii) for sets Λ of level $k+2$ and $\ell = \text{dis}$. The proof of (ii) for $\ell = \text{ord}$ is identical.

A.3 Proof of Lemma 6.3 (iii)

We start with the observation that the weights $K_\ell(\gamma)$, and hence the weights $K'_\ell(\gamma)$ are continuous functions of β . Since the free energies f_ℓ are given in terms of an absolutely convergent power series in the weights $K'_\ell(\gamma)$, they are continuous functions of β as well. Taking into account this continuity, the following lemma immediately implies Lemma 6.3 (iii).

Lemma A.2 *Assume that (A.1) holds.*

(i) *If $a_{\text{ord}} = 0$, then $M(\beta) > 0$.*

(ii) *If $a_{\text{ord}} > 0$, then $M(\beta) = 0$.*

Proof. At this point, the proof of Lemma A.2 is pretty standard. We therefore only sketch the main steps.

First, we note that for $\Lambda = \Lambda_L = \{1, \dots, L\}^d$, the representations (2.1) and (2.2) can be generalized to the model with 1-boundary conditions defined in (1.4). Indeed, let G_+ be the induced graph on $\Lambda_+ = \{0, 1, \dots, L+1\} \subset \mathbb{Z}^d$. The Edwards-Sokal measure $\pi_{\Lambda,1}$ corresponding to $\mu_{\Lambda,1}$ can then be obtained from the measure π_{G_+} by conditioning on $\sigma_x = 1$ for all $x \in \Lambda_+ \setminus \Lambda$ and $xy \in A$ whenever $\{x, y\} \subset \Lambda_+ \setminus \Lambda$. Next, we observe that in the conditional measure $\pi_{G_+}(\cdot | A)$, a spin at a vertex $x \in \Lambda$ has probability $1/q$ of taking the value 1 unless x lies in the same component of (Λ_+, A) as $\Lambda_+ \setminus \Lambda$, in other words, unless $x \in \mathbf{Ext}(A)$. Keeping these two observations in mind, the derivation of the representation (6.20) can easily be adapted to obtain a contour representation for the magnetization. Setting $\mathbf{\Lambda} = (-\frac{1}{4}, L + \frac{5}{4})^d \subset \mathbb{R}^d$, this gives

$$M_\Lambda(\beta) = \left(1 - \frac{1}{q}\right) \frac{1}{Z_{\text{ord}}(\mathbf{\Lambda})} e^{-e_{\text{ord}}|\Lambda_+|} \sum_{\Gamma} \frac{|\Lambda \cap \mathbf{Ext} \Gamma|}{|\Lambda|} \prod_{\gamma \in \Gamma} K_{\text{ord}}(\gamma),$$

where the sum goes over sets of pairwise compatible contours in $\mathbf{\Lambda}$ with external label ord . Note that we have chosen $\mathbf{\Lambda}$ in such a way that all edges in $\Lambda_+ \setminus \Lambda$ lie in $\mathbf{Ext} \gamma$ whenever γ is a contour in $\mathbf{\Lambda}$, corresponding to the above conditioning in π_{G_+} .

If $a_{\text{ord}} = 0$, the weights $K_{\text{ord}}(\gamma)$ are bounded by $e^{-c\beta\|\gamma\|}$ for all γ . As a consequence, we may use a standard Peierls argument to show that the probability that a given point $x \in \Lambda$ lies *not* in $\mathbf{Ext} \Gamma$ is small uniformly in L and $x \in \Lambda_L$, implying that $M(\beta) > 0$, which proves (i).

Assume finally that $a_{\text{ord}} > 0$ (which implies in particular that $a_{\text{dis}} = 0$). We will show that with probability tending to one, $|\Lambda \cap \mathbf{Ext} \Gamma| \leq L^{d-\varepsilon}$. Since the ratio $|\Lambda \cap \mathbf{Ext} \Gamma|/|\Lambda|$ in the definition of the magnetization is bounded uniformly in L , this will show that $M_\Lambda(\beta) \rightarrow 0$ as $L \rightarrow \infty$.

Recall the definition of large contours from the last proof, and let Γ_{large} be the set of external contours in Γ which are large. Then $|\Lambda \cap \mathbf{Ext} \Gamma| \leq |\Lambda \cap \mathbf{Ext} \Gamma_{\text{large}}|$, implying that it will be enough to show that the probability that $|\Lambda \cap \mathbf{Ext} \Gamma_{\text{large}}| \geq L^{d-1/2}$ goes to zero as $L \rightarrow \infty$. To bound this probability, we will prove an upper bound on the sum over contours with $|\Lambda \cap \mathbf{Ext} \Gamma_{\text{large}}| \geq L^{d-1/2}$ and a lower bound on $Z_{\text{ord}}(\mathbf{\Lambda})$.

To obtain the desired upper bound, we proceed as in the proof of Lemma A.1 (ii), leading to the estimate

$$\begin{aligned}
e^{-e_{\text{ord}}|\Lambda_+|} & \sum_{\substack{\Gamma: \\ |\Lambda \cap \mathbf{Ext} \Gamma_{\text{large}}| > L^{d-1/2}}} \prod_{\gamma \in \Gamma} K_{\text{ord}}(\gamma) \leq \\
& \leq e^{-f|\Lambda_+|} e^{2\|\partial\Lambda\|} \max_{\substack{\Gamma_{\text{ext}}: \\ |\Lambda \cap \mathbf{Ext} \Gamma_{\text{ext}}| > L^{d-1/2}}} e^{-\frac{a_{\ell}}{2}|\Lambda \cap \mathbf{Ext} \Gamma_{\text{ext}}|} \prod_{\gamma \in \Gamma_{\text{ext}}} e^{-\frac{c}{2}\beta\|\gamma\|} \quad (\text{A.14}) \\
& \leq e^{-f|\Lambda_+|} e^{2\|\partial\Lambda\|} e^{-\frac{a_{\text{ord}}}{2}L^{d-1/2}}
\end{aligned}$$

where the maximum in the second to last line goes over sets of mutually external contours in Λ which all have external label ord .

To bound $Z_{\text{ord}}(\Lambda)$ from below we restrict the sum in (6.16) to a single term, the term $\Gamma_{\text{ext}} = \{\gamma_0\}$, where γ_0 is the contour $\gamma_0 = \partial[1/4, L + 3/4]^d$. This gives

$$Z_{\text{ord}}(\Lambda) \geq e^{-e_{\text{ord}}|\Lambda_+ \setminus \Lambda|} e^{-\kappa\|\gamma_0\|} Z_{\text{dis}}(\mathbf{Int} \gamma_0) \geq e^{-fL^d} e^{-2d(\kappa+1+O(e^{-\beta}))L^{d-1}},$$

where we used the bound (A.9) and the fact that $a_{\text{dis}} = 0$ in the second step. After extracting the leading contribution e^{-fL^d} , the right hand side falls at most like an exponential in L^{d-1} , while the corresponding decay in (A.14) is exponential in $L^{d-1/2}$. This proves that the ratio of (A.14) and $Z_{\text{ord}}(\Lambda)$ goes to zero as $L \rightarrow \infty$, as desired.