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

Antonio Blanca[1]⋆, David Galvin[2]⋆⋆, Dana Randall[3]⋆⋆⋆, and Prasad Tetali[4]†

[1] Computer Science Division, U.C. Berkeley, Berkeley, CA 94720.
[2] Department of Mathematics, University of Notre Dame, Notre Dame, IN 46656.
[3] School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332-0765.
[4] School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332-0280.

Abstract. The hard-core model has attracted much attention across several disciplines, representing lattice gases in statistical physics and independent sets in discrete mathematics and computer science. On finite graphs, we are given a parameter λ, and an independent set I arises with probability proportional to $\lambda^{|I|}$. On infinite graphs a 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, i.e., existence of multiple Gibbs states. On finite graphs we are interested in determining the mixing time of local Markov chains.

On $\mathbb{Z}^2$ it is conjectured that these problems are related and that both undergo a phase transition at some critical point $\lambda_c \approx 3.79$ [2]. For the question of phase coexistence much of the work to date has focused on the regime of uniqueness, with the best result to date being recent work of Vera et al. [29] showing that there is a unique Gibbs state for all $\lambda < 2.48$. Here we give the first explicit result in the other direction, showing that there are multiple Gibbs states for all $\lambda > 5.3646$. Our proof begins along the lines of the standard Peierls argument, but we add two significant innovations. First, building on the idea of fault lines introduced by Randall [23], we construct an event that distinguishes two boundary conditions and yet always has long contours associated with it, obviating the need to accurately enumerate short contours. Second, we obtain vastly improved bounds on the number of contours by relating them to a new class of self-avoiding walks on an oriented version of $\mathbb{Z}^2$.

The best result for rapid mixing of local Markov chains on boxes of $\mathbb{Z}^2$ is also when $\lambda < 2.48$ [29]. Here we extend our characterization of fault lines to show that local Markov chains will mix slowly when $\lambda > 5.3646$ on lattice regions with periodic (toroidal) boundary conditions and when $\lambda > 7.1031$ with non-periodic (free) boundary conditions. The arguments here rely on a careful analysis that relates contours to taxi walks and represent a sevenfold improvement to the previously best known values of λ [23].

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 to be non-adjacent. Viewing a lattice such as $\mathbb{Z}^d$ as a graph, allowed configurations of molecules naturally correspond to independent sets in the graph.

Given a finite graph G, let Ω be the set of independent sets of G. Given a (fixed) *activity* (or *fugacity*) $\lambda \in \mathbb{R}^+$, 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 Ω, assuming G is finite, as $\mu(I) = w(I)/Z$, where the normalizing constant $Z = Z(G, \lambda) = \sum_{J\in\Omega} w(J)$ 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 that as a parameter of the system is varied – such as the inverse temperature β for the Ising model or the activity λ for the hard-core model – the system undergoes a phase transition at a critical point.

For the classical Ising model, Onsager, in seminal work [20], 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 [3], settling a more than half-a-century old open

⋆ ablanca@cs.berkeley.edu. Supported by a UC Berkeley Chancellor's Fellowship for Graduate Study and an NSF Graduate Research Fellowship.
⋆⋆ dgalvin1@nd.edu. Supported in part by the Simons Foundation and by NSA grant number H98230-13-1-0248.
⋆⋆⋆ randall@cc.gatech.edu. Supported in part by NSF grant CCF-1219202.
† tetali@math.gatech.edu. Supported in part by NSF grant DMS-1101447.

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. [7]). Regardless, a non-rigorous prediction from the statistical physics literature [2] suggests $\lambda_c \approx 3.796$ for $\mathbb{Z}^2$.

Thus, from a statistical physics or probability viewpoint, 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 [30] and Sly [27] have recently 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 it for the hard-core model above a certain critical threshold also turns out to be hard.

Starting with Dobrushin [8] in 1968, physicists have been developing techniques to characterize the regimes on either side of λ_c for the hard-core model. Most attention has focused on establishing ever larger values of λ below which there is always uniqueness of phase. The problem has proved to be a fruitful one for the blending of ideas from physics, discrete probability and theoretical computer science, with improvements to our understanding of the problem having been made successively by Radulescu and Styer [22], van den Berg and Steif [4] and Weitz [30], among others. The state of the art is work by Vera et al. [29], expanding on ideas from Weitz and Restrepo et al. [24], which establishes uniqueness for all $\lambda < 2.48$.

Much less is known about the regime of phase coexistence. Dobrushin [8] established phase coexistence for all $\lambda > C$, but did not explicitly calculate C. Borgs had estimated that $C = 80$ was the *theoretical* lower limit of Dobrushin's argument [5], but a recent computation by the second author suggests that the *actual* consequence of the argument is more like $C \approx 300$.

We are now prepared to state our first main result. Using insights from physics and combinatorics, we establish the first explicit upper bound on λ_c for the hard-core model on $\mathbb{Z}^2$.

Theorem 1. *For all $\lambda > 5.3646$, the hard-core model on $\mathbb{Z}^2$ with activity λ admits multiple Gibbs states.*

From a computational standpoint, there are two natural questions to ask concerning the hard-core model on a finite graph. Can the partition function be approximated, and how easy is it to sample from a given Gibbs distribution? For both questions, a natural and powerful method is offered by Markov chain algorithms – carefully constructed random walks on the space of independent sets of a graph whose equilibrium distributions are the desired Gibbs distributions. One of the most commonly studied families of Markov chains are the *local-update* chains, such as Glauber dynamics, that change the state of a bounded number of vertices at each step (in particular, Glauber dynamics changes the state of a single vertex at each step).

The efficiency of the Markov chain method relies on the underlying chain being rapidly mixing; that is, it must fairly quickly reach a distribution that is close to stationary. For many problems, local chains seem to mix rapidly below some critical point, while mixing slowly above that point. Most notably for the Ising model on $\mathbb{Z}^2$, simple local Markov chains are rapidly mixing (in fact, with optimal rate) for $\beta < \beta_c(\mathbb{Z}^2)$ and slowly mixing for $\beta > \beta_c(\mathbb{Z}^2)$; various mathematical physicists (including Aizenmann, Holley, Martinelli, Olivieri, Schonmann, Stroock and Zegarlinski) have contributed to this work. Recently the Ising picture was completed by Lubezky and Sly [16], who showed polynomial mixing at $\beta = \beta_c(\mathbb{Z}^2)$.

Once again, the known bounds are less sharp for the hard-core model. Luby and Vigoda [17] showed that Glauber dynamics on independent sets is fast when $\lambda \leq 1$ on the 2-dimensional lattice and torus. Weitz [30] reduced the analysis on the grid to the tree, thus establishing that in this same setting Glauber dynamics is fast up to the critical point for the 4-regular tree, in effect for $\lambda < 1.6875$. Again, the best result to date is due to Vera et al. [29]. Using now standard machinery (establishing so-called *strong spatial mixing*), they proved that the natural Glauber dynamics on the space of hard-core configurations on boxes in $\mathbb{Z}^2$ is rapidly mixing for all $\lambda < 2.48$. These results also lead to efficient deterministic algorithms for approximating the hard-core partition function on (finite regions of) $\mathbb{Z}^2$.

As with phase-coexistence, it is believed that there is a critical value λ_c^{mix} across which the Glauber dynamics for sampling from hard-core configurations on a n by n box in $\mathbb{Z}^2$ flips from mixing in time polynomial in n, to exponential in n, and that it coincides with λ_c. Pinning down such a precise transition seems beyond reach at the moment, but we can try to establish a short range of values into which λ_c^{mix} (if it exists) must fall, by finding ever smaller bounds for λ above which the mixing time is exponential.

Borgs et al. [6] 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 establishing a finite constant above which mixing is slow on $\mathbb{Z}^2$. The first effective bound was provided by Randall [23], who showed slow mixing for $\lambda > 50.526$ on boxes with periodical boundary conditions, and for $\lambda > 56.812$ on boxes with free boundary* (but did not address the question of phase coexistence).

Our second main result establishes slow mixing of Glauber dynamics on boxes in $\mathbb{Z}^2$ for values of λ that are an order of magnitude lower than the previously best known bounds.

Theorem 2. *For all $\lambda > 5.3646$, the mixing time for Glauber dynamics for the hard-core model on n by n boxes in $\mathbb{Z}^2$ with periodic boundary conditions and with activity λ is exponential in n. For free boundary conditions, we have the same result with $\lambda > 7.1031$.*

The proofs of Theorems 1 and 2 utilize combinatorial, computational and physical insights. The standard approach to showing multiple Gibbs distributions in a statistical physics model is to consider the limiting distributions corresponding to two different boundary conditions on boxes in the lattice centered at the origin, and find a statistic that separates these two limits. For the hard-core model, it suffices (see [4]) to compare the *even boundary condition* – all vertices on the boundary of a box at an even distance from the origin are occupied – and its counterpart the odd boundary condition, and the distinguishing statistic is typically the occupation of the origin. Under odd boundary condition the origin should be unlikely to be occupied, since independent sets with odd boundary and (even) origin occupied must have a *contour* – a two-layer thick unoccupied loop of vertices separating an inner region around the origin that is in "even phase" from an outer region near the boundary that is in "odd phase". For large enough λ, such an unoccupied layer is costly, and so such configurations are unlikely. This is essentially the *Peierls argument* for phase coexistence, and was the approach taken by Dobrushin [8].

As we will see presently, the effectiveness of the Peierls argument is driven by the number of contours of each possible length – better upper bounds on the number of contours translate directly to better upper bounds on λ_c. Previous (unpublished) work on phase coexistence in the hard-core model on $\mathbb{Z}^2$ had viewed contours as simple polygons in $\mathbb{Z}^2$, which are closely related to the very well studied family of self-avoiding walks. While this is essentially the best possible point of view when applying the Peierls argument on the Ising model, it is far from optimal for the hard-core model. One of the two major breakthroughs of the present paper is the discovery that hard-core contours, if appropriately defined, can be viewed as simple polygons in the oriented *Manhattan lattice* (orient edges of $\mathbb{Z}^2$ that are parallel to the x-axis (resp. y-axis) positively if their y-coordinate (resp. x-coordinate) is even, and negatively otherwise), with the additional constraint that contours cannot make two consecutive turns. The number of such polygons can be understood by analyzing a new class of self-avoiding walks, that we refer to as *taxi walks*. The number of taxi walks turns out to be significantly smaller than the number of ordinary self-avoiding walks, leading to much better bounds on λ_c than could possibly have been obtained previously.

The number $\widetilde{c}_n$ of taxi walks of length n is asymptotically controlled by a single number $\mu_t > 0$, the *taxi walk connective constant*, in the sense that $(\widetilde{c}_n)^{1/n} \sim \mu_t$ as $n \to \infty$. Adapting methods of Alm [1] we obtain good estimates on μ_t, allowing us to understand $\widetilde{c}_n$ for large n. It is difficult to control $\widetilde{c}_n$ for small n, however, presenting a major stumbling block to the effectiveness of the Peierls argument. Using the statistic "occupation of origin" to distinguish the two boundary conditions, one inevitably has to control $\widetilde{c}_n$ for both small and large n. The lack of precise information about the number of short contours leads to discrepancies between theoretical lower limits and actual bounds, such as that between $C = 80$ (theoretical best possible) and $C \approx 300$ for phase coexistence on $\mathbb{Z}^2$, discussed earlier.

The second breakthrough of the present paper is the idea of using an event to distinguish the two boundary conditions that has the property that every independent set in the event has associated with it a *long* contour. This allows us to focus exclusively on the asymptotic growth rate of contours/taxi walks, and obviates the need for an analysis of short contours. The immediate result of this breakthrough is that the actual limits of our arguments agree exactly with their theoretical counterparts. The distinguishing event we use extends the idea of *fault lines*, which we discuss in more detail below in the context of slow mixing.

* Note that stronger bounds were reported in [23], but these were due to a minor error. Specifically, a missing factor of 2 was carried through all of the computations in that reference; see http://www.math.gatech.edu/∼randall/ind-fix.pdf for the corrected version or Section 4.1 here for more details. In any case, the bounds we report here improve on the original claims as well.

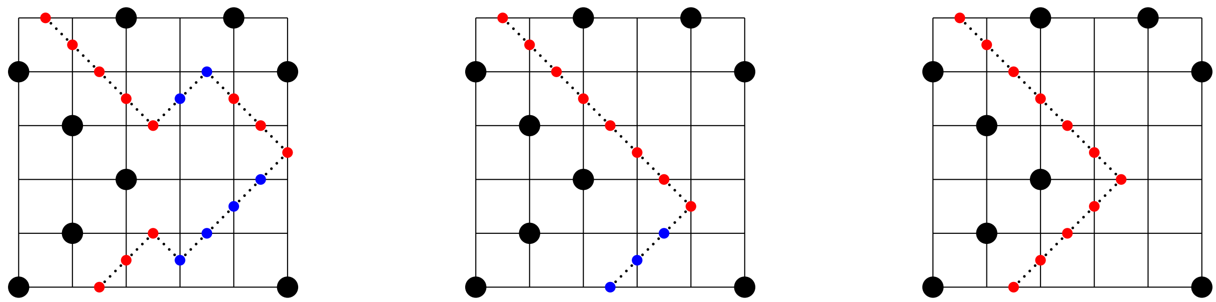

(a) (b) (c)

Fig. 1: Independent sets with (a) a spanning path with four alternation points in $G_\Diamond$, (b) a fault line with one alternation point in $G_\Diamond$, and (c) a fault line with no alternation points in $G_\Diamond$.

The traditional argument for slow mixing 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: 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 by at most 1 in each step, the Markov chain has a bottleneck leading to *torpid* (slow) mixing.

Our work builds on a novel idea from [23] in which the notion of *fault lines* was introduced to establish slow mixing for Glauber dynamics on hard-core configurations for moderately large λ, still improving upon what was best known at that time. Randall [23] gave an improvement by realizing that the state space could be partitioned according to certain *topological obstructions* in configurations, rather than the relative numbers of odd or even vertices. This approach gives better bounds on λ, and 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 without 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, where G^2 connects vertices at distance 2 in 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, allowing for a similar argument.

We improve the argument by refining our consideration of fault lines, which had previously been characterized as (rotated) self-avoiding walks in $\mathbb{Z}^2$. Here we observe that, suitably modified, they are in fact taxi walks and so the machinery developed for phase coexistence can be brought to bear in the mixing context.

We believe that both breakthroughs of the present paper have more general applicability. A natural next step is to study the hard-core model on $\mathbb{Z}^d$ for $d \geq 3$. Establishing bounds on the critical activity for $\mathbb{Z}^3$ is an immediate challenge, as is pinning down how $\lambda_c(d)$ changes with d. The best upper bounds are $\widetilde{O}(d^{-1/4})$ for slow mixing [10] and $\widetilde{O}(d^{-1/3})$ [21] for phase coexistence; the lower bounds for both are $\Omega(d^{-1})$.

The paper is laid out as follows. Section 2 provides the necessary combinatorial background, including the notion of fault lines and crosses, the connection between fault lines and taxi walks, and the characterization of independent sets on finite regions of $\mathbb{Z}^2$. In Section 3 we explain how fault lines and taxi walks can be used to prove Theorem 1 (phase coexistence). In Section 4 we prove Theorem 2 (slow mixing). Finally, in Section 5, we provide the detailed analysis of the number of taxi walks. In this section we also give lower bounds on the number of taxi walks, and show that the present approach cannot give phase coexistence for any λ below 4.33.

2 Combinatorial background

Here we introduce the notions of crosses, fault lines and taxi walks, which are the key ingredients of the proofs of both Theorem 1 and Theorem 2.

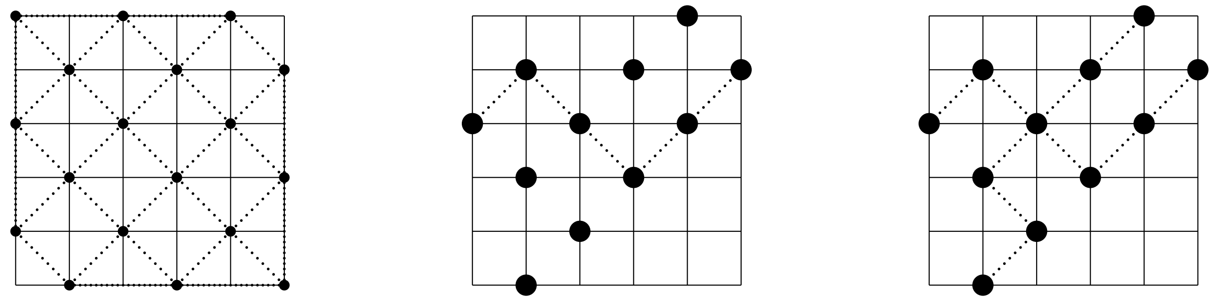

(a) (b) (c)

Fig. 2: The figure shows: (a) G_1, (b) an independent set with an odd bridge in G_1, and (c) two odd bridges forming an odd cross in G_1.

2.1 Crosses and fault lines

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 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 $V_b \subseteq V$ contains all 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 at Hamming distance 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 distinguishing events in our proof of phase coexistence and the partition of the state space for our proofs of slow mixing.

Given an independent set $I \in \Omega$, we say that a simple 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 on p according to whether the corresponding edge in G has an odd or an even vertex to its left as we traverse the path. Specifically, suppose vertex $v \in V_\Diamond$ on the path p bisects an edge $e_v \in E$. Color v *blue* if the odd vertex in e_v is to the left when the path crosses v, and *red* otherwise (note that each $v \in E$ has one odd and one even endvertex). When the color of the vertices along the path changes, we have an *alternation point* (see Fig. 1, parts (a) and (b)). It was shown in [23] that if an independent set has a spanning path, then it must also have one with zero or one alternation points. We call such a path a *fault line*, and let $\Omega_{\mathcal{F}}$ be the set of independent sets in Ω with at least one fault line (see Fig. 1, parts (b) and (c)).

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. See Fig. 2.

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.

We can now partition the state space Ω into three sets, with one separating the other two; this partition is critical to the proofs of both Theorem 1 and Theorem 2. The following lemma was proven in [23].

Lemma 1. *The set of independent sets on G can be partitioned into sets $\Omega_{\mathcal{F}}$, Ω_0 and Ω_1, consisting of configurations with a fault line, an even cross or an odd cross. If $I \in \Omega_0$ and $I' \in \Omega_1$ then $|I \triangle I'| > 1$.*

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 $v_1 = u_1 \pm 1 (\mathrm{mod}\, n)$ and $v_2 = u_2$ or $v_2 = u_2 \pm 1 (\mathrm{mod}\, n)$ 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 symmetric difference of size 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.

Given $I \in \widehat{\Omega}$, we say that I has a *fault* $F = (F_1, F_2)$ if there are a *pair* of vertex-disjoint non-contractible cycles F_1, F_2 in $\widehat{G}_\Diamond$ whose vertices correspond to edges in $\widehat{G}$ where both endpoints are unoccupied, and such that the vertices on each cycle are all red or all blue (i.e., the endpoints in $\widehat{G}$ to one side of either cycle 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 lemma (from [23]) utilizes faults to partition $\widehat{\Omega}$.

Lemma 2. *The set of independent sets on $\widehat{G}$ can be partitioned into sets $\widehat{\Omega}_{\mathcal{F}}, \widehat{\Omega}_0, \widehat{\Omega}_1$, consisting of configurations with a fault, an even cross or an odd cross. If $I \in \widehat{\Omega}_0$ and $I' \in \widehat{\Omega}_1$ then $|I \triangle I'| > 1$.*

2.2 Taxi walks

The strategy for the proofs of phase coexistence and 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 [23] 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 overcount because, as we shall see, this includes all spanning paths with an arbitrary number of alternation points. Instead, we can get much better bounds on the number of fault lines by only counting a subset of 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 (with respect to this edge orientation). 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 in many well-known metropolises.

We now define taxi walks. Let $\mathbb{Z}^2$ be an orientation of $\mathbb{Z}^2$ in which an edge parallel to the x-axis (resp. y-axis) is oriented in the positive x-direction if its y-coordinate is even (resp. oriented in the positive y-direction if its x-coordinate is even), and is oriented in the negative direction otherwise (note that this agrees with the orientation placed on $G_\Diamond$ above). It is common to refer to $\mathbb{Z}^2$ as the *Manhattan lattice*: streets are horizontal, with even streets oriented East and odd streets oriented West, and avenues are vertical, with even avenues oriented North and odd avenues oriented South.

Definition 1. *A* taxi walk *is an oriented walk in $\mathbb{Z}^2$ that begins at the origin, never revisits a vertex, and never takes two left or two right turns in a row.*

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

Lemma 3. *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 a 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. $\square$

The same argument shows that if F is a fault line with an alternation point, then there is a fault line that is the concatenation of two taxi walks (or the reversals of taxi walks). Lemma 3, and the extension just mentioned, are key ingredients in our proofs of both phase coexistence and slow mixing. They allow us to assume, as we do throughout, that all fault lines we work with are essentially taxi walks. For phase coexistence we will also need to understand the connection between *Peierls contours and taxi walks*.

Given an independent set I in $\mathbb{Z}^2$, let $(I^{\mathcal{O}})^+$ be the set of odd vertices in I together with their neighbors. Let R be any finite component of $(I^{\mathcal{O}})^+$, and let W be the unique infinite component of $\mathbb{Z}^2 \setminus R$. Let C be the complement of W (the process of going from R to C is essentially one of "filling in holes" in R). Finally, let γ be the set of edges with one end in W and one in C, and write $\gamma_\Diamond$ for the subgraph of $G_\Diamond$ induced by γ.

Lemma 4. *In $G_\Diamond$, $\gamma_\Diamond$ is a directed cycle that does not take two consecutive turns. Consequently, if an edge is removed from γ, the resulting path in $\mathbb{Z}^2$ (suitably translated and rotated) is a taxi walk.*

Proof: Because γ separates W from its complement, $\gamma_\Diamond$ must include a cycle surrounding a vertex of W, and since γ is in fact a *minimal* edge cutset (W and C are both connected), $\gamma_\Diamond$ must consist of just this cycle. To see both that $\gamma_\Diamond$ is correctly (i.e. uniformly) oriented in $G_\Diamond$, and that it does not take two consecutive turns, note that if either of these conditions were violated then we must have one of the following: a vertex of W (or C) all of whose neighbors are in C (or W), or a unit square in $\mathbb{Z}^2$ with both even vertices in C and both odd vertices in W (an easy case analysis). All of these situations lead to a 4-cycle in $\gamma_\Diamond$, a contradiction since $\gamma_\Diamond$ is a cycle whose length is evidently greater than 4 (in fact it must have length at least 12). $\square$

A critical step in our arguments will be bounding the number of taxi walks. We start by recalling facts about standard self-avoiding walks (which have been studied extensively, although many basic questions remain; see, e.g., [18]). On $\mathbb{Z}^2$, the number c_n of walks of length n grows exponentially with n as $2^n \leq c_n \leq 4 \times 3^{n-1}$, since there are 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. Hammersley and Welsh [14] showed that $c_n = \mu^n \exp(O(\sqrt{n}))$, where $\mu \approx 2.64$ is known as the *connective constant*. It is believed that $\exp(O(\sqrt{n}))$ here can be replaced by $\Theta(n^{11/32})$ (this is supported by considerable experimental and heuristical evidence).

Letting $\widetilde{c}_n$ be the number of taxi walks of length n, we quickly get $2^{n/2} < \widetilde{c}_n < 4 \times 3^{n-1}$. The upper bound here uses $\widetilde{c}_n < c_n$, and for the lower bound we observe that if we take two steps at a time in one direction we can always go East or North. With little extra work, we can make a significant improvement:

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

Proof: At each vertex there are exactly two outgoing edges in $\mathbb{Z}^2$. 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. 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. In this encoding, the condition forbidding consecutive turns forces s to avoid having two 1's 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. $\square$

General considerations (discussed in Section 5) imply that there is a *taxi walk connective constant* $\mu_t > 0$ such that $\widetilde{c}_n = f_t(n)\mu_t^n$, where $f_t(n)$ grows sub-exponentially. A consequence of this is that

$$\text{if } \mu > \mu_t \text{ then for all large } n,\ \widetilde{c}_n < \mu^n. \tag{1}$$

Lemma 5 implies $\mu_t \leq \phi$, and as we shall see from Theorems 4 and 5 below, this is enough to obtain phase coexistence, and slow mixing on the torus, for all $\lambda > \phi^4 - 1 \approx 5.85$. To obtain the stronger Theorems 1 and 2 we use more sophisticated tools, described in Section 5, to improve our bounds on $\widetilde{c}_n$.

Theorem 3. *We have $1.5196 < \mu_t < 1.5884$ and $4.3332 < \mu_t^4 - 1 < 5.3646$.*

3 Proof of Theorem 1 (phase coexistence on $\mathbb{Z}^2$ for large λ)

We work towards the following stronger statement that implies Theorem 1 via Theorem 3.

Theorem 4. *The hard-core model on $\mathbb{Z}^2$ with activity 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, but just say 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, even in the limit as the size of the box grows to infinity. See e.g. [12] for a general treatment, or [4] 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 $\mathcal{A}$ 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. [4]) to be enough to establish 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 1) $\mu_n^{\mathrm{o}}(\mathcal{A}) \geq 2/3$.

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.

3.1 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 (using the notation from the discussion preceding Lemma 4). 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, and so as in the discussion preceding Lemma 4, we can associate to R a cutset γ separating it from the boundary of U_n.

Notice that 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, with all edges from the interior of γ to the exterior that all 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 have a lower bound on γ that is linear in m; in particular, clearly $|\gamma| \geq m$. Note also that by Lemma 4, $\gamma_\Diamond$ must be a closed taxi walk. (See Fig. 3, part (a).)

We now come to the heart of the 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, and we may augment it 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 [6, Lemma 6] or [11, Proposition 2.12] where this is proved in essentially the same setting. Furthermore, from [6, Lemma 5] each of the four possible shift directions free up exactly $|\gamma|/4$ vertices that can be added to the modified independent set.

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, which we can assume by Lemma 3 is a taxi walk or the concatenation of two taxi walks. Whether it has zero or one alternation points, we can find a path $P = u_1 u_2 \dots u_k$ in $\mathbb{Z}^2$ with k linear in m, with u_1 and u_k both odd, with 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 (see Fig. 3, part (b)). 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 line (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, 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 standard strategies outlined in [6] and [11] can easily be used to derive the properties in this case. 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 $m/2$).

3.2 The Peierls argument

For $J \in \mathcal{J}_n^{\mathrm{e}}$ set $w(J) = \lambda^{|J \cap U_n|}$. Our aim is to show that $w(\mathcal{A}_n^{\mathrm{e}})/w(\mathcal{J}_n^{\mathrm{e}}) \leq 1/3$. 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 $|\gamma|/4$

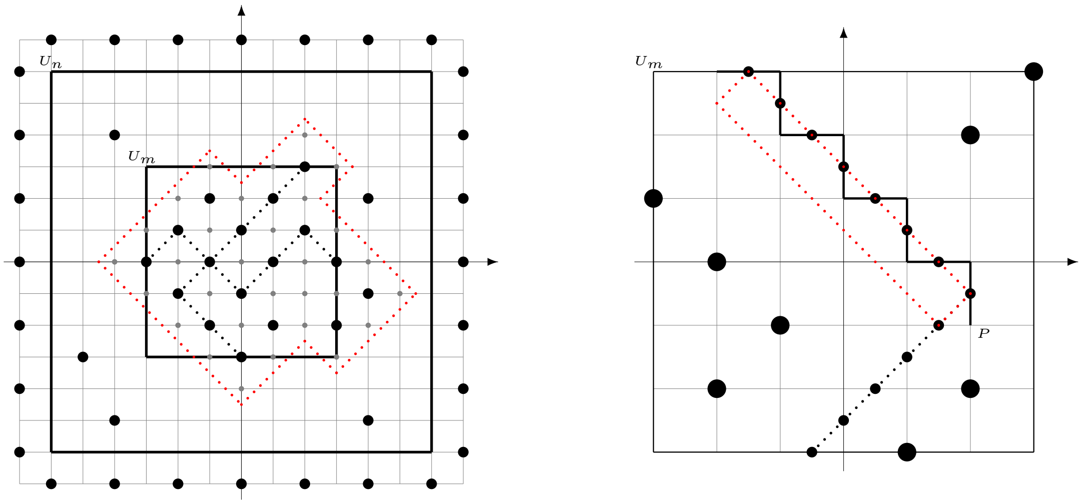

(a) (b)

Fig. 3: Independent sets in $\mathcal{A}_n^{\mathrm{e}}$ with (a) an odd cross in U_m with the corresponding C and $\gamma_\Diamond$ and (b) a fault line in U_m with the corresponding P and $\gamma_\Diamond$.

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

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

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 reconstructed from J and γ, since the set S of added vertices can easily be identified; cf. [11, Section 2.5]). It follows that we may bound $f(\ell)$ by the number of contours of length 4ℓ with a vertex of U_m in their interiors.

Fix $\mu > \mu_t$. By the properties of contours we have established, up to translations of contours this number is at most the maximum of $\mu^{4\ell}$ and $\sum_{j+k=4\ell:\ j,k \ge m/2} \mu^j \mu^k = 4\ell\mu^{4\ell}$ (for all large m, here using (1). The restriction of $G_\Diamond$ to U_m has at most $4(2m+1)^2 \le 17m^2$ edges, so there are at most this many translates of a contour that can have a vertex of U_m in its interior. We may bound $f(\ell)$ by $68m^2\ell\mu^{4\ell}$ and so the sum in (2) by $\sum_{\ell \ge m} 68m^2\ell(\mu^4/(1+\lambda))^\ell$. 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.

4 Proof of Theorem 2 (slow mixing of Glauber dynamics on finite regions)

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 $\pi(I) = \lambda^{|I|}/Z$, where $Z = \sum_{I' \in \Omega} \lambda^{|I'|}$. *Glauber dynamics* is a local Markov chain that connects two independent sets if they have symmetric difference of size one. The Metropolis probabilities [19] 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 [26], is a good measure of a chain's mixing rate. Let

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

where $\pi(S) = \sum_{x\in S}\pi(x)$. From [26] we know that $\frac{\Phi^2}{2} \le Gap(P) \le 2\Phi$, where $Gap(P)$ is the spectral gap of the transition matrix. The spectral gap is well-known to be a measure of the mixing rate of a Markov chain (see, e.g., [25]), so a partition of the state space witnessing exponentially small conductance is sufficient to show slow mixing. The machinery of Section 2.1 provides such a partition.

4.1 Glauber dynamics on the 2-d torus

Let n be even, and let $\widehat{G} = \{0,\dots,n-1\}\times\{0,\dots,n-1\}$ be the $n\times n$ lattice region with toroidal boundary conditions. Let $\widehat{\Omega}$ be the set of independent sets on $\widehat{G}$, and $\widehat{\pi}$ the Gibbs distribution. Lemma 2 shows that $\widehat{\Omega}$ may be partitioned into $\widehat{\Omega}_{\mathcal{F}}$ (independent sets with a fault), $\widehat{\Omega}_0$ (independent sets with an even cross) and $\widehat{\Omega}_1$ (independent sets with an odd cross), and that furthermore $\widehat{\Omega}_0$ and $\widehat{\Omega}_1$ are not directly connected by moves in the chain. It remains to show that $\widehat{\pi}(\widehat{\Omega}_{\mathcal{F}})$ is exponentially smaller than both $\widehat{\pi}(\widehat{\Omega}_0)$ and $\widehat{\pi}(\widehat{\Omega}_1)$. (Clearly $\widehat{\pi}(\widehat{\Omega}_0) = \widehat{\pi}(\widehat{\Omega}_1)$ by symmetry.) Notice that on the torus we may assume that fault lines have no alternation points; since they start and end at the same place, the number of alternation points must be even.

For an independent set $I\in\widehat{\Omega}_{\mathcal{F}}$ with fault $F=(F_1,F_2)$, partition 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 F to be the number of edges (in $G_\Diamond$) on F_1 and F_2. Note that if F_1 has no alternation points then it has length $N = 2n+2\ell$ for some positive integer ℓ.**

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 $F' := (F_1', F_2')$. 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 its *first* fault. The following results are modified from [23] and rely on the new characterization of faults as taxi walks.

Lemma 6. *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 $4n+4\ell$. Then $\pi(\widehat{\Omega}_F)\le(1+\lambda)^{-(n+\ell)}$.*

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

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

□

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

Proof: Fix μ satisfying $\lambda>\mu^4-1>\mu_t^4-1$, where μ_t is the taxi walk connective constant. Summing over possible locations for the two faults F_1 and F_2 and using Lemma 6,

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

** In [23, Section 2.2] this is erroneously presented as $N = n+2\ell$, and the missing factor 2 remains absent for all the remaining calculations; the corrected calculations lead to the weaker bounds quoted in the introduction. It should be stressed that this missing factor *only* affects the numerical calculations in [23], and not any of the preceding theoretical discussion of fault lines and crosses.

The second inequality here uses Theorem 3. By our choice of μ 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. $\square$

From Section 2.2 we know that $\mu_t^4 - 1 < 5.3646$. From Theorem 5, we thus get the first part of Theorem 2, as well as the following stronger result.

Corollary 1. *Fix $\lambda > \mu_t^4 - 1$. 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 $c > 0$ (depending on λ).*

Proof: We will bound the conductance by considering $S = \widehat{\Omega}_0$. It is clear that $\widehat{\pi}(S) \leq 1/2$ since $\overline{S} = \widehat{\Omega}_{\mathcal{F}} \cup \widehat{\Omega}_1$ and $\widehat{\pi}(\widehat{\Omega}_0) = \widehat{\pi}(\widehat{\Omega}_1)$. Thus,

$$\Phi \leq \frac{\sum_{s \in \widehat{\Omega}_0, t \in \widehat{\Omega}_{\mathcal{F}}} \widehat{\pi}(s) P(s,t)}{\widehat{\pi}(\widehat{\Omega}_0)} = \frac{\sum_{s \in \widehat{\Omega}_0, t \in \widehat{\Omega}_{\mathcal{F}}} \widehat{\pi}(t) P(t,s)}{\widehat{\pi}(\widehat{\Omega}_0)} \leq \frac{\sum_{t \in \widehat{\Omega}_{\mathcal{F}}} \widehat{\pi}(t)}{\widehat{\pi}(\widehat{\Omega}_0)} = \frac{\widehat{\pi}(\widehat{\Omega}_{\mathcal{F}})}{\widehat{\pi}(\widehat{\Omega}_0)}.$$

Given Theorem 5, it is trivial to show that $\widehat{\pi}(\widehat{\Omega}_0) > 1/3$, thereby establishing that the conductance is exponentially small. It follows that Glauber dynamics takes exponential time to converge. $\square$

4.2 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.

Theorem 6. *Fix λ satisfying*

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

where μ_t is the taxi walk connective constant. Glauber dynamics for independent sets on the $n \times n$ grid G takes time at least e^{cn} to mix, for some constant $c = c(\lambda) > 0$.

Simple algebra reveals that the second condition is satisfied whenever $\lambda^2 + (2 - \mu_t^2 - \mu_t^4)\lambda + (1 - \mu_t^2) > 0$. Using $\mu_t < 1.5884$ we find that both of these conditions are met when $\lambda > 7.1031$, as claimed in Theorem 2.

As before, we partition the state space Ω into three sets, namely $\Omega_{\mathcal{F}}$ (independent sets with a fault line), Ω_0 (independent sets with an even cross), and Ω_1 (independent sets with an odd cross). From Lemma 1 we know that these sets do indeed partition Ω, and that furthermore Ω_0 and Ω_1 are not connected by moves in P. It remains to show that $\pi(\Omega_{\mathcal{F}})$ is exponentially smaller than $\pi(\Omega_0)$ and $\pi(\Omega_1)$.

Let $I \in \Omega_{\mathcal{F}}$ be an independent set with vertical fault line F. The fault F 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 points, and the edges form a path in $G_\diamondsuit$. We will represent length of the path as $N = n + 2\ell$, for some $\ell \in \mathbb{Z}$. Strictly speaking this is only correct for fault lines with zero alternation points; but this slight abuse 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 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'.

We use the following result from [23].

Lemma 7. *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'.*

Moreover, F and F' are taxi walks or the concatenation of two taxi walks at an alternation point.

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; 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. Let the length of F be $n + 2\ell$, 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 8. *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 $n + \ell - \delta$ and let $|r|$ denote the number of bits set to 1, where $|r| \leq n + \ell$. The main step is to define an injective map $\phi_{F,J} : \Omega_{F,J} \times \{0,1\}^{n+\ell} \to \Omega$ such that, for any $I \in \Omega_{F,J}$,

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

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 7 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 9. *Let $G_{1,n}$ be a $1 \times n$ strip, and let Ω_n be the set of independent sets in $G_{1,n}$. Then*

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

for some constant c.

Proof: Let Ω_i be the set of independent sets in $G_{1,i}$ and let $T_i = \sum_{J\in\Omega_i}\lambda^{|J|}$. 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 7. *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 λ satisfies the hypothesis of Theorem 6.

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. Using Lemma 9 for the third inequality and Theorem 3 for the fourth, we have

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

This means that we will have $\pi(\Omega_{\mathcal{F}}) \leq p(n)e^{-c'n}$, for some polynomial $p(n)$, if both

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

hold. Taking μ arbitrarily close to μ_t, the result follows. □

Theorem 6 follows exactly as Corollary 1 in Section 4.1.

5 Taxi walks: Bounds and Limits

We conclude by justifying the upper bound on the number of taxi walks given in Theorem 3, as well as providing a lower bound on μ_t. It is necessary to first establish the submultiplicativity of $\widetilde{c}_n$ (or, equivalently, the subadditivity of $\log \widetilde{c}_n$).

Lemma 10. *Let $\widetilde{c}_n$ be the number of taxi walks of length n and let $1 \leq i \leq n-1$. Then $\widetilde{c}_n \leq \widetilde{c}_i\, \widetilde{c}_{n-i}$.*

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 $\widetilde{c}_n \leq \widetilde{c}_i\, \widetilde{c}_{n-i}$. □

It follows from Lemma 10 that $a_n = \log \widetilde{c}_n$ is subadditive, i.e., $a_{n+m} \leq a_n + a_m$. By Fekete's Lemma (see, e.g., [28, Lemma 1.2.2]) we know that $\lim_{n\to\infty} a_n/n$ exists and

$$\lim_{n\to\infty} \frac{a_n}{n} = \inf \frac{a_n}{n}. \tag{3}$$

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.

Subadditivity gives us a strategy for getting a better bound on μ_t. From (3) we see that for all n, $\log \widetilde{c}_n/n$ is an upper bound for $\log \mu_t$. We exactly enumerated taxi walks of length n, for $n \leq 60$; see `http://nd.edu/~dgalvin1/TD/` for this and other data. Using $c_{60} = 2189670407434$ gives a bound of $\mu_t < 1.6058$. Note that exact counts for larger n will immediately improve our bounds on both μ_t and λ_c.

The connective constant for ordinary self-avoiding walks has been well studied, and some of the methods used to obtain bounds there can be adapted to deal with taxi walks. In particular, a method of Alm [1] is useful. Fix $n > m > 0$.

Construct a square matrix $A(m,n)$ whose ij entry counts the number of taxi walks of length n that begin with the ith taxi walk of length m, and end with the jth taxi walk of length m, for some fixed ordering of the walks of length m. To make sense of this, it is necessary to choose, for each $v \in \mathbb{Z}^2$, an orientation preserving map f_v of $\mathbb{Z}^2$ that sends the origin to v; saying that a walk of length n ends with the jth walk of length m means that if the length m terminal segment of the walk is transformed by f_v^{-1} to start at the origin, where v is the first vertex of the terminal segment, then the result is the jth walk of length m. Then a theorem of Alm [1] may be modified to show that μ_t is bounded above by $\lambda_1(A(m,n))^{1/(n-m)}$, where λ_1 indicates the largest positive eigenvalue. (Note that when $m=0$ this recovers the subadditivity bound discussed earlier).

We have calculated $A(20,60)$. This is a square matrix of dimension 20114, and a simple symmetry argument reduces the dimension by a factor of 2. Using MATLAB, we could estimate the largest eigenvalue of this reduced matrix to obtain $\mu_t < 1.5884$ and $\mu_t^4 - 1 < 5.3646$.

A similar strategy can be used to derive lower bounds on μ_t in order to determine the theoretical limitations of our approach of characterizing contours by taxi walks. We have already given the trivial lower bound $\mu_t \geq \sqrt{2}$. To improve this, we consider *bridges* (introduced for ordinary self-avoiding walks by Kesten [15]). A *bridge*, for our purposes, is a taxi walk that begins by moving from the origin $(0,0)$ to the point $(1,0)$, never revisits the y-axis, and ends by taking a step parallel to the x-axis to a point on the walk that has maximum x-coordinate over all points in the walk (but note that this maximum does not have to be uniquely achieved at the final point).

Let b_n be the number of bridges of length n. Then bridges are supermultiplicative, i.e., $b_n \geq b_i b_{n-i}$ (and $\log b_n$ is superadditive). To see this, note that if β_1 and β_2 are bridges, then they both begin and end at vertices whose y-coordinates are even because they are taking steps to the East. If the parities of the x-coordinates of the first vertices in β_1 and β_2 agree, then the concatenation of β_1 and an appropriate translation of β_2 is also a bridge; if the parities are different then concatenation of β_1 with a translation of β_2 after reflecting horizontally will be a valid bridge. Notice that the parity of the x-coordinate of the two pieces allows us to recover whether a reflection was necessary to keep the walk on the directed Manhattan lattice, so bridges are indeed supermultiplicative. It similarly follows that there are at least b_n^k taxi walks of length kn (just concatenate k length n bridges), so that

$$\mu_t = \lim_{m\to\infty} \tilde{c}_m^{1/m} \geq \lim_{k\to\infty} (b_n^k)^{1/nk} = b_n^{1/n}.$$

We have enumerated bridges of length up to 60, in particular discovering that $b_{60} = 80312795498$, leading to $\mu_t > 1.5196$ and $\mu_t^4 - 1 > 4.3332$.

A consequence of our lower bound on μ_t is that our present approach to phase coexistence cannot give anything better than $\lambda_c \leq 4.3332$; this tells us that new ideas will be needed to reach the value of 3.796 suggested by computations as the true value of λ_c.

References

1. S. Alm. Upper bounds for the connective constant of self-avoiding walks. *Combinatorics, Probability & Computing* **2**: 115–136, 1993.
2. R.J. Baxter, I.G. Entig, and S.K. Tsang. Hard-square lattice gas. *J. of Statistical Physics*, **22**: 465–489, 1980.
3. V. Beffara and H. Duminil-Copin. The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. *Probability Theory and Related Fields*, **153**: 511–542, 2012.
4. J. van den Berg and J.E. Steif. Percolation and the hard-core lattice model. *Stochastic Processes and their Applications*, **49**: 179–197, 1994.
5. C. Borgs. Personal communication.
6. 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 (FOCS)*, 218–229, 1999.
7. G. R. Brightwell, O. Häggström, and P. Winkler. Non- monotonic behavior in hard-core and Widom-Rowlin- son models. *J. of Statistical Physics*, **94**: 415–435, 1999.
8. R.L. Dobrushin. The problem of uniqueness of a Gibbs random field and the problem of phase transitions. *Functional Analysis and its Applic.*, **2**: 302–312, 1968.
9. 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*, Birkhäuser, Boston, 347–370, 1985.
10. D. Galvin. Sampling independent sets on the discrete torus. *Random Str. & Algorithms,* **33**: 356–376, 2008.
11. D. Galvin and J. Kahn. On phase transitions in the hard-core model on Z^d. *Combinatorics, Probability & Computing*, **13**: 137–164, 2004.

12. H.-O. Georgii, Gibbs Measures and Phase Transitions. de Gruyter, Berlin, 1988.
13. R. Glauber. Time dependent statistics of the Ising model. *J. Mathematical Physics,* **4**: 294–307, 1963.
14. 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.
15. H. Kesten. On the number of self-avoiding walks. *J. Math. Phys.,* **4**: 960–969, 1963.
16. E. Lubezky and A. Sly. Critical Ising on the square lattice mixes in polynomial time. To appear in *Communications in Mathematical Physics*.
17. M. Luby and E. Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets. *Random Structures & Algorithms*, **15**: 229–241, 1999.
18. N. Madras and G. Slade. The Self-Avoiding Walk. Birkhäuser, Boston, 1993.
19. 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.
20. L. Onsager. Crystal statistics. I. A two-dimensional model with an order-disorder transition. *Physics Review Letters*, **65**: 117–149, 1944.
21. R. Peled and W. Samotij. Odd cutsets and the hard-core model on $\mathbb{Z}^d$. To appear in *Annales de l'Institut Henri Poincaré,* 2013.
22. D. C. Radulescu and D. F. Styer. The Dobrushin-Shlosman phase uniqueness criterion and applications to hard squares. *Journal of Statistical Physics*, **49**: 281–295, 1987.
23. D. Randall. Slow mixing of Glauber dynamics via topological obstructions. *Proc 17th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 870–879, 2006.
24. R. Restrepo, J. Shin, P. Tetali, E. Vigoda and L. Yang. Improved mixing condition on the grid for counting and sampling independent sets. *Proc. 52nd IEEE Symp. on Foundations of Computer Science (FOCS)*, 140–149, 2011.
25. A. Sinclair. Counting and generating combinatorial structures: a Markov chain approach. Birkhäuser, 1993.
26. A.J. Sinclair and M.R. Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains *Information and Computation* **82**: 93–133, 1989.
27. A. Sly. Computational Transition at the Uniqueness Threshold. *Proc. 51st IEEE Symp. on Foundations of Computer Science (FOCS)*, 287–296, 2010.
28. J.M. Steele. Probability Theory and Combinatorial Optimization. SIAM, 1997.
29. J. Vera, E. Vigoda and L. Yang. Improved bounds on the phase transition for the hard-core model in 2-dimensions. To appear in *Proc. 17th Workshop on Randomization and Computation (RANDOM)*, 2013.
30. D. Weitz. Counting independent sets up to the tree threshold. *Proc. 38th ACM Symp. on the Theory of Computing (STOC)*, 140–149, 2006.