Phase Coexistence and Slow Mixing for the Hard-Core Model on $\mathbb{Z}^2$

Antonio Blanca[*] David Galvin[†] Dana Randall[‡] Prasad Tetali[§]

Abstract

The hard-core model has attracted much attention across several disciplines, representing lattice gases in statistical physics and independent sets in the discrete setting. On finite graphs, we are given a parameter λ, and each independent set I arises with probability proportional to $\lambda^{|I|}$. On infinite graphs the Gibbs distribution is defined as a suitable limit with the correct conditional probabilities. In the infinite setting we are interested in determining when this limit is unique and when there is phase coexistence – existence of multiple Gibbs states. In the finite setting, for example on finite regions of the square lattice $\mathbb{Z}^2$, we are interested in determining when local Markov chains are rapidly mixing. These problems are believed to be related and it is conjectured that both undergo a phase transition at some critical point $\lambda = \lambda_c \approx 3.79$ [1]. It remains open whether there is a single critical point, although it was recently shown that on general graphs of maximum degree Δ, the computational complexity of computing the partition function (namely, the λ-weighted count of independent sets) undergoes a phase transition at the unique well-known critical point $\lambda_c(\mathbb{T}_\Delta)$ at which the Δ-regular infinite tree $\mathbb{T}_\Delta$ undergoes a transition from uniqueness to having multiple Gibbs states [25, 27].

On $\mathbb{Z}^2$, Restrepo et al. [22] recently showed that there is a unique Gibbs state and strong spatial mixing as long as $\lambda < 2.3882$, building on breakthrough ideas of Weitz [27]. It has been shown that there are finite values for λ above which the mixing time of an associated local Markov chain is slow [5, 21], and where there will be phase-coexistence [7], although these bounds are far from the conjectured threshold. We greatly improve upon these bounds by showing that local Markov chains will be slow when $\lambda > 5.68014$ on lattice regions with periodic (toroidal) boundary conditions and when $\lambda > 7.439$ with non-periodic (free) boundary conditions. Our arguments build on the idea of fault lines introduced by Randall [21] and use a careful analysis of a new family of self-avoiding walks in the two-dimensional lattice to get improved bounds. In addition, we extend these arguments to the infinite setting to show phase coexistence when $\lambda > 5.68014$. This is the first time fault lines have been used in the context of non-uniqueness. The arguments here represent a more than tenfold improvement to the best value of λ that could possibly be obtained using previously known methods, such as those described in [5].

[*]Computer Science Division, University of California at Berkeley, Berkeley, CA 94720.
[†]Department of Mathematics, University of Notre Dame, Notre Dame, IN 46556
[‡]College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0760.
[§]School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332-0280.

1 Introduction

The hard-core model was introduced in statistical physics as a model for lattice gases, where each molecule occupies non-trivial space in the lattice, requiring occupied sites in the lattice to be non-adjacent. When a (discrete) lattice such as $\mathbb{Z}^d$ is viewed as a graph, the allowed configurations of molecules naturally correspond to independent sets in the graph.

Given a graph G, let $\mathcal{I}$ be the set of independent sets of G. Given a (fixed) *fugacity (or activity)* $\lambda \in \mathbb{R}^+$, the weight associated with each independent set I is $w(I) = \lambda^{|I|}$. The associated *Gibbs (or Boltzmann) distribution*, $\mu = \mu_{G,\lambda}$ is defined on $\mathcal{I}$, assuming G is finite, as $\mu(I) = w(I)/Z$, where the normalizing constant $Z = Z(G, \lambda) = \sum_{J \in \mathcal{I}} w(J)$ is commonly called the *partition function*. Physicists are interested in the behavior of models on an infinite graph (such as the integer lattice $\mathbb{Z}^d$), where the Gibbs measure is defined as a certain weak limit with appropriate conditional probabilities. For many models it is believed, as a parameter of the system is varied – such as the inverse temperature β for the Ising model or the fugacity λ for the hard-core model – that the system undergoes a phase transition at a critical point.

For the classical Ising model, Onsager, in seminal work [19], established the precise value of the critical temperature $\beta_c(\mathbb{Z}^2)$ to be $\log(1+\sqrt{2})$. Only recently have the analogous values for the (more general) q-state Potts model been established in breakthrough work by Beffara and Duminil-Copin [2], settling a more than half-a-century old open problem. Establishing such a precise value for the hard-core model with currently available methods seems nearly impossible. Even the *existence* of such a (unique) critical activity λ_c, where there is a transition from a unique Gibbs state to the coexistence of multiple Gibbs states remains conjectural for $\mathbb{Z}^d$ ($d \geq 2$; it is folklore that there is no such transition for $d = 1$), while it is simply untrue for general graphs (even general trees, in fact, thanks to a result of Brightwell et al. [6]). Regardless, a non-rigorous prediction from the statistical physics literature [1] suggests $\lambda_c \approx 3.796$ for $\mathbb{Z}^2$.

Thus, from a statistical physics or probability point of view, understanding the precise dependence on λ for the existence of unique or multiple hard-core Gibbs states is a natural and challenging problem. Moreover, breakthrough works of Weitz [27] and Sly [25] in recent years identified $\lambda_c(\mathbb{T}_\Delta)$ – the critical activity for the hard-core model on an infinite Δ-regular tree – as a *computational* threshold where estimating the hard-core partition function on general Δ-regular graphs undergoes a transition from being in P to being NP-hard (specifically, there is no PTAS unless $NP = RP$), further motivating the study of such (theoretical) physical transitions and their computational implications. While it is not surprising that for many fundamental problems computing the partition function *exactly* is intractable, it is remarkable that even approximating the partition function of the hard-core model above a certain critical threshold also turns out to be hard.

Our current work is inspired by these striking developments as well as a recent improvement due to Restrepo et al. [22] for the hard-core model on $\mathbb{Z}^2$, building on novel arguments introduced by Weitz [27] and establishing uniqueness for all $\lambda < 2.3882$. Here we establish, inter alia, phase coexistence for the hard-core model on $\mathbb{Z}^2$ for $\lambda > 5.68014$, shortening the interval (for λ_c to exist) significantly by providing more than a tenfold improvement to the bounds obtainable using the best previously known methods [5].

Returning to the issue of approximating the partition function or sampling from the desired Gibbs distribution, Markov chain algorithms offer a natural and powerful method. But for the method to be efficient, the underlying Markov chain must be rapidly mixing. For many problems, local Markov chains, known as Glauber dynamics, seem to be rapidly mixing below some critical point, while mixing quite slowly above it. Most notably for the Ising model on $\mathbb{Z}^2$, simple local Markov chains are known to be rapidly mixing (in fact, with optimal rate) for $\beta < \beta_c(\mathbb{Z}^2)$ and slowly mixing beyond that point, thanks to a series of papers by various mathematical physics experts (including Aizenmann, Holley, Stroock, Zegarlinski, Martinelli, Olivieri, Schonmann).

Once again, the known bounds are less sharp for the hard-core model. In the following we provide further details and mention precise bounds.

1.1 Previous bounds for slow-mixing and phase coexistence

Starting with Dobrushin [7] in 1968, physicists have been developing techniques to systematically characterize the regimes on either side of λ_c for the hard-core model. The local Markov chain known as Glauber dynamics connects pairs of configurations with Hamming distance one, with transition probabilities defined so that the unique stationary measure is the Gibbs distribution. Luby and Vigoda [16] showed that Glauber dynamics on independent sets is fast when $\lambda \leq 1$ on the 2-dimensional lattice and torus. Weitz [27] showed how to reduce the analysis on the grid to the tree, also establishing that Glauber dynamics is fast up to the critical point on the 4-regular tree, in effect for $\lambda < 1.6875$. Besides building on the work of Weitz, Restrepo et al. [22] made crucial use of the properties of the square lattice in achieving their improvement for the uniqueness regime. Using now standard machinery (by way of establishing the so-called *strong spatial mixing*), they also proved that the natural Glauber dynamics on the space of hard-core configurations is rapidly mixing for the same range of $\lambda < 2.3882$. These results also lead to efficient deterministic algorithms for approximating the partition function for independent sets on $\mathbb{Z}^2$.

On the other hand, Borgs et al. [5] showed that Glauber dynamics is slow on toroidal lattice regions in $\mathbb{Z}^d$ (for $d \geq 2$), when λ is sufficiently large (in particular, growing with d). Although the bound on λ for $\mathbb{Z}^2$ remains unpublished, it is known that $\lambda > 80$ is the best possible using their methods. Informally the argument is based on the observation that when λ is large, the Gibbs distribution favors dense configurations, and Glauber dynamics will take exponential time to converge to equilibrium. The slow convergence arises because the Gibbs distribution is bimodal: the dense configurations lie predominantly on either the odd or the even sublattice, 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 independent set by at most 1 in each step, the Markov chain has a bottleneck leading to *torpid* (slow) mixing.

Our present work builds on a novel idea from [21] in which the notion of *fault lines* was introduced to establish slow mixing for the Glauber dynamics on hard-core configurations for moderately large λ, still improving upon what was best known at that time. Randall [21] gave an improvement by realizing that the state space could be partitioned according to certain *topological obstructions* in the configurations, rather than the relative numbers of odd or even vertices. Not only does this approach give better bounds on λ, but it also greatly simplifies the calculations. First consider an $n \times n$ lattice region G with free (non-periodic) boundary conditions. A configuration I is said to have a *fault line* if there is a width two path of unoccupied vertices in I from the top of G to the bottom or from the left boundary of G to the right. Configurations that do not have a fault line must have a *cross* of occupied vertices in either the even or the odd sublattices forming a connected path in G^2 from both the top to the bottom and from the left to the right of G. Roughly speaking the set of configurations that have a fault line forms a cut set that must be crossed to move from a configuration that has an odd cross to one with an even cross, and it was shown that fault lines are exponentially unlikely when λ is large. Likewise, if $\widehat{G}$ is an $n \times n$ region with periodic boundary conditions, it was shown that either there is an odd or an even cross forming non-contractible loops in two different directions or there is a pair of non-contractible fault lines, and a similar argument can be made. Using these arguments it can be shown that Glauber dynamics is slowly mixing on $\widehat{G}$ when $\lambda > 50.59$ and on G when $\lambda > 56.88$. (Better bounds were originally reported in [21] due to a minor error, although our current results improve on the original claims as well.)

1.2 Our results

In the present work, we establish that local Markov chains will be slow when $\lambda > 5.68014$ on lattice regions with periodic (toroidal) boundary conditions and when $\lambda > 7.439$ with non-periodic (free) boundary conditions. Building on the idea of fault lines, we use a more careful analysis to define *taxi walks*, a new family of self-avoiding walks in the two-dimensional lattice. The previous bounds just used the fact that fault lines are self-avoiding walks on a rotated grid. We observe here that they in fact lie on an *oriented* version of $\mathbb{Z}^2$ where there are at most two ways to extend a walk at each step instead of 3. In addition, we show that if there is a fault line, there is always one that avoids taking 2 turns in a row, further reducing their cardinality. We show that with this characterization, the number of fault lines of length n is at most the nth Fibonacci number. Capitalizing on the fact that the walks are also self-avoiding, we get an additional improvement further reducing their number. While we cannot enumerate walks exactly, we use the fact that the log of the number of walks is subadditive to derive bounds on the total number. This leads to improvements on the bounds for λ.

Finally, we extend these arguments to the infinite setting to show phase coexistence when $\lambda > 5.68014$. This is the first time fault lines have been used in the context of non-uniqueness. The arguments here represent a significant improvement to the best value of λ that could possibly be obtained using previously known methods, such as those described in [5].

We believe that using fault lines of the type introduced in [21] has more general applicability; a natural next step is to study the hard-core model on $\mathbb{Z}^d$ for $d \geq 3$. Establishing reasonable bounds on the critical activity for $\mathbb{Z}^3$ is a challenging next step, as is pinning down how the critical value changes with d. The best upper bounds are $\widetilde{O}(d^{-1/4})$ for slow mixing [9] and $\widetilde{O}(d^{-1/3})$ [20] for phase coexistence; the best known lower bounds in both cases are $\Omega(d^{-1})$.

The rest of this manuscript is structured as follows. Section 2 provides much of the background material, including the precise definition of the relevant Markov chain and the characterization of independent sets on finite regions of $\mathbb{Z}^2$. In Section 3.1 we introduce *taxi walks* and derive bounds on their cardinality. In the remainder of Section 3 we use a characterization based on fault lines to characterize a bad cut in the state space, thereby showing that the local Markov chain requires exponential time to reach equilibrium. Finally, in Section 4, we explain how to use fault lines on the infinite lattice $\mathbb{Z}^2$ in order to show phase coexistence above a certain fugacity λ.

2 Background: Markov chains and fault lines

2.1 Glauber dynamics on independent sets on $\mathbb{Z}^2$

Let $G \subset \mathbb{Z}^2$ be an $n \times n$ lattice region and let Ω be the set of independent sets on G. Our goal is to sample from Ω according to the Gibbs distribution, where each $I \in \Omega$ is assigned probability

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

and $Z = \sum_{I' \in \Omega} \lambda^{|I'|}$ is the normalizing constant known as the *partition function*.

Glauber dynamics is a local Markov chain that connects two independent sets if they have Hamming distance one. The Metropolis probabilities [18] that force the chain to converge to the Gibbs distribution are given by

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

The conductance, introduced by Jerrum and Sinclair [24], is a good measure of the mixing rate of a chain. 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 classical theorem provides the connection between low conductance and slow mixing.

Theorem 2.1. [24] *For any Markov chain with conductance Φ we have $\frac{\Phi^2}{2} \leq Gap(P) \leq 2\Phi$, where $Gap(P)$ is the spectral gap of the adjacency matrix.*

The spectral gap is well-known to be a measure of the mixing rate of a Markov chain (see, e.g., [23]), so an exponentially small conductance is sufficient to show slow mixing. Using Theorem 2.1, our goal is therefore to define a partition that has exponentially small conductance. Section 2.2 introduces the notation that we will use to characterize this partition.

2.2 Crossings and obstructions

We begin by defining some useful graph structures. Let $G = (V, E)$ be a simply connected region in $\mathbb{Z}^2$, say the $n \times n$ square. We define the graph $G_\Diamond = (V_\Diamond, E_\Diamond)$ as follows. The vertices $V_\Diamond$ are associated with the midpoints of edges in E. Vertices u and v in $V_\Diamond$ are connected by an edge in $E_\Diamond$ if and only if they are the midpoints of incident edges in E that are perpendicular. Notice that $G_\Diamond$ is a region in a smaller Cartesian lattice that has been rotated by 45 degrees.

We will also make use of the *even and odd subgraphs* of G. For $b \in \{0, 1\}$, let $G_b = (V_b, E_b)$ be the graph whose vertex set is the set $V_b \subseteq V$ of vertices with parity b (i.e., the sum of their coordinates has parity b), with $(u, v) \in E_b$ if u and v are connected in G^2. We refer to G_0 and G_1 as the even and odd subgraphs. The graphs $G_\Diamond, G_0$ and G_1 play a central role in defining the features of independent sets that determine the partition of the state space for our proofs of slow mixing.

Given an independent set $I \in \Omega$, we say that a path p in $G_\Diamond$ is *spanning* if it extends from the top boundary of $G_\Diamond$ to the bottom, or from the left boundary to the right, and each vertex in p corresponds to an edge in G such that both endpoints are unoccupied in I. It will be convenient to color the vertices in $V_\Diamond$ along a spanning path using the parity of the vertex to the "left" (or "top") of the path in V. In particular, recall that each vertex $v \in V_\Diamond$ on the path p bisects an edge $e_v \in E$. Each edge in E has an odd and an even endpoint, and we color v *blue* if the odd vertex in e_v is to the left when the path crosses v, and *red* otherwise. Every time the color of the vertices along the path changes, we have an *alternation point*. It was shown in [21] that if an independent set has a spanning path, then it must also have one with zero or one alternation points. We call this path a *fault line*, and we let $\Omega_\mathcal{F}$ be the set of configurations in Ω that contain at least one fault line.

We say that $I \in \Omega$ has an *even bridge* if there is a path from the left to the right boundary or from the top to the bottom boundary in G_0 consisting of occupied vertices in I. Similarly, we say it has an *odd bridge* if it traverses G_1 in either direction. We say that I has a *cross* if it has both left-right and a top-bottom bridges.

Notice that if an independent set has an even top-bottom bridge it cannot have an odd left-right bridge, so if it has a cross, both of its bridges must have the same parity. We let $\Omega_0 \subseteq \Omega$ be the set of configurations that contain an even cross and let $\Omega_1 \subseteq \Omega$ be the set of those with an odd cross.

These definitions provide a useful characterization that partitions the state space $\mathcal{I}$ into three sets with one separating the other two. The following lemmas were proven in [21].

Lemma 2.2. *The sets $\Omega_\mathcal{F}$, Ω_0 and Ω_1 consisting of configurations with a fault line, an even cross or an odd cross, form a partition of the state space Ω.*

Lemma 2.3. *If I and I' are two independent sets on G such that I has an even cross and I' has an odd cross, then $P(I, I') = 0$.*

It will be useful to extend these definitions to the torus as well. Let n be even, and let $\widehat{G}$ be the $n \times n$ toroidal region $\{0, \ldots, n-1\} \times \{0, \ldots n-1\}$, where $v = (v_1, v_2)$ and $u = (u_1, u_2)$ are connected if $v_1 = u_1 \pm 1 (\text{mod } n)$ and $v_2 = u_2$ or $v_2 = u_2 \pm 1 (\text{mod } n)$ and $v_1 = u_1$. Let $\widehat{\Omega}$ be the set of independent sets on $\widehat{G}$ and let $\widehat{\pi}$ be the Gibbs distribution. As before, we consider Glauber dynamics that connect configurations with Hamming distance one.

We define $\widehat{G}_\Diamond, \widehat{G}_0$ and $\widehat{G}_1$ as above to represent the graph connecting the midpoints of perpendicular edges (including the boundary edges), and the odd and even subgraphs. As with $\widehat{G}$, all of these have toroidal boundary conditions.

Let $I \in \widehat{\Omega}$ be an independent set on $\widehat{G}$. We say that I has a fault if there are a *pair* of vertex-disjoint non-contractible cycles in $\widehat{G}_\Diamond$ whose vertices correspond to edges in $\widehat{G}$ and whose endpoints are both unoccupied and are all red or are all blue (i.e., the vertices in G to the left (top) of the path all have the same parity). We say that I has a *cross* if it has at least two non-contractible cycles of occupied sites in I with different winding numbers.

The next two lemmas from [21] characterize the partition of the state space and the cut set $\widehat{\Omega}_{\mathcal{F}}$.

Lemma 2.4. *If $\widehat{G}$ is a lattice region with toroidal boundary conditions and $\widehat{\Omega}$ is the set of independent sets on $\widehat{G}$, then $\widehat{\Omega}$ can be partitioned into sets $\widehat{\Omega}_{\mathcal{F}}, \widehat{\Omega}_0, \widehat{\Omega}_1$ where $\widehat{\Omega}_{\mathcal{F}}$ is the set of configurations with a fault and $\widehat{\Omega}_b$ is the set of configurations with a cross with parity b, for $b \in \{0, 1\}$.*

Lemma 2.5. *If I and I' are two independent sets on $\widehat{G}$ such that I has an even cross in $\widehat{G}_0$ (consisting of even vertices) and I' has an odd cross in $\widehat{G}_1$ (consisting of odd vertices), then $P(I, I') = 0$.*

3 Lower bounds on the mixing time

Here we bound the mixing time of Glauber dynamics by showing that the conductance is exponentially small. We start by defining "taxi walks" since they play a critical role in all that follows.

3.1 Taxi walks

The strategy for the proofs of slow mixing will be to use a "Peierls argument" to define a map from $\Omega_{\mathcal{F}}$ to Ω that takes configurations with fault lines to ones with exponentially larger weight. The map is not injective, however, so we need to be careful about how large the pre-image of a configuration can be, and for this it is necessary to get a good bound on the number of fault lines. In [21] the number of fault lines was bounded by the number of self-avoiding walks in $G_\Diamond$ (or $\widehat{G}_\Diamond$ on the torus). However, this is a gross over count because, as we shall see, this includes all spanning paths with an arbitrary number of alternation points. We can get much better bounds on the number of fault lines by only counting self-avoiding walks with zero or one alternation points.

To begin formalizing this idea, we put an orientation on the edges of $G_\Diamond$. Each edge $(u, v) \in E_\Diamond$ corresponds to two edges in E that share a vertex $w \in V$. We orient the edge "clockwise" around w if w is even and "counterclockwise" around w if w is odd. For paths with zero alternation points, all of the edges must be oriented in the same direction. If we rotate $G_\Diamond$ so that the edges are axis aligned, then this simply means that the horizontal (resp. vertical) edges alternate direction according to the parity of the y- (resp. x-) coordinates, like many common metropolises.

We can now define "taxi walks". Let $\vec{\mathbb{Z}}$ be a directed grid region where streets are horizontal, with even numbered streets oriented East and odd numbered streets oriented West and avenues are vertical, with even numbered avenues oriented North and odd numbered avenues oriented South.

Definition 3.1. *A* taxi walk *is an oriented walk in $\vec{\mathbb{Z}}$ that never revisits any vertex (and thus is self-avoiding) and never takes two left or two right turns in a row.*

We call these taxi walks because violation of either restriction during a taxi ride would cause suspicion among savvy passengers.

Lemma 3.1. *If an independent set I has a fault line F with no alternations, then it also has a fault line F' so that either F' or F'^R (the reversal of F') is a taxi walk.*

Proof. It is straightforward to see that if I has a fault line F with no alternation points, then it must have all of its edges oriented the same way (in $G_\Diamond$) and it must be self-avoiding. Suppose F is the minimal length fault line in I without any alternations, and suppose that F has two successive turns. Because of the parity constraints, the vertices immediately before and after these two turns must both connect edges that are in the same direction, and these five edges can be replaced by a single edge to form a shorter fault line without any alternations. This is a contradiction to F being minimal, completing the proof. □

The critical step will be bounding the number of taxi walks. We start by recalling what we know about the number of standard self-avoiding walks. Self-avoiding walks have been studied extensively, although many basic questions remain (see, e.g., [17]). It is easy to see that in $\mathbb{Z}^2$, the number c_n of walks of length n, grows exponentially with n as $2^n \leq c_n \leq 4 * 3^{n-1}$, since there are always at most 3 ways to extend a self-avoiding walk of length $n-1$ and walks that only take steps to the right or up can always be extended in 2 ways. It is believed that in $\mathbb{Z}^2$ the number of walks grows as $\mu^n f(n)$, where μ is known as the *connective constant* and $f(n) = \Theta(n^{11/32})$. In 1962 Hammersley and Welsh [13] showed that $c_n \approx c\mu^n \exp(O(\sqrt{n}))$, for some constant c, although there is a lot of experimental and heuristical evidence to suggest $f(n)$ grows as a small polynomial.

Let $\widetilde{c}_n$ be the number of taxi walks of length n. It is easy to see that $2^{n/2} < \widetilde{c}_n < 4 * 3^{n-1}$ by observing that we can always take pairs of steps East or North and since $\widetilde{c}_n < c_n$, the number of standard self-avoiding walks.

Lemma 3.2. *Let $\widetilde{c}_n$ be the number of taxi walks of length n. Then $\widetilde{c}_n = O((1+\sqrt{5})/2)^n$.*

Proof. Notice that at each vertex there are exactly two outgoing edges in $\vec{\mathbb{Z}}$. If we arrive at v from u, then one of the outgoing edges continues the walk in the same direction and the other is a turn. Notice that the two allowable directions are determined by the parity of the coordinates of v, so we can encode each walk as a bitstring $s \in \{0,1\}^{n-1}$. If $s_1 = 0$ then the walk starts by going East (along a street) and if $s_0 = 1$ the walk starts North along an avenue. For all $i > 1$, if $s_i = 0$ the walk continues in the same direction as the previous step, while if $s_i = 1$ then the walk turns in the permissible direction. Using this encoding, it is easy to see that a walk s will never have two ones in a row, and hence $\widetilde{c}_n \leq f_n = O(\phi^n)$, where f_n is the nth Fibonacci number and $\phi = (1+\sqrt{5})/2 \approx 1.618$ is the golden ratio. □

We can get even better upper bounds on $\widetilde{c}_n$ from subadditivity.

Lemma 3.3. *Let $\widetilde{c}_n$ be the number of taxi walks of length n and let $1 \leq i \leq n-1$. Then $c_n \leq c_i\ c_{n-i}$.*

Proof. As with traditional self-avoiding walks, the key is to recognize that if we split a taxi walk of length n into two pieces, the resulting pieces are both self-avoiding. Let $s = s_1, \ldots, s_n$ be a taxi walk of length n and let $1 \leq i \leq n-1$. Then the initial segment of the walk $s_I = s_1, \ldots, s_{i+1}$ is a taxi walk of length i. Let $p = (x, y)$ be the ith vertex of the walk s. Let s_F be the final $n-i$ steps of the walk s starting at p. We define $f(s_F)$ by translating the walk so that $f(p)$ is the origin, reflecting horizontally if p_x is odd and reflecting vertically if p_y is odd. Notice that this always produces a valid taxi walk of length $n-i$ and the map f is invertible given p. Therefore $c_n \geq c_i\ c_n - i$. □

It follows from Lemma 3.1 that $a_n = \log c_n$ is subadditive, i.e., $a_{n+m} \le a_n + a_m$. Fekete's Lemma states that for every subadditive sequence $\{a_n\}_{n=1}^{\infty}$, the limit $\lim_{n\to\infty} a_n/n$ exists and

$$\lim_{n\to\infty} \frac{a_n}{n} = \inf \frac{a_n}{n} \tag{1}$$

(see, for example, [26, Lemma 1.2.2]). Thus, we can write the number of taxi walks as $\widetilde{c}_n = \mu_t^n f_t(n)$, where μ_t is the connective constant associated with taxi walks and $f_t(n)$ is subexponential in n. Frequently we will use $\widetilde{c}_n \le \mu^n$ for any fixed $\mu > \mu_t$ (valid for all large n).

Subadditivity gives us a strategy for getting a better bound on μ_t. From (1) we see that for all n, $\log c_n/n$ is an upper bound for $\log \mu_t$. The number $\widetilde{c}_n$ of taxi walks of length n, with $40 \le n \le 60$, were enumerated on a super-computer using 240 cores [http://www.nersc.gov/systems/hopper-cray-xe6/]. Using $c_{60} = 2189670407434$ gives a bound of $\mu_t < 1.6057317$. Note that exact counts for larger n will improve the bound on the connective constant as well as our bound on λ for independent sets.

n	$\widetilde{c}_n$	Estimate of $\mu_t^4 - 1$
40	219324398	5.825095
41	348109128	5.812995
42	552582790	5.801572
43	877163942	5.790699
44	1389806294	5.779188
45	2204289314	5.768817
46	3496483316	5.758977
47	5546212122	5.749573
48	8783360626	5.739666
49	13922238632	5.730664
50	22069957494	5.722087
51	34986181158	5.713860
52	55383388278	5.705232
53	87740467384	5.697333
54	139014623272	5.689781
55	220254102104	5.682515
56	348536652664	5.674924
57	551914140382	5.667929
58	874039817792	5.661222
59	1384184997874	5.654751
60	2189670407434	5.648014

3.2 Glauber dynamics on the 2-d torus

We are now ready to complete the proof of slow mixing, starting first with the two-dimensional torus. Let n be even, and let $\widehat{G} = \{0, \ldots, n-1\} \times \{0, \ldots, n-1\}$ be the $n \times n$ lattice region with toroidal boundary conditions. We take $\widehat{\Omega}$ to be the set of independent sets on $\widehat{G}$ and let $\widehat{\pi}$ be the Gibbs distribution. Our strategy is to partition the state space $\widehat{\Omega}$ into three sets: $\widehat{\Omega}_{\mathcal{F}} \cup \widehat{\Omega}_0 \cup \widehat{\Omega}_1$. The set $\widehat{\Omega}_0$ contains all independent sets with an even cross, $\widehat{\Omega}_1$ contains all independent sets with an odd cross, and $\widehat{\Omega}_{\mathcal{F}}$ is the set of independent sets containing a fault. We have shown that these three sets form a pairwise disjoint partition of the state space, and furthermore $\widehat{\Omega}_0$ and $\widehat{\Omega}_1$ are not directly connected by moves in the chain P. Our last remaining step is showing that $\widehat{\pi}(\widehat{\Omega}_{\mathcal{F}})$ is exponentially smaller than both $\widehat{\pi}(\widehat{\Omega}_0)$ and $\widehat{\pi}(\widehat{\Omega}_1)$. (Clearly we know that $\widehat{\pi}(\widehat{\Omega}_0) = \widehat{\pi}(\widehat{\Omega}_1)$ by symmetry.)

Let $I \in \widehat{\Omega}_{\mathcal{F}}$ be an independent set with fault $F = (F_1, F_2)$. The fault lines partition the vertices of I into two sets, I_A and I_B, depending on which side of F_1 and F_2 they lie. Define the length of the fault to be the total number of edges on the path F_1 and F_2 in $G_\Diamond$. All fault lines have length $N = n + 2\ell$, for some integer ℓ, since they all have the same parity.

Let $I' = \sigma(I, F)$ be the configuration formed by shifting I_A one to the right. Let $F_1' = \sigma(F_1)$ and $F_2' = \sigma(F_2)$ be the images of the fault under this shift. We define the points that lie in $F_1 \cap F_1'$ and $F_2 \cap F_2'$ to be the points that fall "in between" F and $F' := (F_1', F_2')$.

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 $\widehat{\pi}(\widehat{\Omega}_1)$. It will be convenient to order the set of possible fault lines so that given a configuration $I \in \widehat{\Omega}_{\mathcal{F}}$ we can identify the *first* fault it contains. The following lemmas are modified from [21].

Lemma 3.4. *Let $\widehat{\Omega}_F$ be the configurations in $\widehat{\Omega}_{\mathcal{F}}$ with "first" fault $F = (F_1, F_2)$. Write the length of F as $2n + 2\ell$. 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 $\widehat{G}$ 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

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

□

Theorem 3.5. *Let $\widehat{\Omega}$ be the set of independent sets on $\widehat{G}$ weighted by $\widehat{\pi}(I) = \lambda^{|I|}/Z$, where $Z = \sum_{I \in \widehat{\Omega}} \lambda^{|I|}$. Let $\Omega_{\mathcal{F}}$ be the set of independent sets on $\widehat{G}$ with a fault. Then for any $\lambda > \mu_t^4 - 1$ there is a constant $c > 0$ such that*

$$\widehat{\pi}(\Omega_{\mathcal{F}}) \leq e^{-cn}.$$

Proof. Summing over possible locations for the two faults F_1 and F_2 and using Lemma 3.4, we have

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

Choosing μ above to satisfy $\lambda > \mu^4 - 1 > \mu_t^4 - 1$ we get (for large n) $\pi(\Omega_{\mathcal{F}}) \leq e^{-cn}$ for some constant $c > 0$; and we can easily modify this constant to deal with all smaller values of n. □

From Section 3.1 we know that $\mu_t < 1.6057317$ and so $\mu_t^4 - 1 < 5.648014$. Combining Theorems 2.1 and 3.5, we get the following corollary as an immediate consequence.

Corollary 3.6. *Glauber dynamics for sampling independent sets on the $n \times n$ torus $\widehat{G}$ takes time at least e^{cn} to mix, for some constant $c > 0$, when $\lambda > 5.648014$.*

3.3 Grid regions with non-periodic boundary conditions

For regions with non-periodic boundary conditions we also employ a weight-increasing map from configurations with fault lines by performing a shift and adding vertices. In this setting, however, we not only have to reconstruct the position of the fault line, but we must also encode the part of the configuration lost by the shift due to the finite boundary. In this section we give the proof of the following result.

Theorem 3.7. *Glauber dynamics for independent sets on the $n \times n$ grid G takes time at least e^{cn} to mix, for some constant $c > 0$, when $\lambda > 7.439$.*

As before, we partition the state space Ω into three sets: $\Omega_{\mathcal{F}} \cup \Omega_0 \cup \Omega_1$. The set Ω_0 contains all independent sets with an even cross, Ω_1 contains all independent sets with an odd cross, and $\Omega_{\mathcal{F}}$ is the set of independent sets containing a fault line. We 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 $\pi(\Omega_1)$.

Let $I \in \Omega_{\mathcal{F}}$ be an independent set with a vertical fault line F. The fault line partitions the vertices of G into two sets, Right(F) and Left(F), depending on the side of the fault on which they lie. Recall that a fault has zero or one alternation point, and the edges form a path (or pair of paths) in $G_\Diamond$. Define the length of a fault to be the total number of edges on this path (or paths) in $G_\Diamond$. 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 Right(F) one to the right. We will not be concerned right now if some vertices "fall off" the right side of the region G. Let $F' = \sigma(F)$ be the F shifted one to the right. We define the points that lie in Right$(F) \cap$ Left(F') to be the points that fall "in between" F and F'.

The following lemmas are modified from [21].

Lemma 3.8. *Let I be an independent set with a fault line F. Let $I' = \sigma(F, I)$ and $F' = \sigma(F)$ be defined as above.*

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'.*

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 G 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. The length of the fault is $n + 2\ell$, for some integer ℓ.

Let $G_{1,n}$ be the $1 \times n$ lattice representing the last column of G, and let J be any independent set on $G_{1,n}$. 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 $G_{1,n}$.

Lemma 3.9. *Let F be a fault in G with length $n + 2\ell$ and let δ equal the number of alternation points on F (so $\delta = 0$ or 1). Let J be an independent set on $G_{1,n}$. 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 $\phi_{F,J} : \Omega_{\mathcal{F}} \times \{0,1\}^{n+\ell} \to \Omega$ 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. For any $I \in \Omega_{F,J}$, we delete the last column (which is equal to J). Next, recalling that any fault line partitions G 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 3.8 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 3.10. *Let $G_{1,n}$ be a $1\times n$ strip, and let Ω_r be the set of independent sets on $G_{1,n}$. Then*

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

for some constant c.

Proof. Let S_i be the set of independent sets on $G_{1,n}$ 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 3.11. *Let Ω be the set of independent sets on the $n\times n$ lattice G 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 G with a fault line. Then*

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

for some polynomial $p(n)$ and constant $c' > 0$, whenever $\lambda > 7.439$.

Proof. We will make use of the injective map $\phi_{F,J} : \Omega_{F,J} \times \{0,1\}^N \to \Omega$, where $N = n + 2\ell$ is the length of the fault line. We also use our bound for the number of taxi walks from Section 3.1: for any $\mu \geq 1.6057317$, the number of walks of length N is at most μ^N for all large N.

We now have

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

where the third equality follows from Lemma 3.10. This means that we will have $\pi(\Omega_{\mathcal{F}}) \leq p(n)e^{-c'n}$, for some polynomial $p(n)$, if

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

where we may take μ to be anything at least as large as 1.6057317. Simple algebra reveals that the second condition is satisfied whenever $\lambda^2 + (2-\mu^2-\mu^4)\lambda + (1-\mu^2) > 0$. Taking $\mu = 1.6057317$, we find that both of these conditions are met when $\lambda > 7.439$. □

Finally, we show how Theorem 3.7 follows as an immediate consequence.

Proof of Theorem 3.7. We will bound the conductance by considering the cut $S = \Omega_0$. It is clear that $\pi(S) \leq 1/2$ since $\overline{S} = \Omega_{\mathcal{F}} \cup \Omega_1$ 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 3.11, it is trivial to show that $\pi(\Omega) > 1/3$, thereby establishing that the conductance is exponentially small. It follows from Theorem 2.1 that Glauber dynamics takes exponential time to converge. □

4 Phase coexistence

Here we prove the following (which implies multiple Gibbs states for all $\lambda > 5.648014$).

Theorem 4.1. *The hard-core model on $\mathbb{Z}^2$ with fugacity λ admits multiple Gibbs states for all $\lambda > \mu_t^4 - 1$, where μ_t is the connective constant of taxi walks.*

We will not review the theory of Gibbs states, contenting ourselves with saying informally that an interpretation of the existence of multiple Gibbs states is that the local behavior of a randomly chosen independent set in a box can be made to depend on a boundary condition imposed on the box, even in the limit as the size of the box grows to infinity. See e.g. [11] for a very general treatment, or [3] for a treatment specific to the hard-core model on the lattice.

Let U_n be the box $[-n, +n]^2$, and I^{e} the independent set consisting of all even vertices of $\mathbb{Z}^2$. Let $\mathcal{J}_n^{\mathrm{e}}$ be the set of independent sets that agree with I^{e} off U_n, and μ_n^{e} the distribution supported on $\mathcal{J}_n^{\mathrm{e}}$ in which each set is selected with probability proportional to $\lambda^{|I \cap U_n|}$. Define μ_n^{o} analogously (with "even" everywhere replaced by "odd"). We will exhibit an event $\mathcal{A}$ that depends only on finitely many vertices, with the property that for all large n, $\mu_n^{\mathrm{e}}(\mathcal{A}) \leq 1/3$ and $\mu_n^{\mathrm{o}}(\mathcal{A}) \geq 2/3$. This is well known (see e.g. [3]) to be enough to establish the existence of multiple Gibbs states.

The event $\mathcal{A}$ depends on a parameter $m = m(\lambda)$ whose value will be specified later. Specifically, $\mathcal{A}$ consists of all independent sets in $\mathbb{Z}^2$ whose restriction to U_m contains either an odd cross or a fault line. We will show that $\mu_n^{\mathrm{e}}(\mathcal{A}) \leq 1/3$ for all sufficiently large n; reversing the roles of odd and even throughout, the same argument gives that under μ_n^{o} the probability of U_m having either an *even* cross or a fault line is also at most 1/3, so that (by Lemma 2.2) $\mu_n^{\mathrm{odd}}(\mathcal{A}) \geq 2/3$.

Write $\mathcal{A}_n^{\mathrm{e}}$ for $\mathcal{A} \cap \mathcal{J}_n^{\mathrm{e}}$; note that for all large n we have $\mu_n^{\mathrm{e}}(\mathcal{A}) = \mu_n^{\mathrm{e}}(\mathcal{A}_n^{\mathrm{e}})$. To show $\mu_n^{\mathrm{e}}(\mathcal{A}_n^{\mathrm{e}}) \leq 1/3$ we will use the fact that $I \in \mathcal{A}_n^{\mathrm{e}}$ is in "even phase" (predominantly even-occupied) outside U_n, but because of either the odd cross or the fault line in U_m it is not in even phase close to U_m; so there must be a "contour" marking the furthest extent of the even phase inside U_n. We will modify I inside the contour via a weight-increasing map, showing that an odd cross or fault line is unlikely.

4.1 The contour and its properties

Fix $I \in \mathcal{A}_n^{\mathrm{e}}$. If I has an odd cross in U_m, we proceed as follows. Let $(I^{\mathcal{O}})^+$ be the set of odd vertices of I, together with their neighbors. Let R be the component of $(I^{\mathcal{O}})^+$ that includes a particular odd cross. Note that because I agrees with I^{e} off U_n, R does not reach the boundary of U_n (the vertices in U_n with a neighbor outside U_n). Let W be the component of the boundary of U_n in the complement of R. Finally, let C be the complement of W in U_n and let $\gamma = \gamma(I)$ be the set of edges with one end in W and one end in C. Write $\gamma_\Diamond$ for the subgraph of $G_\Diamond$ induced by γ.

Evidently γ is an edge cutset in U_n separating an interior connected region that meets U_m from an exterior connected region that includes the boundary of U_n. Also it is evident that all edges from the interior of γ to the exterior go from an unoccupied even vertex to an unoccupied odd vertex. This implies that $|\gamma|$, the number of edges in γ, is a multiple of 4; specifically four times the difference between the number of even and odd vertices in the interior of γ. Because the interior includes two points of the odd cross that are at distance at least $2m+1$ from each other in U_m, we can put a lower bound on γ that is linear in m; for example, it is certainly true that $|\gamma| \geq m$.

We now come to the heart of the so-called Peierls argument. If we modify I by shifting it by one axis-parallel unit (positively or negatively) in the interior of γ and leaving it unchanged elsewhere, then the resulting set is still independent. Moreover, we may augment the shifted independent set with any vertex in the interior whose neighbor in the direction opposite to the shift is in the exterior. This is a straightforward verification; see [5, Lemma 6] or [10, Proposition 2.12] where this is proved in essentially the same setting. Furthermore, from [5, Lemma 5] each of the four possible shift directions free up exactly $|\gamma|/4$ vertices that can be added to the modified independent set.

The property of γ that allows us to control the number of contours that can occur is that $\gamma_\Diamond$ is a cycle in $G_\Diamond$, and the removal of any edge makes it a taxi walk. That $\gamma_\Diamond$ is a cycle follows quickly from the fact that γ is a *minimal* edge cutset. For the taxi walk claim, note that there are two ways that we might have consecutive turns in the same direction in $\gamma_\Diamond$. Either there is a vertex of $\mathbb{Z}^2$ with all four of its incident edges in γ, contradicting the connectivity of the interior and exteriors of γ (note that these both have more than one vertex); or there is a 4-cycle in $\gamma_\Diamond$ sitting inside a 4-cycle in $\mathbb{Z}^2$, contradicting the minimality of γ. The appropriate orientation of the walk is ensured by the fact that all interior endvertices of edges in γ have the same parity.

We now describe the contour if I has a fault line in U_m. If there happens to be an odd occupied vertex in U_m then we construct γ as before, starting with some arbitrary component of $(I^{\mathcal{O}})^+$ that meets U_m in place of the component of an odd cross. If the resulting γ has a fault line in its interior, then γ and its associated $\gamma_\Diamond$ satisfy all the previously established properties immediately.

Otherwise, choose a fault line. Whether it has zero or one alternation points, we can find a path $P = u_1 u_1 \ldots u_k$ in $\mathbb{Z}^2$ with k linear in m, u_1 and u_k both odd, no two consecutive edges parallel, and with the midpoints of the edges of the path inducing an alternation-free sub-path of the chosen fault line (essentially we are just taking a long piece of the fault line, on an appropriately chosen side of the alternation point, if there is one). This sub-path F_1 is a taxi walk. Next, we find a second path in $G_\Diamond$, disjoint from F_1, that always bisects completely unoccupied edges, and that taken together with F_1 completely encloses P. If there are no occupied odd vertices adjacent to even vertices of P, such a path is easy to find: we can shift F_1 one unit in an appropriate direction, and close off with an additional edge at each end. If there are some odd occupied vertices adjacent to some even vertices of P, then this translate of F_1 has to be looped around the corresponding components of $(I^{\mathcal{O}})^+$. Such a looping is possible because $(I^{\mathcal{O}})^+$ does not reach the boundary of U_n, nor does it enclose the fault lines (if it did, we would be in the case of the previous paragraph).

This second path we have constructed may not be a taxi walk; however, following the proof of Lemma 3.1, we see that a *minimal* path F_2 satisfying the conditions of our constructed path is indeed a taxi walk. We take the concatenation of F_1 and F_2 to be $\gamma_\Diamond$ in this case, and take γ to be the set of edges that are bisected by vertices of $\gamma_\Diamond$. The contours in this case satisfy all the properties of those in the previous case. The evident properties remain evident, and those from [5] and [10] can be derived in this case using the methods of those references. The one difference is that now $\gamma_\Diamond$ may not be a closed taxi walk; but at worst it is the concatenation of two taxi walks, both of length linear in m (and certainly it can be arranged that each has length at least $m/2$).

4.2 The Peierls argument for phase coexistence

For $J \in \mathcal{J}_n^{\mathrm{e}}$ set $w(J) = \lambda^{|J \cap U_n|}$; we must show $w(\mathcal{A}_n^{\mathrm{e}})/w(\mathcal{A}_n^{\mathrm{e}}) \leq 1/3$. For $I \in \mathcal{A}_n^{\mathrm{e}}$, let $\varphi(I)$ be the set of independent sets obtained from I by shifting in the interior parallel to $(1,0)$ and adding all subsets of the $|\gamma|/4$ vertices by which the shifted independent set can be augmented. For $J \in \varphi(I)$, let S denote the set of added vertices. Define a bipartite graph on partite sets $\mathcal{A}_n^{\mathrm{e}}$ and $\mathcal{J}_n^{\mathrm{e}}$ by joining $I \in \mathcal{A}_n^{\mathrm{e}}$ to $J \in \mathcal{J}_n^{\mathrm{e}}$ if $J \in \varphi(I)$. Give edge IJ weight $w(I)\lambda^{|S|} = w(J)$ (where S is the set of vertices added to I to obtain J).

The sum of the weights of edges out of those $I \in \mathcal{A}_n^{\mathrm{e}}$ with $|\gamma(I)| = 4\ell$ is $(1+\lambda)^\ell$ times the sum of the weights of those I. For each $J \in \mathcal{J}_n^{\mathrm{e}}$, the sum of the weights of edges into J from this set of

I's is $w(J)$ times the degree of J to the set. If $f(\ell)$ is a uniform upper bound on this degree, then

$$\frac{w(\mathcal{A}_n^{\mathrm{e}})}{w(\mathcal{J}_n^{\mathrm{e}})} \leq \sum_{\ell \geq m/4} \frac{f(\ell)}{(1+\lambda)^\ell}. \tag{2}$$

The lower bound on ℓ here is crucial. The standard Peierls argument takes $\mathcal{A}$ to be the event that a fixed vertex is occupied, and the analysis of probabilities associated with this event requires dealing with short contours, leading to much weaker bounds than we are able to obtain.

To control $f(\ell)$, observe that for each $J \in \mathcal{J}_n^{\mathrm{e}}$ and contour γ of length 4ℓ there is at most one I with $\gamma(I) = \gamma$ such that $J \in \varphi(I)$ (I can be uniquely reconstructed from J and γ, since the set S of added vertices can easily be identified; cf [10, Section 2.5]). It follows that we may bound $f(\ell)$ by the number of contours of length 4ℓ that have a vertex of U_m in their interiors.

Let μ be any upper bound on μ_t. By the properties of contours we have established, up to translation this number is at most the maximum of $\mu^{4\ell}$ and $\sum_{j+k=4\ell:\ j,k\geq m/2} \mu^j\mu^k = 4\ell\mu^{4\ell}$ (at least for all large m). The restriction of $G_\Diamond$ to U_m has at most $4(2m+1)^2 \leq 17m^2$ edges, so there are at most this many translates of any particular contour that can have a vertex of U_m in its interior. If follows that we may bound $f(\ell)$ by $68m^2\ell\mu^{4\ell}$ and so the sum in (2) by $\sum_{\ell\geq m} 68m^2\ell(\mu^4/(1+\lambda))^\ell$. For any fixed $\lambda > \mu^4 - 1$, there is an m large enough so that this sum is at most $1/3$; we take any such m to be $m(\lambda)$, completing the proof of phase coexistence.

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] V. Beffara and H. Duminil-Copin. The self-dual point of the two-dimensional random cluster model is critical above $q = 1$. To appear in *Probability Theory and Related Fields.* Preprint available from the arXiv at: http://arxiv.org/abs/1006.5073

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

[4] C. Borgs. Polymers, contours and their applications in discrete mathematics. Monograph in preparation, 2004.

[5] 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.

[6] G. R. Brightwell, O. Häggström, and P. Winkler. Nonmonotonic behavior in hard-core and Widom-Rowlinson models. *J. of Statistical Physics*, **94** (3): 415–435, 1999.

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

[8] 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, Birhauser, Boston, 1985.

[9] D. Galvin. Sampling independent sets on the discrete torus. *Random Structures Algorithms* **33**, 356–376, 2008.

[10] D. Galvin and J. Kahn. On phase transitions in the hard-core model on Z^d. *Combin. Probab. Comput.* **13** 137–164, 2004.

[11] H.-O. Georgii, Gibbs Measures and Phase Transitions. de Gruyter, Berlin, 1988.

[12] R. Glauber. Time dependent statistics of the Ising model. *Journal of Mathematical Physics* **4**: 294-307, 1963.

[13] J.M. Hammersley and D.J.A. Welsh. Further results on the rate of convergence to the connective constant of the hyper cubic lattice. *Quarterly Journal of Mathematics* **2**, 108–110, 1962.

[14] 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.

[15] E. Lubetzky and A. Sly. Critical Ising on the square lattice mixes in polynomial time. To appear in *Communications in Mathematical Physics.*

[16] M. Luby and E. Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets. *Random Structures and Algorithms*, 229–241, 1999.

[17] N. Madras and G. Slade. The Self-Avoiding Walk. Birkhäuser, Boston, 1993.

[18] 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.

[19] L. Onsager. Crystal statistics. I. A two-dimensional model with an order-disorder transition. *Physics Review Letters*, **65** (3–4): 117–149, 1944.

[20] R. Peled and W. Samotij. Odd cutsets and the hard-core model on $\mathbb{Z}^d$. Preprint available from the arXiv at: http://arxiv.org/abs/1106.3594.

[21] D. Randall. Slow mixing of Glauber dynamics via topological obstructions. *Proc 17th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 870–879, 2006.

[22] R. Restrepo, J. Shin, P. Tetali, E. Vigoda and L. Yang. Improved mixing condition on the grid for counting and sampling independent sets. *Foundations of Computer Science (FOCS)*, 140–149, 2011.

[23] A. Sinclair. Counting and generating combinatorial structures: a Markov chain approach. Birkhäuser, Boston, 1993.

[24] A.J. Sinclair and M.R. Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains *Information and Computation* **82:**93–133, 1989.

[25] A. Sly. Computational Transition at the Uniqueness Threshold. *Proc. of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 287-296, 2010.

[26] J.M. Steele. Probability Theory and Combinatorial Optimization. SIAM, 1997.

[27] D. Weitz. Counting independent sets up to the tree threshold. *Proc. 38th ACM Symposium on the Theory of Computing (STOC)*, 140–149, 2006.