

Slow Mixing of Glauber Dynamics Via Topological Obstructions

Dana Randall*

Abstract

Many local Markov chains based on Glauber dynamics are known to undergo a phase transition as a parameter λ of the system is varied. For independent sets on the 2-dimensional Cartesian lattice, the Gibbs distribution assigns each independent set a weight $\lambda^{|I|}$, and the Markov chain adds or deletes a single vertex at each step. It is believed that there is a critical point $\lambda_c \approx 3.79$ such that for $\lambda < \lambda_c$, local dynamics converge in polynomial time while for $\lambda > \lambda_c$ they require exponential time. We introduce a new method for showing slow mixing based on the presence or absence of certain topological obstructions in the independent sets. Using elementary arguments, we show that Glauber dynamics will be slow for sampling independent sets in 2 dimensions when $\lambda \geq 56.812$, improving on the best known bound while using simpler arguments. We also show they are slow on the torus when $\lambda \geq 50.5253$.

1 Introduction

Glauber dynamics [9], or Markov chains employing local moves, can be very useful for sampling from large combinatorial state spaces. We consider the behavior of these Markov chains in the context of sampling independent sets on 2-dimensional lattice regions. Let Λ be an $n \times n$ region on the Cartesian lattice and consider the set of independent sets Ω on Λ . The Gibbs distribution on Ω depends on a parameter $\lambda > 0$, and is given by

$$\pi(I) = \lambda^{|I|}/Z,$$

where $|I|$ is the size of the independent set and $Z = \sum_{I \in \Omega} \lambda^{|I|}$ is the normalizing constant known as the partition function. The Glauber dynamics connects pairs of configurations with Hamming distance one, with transition probabilities defined so that the unique stationary measure is the Gibbs distribution.

An interesting phenomenon occurs as the parameter λ is varied. When λ is sufficiently small, the Gibbs distribution favors configurations that are sparse, and Glauber dynamics will be efficient. However, when λ is sufficiently large, the distribution favors dense configurations, and Glauber dynamics will take an exponential amount of time to converge to equilibrium. The slow convergence arises because the Gibbs distribution is bimodal; configurations that are extremely dense lie predominantly on the odd or the even lattice, while configurations that are roughly half odd and half even have much smaller probability. Since Glauber dynamics changes the relative numbers of even and odd vertices in an

*College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280. Supported in part by NSF grants CCR-0515105 and DMS-0505505.

independent set by at most 1 in each step, the Markov chain has a bottleneck prohibiting fast mixing.

This dichotomy in mixing times is closely related to whether there is a unique limiting distribution for the Gibbs distribution on the infinite lattice, known as a *Gibbs state*. It is believed there is a critical value λ_c such that for $\lambda < \lambda_c$ there is a unique Gibbs state, while for $\lambda > \lambda_c$, there are multiple Gibbs states. Starting with Dobrushin in 1968 [5], physicists have been developing techniques to systematically characterize these two regimes. For independent sets on the two-dimensional lattice, λ_c is believed to be approximately 3.79 [1].

The methods used to show fast and slow mixing of Markov chains build on methods developed in statistical physics in the context of uniqueness of the Gibbs state. Luby and Vigoda [11] showed that Glauber dynamics on independent sets is fast when $\lambda \leq 1$ on the 2-dimensional lattice and torus. Van den Berg and Steif [2] showed that mixing is fast even for λ a little bigger than one by relating mixing times to rigorous bounds on the critical point for site percolation on the grid. On the other hand, Borgs et al. [4] showed that mixing of local Markov chains is slow on toroidal lattice regions for sufficiently large λ . Although the bound on λ remains unpublished, the authors claim that their methods yield slow mixing in two dimensions for $\lambda > 80$ [7, 16]. All of these bounds are far from the believed critical point, which is also 3.79 in the context of mixing.

In this paper, we introduce a new method that shows slow mixing of local dynamics and we demonstrate the method in the context of independent sets. The set up is largely the same as previous approaches: we show that the state space can be partitioned into three pieces, two exponentially larger than the third, such that moving from one of the two large pieces to the other with steps of the Markov chain requires passing through the small one in the middle. Such a partition can be used to show that the conductance of the chain is exponentially small, thereby proving that the chain takes exponential time to converge (see [15]). We employ a common technique known as a *Peierls argument*, used in statistical physics to show non-uniqueness and in computer science to show slow mixing. The Peierls argument constructs injections from one part of the state space to another in a way that allows one to infer that one piece of the partition is exponentially smaller than the other two.

What is new in our approach is the way in which we define the partition of the state space Ω into three pieces. Traditionally this partition is based on the contours around components of the independent sets whose boundaries lie on either the even or the odd sublattice. Following this approach, if we take the middle (small) piece in the partition to be the independent sets that are half even and half odd, then we are faced with the daunting task of bounding the number of such configurations. Borgs et al. [4] built on a sophisticated technique from statistical physics known as *Pirigov-Sinai theory*, as well as powerful combinatorial tools for counting contours, in order to carry out their proofs of slow mixing.

Here, instead, we characterize the three pieces of the partition using simple *topological obstructions* in the independent sets. These are defined using the vertices that are unoccupied in the independent set rather than the vertices that are present. In other words, instead of looking at (potentially) multiple contours around components of points in the independent set, we concentrate on a single connected feature in the complement that must be present whenever an independent set is sufficiently sparse. We call this a

fault line, and, roughly speaking, it is a width 2 path from the top of Λ to the bottom (or from the left side to the right) that is unoccupied (see Figure 2). The presence of a fault line in an independent set characterizes our middle (small) piece $\Omega_{\mathcal{F}}$ in the partition of the state space. If there is no fault line, then there must be *monochromatic bridges* in each direction, i.e., paths of points in the independent set that lie solely on one of the two sublattices and extend from one side of Λ to the other. We assign all the independent sets that do not have a fault line to Ω_0 or Ω_1 depending on whether these monochromatic bridges fall on the even or odd sublattices — indeed all the bridges in an independent set must fall on the same sublattice if there is no fault line!

This characterization tremendously simplifies the proofs of slow mixing. Moreover, it yields much stronger bounds on λ : we show that Glauber dynamics is slow on the two-dimensional grid when $\lambda < 56.812$ and on the two-dimensional torus when $\lambda < 50.5253$, resulting in an improvement over the best bounds to date while employing much simpler arguments. While we demonstrate this approach only in the context of independent sets on two-dimensional lattice regions, we note that the technique is far more general and is likely to yield improved bounds for slow mixing in the context of other sampling problems.




Figure 1: The graph H and a fault line F in $H_U(I)$, shown dotted in part (c).

2 Glauber dynamics on the 2-dimensional grid

Let Λ be an $n \times n$ region on the 2-dimensional Cartesian lattice and let Ω be the set of independent sets on Λ . Our goal is to sample from Ω according to the Gibbs distribution in which each $I \in \Omega$ is assigned probability

$$\pi(I) = \lambda^{|I|}/Z,$$

where $Z = \sum_{I' \in \Omega} \lambda^{|I'|}$ is the normalizing constant.

Glauber dynamics is a local Markov chain that connects two independent sets if they have Hamming distance one. The Metropolis probabilities [13] that force the chain to converge to the Gibbs distribution are given by

$$P(I, I') = \begin{cases} \frac{1}{2n} \min\left(1, \lambda^{|I'| - |I|}\right), & \text{if } I \oplus I' = 1, \\ 1 - \sum_{J \sim I} P(I, J), & \text{if } I = I', \\ 0, & \text{otherwise.} \end{cases}$$

Our concern is determining how quickly this chain converges to its stationary distribution, as this controls the efficiency of the sampling algorithm. For this we study

the *mixing time* τ , which measures the chain's distance to stationarity in terms of the total variation distance:

$$\tau(\varepsilon) = \min\{t : \|P^t, \pi\|_{tv} \leq \varepsilon, \forall t' \geq t\}.$$

The *spectral gap* of the transition matrix $\text{Gap}(P) = 1 - |\lambda_1|$, where $\lambda_0 = 1 \geq |\lambda_1| \geq \dots, |\lambda_{|\Omega|-1}|$ are the eigenvalues of the transition matrix P , is known to be inversely related to the mixing time. However, it is typically prohibitively difficult to calculate.

Alternatively, the conductance provides another bound on the mixing time that can be more amenable to analysis. Let

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

where $\pi(S) = \sum_{x \in S} \pi(x)$ is the weight of the cutset S . The following theorem, due to Jerrum and Sinclair [15] establishes the connection between small conductance and slow mixing.

THEOREM 2.1. *For any Markov chain with conductance Φ ,*

$$\frac{\Phi^2}{2} \leq \text{Gap}(P) \leq 2\Phi.$$

Hence, exponentially small conductance implies that a chain requires exponential time to converge to equilibrium.

Using this theorem, our goal now is to define a partition of the state space that has very small conductance, thereby showing that Glauber dynamics mixes slowly. Section 2.1 introduces the notation that we will use to characterize this partition and in Section 2.2 we show how to use these definitions to bound the conductance.

2.1 Topological obstructions Given a lattice region $\Lambda = (V, E)$, we first define a graph on the edges of Λ whose adjacencies capture those pairs of edges that are incident and perpendicular. In particular, let $H = (V_H, E_H)$ be defined so that $V_H = E$ and $(e_1, e_2) \in E_H$ iff e_1 and e_2 share one vertex and their other vertices have no coordinates in common. In other words, if $e_1 = (u_1, v_1)$ and $e_2 = (u_2, v_2)$ where $u_1 = v_1$, then $e_1 \sim e_2$ if $v_1 = v_2 + \alpha$, where $\alpha = (\pm 1, \pm 1)$. Hence an edge $e = (u, v) \in \Lambda$ corresponds to a vertex V_H with at most 4 neighbors in H , two arising from edges (in Λ) incident to u , and two incident to v . In what follows it will sometimes be convenient to identify vertices e in $H_U(I)$ with the midpoint of the corresponding edge in Λ . (See Figure 1.)

The graph H plays a crucial role in defining the features of independent sets that determine the partition of the state space for our proof of slow mixing. Given an independent set $I \in \Omega$, we first focus on the set of unoccupied vertices $\bar{I} = V \setminus I$. Let $H_U(I)$ be the induced subgraph of H whose vertices correspond to edges in Λ whose endpoints are both unoccupied. A couple of observations about the graph $H_U(I)$ will be very useful later on. If I is an independent set, then every vertex v in $H_U(I)$ that lies in the interior of H must have degree at least two in $H_U(I)$. We can see this by noticing that such a vertex corresponds to an edge $e = (u, v)$ in Λ with $u, v \in \bar{I}$. Therefore the two unit squares in Λ bounded on one side by e each contain at most one vertex in I . In fact, there is at least

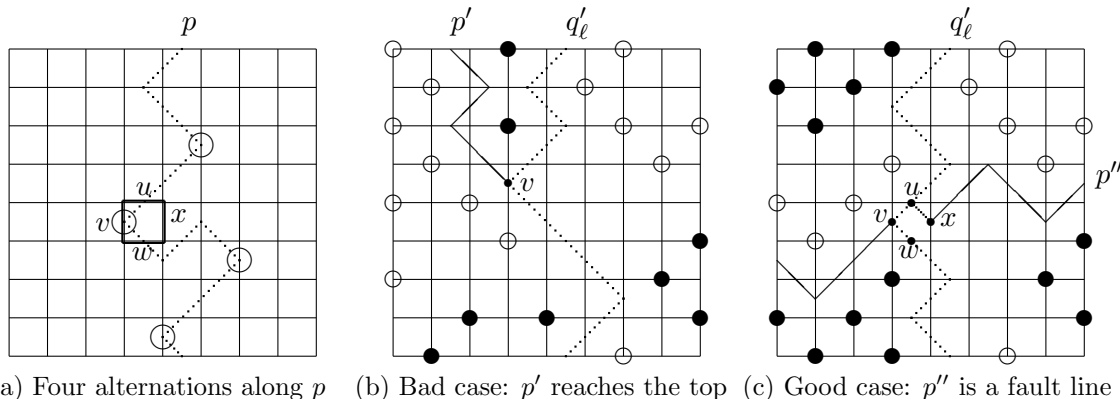


Figure 2: Reducing the number of alternations on a spanning path.

one edge incident to v arising from unoccupied vertices in each of these two squares; we make use of this fact in lemma 2.1.

The following definition begins our characterization of independent sets in Ω .

DEFINITION 2.1. *We say that $H_U(I)$ spans Λ if it contains a connected path p from the top boundary of H to the bottom, or from the left side to the right side. We call p a spanning path if it is minimal, meaning that it hits the top and bottom boundaries, or left and right, exactly once.*

Any spanning path p in $H_U(I)$ is a topological obstruction that partitions the grid Λ into two regions, left and right if p is vertical, or bottom and top if p is horizontal. These regions are the components formed when we remove all the edges in Λ corresponding to vertices on p .

Consider superimposing the path p and the lattice Λ so that a vertex on p bisects its corresponding edge. In some case it turns out that the path is only tangent to an edge e in Λ , so that cutting e is not necessary to disconnect the lattice. If the path touches e but does not cross it, then we say v is an *alternation* (see Figure 2, part (a)).

Notice that if p does not have any alternations, then it is nicely behaved in the sense that p cuts edges in such a way that the vertices on either side of the path all lie on the same sublattice of Λ . Similarly, if p has one alternation at v and we cut along p , then the vertices of Λ that fall just to one side of p have one parity until v and then have the other parity, with the complementary parities on the other side of the cut.

It turns out that the occurrence of an alternation on a spanning path might be unavoidable. Imagine the top left and bottom right “quadrants” of a square lattice region are densely packed with vertices on odd points, and the other two quadrants are densely packed with vertices on even points, allowing some space between these as necessary to ensure that this is an independent set. Then, by parity, any spanning path from the top to the bottom must contain an alternation, and similarly for paths from the left to the right of Λ . The following definition and lemma shows that this example is the worst case possible if we are interested in minimizing the occurrence of alternations.

DEFINITION 2.2. *We call a spanning path a fault line if it has zero or one alternation point(s).*

LEMMA 2.1. *Let I be an independent set such that $H_U(I)$ has a spanning path. Then it has a fault line.*

Proof. Let I be an independent set such that $H_U(I)$ has at least one spanning path and assume without loss of generality that it is vertical (since if it isn't we could rotate the square by 90° for the purpose of this argument). Just as every spanning path p partitions Λ into two pieces, it also naturally partitions the vertices of H that are not on the path p into two sets; we call these the *left side* $L(p)$ and the *right side* $R(p)$. There is a partial order on vertical spanning paths, where we say $p \leq p'$ if $L(p) \subseteq L(p')$. This partial order gives rise to a natural lattice structure on paths in $H_U(I)$, and there is a vertical spanning path p_ℓ that is minimal in the sense that $L(p_\ell) \subseteq L(p)$ for all vertical spanning paths p . We call p_ℓ the *leftmost* spanning path.

It will be convenient to restrict our attention to “clean paths” that do not have two consecutive alternations since we can always remove these. If a spanning path has such a pair of alternation points, then there are four consecutive vertices v_1, v_2, v_3, v_4 on the path (defining three consecutive edges) that correspond to the four sides of a square in Λ . Removing these three edges and instead adding the edge (v_1, v_4) results in a path with fewer alternation pairs, and we can iterate this process until this is a clean path. If we perform this modification to the leftmost path p_ℓ , forming q_ℓ , then we do in fact get the leftmost clean path (i.e., the clean path that is left of all other clean paths). To see this, observe that $q_\ell > p_\ell$ by the minimality of p_ℓ . If there were any paths that contained either alternation point from a pair we removed but not the other, then p_ℓ again could not have been minimal. Therefore q_ℓ is the leftmost clean path.

If q_ℓ does not contain alternation points, then we have identified a fault line, so we assume that it contains at least one alternation. Let v be any such alternation point on q_ℓ . Recall that for v to be an alternation, q_ℓ must hit v but not cross it (see Figure 2). This means that if u and w are the two neighbors of v along this path, then u, v, w correspond to three of the four edges around some unit square in Λ , and all four corners of the square must be unoccupied in I . Let x be vertex in $H_U(I)$ corresponding to the remaining side of this square. The point x cannot fall on the left side of q_ℓ or the new path q' formed by replacing v with x along q_ℓ , which is also clean, would violate the minimality condition of q_ℓ . In addition, x cannot fall on the path q_ℓ because the presence of either edge (u, x) or (w, x) would mean the path is not clean; likewise, if x lies on the path but is not adjacent to u or w on p , then, adding the edge (u, x) and removing the segment of p between u and x would result in a path left of q_ℓ . Therefore x lies to on right side of q_ℓ . In addition, we already observed that every vertex v in $H_U(I)$ that lies in the interior of H has degree at least two, and there must be one neighbor representing a side of each of the two unit squares containing v . The two neighbors of v along the path q_ℓ both come from the same square since v is an alternation, so v must have at least one additional neighbor y in $H_U(I)$. If y were to lie on the path q_ℓ then we could shorten the path by connecting v and y directly, and this new path would violate the minimality of q_ℓ . Therefore y lies on the left side of q_ℓ .

We will now argue that there must be a path in $H_U(I)$ extending from v to the left boundary of H . We defined q_ℓ so that v cannot lie on the top or bottom boundary of H . Similarly, it cannot lie on the right boundary since this would prohibit x from lying to the right of q_ℓ . So v either lies on the left boundary of H , in which case we have identified

a trivial path from v to this boundary, or it lies on the interior of H . Considering the remaining case in which v lies in the interior, let $z_0 = v$ and let $z_1 = y$ be the neighbor of v that lies to the left of q_ℓ . We construct the rest of this path recursively. Given a vertex z_i , we define z_{i+1} to be any neighbor of z_i that is on the other side of z_i from z_{i-1} (i.e., mapping back to the edges in Λ , the unit square including z_i and z_{i+1} is distinct from the one defined by z_i and z_{i-1}). This construction continues until we hit a point z_i that is on the boundary of H or the path q_ℓ . However, since q_ℓ is the leftmost vertical path, it is impossible for z_i to lie on the path q_ℓ or we would be able to find a new path which is clean and to the left of q_ℓ ; we simply replace the segment of q_ℓ from v to z_i with this new path we just constructed, removing any alternation pairs if necessary. Likewise, the other potentially bad case in which z_i lies on the top or bottom boundary, depicted in part (b) of Figure 2, is also impossible if q_ℓ is the leftmost clean path. Therefore $v = z_0, z_1, \dots, z_i$ must be a path in $H_U(I)$ from v to the left boundary of H . Moreover, by construction this path has no alternation points.

Our last step is constructing a second path from v to the right boundary of H , since together these two paths from v to the boundary will form a spanning path. As before, let $x \in G_U(I)$ be the final vertex in the square defined by the two edges (u, v) and (w, v) on q_ℓ incident to v . We have already argued that x falls to the right of q_ℓ . Let q_r be the rightmost vertical spanning path that includes v but does not include x . The path q_ℓ is an example of a such a path, so this is always defined. Because there is a lattice structure on spanning paths, we know that $q_r \geq q_\ell$, so q_r must include vertices u and w . Replace v with x along this path. Now x is an alternation point with v to the left of the path so, by a construction similar to the one above, there must be a path from x to the right boundary of H that does not contain any alternation points. Adding edges (v, u) and (u, x) to these two paths we constructed yields a horizontal spanning path with at most one alternation point at u , as depicted in part (c) of Figure 2. Hence this spanning path is indeed a fault line in I .

The independent sets with spanning paths, and therefore fault lines, form one piece in the partition of Ω that we will use to show slow mixing. To complete this characterization, we now consider the remaining independent sets that do not have spanning paths.

In what follows, it will be convenient to refer to the two natural sublattices of the bipartite lattice $\Lambda = (V, E)$. Let V_0 be the even vertices in Λ , i.e., points (x, y) such that $x + y$ is even, and let V_1 be the odd vertices $V \setminus V_0$. The sublattices denoted $\Lambda_0 = (V_0, E_0)$ and $\Lambda_1 = (V_1, E_1)$ connect vertices on V_0 and V_1 respectively that have distance exactly two in Λ . Thus, each vertex in each of these sublattices has at most eight neighbors. We use $\Lambda^2 = \Lambda_0 \cup \Lambda_1$ to represent the union of these two sublattice, when convenient, so Λ^2 has two distinct connected components, one on the odd vertices and one on the even vertices. The *boundary* of Λ_0 is the set of even vertices that are either on the boundary of Λ or adjacent to an odd vertex on the boundary of Λ ; we define the boundary of Λ_1 similarly. Finally, the boundary of Λ^2 refers to the union of these two boundaries and is precisely the union of vertices on the boundary of Λ with those adjacent to a vertex on the boundary of Λ .

DEFINITION 2.3. *A monochromatic bridge in an independent set I is a path q in Λ_0 or Λ_1 that includes only occupied vertices and crosses from the left boundary to the right boundary of Λ^2 , or from the top to the bottom. If it crosses from left to right we call it a*

horizontal bridge and if it crosses from top to bottom, we call it a vertical bridge. We say that I has a monochromatic cross if it contains both a horizontal and a vertical bridge.

It is straightforward to see that any independent set that has a horizontal bridge on one sublattice cannot have a vertical bridge on the other sublattice. Therefore all bridges in an independent set with a monochromatic cross must have the same parity.

LEMMA 2.2. *If I is an independent set then it has either a monochromatic cross or a fault line, but not both.*

Proof. It is fairly easy to see that if I has a fault line it cannot also have a monochromatic cross. This is because every point along a fault corresponds to two unoccupied vertices in L , one on each sublattice. Removing these vertices disconnects the left (or top) boundary of Λ^2 from the right (resp. bottom) boundary on each of the two sublattices.

For the remainder of the proof we concentrate on verifying the converse. Suppose I does not have a monochromatic cross, and in particular assume without loss of generality that I does not have a horizontal bridge. We will show that I must have a vertical spanning path, and therefore a fault line.




Figure 3: The shift operation, and the points falling “between” F and F' .

Let ∂_L be the set of vertices in I that lie on the left boundary of Λ or have a neighbor on the left boundary of Λ (as these are precisely the boundary vertices of Λ^2). Let $C(v)$ be the connected component in Λ^2 containing v , together with all of the neighbors of these vertices in Λ (i.e., lying in the other sublattice). Now define $I_L = \cup_{v \in \partial_L} C(v)$ to be the union of all of these components that contain at least one vertex in ∂_L . We assume that I does not have a horizontal bridge, so I_L cannot contain any vertex on the right boundary of Λ^2 .

We now restrict our attention to the complement of this set, $\overline{I_L}$, and let $C_R \subseteq \overline{I_L}$ be the connected component in the induced subgraph on $\overline{I_L}$ containing the right boundary of Λ . Borrowing terminology from [8], we define the *inner boundary* of C_R to be the set of vertices in C_R that have neighbors in I_L , together with any vertices in C_R that lie on

the left boundary of Λ . The *outer boundary* of C_R are the vertices in I_L with neighbors in C_R . The union of these two sets contains a connected path p in Λ extending from the top of the lattice to the bottom that does not partition I_L into two or more pieces, so p roughly follows the boundary of $I_L \cup \partial_L$. All of the vertices on p are unoccupied; if they are on the inner boundary then they are the unoccupied neighbors of a component $C(v)$ and if they are on the outer boundary then they must be unoccupied or they would have been part of a connected component $C(v)$.

We can now infer that the edges on p all correspond to vertices in $H_U(I)$. Let S_p be the possibly larger set of vertices in $H_U(I)$ that correspond to edges in Λ with at least one endpoint on p . We will show that the set S_p must contain a vertical spanning path. For the sake of contradiction, assume S_p is not connected. Since all edges along p give rise to a vertex in S_p , if S_p is not connected then p must contain a set of three colinear vertices v_1, v_2, v_3 such that v_2 has only occupied neighbors. If v_2 were on the left boundary and had an occupied neighbor, then it would have been in I_L . It also cannot lie on any other boundary since I does not have a horizontal bridge and the path p is a minimal path hitting the top and bottom boundaries once. Finally, if v is on the interior of Λ and both of its neighbors are occupied, then both lie in I_L so v_1, v_2, v_3 would disconnect I_L , also a contradiction. Therefore it must be that S is connected in $H_U(I)$. In addition it spans the H from the top to the bottom since p spans Λ , so we can conclude that I contains a fault line.

The following lemma finishes our characterization of independent sets. It confirms that in order to move from a configuration with an even monochromatic cross to one with an odd monochromatic cross using steps of the Markov chain, we must pass through a configuration with a fault line.

LEMMA 2.3. *If I and I' are two independent sets on Λ such that I has a monochromatic cross in Λ_0 (consisting of even vertices) and I' has a monochromatic cross in Λ_1 (consisting of odd vertices), then $P(I, I') = 0$.*

Proof. We first show that $I \neq I'$ by proving that a configuration cannot simultaneously contain a monochromatic cross of each parity. Suppose that I has both an odd and an even monochromatic cross, and consider the even vertical bridge and the odd horizontal bridge. Since each bridge connects vertices at distance 2 apart in Λ , there must be a vertex v on one of the bridges that is within distance 1 of a point w on the other bridge. But v and w cannot both be in I since they are nearest neighbors, hence we have a contradiction.

Now consider the case where $P(I, I') > 0$ and $I \neq I'$. We can assume without loss of generality that $I = I' \cup \{v\}$, for some $v \notin I$. But if I has an odd cross in Λ_0 and we add a vertex to form I' containing an even cross, then I' must in fact have a monochromatic cross of each parity. From the previous argument this is impossible, so I and I' must be at distance at least two in the Markov chain P .

2.2 Mixing times The purpose of the previous section was to identify a useful partition of the state space into three sets: $\Omega = \Omega_0 \cup \Omega_{\mathcal{F}} \cup \Omega_1$. The set Ω_0 contains all independent sets with an even monochromatic cross, Ω_1 contains all independent sets with an odd monochromatic cross, and $\Omega_{\mathcal{F}}$ is the set of independent sets containing a fault line. We

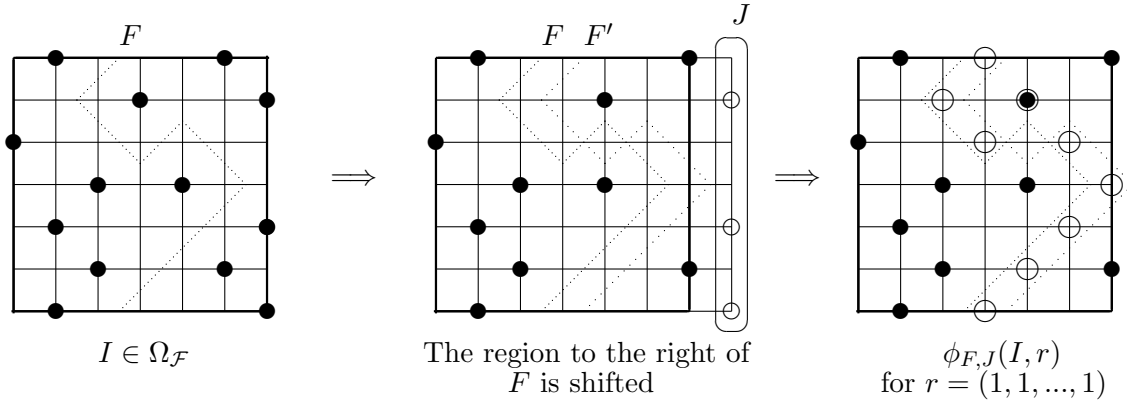


Figure 4: The injective map $\phi_{F,J}$.

have shown that these three sets form a pairwise disjoint partition of the state space, and furthermore Ω_0 and Ω_1 are not connected by moves in P . Our last remaining step is showing that $\pi(\Omega_{\mathcal{F}})$ is exponentially smaller than both $\pi(\Omega_0)$ and $\pi(\Omega_1)$. Clearly we know that $\pi(\Omega_0) = \pi(\Omega_1)$ by symmetry.

The key ingredient for bounding $\pi(\Omega_{\mathcal{F}})$ is an injective map from $\Omega_{\mathcal{F}}$ into Ω . In order to define this injection, we need to understand a “shift operation” on elements of $\Omega_{\mathcal{F}}$. Let $I \in \Omega_{\mathcal{F}}$ be an independent set with a vertical fault line F . The fault line partitions the vertices in Λ into two sets, $\text{Right}(F)$ and $\text{Left}(F)$, depending on the side of the fault on which they lie. Recall that a fault has at most one alternation point, and that all other points along the path in $H_U(I)$ bisect an edge in Λ . Define the length of a fault to be the number of edges in Λ that are bisected by vertices in F . Notice that all fault lines with zero alternation points have length $N = n + 2\ell$, for some integer ℓ , since they all have the same parity. We will use this representation even if there is a single alternation point; this will affect the analysis of what follows by only a constant factor.

Let $I' = \sigma(I, F)$ be the configuration formed by shifting $\text{Right}(F)$ one to the right. We will not be concerned right now if some vertices “fall off” the right side of the region Λ . Let $F' = \sigma(F)$ be the F shifted one to the right. We define the points that lie in $\text{Right}(F) \cap \text{Left}(F')$ to be the points that fall “in between” F and F' . (See Figure 3.)

Note that F' also divides Λ into two pieces, $\text{Left}(F')$ and $\text{Right}(F')$. We say that points in $\text{Right}(F) \cap \text{Left}(F')$ are *in between* F and F' .

LEMMA 2.4. *Let I be an independent set with a fault line F . Let $I' = \sigma(F, I)$ and $F' = \sigma(F)$ be defined as above. Then*

1. F and F' are both fault lines in I' .
2. If we form I'' by adding all the points that lie in between F and F' to I' (except the unique odd point incident to the alternation point, if it exists), then I'' will be an independent set.
3. If $|F| = n + 2\ell$, then there are exactly $n + \ell$ points that lie in between F and F' .

Proof. 1. Points in $\text{Left}(F)$ are not shifted, so any of these adjacent to F remain unoccupied after the shift. Points in $\text{Right}(F) \cap \text{Left}(F')$ were unoccupied before the shift, and therefore remain unoccupied when $\text{Right}(F)$ is shifted to the right. So F is a

fault in I' . All points that are adjacent to F' are the image of points adjacent to F under the shift, so these are all unoccupied in I' and F' is also a fault.

2. If $u \in \text{Right}(F) \cap \text{Left}(F')$, then any occupied neighbor of u must be in $\text{Right}(F)$ and has been shifted so that its preimage is now unoccupied. Moreover, only points in $\text{Left}(F')$ can be shifted to a neighbor of u , so these remain unoccupied after the shift. Therefore all neighbors of u in I' are unoccupied. Finally, notice that if there are no alternations, then all points in $\text{Right}(F) \cap \text{Left}(F')$ lie on the same sublattice, so they can all be added simultaneously. If there is an alternation $v \in H_U(I)$, then the only two possible neighbors in $\text{Right}(F) \cap \text{Left}(F')$ with different parity are incident to the edge corresponding to v . We do not add the odd one, so again I'' will be an independent set.

3. Consider any horizontal line across Λ . Each such line crosses F an odd number of times. The point immediately to the right of the first crossing is in $\text{Right}(F) \cap \text{Left}(F')$, the next is in $\text{Left}(F)$, and this alternation continues. If the number of crossings is $2k + 1$, then there are exactly $k + 1$ that are in $\text{Right}(F)$. Summing over the n rows gives $n + \ell$ points in $\text{Right}(F) \cap \text{Left}(F')$, where $n + 2\ell$ is the total length of the fault F .

We are now ready to prove the main lemma we will use to show that $\pi(\Omega_{\mathcal{F}})$ is very small. Let $I \in \Omega_{\mathcal{F}}$ be an independent set with a fault line, which we assume is vertical. (If I only has horizontal fault lines, we can rotate Λ so that it is vertical for the purpose of this argument; the net effect of ignoring these independent sets is at most a factor of 2 in the upper bound on $\pi(\Omega_{\mathcal{F}})$, and this will get incorporated into other constant factors.) Let $F = F(I)$ be the leftmost fault line. To span the $n \times n$ box, the length of the fault must be at least $2n$ and it will have even length, so we write the length as $2n + 2\ell$, for some non-negative integer ℓ .

Let $\Lambda_{1,n}$ be the $1 \times n$ lattice representing the last column of Λ , and let J be any independent set on $\Lambda_{1,n}$. We further partition $\Omega_{\mathcal{F}}$, into $\cup_{F,J} \Omega_{F,J}$, where $I \in \Omega_{F,J}$ if it has leftmost fault line F and is equal to J when restricted to the last column $\Lambda_{1,n}$.

LEMMA 2.5. *Let F be a fault in Λ with length $2n + 2\ell$ and let δ equal the number of alternation points on F (so $\delta = 0$ or 1). Let J be an independent set on $\Lambda_{1,n}$. With $\Omega_{F,J}$ defined as above, we have*

$$\pi(\Omega_{F,J}) \leq \lambda^{|J|} (1 + \lambda)^{-(n+\ell-\delta)}.$$

Proof. Let $r \in \{0, 1\}^{n+\ell-\delta}$ be any binary vector of length $n + \ell$ and let $|r|$ denote the number of bits set to 1, where $|r| \leq n + \ell$. The main step in this proof is to define an injective map $\phi_{F,J} : \Omega_{\mathcal{F}} \times \{0, 1\}^{n+\ell} \rightarrow \Omega$ such that, for any $I \in \Omega_{\mathcal{F}}$,

$$\pi(\phi_{F,J}(I, r)) = \pi(I) \lambda^{-|J|+|r|}.$$

Given this map, we have

$$\begin{aligned}
1 = \pi(\Omega) &\geq \sum_{I \in \Omega_{F,J}} \sum_{r \in \{0,1\}^{n+\ell-\delta}} \pi(\phi_{F,J}(I, r)) \\
&= \sum_{I \in \Omega_{F,J}} \sum_{r \in \{0,1\}^{n+\ell-\delta}} \pi(I) \lambda^{-|J|+|r|} \\
&= \sum_{I \in \Omega_{F,J}} \pi(I) \lambda^{-|J|} \sum_{r \in \{0,1\}^{n+\ell-\delta}} \lambda^{|r|} \\
&= \sum_{I \in \Omega_{F,J}} \pi(I) \lambda^{-|J|} (1 + \lambda)^{n+\ell-\delta} \\
&= \lambda^{-|J|} (1 + \lambda)^{n+\ell-\delta} \pi(\Omega_{F,J}).
\end{aligned}$$

We define the injective map $\phi_{F,J}$ in stages, referring the reader to Figure 4. For any $I \in \Omega_{F,J}$, we delete the last column (which is equal to J). Next, recalling that any fault line partitions Λ into two pieces, we identify all points in I that fall on the right half and shift these to the right by one using the map $\sigma(I, F)$. From Lemma 2.4 we know that the number of points that fall between these two fault lines is $n + \ell$, where $n + 2\ell$ is the length of the fault. The final step defining the map is to insert new points into the independent set along this strip between the two faults using the vector r , thereby adding $|r|$ new points. The new independent set $\phi_{F,J}(I, r)$ has $|I| - |J| + |r|$ points, and hence has weight $\pi(I) \lambda^{-|J|+|r|}$.

LEMMA 2.6. *Let $\Lambda_{1,n}$ be a $1 \times n$ strip, and let Ω_r be the set of independent sets on $\Lambda_{1,n}$. Then*

$$\sum_{J \in \Omega_r} \lambda^{|J|} \leq c_1 \left(\frac{1 + \sqrt{1 + 4\lambda}}{2} \right)^n,$$

for some constant c_1 .

Proof. Let S_i be the set of independent sets on $\Lambda_{1,n}$ and let $T_i = \sum_{J \in S_i} \lambda^{|J|}$. Then $T_0 = 1, T_1 = 1 + \lambda$, and

$$T_i = T_{i-1} + \lambda T_{i-2}.$$

Solving this Fibonacci-like recurrence yields the lemma.

THEOREM 2.2. *Let Ω be the set of independent sets on the $n \times n$ lattice Λ weighted by $\pi(I) = \lambda^{|I|}/Z$, where $Z = \sum_{I \in \Omega} \lambda^{|I|}$ is the normalizing constant. Let $\Omega_{\mathcal{F}}$ be the set of independent sets on Λ with a fault line. Then*

$$\pi(\Omega_{\mathcal{F}}) \leq p(n) e^{-c_2 n},$$

for some polynomial $p(n)$ and constant $c_2 > 0$, whenever $\lambda \geq 56.812$.

Proof. We will make use of the injective map $\phi_{F,J} : \Omega_{F,J} \times \{0,1\}^N \rightarrow \Omega$, where $N = n + 2\ell$ is the length of the fault line. Our final ingredient will be to bound the number of fault lines of length N . Notice that fault lines are self-avoiding walks in a lattice region that has been rotated by 45 degrees. A crude upper bound for the number of walks of length

N is $n3^N$, since we have n choices for our starting point along the top border of Λ and at most 3 choices for the next point along the walk. There has been a lot of work deriving better bounds on the number of self-avoiding walks, and the best known rigorous upper bound for the so-called “connective constant” is $\mu < 2.6792$, where the number of walks is less than μ^N [14].

We now have

$$\begin{aligned}
\pi(\Omega_{\mathcal{F}}) &= \sum_{F,J} \pi(\Omega_{F,J}) \\
&\leq \sum_{F,J} \lambda^{|J|} (1+\lambda)^{-(n+\ell-\delta)} \\
&\leq \lambda \sum_F (1+\lambda)^{-(n+\ell)} \sum_{J \in \Omega_r} \lambda^{|J|} \\
&\leq \lambda c_1 \sum_F (1+\lambda)^{-(n+\ell)} \left(\frac{1+\sqrt{1+4\lambda}}{2} \right)^n \\
&\leq \lambda c_1 \sum_{i=0}^{n^2} n \mu^{2n+2i} (1+\lambda)^{-(n+i)} \\
&\quad \times \left(\frac{1+\sqrt{1+4\lambda}}{2} \right)^n \\
&= \lambda c_1 n \sum_i \left(\frac{\mu^2}{(1+\lambda)} \right)^i \left(\frac{\mu^2(1+\sqrt{1+4\lambda})}{2(1+\lambda)} \right)^n,
\end{aligned}$$

where the third equality follows from Lemma 2.6. This means that we will have $\pi(\Omega_{\mathcal{F}}) \leq p(n)e^{-c_2 n}$, for some polynomial $p(n)$, if

1. $(1+\lambda) > \mu^2$ and
2. $2(1+\lambda) > \mu(1+\sqrt{1+4\lambda})$.

Simple algebra reveals that the second condition is satisfied whenever $\lambda^2 + (2 - \mu^2 - \mu^4)\lambda + (1 - \mu^2) > 0$. Taking $\mu = 2.6792$, we find that both of these conditions are met when $\lambda > 56.812$.

We get the following corollary as an immediate consequence.

COROLLARY 2.1. *Glauber dynamics for independent sets on the $n \times n$ grid Λ takes time at least $e^{c_2 n}$ to mix, for some constant $c_2 > 0$, when $\lambda > 56.812$.*

Proof. We will bound the conductance by considering the cut $S = \Omega_0$. It is clear that

$\pi(S) \leq 1/2$ since $\bar{S} = \Omega_{\mathcal{F}} \cup \Omega_1$ and $\pi(\Omega_0) = \pi(\Omega_1)$. Thus,

$$\begin{aligned} \Phi &\leq \Phi_S = \frac{\sum_{s \in \Omega_0, t \in \Omega_{\mathcal{F}}} \pi(s) P(s, t)}{\pi(\Omega_0)} \\ &= \frac{\sum_{s \in \Omega_0, t \in \Omega_{\mathcal{F}}} \pi(t) P(t, s)}{\pi(\Omega_0)} \\ &\leq \frac{\sum_{t \in \Omega_{\mathcal{F}}} \pi(t)}{\pi(\Omega_0)} \\ &= \frac{\pi(\Omega_{\mathcal{F}})}{\pi(\Omega_0)}. \end{aligned}$$

Given theorem 2.2, it is trivial to show that $\pi(\Omega) > 1/3$, thereby establishing that the conductance is exponentially small. It follows from theorems 2.1 that Glauber dynamics takes exponential time to converge.

Remark: Note that the only part of the calculations leading to the proof of Theorem 2.2 that builds on nontrivial bounds is the estimate $\mu < 2.6792$. Using the trivial bound of $\mu \leq 3$ instead, our result becomes completely self-contained, showing instead that Glauber dynamics is slowly mixing when $\lambda > 10.1962$.

3 Glauber dynamics on the 2-d torus

Let n be even, and let $\widehat{\Lambda} = \mathbb{Z}/n \times \mathbb{Z}/n$ be the $n \times n$ lattice region Λ with toroidal boundary conditions. We take $\widehat{\Omega}$ to be the set of independent sets on $\widehat{\Lambda}$ and $\widehat{\pi}$ to be the Gibbs distribution. As before, we consider Glauber dynamics that connect configurations with Hamming distance one.

Most of the definitions are analogous to the grid case, only now we do not have to worry about alternations at all. Given I , we define $H_U(I)$ as before by including edges whose endpoints are unoccupied in I .

DEFINITION 3.1. (1) We say I has a spanning cycle if there is a non-contractable cycle in $H_U(I)$. The independent set I has a fault $F = (F_1, F_2)$ if there are two vertex disjoint spanning cycles F_1 and F_2 that do not have alternation points. (By parity, if there is one there must be two.)

We say I has a monochromatic cross if there are two cycles in $\widehat{\Lambda}^2$ with different winding numbers, where $\widehat{\Lambda}^2$ is the graph on the even (or odd) sublattice of $\widehat{\Lambda}$ that connects points at distance 2.

We now partition $\widehat{\Omega}$ into $\widehat{\Omega} = \widehat{\Omega}_0 \cup \widehat{\Omega}_{\mathcal{F}} \cup \widehat{\Omega}_1$, where $\widehat{\Omega}_0$ and $\widehat{\Omega}_1$ are the sets that have a monochromatic cross of the appropriate parity and $\widehat{\Omega}_{\mathcal{F}}$ is the set of configurations that have a fault. It is not difficult to show that these three sets form a partition of the state space and that it is impossible to move from $\widehat{\Omega}_0$ to $\widehat{\Omega}_1$ using steps of the Markov chain without passing through $\widehat{\Omega}_{\mathcal{F}}$. We defer the details for the final version.

In order to show slow mixing of Glauber dynamics, it is now enough to show that $\widehat{\pi}(\widehat{\Omega}_{\mathcal{F}})$ is exponentially smaller than $\widehat{\pi}(\widehat{\Omega}_0)$ and $\widehat{\pi}(\widehat{\Omega}_1)$. Note that if $F = (F_1, F_2)$ is a fault, then each of F_1 and F_2 must have at least $2n$ steps, and each path has even length, so the total length can be written as $4n + 2\ell$ for some non-negative integer ℓ .

LEMMA 3.1. Let I be an independent set on Λ^2 and let $F = (F_1, F_2)$ be a fault in I with total length $4n + 2\ell$. Let $\widehat{\Omega}_F$ be the configurations in $\widehat{\Omega}_{\mathcal{F}}$ with “first” fault F . Then

$$\pi(\widehat{\Omega}_F) \leq (1 + \lambda)^{-(n+\ell)}.$$

Proof. We define an injection $\phi_F : \widehat{\Omega}_F \times \{0, 1\}^{n+\ell} \hookrightarrow \Omega$ so that

$$\widehat{\pi}(\phi_F(I, r)) = \widehat{\pi}(I)\lambda^{|r|}.$$

The injection is formed by cutting the torus Λ^2 along F_1 and F_2 and shifting one of the two connected pieces in any direction by one unit. There will be exactly $n + \ell$ unoccupied points near F that are guaranteed to have only unoccupied neighbors. We add a subset of the vertices in this set to I according to bits that are one in the vector r .

Given this map, we have

$$\begin{aligned} 1 &= \widehat{\pi}(\widehat{\Omega}) \geq \sum_{I \in \widehat{\Omega}_F} \sum_{r \in \{0,1\}^{n+\ell}} \widehat{\pi}(\phi_F(I, r)) \\ &= \sum_{I \in \widehat{\Omega}_F} \widehat{\pi}(I) \sum_{r \in \{0,1\}^{n+\ell}} \lambda^{|r|} \\ &= \widehat{\pi}(\widehat{\Omega}_F) (1 + \lambda)^{n+\ell}. \end{aligned}$$

THEOREM 3.1. Let $\widehat{\Omega}$ be the set of independent sets on Λ^2 weighted by $\widehat{\pi}(I) = \lambda^{|I|}/Z$, where $Z = \sum_{I \in \widehat{\Omega}} \lambda^{|I|}$ is the normalizing constant. Let $\Omega_{\mathcal{F}}$ be the set of independent sets on Λ^2 with a fault. Then

$$\widehat{\pi}(\Omega_{\mathcal{F}}) \leq p(n) e^{-c_3 n},$$

for some polynomial $p(n)$ and constant $c_3 > 0$, whenever $\lambda \geq 50.5253$.

Proof. Summing over possible locations for the two faults F_1 and F_2 and appealing to Lemma 3.1, we have

$$\begin{aligned} \widehat{\pi}(\Omega_{\mathcal{F}}) &= 2 \sum_F \widehat{\pi}(\widehat{\Omega}_F) \\ &\leq 2 \sum_F (1 + \lambda)^{-(n+\ell)} \\ &\leq 2 \sum_{i=0}^{(n^2-2n)/2} \binom{n}{2} \mu^{4n+2i} (1 + \lambda)^{-(n+i)} \\ &< n^2 \sum_i \left(\frac{\mu^4}{(1 + \lambda)} \right)^n \left(\frac{\mu^2}{(1 + \lambda)} \right)^i. \end{aligned}$$

Hence, we have $\pi(\Omega_{\mathcal{F}}) \leq p(n) e^{-c_3 n}$ for some polynomial $p(n)$ and constant $c_3 < 0$ whenever $(1 + \lambda) \geq \mu^2$. Taking the known bound $\mu < 2.6792$ gives $\lambda \geq 50.5253$.

Combining Theorems 2.1 and 3.1, we get the following corollary as an immediate consequence.

COROLLARY 3.1. *Glauber dynamics on independent sets on the $n \times n$ torus Λ^2 takes time at least $e^{c_3 n}$ to mix, for some constant $c_3 > 0$, when $\lambda \geq 50.5253$.*

Acknowledgments: This work benefitted from conversations with David Galvin and Prasad Tetali.

References

- [1] R.J. Baxter, I.G. Entig, and S.K. Tsang. Hard-square lattice gas. *Journal of Statistical Physics*, **22**: 465–489, 1989.
- [2] J. van den Berg and J.E. Steif. Percolation and the hard-core lattice model. *Stochastic Processes and their Applications*, **49**: 179–197, 1994.
- [3] C. Borgs. Polymers, contours and their applications in discrete mathematics. Monograph in preparation, 2004.
- [4] C. Borgs, J.T. Chayes, A. Frieze, J.H. Kim, P. Tetali, E. Vigoda, and V.H. Vu. Torpid mixing of some MCMC algorithms in statistical physics. *Proc. 40th IEEE Symp. on Foundations of Computer Science*, 218–229, 1999.
- [5] R.L. Dobrushin. The problem of uniqueness of a Gibbsian random field and the problem of phase transitions. *Functional Analysis and its Applications*, **2**: 302–312, 1968.
- [6] R.L. Dobrushin and S.B. Shlosman. Constructive criterion for the uniqueness of Gibbs fields. In *Statistical physics and dynamical systems*, vol. 10 of *Progress in Physics*, 347–370, Birkhauser, Boston, 1985.
- [7] D. Galvin. Personal communication, 2004.
- [8] D. Galvin and J. Kahn. On phase transitions in the hard-core model on Z^d . To appear in *Combinatorics, Probability and Computing*.
- [9] R. Glauber. Time dependent statistics of the Ising model. *Journal of Mathematical Physics* **4**: 294–307, 1963.
- [10] R. Kannan, M.W. Mahoney, and R. Montenegro. Rapid mixing of several Markov chains for a hard-core model. *14th Annual ISAAC*, 663–675, 2003.
- [11] M. Luby and E. Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets. *Random Structures and Algorithms*, 229–241, 1999.
- [12] N. Madras and G. Slade. *The Self-Avoiding Walk*. Birkhäuser, Boston, 1993.
- [13] N. Metropolis, A. W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, **21**: 1087–1092, 1953.
- [14] Pönitz and Tittman. Improved upper bounds for self-avoiding walks in Z^d . *Electronic Journal of Combinatorics*, **7**, 2000.
- [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] P. Tetali. Personal communication, 2004.