

The Effect of Boundary Conditions on Mixing Rates of Markov Chains

Nayantara Bhatnagar^{1*}, Sam Greenberg^{2*}, and Dana Randall^{1*}

¹ College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280.

² School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332-0160.

Abstract. Many natural Markov chains undergo a phase transition as a temperature parameter is varied; a chain can be rapidly mixing at high temperature and slowly mixing at low temperature. Moreover, it is believed that even at low temperature, the rate of convergence is strongly dependent on the environment in which the underlying system is placed. It is believed that the boundary conditions of a spin configuration can determine whether a local Markov chain mixes quickly or slowly, but this has only been verified previously for models defined on trees. We demonstrate that the mixing time of Broder's Markov chain for sampling perfect and near-perfect matchings does have such a dependence on the environment when the underlying graph is the square-octagon lattice. We show the same effect occurs for a related chain on the space of Ising and “near-Ising” configurations on the two-dimensional Cartesian lattice.

1 Introduction

Boundary conditions play a crucial role in statistical physics for determining the uniqueness of Gibbs states, or the limiting distributions of families of configurations on the infinite lattice. Consider the Ising model on the $n \times n$ Cartesian lattice, a fundamental physical model for ferromagnetism. Each configuration σ in the state space $S = \{+, -\}^{n^2}$ consists of an assignment of a $+$ or $-$ spin to each of the vertices, and the *Gibbs distribution* assigns weight

$$\pi(\sigma) = \lambda^{-D(\sigma)} / Z,$$

where $D(\sigma) = |\{(i, j) \in E \mid \sigma(i) \neq \sigma(j)\}|$ and Z is the normalizing constant or *partition function*. In the classical description of the Ising model, $\lambda = e^{2\beta}$, where $\beta > 0$ is inverse temperature.

To characterize when there is a phase transition in a physical model, physicists study whether there is a unique limiting distribution as $n \rightarrow \infty$. The vertices on the boundary of an $n \times n$ grid are hard-wired to be $+$ in one case and $-$ in another. The Gibbs measure on the interior is defined as the limiting distribution conditioned on the boundary. It is well known that there is a critical value λ_c such that, for $\lambda < \lambda_c$, the limiting distribution is unique, yet for $\lambda > \lambda_c$,

* Supported in part by NSF grants CCR-0515105 and DMS-0505505.

correlations between the spins of vertices inside a finite region and the spins on the boundary of that region persist over long distances and there are multiple limiting distributions (see, e.g., [3]). A related effect has been observed in the context of mixing times of local chains on finite regions. The *mixing time* of a chain, i.e., the number of steps required so that probabilities of reaching each configuration is close to the stationary distribution, undergoes a similar phase change. When λ is sufficiently small, local dynamics are efficient, while when λ is large, local chains require exponential time to converge to equilibrium [16]. This is because at low enough temperature the Gibbs distribution strongly favors configurations that are predominantly one spin; it takes exponential time to move between mostly + and mostly – states using local chains [8, 9].

A natural question that integrates these two perspectives is: *Can one type of boundary condition cause a Markov chain to mix slowly, while another causes the same chain to mix rapidly?* Martinelli, Sinclair, and Weitz [10,11] answered this question in the affirmative in the context of spin systems on trees. The question remains unresolved when configurations are defined on lattices, although the same effect is believed to occur. Martinelli [8] showed that mixing times of Glauber (local) dynamics on Ising configurations of the 2-dimensional lattice can vary by an exponential factor, though the mixing time for both of his boundary conditions were shown to be exponential – in reality the boundary that leads to faster mixing is believed to converge in polynomial time.

Models and Results: The first problem we consider is sampling matchings on finite regions of the square-octagon lattice. This is the lattice formed by tightly packing octagons so that the uncovered space forms smaller squares (see Figure 1). For certain finite regions R of this lattice, there is an ergodic Markov chain on the set of perfect matchings; it starts at any matching and repeatedly does the following: choose a square or octagonal face uniformly, and if the matching alternates edges around this face, then “rotate” to the other matching. Propp [12] used coupling-from-the-past [13] on certain regions to generate so-called “diabolo tilings of fortresses,” and conjectured that the chain mixes slowly. The only proof of slow mixing for this model requires “activities” on the edges that weigh matchings according to the number of edges that bound squares on the lattice [4]. It has been conjectured that there is a region such that one boundary condition will cause this local Markov chain to mix quickly while another will mix slowly; however, like the Ising model at sufficiently low temperature, it remains a challenge to show fast mixing for such a contour model, even though there is a boundary for which the chain is believed to mix rapidly.

In this paper, we consider instead the Broder-chain on the set of perfect and near-perfect matchings on the square-octagon lattice. For a finite, simply-connected region R on this lattice, let the *boundary* of R be the set of vertices that have neighbors both inside and outside R on the infinite lattice. The boundary condition is defined by specifying, for each vertex on the boundary, whether it is to be included in the matchings or not. We hardwire a boundary condition and start with a perfect matching on the remaining region. The Broder-chain successively chooses an edge and this edge is added, deleted, or exchanged with

another edge (if exactly one endpoint was matched). The chain converges to the uniform distribution on perfect and near-perfect matchings. We show that, for a family of regions, there are two types of boundary conditions, one that causes the Broder-chain to mix slowly and another that causes the chain to mix rapidly. Remarkably, these two boundary conditions differ by the deletion of only four vertices. This is the first proof of slow mixing for matchings on the square-octagon lattice without activities on the edges.

The second model we consider is the Ising model on \mathbb{Z}^2 . It is strongly believed that Glauber dynamics are very sensitive to boundary conditions, even at low temperature, and that they will be fast for the all plus boundary and slow for two sides fixed to plus and two to minus. It is useful to view the Ising model in terms of contours. Given an Ising configuration, take the union of all edges that separate a + spin from a - spin. For the all-plus boundary condition, these edges form an even degree subgraph that can be thought of as sets of closed contours. Glauber dynamics perform local changes to these contours and, at low temperature, short contours are thermodynamically favorable. Fernandez, Ferrari, and Garcia [2] proposed a chain that, in one step, moves between two configurations that differ by a single contour. They show that at sufficiently low temperature, this chain converges quickly to stationarity, but unfortunately there is no efficient way to perform a step of the chain.

Instead, we consider a Broder-type chain on the set of Ising and “near-Ising” configurations on finite regions of the Cartesian lattice. For the all-plus boundary condition, for example, near-Ising configurations allow exactly one contour to be open while the others must be closed. A step of the chain removes or adds an edge, with transition probabilities chosen so that we converge to the Gibbs distribution on Ising and near-Ising configurations. We show that, at sufficiently low temperature, there are two boundary conditions so that the chain mixes slowly with one and quickly with the other. The fast mixing results are of independent interest because they demonstrate how to extend the result by Fernandez et al. to define a rapidly mixing chain at low temperature that can be efficiently implemented. Moreover, this gives a much more efficient algorithm for sampling Ising configurations at low temperature than the only other rigorous method known previously [14].

Techniques: A key fact underlying our results is that both the matching and Ising models considered here can be reformulated as contour models. Contours for the Ising model are unrestricted, while the contours arising from matchings on the square-octagon lattice are required to turn left or right at every step. The contours arising from these matchings behave similarly to Ising contours at low temperature where long contours are penalized.

In the first (fast) cases, boundaries are defined so that initially all vertices, including the boundary, have even degree in the contour representation. During the simulation, configurations must have zero or one open path. We show that the weight of perfect and near-perfect configurations are polynomially related. Following the canonical path technique introduced by Jerrum and Sinclair [5], this suffices to show polynomial mixing for both models.

In the second (slow) cases, our proofs are based on a *Peierls argument* from statistical physics (see, e.g., [3]). We introduce four designated vertices on the boundary that initially have odd degree in the contour representation, and thus are the endpoints of two paths. At any point during the simulation of the Broder-chains, at most two vertices in the contour representation have changed parity. Once one of the initial paths is disconnected, the two pieces tend to shrink. We prove, via sensitive injections, that it will take exponential time for these to reconnect — this is sufficient to show slow mixing. While this is an artifact of the extended state space, these models do correctly capture the effects of the two boundaries for models where the Gibbs measure favors short contours. On the square-octagon lattice these bounds are especially sensitive because we cannot modify the temperature to establish the required inequalities.

2 Preliminaries

Let \mathcal{M} be an ergodic (i.e., irreducible and aperiodic), reversible Markov chain with finite state space S , transition probability matrix P , and stationary distribution π . Let $P^t(x, y)$ be the t -step transition probability from x to y and let $\|\cdot, \cdot\|$ denote the total variation distance.

Definition 1. For $\varepsilon > 0$, the mixing time $\tau = \min\{t : \|P^t, \pi\| \leq 1/4, \forall t' \geq t\}$.

A Markov chain is *rapidly mixing* if the mixing time is bounded above by a polynomial in n , the size of each configuration in the state space. If the mixing time is exponential in n , the chain is *slowly mixing*. Jerrum and Sinclair defined the conductance of a chain and showed that it bounds mixing time [5].

Definition 2. If a Markov chain has stationary distribution π , we define the conductance Φ as

$$\Phi = \min_{S: \pi(S) \leq 1/2} \frac{\sum_{x \in S, y \notin S} \pi(x)P(x, y)}{\pi(S)}.$$

Theorem 1. An ergodic, reversible chain with conductance Φ is rapidly mixing if and only if $\Phi > 1/p(n)$ for some polynomial $p(\cdot)$.

Jerrum and Sinclair use this theorem to analyze a natural Markov chain on matchings due to Broder [1]. Here we consider the chain on the state space consisting of perfect and near-perfect matchings. At each step, **the Broder-chain** \mathcal{M}_B does the following:

- Choose an edge e uniformly at random.
- ◊ If the endpoints of e are unmatched, add e to the matching.
- ◊ If e is in the matching and the matching is perfect, remove e .
- ◊ If exactly one endpoint of e is matched, remove the matched edge and add e .
- ◊ Otherwise, do nothing.

This chain converges to the uniform distribution on matchings and near-perfect matchings. Jerrum and Sinclair [5, 15] find the following characterization for when the Broder-chain can be used to efficiently sample perfect matchings.




Fig. 1. For $n = 5$, (a) \mathbf{L} , (b) \mathbf{L}' , (c) a near-perfect matching, (d) that matching's contraction.

Theorem 2. Let $\mathbf{S}_{\mathcal{P}}$ be the set of perfect matchings in \mathbf{S} and $\mathbf{S}_{\mathcal{N}}$ be the set of near-perfect matchings. If $|\mathbf{S}_{\mathcal{N}}| \leq p(n)|\mathbf{S}_{\mathcal{P}}|$ for some polynomial $p(\cdot)$, then Φ is at least inverse-polynomial.

Theorem 2 was proven using the *canonical path* technique. The key idea in this proof is to define paths between every pair of states $(I, F) \in S \times S$ in the transition graph of the chain. If not too many paths go through any specific transition then there cannot be a bottleneck in the transition graph. The following summary of their method will be useful in Section 4.

Theorem 3. Suppose there exists a function η such that, for a transition $T = (G, G')$ along the canonical path from I to F ,

1. given T and $\eta(T, I, F)$, we can reconstruct both I and F
2. $\mu(I)\mu(F) \leq \mu(G)\mu(\eta(T, I, F))P(G, G')$.

Then $\Phi = \Omega(n^{-c})$ for some constant c .

This theorem is the key ingredient that establishes our fast mixing results. For our slow mixing results, we show that the conductance is exponentially small. For this, it suffices to identify a bad cut (S, \bar{S}) in the state space.

3 Perfect matchings in the square-octagon lattice

Let \mathbf{L} to be the square-octagon lattice with n squares on a side, for n odd (see Figure 1a). Define \mathbf{L}' to be the same lattice, but with one vertex missing from each of the corner squares (Figure 1b). Regions \mathbf{L} and \mathbf{L}' capture the two boundary conditions we study.

We let \mathbf{S} (resp. \mathbf{S}') be the set of perfect and near-perfect matchings on \mathbf{L} (resp. \mathbf{L}') and let $\mathcal{M}_{\mathbf{B}}$ be the Broder-chain on these state spaces, as described in Section 2. Our first main result is the following:

Theorem 4. There exist constants $c_1, c_2 > 1$ such that the mixing time of $\mathcal{M}_{\mathbf{B}}$ on \mathbf{S} is $O(n^{c_1})$ while the mixing time of $\mathcal{M}_{\mathbf{B}}$ on \mathbf{S}' is $\Omega(e^{c_2 n})$.

3.1 Contraction to contours

Before presenting the proof of Theorem 4, it will be convenient to define a bijection between perfect matchings and a related contour representation. We “contract” the lattice regions \mathbf{L} and \mathbf{L}' by replacing each of the squares with vertices so that only the edges bounded by octagons on each side survive. The result is isomorphic to a subregion of the Cartesian lattice (see Figure 1c and 1d). We use bold script (\mathbf{L} , \mathbf{S} , \mathcal{M}_B , etc.) when referring to the square-octagon graphs, and normal script (L , S , \mathcal{M}_B , etc.) for the contracted case. To that end, define L and L' to be the integer lattice with n vertices on each side.

Consider the effect of this contraction on a perfect matching of \mathbf{L} . The contraction is an even degree subgraph where all vertices of degree 2 are incident to edges that bound two sides of a unit square; if the endpoints were collinear, the corresponding square in \mathbf{L} would have two unmatched vertices. For \mathbf{L}' , we get an even degree subgraph with this turning property, only now there are four vertices of odd degree in the corners, as they correspond squares in \mathbf{L}' with only three vertices present. Finally, near-perfect matchings of \mathbf{L} and \mathbf{L}' contract as above, except two vertices might have the opposite parity (corresponding to the square(s) in the lattice containing the two unmatched vertices).

We define new ground-states S and S' with this in mind. We call a subgraph of L a *turning graph* if all vertices have even degree and vertices of degree 2 “turn corners.” We call a subgraph of L' a *turning graph* if this applies to all but the corner vertices. We call a subgraph of L or L' a *near-turning graph* if all vertices of degree 2 turn, and exactly two vertices have parity different from what was prescribed for turning graphs. Let $S_{\mathcal{T}}$ and $S_{\mathcal{N}}$ (resp. $S'_{\mathcal{T}}$ and $S'_{\mathcal{N}}$) be the turning and near-turning graphs of L (resp. L'). Let S and S' be the union of the turning and near-turning graphs in each case.

Notice that the contraction map is not one-to-one. Let $G \in S$ and let $v \in L$. If the degree of v in G is 2, 3, or 4, there is a unique way to expand v to a square face of $\mathbf{G} \in \mathbf{S}$, but if $d(v) = 0$ or 1, then there are two ways to recover the matched edge(s) that were deleted from the corresponding square, illustrated in Figure 1. For $G \in S \cap S'$, the number of matchings which contract to G is to within a small polynomial factor $q(n)$ of $2^{|V(L)| - |V(G)|}$.

Our first goal is to show that the Broder-chain is fast on \mathbf{L} , and for this we need to show a polynomial relationship between the numbers of near-perfect and perfect matchings. For simplicity, we instead consider their contracted versions in L and introduce a weight μ on turning and near-turning contours as follows: for $G \in S \cup S'$ let $\mu(G) := 2^{-|V(G)|}$. After a normalization, $\mu(G)$ is within $q(n)$ of the number of matchings which contract to G . It then follows that showing a polynomial relationship between $\sum_{G \in S_{\mathcal{T}}} \mu(G)$ and $\sum_{G \in S_{\mathcal{N}}} \mu(G)$ is sufficient to establish fast mixing of the Broder-chain.

The following lemma will be crucial in our proof of slow mixing. It shows that we can encode near-turning contours as a function of the number of vertices it hits, rather than its total length.

Lemma 1. *For $G \in S \cup S'$, let $\mathcal{N}_a(G)$ be the set of near-turning components A , edge-disjoint from G , such that $|V(G \cup A)| = |V(G)| + a$. Then $|\mathcal{N}_a(G)| \leq 4n^{42a}$.*

Proof. Using a standard “Euler-like” decomposition, any A can be described as a single turning path between the odd-degree vertices. If we are given the coordinates of the first odd vertex and the initial direction of the path, we can encode the rest in a binary string, with 0 representing a left-turn and 1 representing a right-turn. This is a natural encoding of the turning-path, but requires $|E(A)|$ bits, possibly more than a . It will be necessary to define an encoding which focuses on vertices instead of edges.

Fortunately, not all turns need to be encoded; whenever the path touches the graph G , either because it turns back on itself or because it touches $G \setminus A$, the next turn is forced. We therefore only encode those turns when our path hits a previously-empty vertices, creating a bitstream of length $a - 1$. This encoding is not completely unique. After the last recorded turn, the path might proceed for any number of forced-turns; our encoding would fail to represent how many. However, knowing the length of the turning-path determines this uniquely, and we can upper bound the length by $2n^2$ edges. Hence each binary string corresponds to at most a polynomial number of turning-paths.

The total size of $\mathcal{N}_a(G)$ is therefore upper bounded by the number of possible lengths times the number of starting vertices and directions, times the number of binary strings of length $a - 1$. This gives us $|\mathcal{N}_a(G)| \leq 2n^2 \cdot n^2 \cdot 4 \cdot 2^{a-1}$. \square

3.2 Fast mixing of \mathcal{M}_B on \mathbf{S}

For the rapid mixing result of Theorem 4, it is sufficient to show a polynomial relationship between the size of $\mathbf{S}_{\mathcal{P}}$ and the size of $\mathbf{S}_{\mathcal{N}}$. Focusing on the contracted representations, this is equivalent to the following.

Lemma 2. *For some polynomial $p(\cdot)$, $\mu(S_{\mathcal{N}}) \leq p(n) \cdot \mu(S_{\mathcal{T}})$.*

Proof. We define a function $f : S_{\mathcal{N}} \rightarrow S_{\mathcal{T}}$. For $G' \in S_{\mathcal{T}}$, define the pre-image of G' to be $f^{-1}(G') = \{G \in S_{\mathcal{N}} : f(G) = G'\}$. We define f in such a way that, although $f^{-1}(G')$ contains many graphs, their total weight is within a polynomial factor of the weight of G' .

Let $A(G)$ be the component of G containing the two odd vertices and let $f(G) = G \setminus A(G)$. partitioning according to the size of $A(G)$, for $G' \in \text{Img}(f)$,

$$\mu(f^{-1}(G')) = \sum_{a=1}^{n^2} \sum_{A \in \mathcal{N}_a(G')} \mu(G' \cup A) = \sum_{a=1}^{n^2} |\mathcal{N}_a(G')| \cdot \mu(G') \cdot 2^{-a} \leq 16n^6 \mu(G').$$

where the last inequality is due to Lemma 1. Then

$$\mu(S_{\mathcal{N}}) = \sum_{G' \in \text{Img}(f)} \mu(f^{-1}(G')) \leq 16n^6 \sum_{G' \in \text{Img}(f)} \mu(G') \leq 16n^6 \mu(S_{\mathcal{T}}). \quad \square$$

Hence, by Theorem 2, the Broder-chain is fast on \mathbf{L} .




Fig. 2. Stages of f : (a) W , (b) removing $B(W)$ and shifting, (c) adding squares along the wall to obtain U .

3.3 Slow mixing of \mathcal{M}_B on S'

We turn now to the behavior of the Broder-chain on L' . Define a *bridge* to be a turning path connecting two corner vertices of L' . Perfect matchings of L' contract to turning subgraphs of L' with two distinct bridges, while near-perfect matchings of L' might map to a graph with only one bridge.

Our strategy will be to show that there is a bad cut in the state space. Let σ, τ be two configurations that each have only one bridge, and suppose that these bridges connect different pairs of vertices on the corners of the boundary. To move from σ to τ it is necessary to pass through a configuration with two bridges. We will show that the set of configurations with two bridges is exponentially smaller than configurations with any one bridge, and so this will establish slow mixing of the Broder-chain. We must define a very sensitive map from configurations with two bridges to those with one. This is accomplished by the following lemma.

Lemma 3. *Let \mathcal{W} be the set of graphs in S' with two bridges and \mathcal{U} be the graphs with only one. There is a constant $c > 1$ such that $\mu(\mathcal{W}) \leq c^{-n}\mu(\mathcal{U})$.*

Proof. We define a function $f : \mathcal{W} \rightarrow \mathcal{U}$ in such a way that $f^{-1}(U)$ is exponentially smaller than the weight of U . Informally, first remove the larger of the two bridges in W ; then shift all components between that bridge and the wall by 1 unit (away from the wall). (See Figure 2.) This allows us to use cells adjacent to the wall to encode the initial part of the bridge, which is crucial to the result.

More precisely, for $W \in \mathcal{W}$, let $B(W)$ be the maximal bridge (with respect to vertices) in W . Let $\mathcal{W}_{\mathcal{T}}$ be the set of turning graphs of S' and let $\mathcal{W}_{\mathcal{N}}$ be the near-turning graphs. If $W \in \mathcal{W}_{\mathcal{T}}$, we can remove $B(W)$ leaving a graph in \mathcal{U} . If $W \in \mathcal{W}_{\mathcal{N}}$, we remove both $B(W)$ and the near-turning component $A(W)$. Suppose $B(W)$ connects the upper-left and lower-left corners of L' . Shift all of the components between $B(W)$ and the left wall one square to the right. This allows us to add edges along this left wall. It would be convenient to be able to add a cycle to any face along the wall; this is not always possible, as the original components might obstruct the addition, even after shifting. However, it can be seen that at least $n/2$ of the faces *do* allow such an addition after shifting. We use these $n/2$ positions to encode the initial segment of the deleted bridge. Break

$n/2$ of these positions into groups of 2^5 . We will add a 4-cycle to exactly one face in each group, thus encoding the first 5 bits of the bit string. (For instance, if the first five bits of $B(W)$ are 01011, we add a cycle on the $1 + 2 + 8 = 11$ th face in the group.) We do this in each of the $\frac{n}{2 \cdot 32}$ groups. In this way, we can encode the first $5n/64$ bits of B at a cost of only $4n/64$ additional edges (4 per face). Let the graph in \mathcal{U} thus obtained be U .

We partition $f^{-1}(U)$ based on the size of the bridge removed. For each U , let $E(U)$ be the extra encoding added along the left wall and let $\mathcal{B}_b(U)$ be the set of bridges B that add b vertices ($|U \cup B| = |U| + b$) and match those $5n/64$ bits encoded in $E(U)$. Then, partitioning according to b ,

$$\begin{aligned} \mu(f^{-1}(U) \cap \mathcal{W}_{\mathcal{T}}) &= \sum_{b=1}^{n^2} \sum_{B \in \mathcal{B}_b(U)} \mu(U \setminus E(U) \cup B) \\ &\leq \sum_{b=1}^{n^2} |\mathcal{B}_b(U)| \cdot \mu(U) \cdot 2^{\frac{4n}{64} - b} \leq 32n^6 2^{\frac{-n}{64}} \cdot \mu(U). \end{aligned}$$

For $f^{-1}(U) \cap \mathcal{W}_{\mathcal{N}}$, we must further partition according to size.

$$\begin{aligned} \mu(f^{-1}(U) \cap \mathcal{W}_{\mathcal{N}}) &= \sum_{b=1}^{n^2} \sum_{B \in \mathcal{B}_b(U)} \sum_{a=1}^{n^2} \sum_{A \in \mathcal{N}_a(U \cup B)} \mu(U \setminus E(U) \cup A \cup B) \\ &\leq \sum_{b=1}^{n^2} \sum_{B \in \mathcal{B}_b(U)} \sum_{a=1}^{n^2} |\mathcal{N}_a(U \cup B)| \mu(U) 2^{\frac{4n}{64} - b - a} \\ &\leq \sum_{b=1}^{n^2} \sum_{a=1}^{n^2} |\mathcal{B}_N(U)| 16n^4 2^a \cdot \mu(U) 2^{\frac{4n}{64} - b - a} \leq 64n^{12} 2^{\frac{-n}{64}} \mu(U), \end{aligned}$$

where the last two inequalities come from Lemma 1. Then, for some $c > 1$,

$$\mu(\mathcal{W}) = \sum_{U \in \text{Img}(f) \cap \mathcal{W}_{\mathcal{T}}} \mu(f^{-1}(U)) + \sum_{U \in \text{Img}(f) \cap \mathcal{W}_{\mathcal{N}}} \mu(f^{-1}(U)) < c^n \sum_{U \in \text{Img}(f)} \mu(U) \leq c^{-n} \mu(\mathcal{U}). \quad \square$$

This establishes an exponentially small cut in S' because the set of turning graphs with a bridge from upper-left to upper-right has the same weight as the set with bridges from lower-left to lower-right, by symmetry. This upper bounds conductance, and verifies the slow mixing result in Theorem 4.

4 Ising model

There is an analogous dichotomy for the mixing time of a Broder-type Markov chain defined for the Ising model. In the Ising model, each face of an $n \times n$ region on the Cartesian lattice is assigned one of two spins, $+$ or $-$. (Traditionally, the spins are assigned to the vertices; we use the equivalent model on faces for




Fig. 3. Paths arising from the plus-minus boundary.

reasons that will soon become clear.) Given a fixed assignment of spins to the faces just outside the boundary, our goal is to sample from the set of possible configurations with that boundary according to the Gibbs measure. We will consider two boundaries, the all-plus boundary β , and the plus-minus boundary β' where we fix $+$ along the horizontal sides and $-$ along the vertical sides.

Given any Ising configuration with a prescribed boundary, we can uniquely reconstruct the spins on the interior from the set of edges that separate faces with unequal spins. We will concentrate on this contour representation of Ising configurations. The Ising model is defined so that the Gibbs measure of a configuration mapping to a subgraph G is proportional to the weight $\mu(G) = \lambda^{-|E(G)|}$. Notice that an Ising configuration with boundary β maps to a graph that can be decomposed into edge-disjoint contours, while an Ising configuration with boundary β' maps to a graph that can be decomposed into a set of contours and two paths connecting the four corners, as in Figure 3. These contours are no longer forced to turn, but the setting is otherwise reminiscent of Section 3.

Following Section 3, we first enlarge our state-space. Let Λ be the set of contours and near-contours of L , where every vertex has even degree except possibly two vertices. (We can think of near-contours representing “near-Ising” configurations, although this does not have a natural interpretation in the spin representation.) Let Λ' be the set of contours and near-contours of L' , where the contours contain only vertices of even degree except in the four corners, and near-contours have this parity everywhere except at two vertices.

We now define a Markov chain \mathcal{M}_I on Λ . Given G in Λ (or Λ'), choose an edge e uniformly at random in L . If $e \in G$, let $G' = G \setminus e$. If $e \notin G$, let $G' = G \cup e$. Then, if G' is in Λ (or Λ'), \mathcal{M}_I sends G to G' with probability $\min(\frac{\mu(G')}{\mu(G)}, 1)$ and does nothing otherwise. Our second main theorem establishes that the mixing time of \mathcal{M}_I is also very sensitive to the boundary.

Theorem 5. *For any $\lambda > 3$, there exists constants $c_1, c_2 > 1$ such that the mixing-time of \mathcal{M}_I on Λ is $O(n^{c_1})$, but the mixing-time of \mathcal{M}_I on Λ' is $\Omega(e^{c_2 n})$.*

4.1 Fast mixing of \mathcal{M}_I on Λ

To prove the fast mixing part of Theorem 5, we again bound conductance. This proof is almost identical to the arguments underlying the fast mixing of the Broder-chain on the square-octagon lattice.

Let $\Lambda_{\mathcal{T}}$ be the set of contours of Λ and let $\Lambda_{\mathcal{N}}$ be the set of near-contours. Our canonical paths map I and F to the closest contour graphs (if they are near-contours) and then define a traditional canonical path between these contour graphs. If $G \in \Lambda_{\mathcal{T}}$, then $\overline{G} = G$. However, if $G \in \Lambda_{\mathcal{N}}$, we again let $A(G)$ be the component of G containing the odd vertices and let $\overline{G} = G \setminus A(G)$. In the canonical path from I to \overline{I} , place a fixed ordering on all vertices, perhaps by vertical and then horizontal location. Decompose $A(I)$ as an Eulerian path, and, starting from the first of the odd vertices, remove this path one edge at a time. The path from \overline{F} to F is defined to be the inverse of the path from F to \overline{F} .

For any contour G , the set of near-contours mapped to G in this manner has small weight. To show this, we use the trivial bound that the number of self-avoiding walks of length ℓ starting at a particular vertex is bounded by c^ℓ , where $c < 3$, see [7]. Observe that the set of H s.t. $\overline{H} = G$ has weight

$$\mu(\{H \in \Lambda_{\mathcal{N}} : \overline{H} = G\}) = \sum_{l=1}^{n^2} \sum_{\substack{A \\ |A|=l}} \mu(G \cup A) < \sum_{l=1}^{n^2} n^2 3^l \mu(G) \lambda^{-l} < n^4 \mu(G)$$

We define a canonical path between contours I and $F \in \Lambda_{\mathcal{T}}$ by “unwinding” the cycles and paths in their symmetric difference. Let $C = I \oplus F$. Order the components $\{C_i\}$ of C by $\{c_i\}$, where c_i is the earliest vertex in C_i . Each C_i is a contour, so it can be written as cycle starting from c_i . The path from I to F unwind each of the components in turn, starting at c_i and complementing each edge in turn. For a transition $T = (G, G')$, define the graph $\eta(T, I, F) := I \oplus F \oplus (G \cup G')$. Given $\eta(T, I, F)$ we will be able to reconstruct I and F given T . To reconstruct I and F from $\eta(T, I, F)$ and G , let $C = \eta(T, I, F) \oplus (G \cup G')$, and divide the edges of G and $\eta(T, I, F)$ according to the components of C . Note that every edge in both I and F is in both G and $\eta(T, I, F)$. Any edge in only one of I or F is in only one of G or $\eta(T, I, F)$. This implies that the weight $\mu(I)\mu(F) \leq \mu(\eta(T, I, F))\mu(G)$. Using Theorem 3, this is sufficient to prove that the conductance is at least $p(n)$, for some polynomial $p(\cdot)$, and thus it implies that the chain is rapidly mixing when we have the all-plus boundary.

4.2 Slow mixing of \mathcal{M}_I on Λ'

We proceed as in Section 3.3 by finding a cut set \mathcal{W} that has exponentially smaller weight than the parts of the state space it separates. This implies that the conductance is exponentially small and that the chain is slowly mixing.

As before, define a *bridge* to be a path connecting corners of L' . (Recall they need no longer turn.) We define \mathcal{W} to be the set of graphs in Λ' with two bridges and \mathcal{U} to be the set of graphs with only one. For $W \in \mathcal{W}$, let $B(W)$ be the maximal (in terms of vertices) bridge in W . If W is a near-contour, let $A(W)$ be the component with internal odd vertices.

We define the function $f : \mathcal{W} \rightarrow \mathcal{U}$ such that, for $W \in \mathcal{W}$, f removes $B(W)$ (and $A(W)$ if one exists). By the choice of λ and the bound on the number of

self-avoiding walks, there exists $c > 1$ such that, for any $U \in \text{Img}(f)$,

$$\begin{aligned} \mu(f^{-1}(U)) &= \sum_A \sum_B \mu(U \cup A \cup B) = \sum_{a=1}^{n^2} \sum_{b=2n}^{n^2} \sum_{\substack{A,B \\ |A|=a \\ |B|=b}} \mu(U) \mu(A) \mu(B) \\ &< \sum_{a=1}^{n^2} \sum_{b=2n}^{n^2} 3^a 3^b \mu(U) \lambda^{-a} \lambda^{-b} < c^n \mu(U). \end{aligned}$$

Then, $\mu(W) = \sum_{U \in \text{Img}(f)} \mu(f^{-1}(U)) < c^{-n} \sum \mu(U) = c^{-n} \mu(U)$ which, by Theorem 1, shows slow mixing.

References

1. A.Z. Broder. How hard is it to marry at random? (On the approximation of the permanent). *Proc. 8th ACM Symposium on Theory of Computing*, 50–58, 1986.
2. R. Fernandez, P.A. Ferrari, and N.L. Garcia. Loss network representation of Ising contours. *Annals of Probability* **29**: 902–937, 2001.
3. H.O. Georgii. *Gibbs measures and phase transitions*. de Gruyter Studies in Mathematics **9**, Walter de Gruyter & Co., Berlin, 1988.
4. S. Greenberg and D. Randall. Slow mixing of Glauber dynamics on perfect matchings of the square-octagon lattice. Preprint, 2006.
5. M.R. Jerrum and A.J. Sinclair. Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation* **82**: 93–133, 1989.
6. M.R. Jerrum and A.J. Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing* **22**: 1087–1116, 1993.
7. N. Madras and G. Slade. *The Self-Avoiding Walk*. Boston, MA, Birkhuser, 1993.
8. F. Martinelli. Lectures on Glauber dynamics for discrete spin models. *Lectures on Probability Theory and Statistics (Saint-Flour, 1997)*, Lecture notes in Mathematics **1717**: 93-191, Springer, Berlin, 1998.
9. F. Martinelli, E. Olivieri and R. Schonmann. For 2-D lattice spin systems weak mixing implies strong mixing. *Comm. Mathematical Physics* **165**: 33–47, 1994.
10. F. Martinelli, A. Sinclair, D. Weitz, The Ising model on trees: Boundary conditions and mixing time. *Comm. Mathematical Physics* **250**: 301–334, 2004.
11. F. Martinelli, A. Sinclair, D. Weitz, Fast mixing for independent sets, colorings and other models on trees. *Proc. 15th ACM/SIAM Symposium on Discrete Algorithms*, 449–458, 2004.
12. J. Propp. Diabolo Tilings of Fortress. Talk given at the MIT Combinatorics Seminar, 1998 (<http://www.math.wisc.edu/~propp/diabolo.ps.gz>).
13. J. Propp and D.B. Wilson. Exact Sampling with Coupled Markov Chains and Applications to Statistical Mechanics. *Random Structures and Algorithms* **9**: 223–252, 1996.
14. D. Randall and D.B. Wilson. Sampling Spin Configurations of an Ising System. *Proc. 10th ACM/SIAM Symposium on Discrete Algorithms*, S959–960, 1999.
15. A.J. Sinclair and M.R. Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation* **82**: 93–133, 1989.
16. L. Thomas. Bound on the mass gap for the finite volume stochastic Ising models at low temperature. *Comm. Mathematical Physics* **126**: 1–11, 1989.