

Perfect Sampling in Infinite Spin Systems via Strong Spatial Mixing*

Konrad Anand[†] and Mark Jerrum[‡]

School of Mathematical Sciences
Queen Mary, University of London
Mile End Road, London E1 4NS
United Kingdom

Abstract

We present a simple algorithm that perfectly samples configurations from the unique Gibbs measure of a spin system on a potentially infinite graph G . The sampling algorithm assumes strong spatial mixing together with subexponential growth of G . It produces a finite window onto a perfect sample from the Gibbs distribution. The run-time is linear in the size of the window, in particular it is constant for each vertex.

1 Introduction

Many interesting computational problems have come out of statistical physics, in particular from the study of spin systems. A wide variety of systems used and studied in probability, physics, machine learning, and theoretical computer science fall under this domain. In brief, a spin system is defined by a graph G , a finite (in our case) set of spins, and a local ‘specification’. A configuration of the system is an assignment of spins to the vertices of the graph. The local specification defines a probability distribution on the configurations of the system. Although the specification is local, this Gibbs distribution can exhibit long-range effects. A primary algorithmic goal in this area is to sample configurations from the Gibbs distribution.

Spin systems have been an active area of research from a computational perspective in the past few decades, with many algorithms proposed and many results on hardness achieved. Many of the most successful algorithms have been based on Markov chain simulation. Particularly prominent is the familiar Glauber dynamics [Gla63] whose mixing time has remained an active area of research, with wonderful new results coming in recent years [CLV20a; CLV20b], prompted by pioneering work on high dimensional expanders [KO20; ALGV18].

*Published in SICOMP

[†]Supported by a studentship from the School of Mathematical Sciences.

[‡]Supported by grant EP/S016694/1 ‘Sampling in hereditary classes’ from the Engineering and Physical Sciences Research Council (EPSRC) of the UK.

Deterministic approaches have also been proposed, based on decay of correlations [Wei06] or Taylor expansion in a zero-free region of the partition function [PR17; BS17]. As with Markov chain simulation, these approaches provide a sample from an approximation to the desired distribution. Another strand of research has been into perfect samplers. Aside from the direct algorithmic challenge of producing perfect samples, there is another, more practical, motivation. This motivation arises not from the exactness of the sampling distribution as such; after all, the t -step distribution of a Markov chain converges exponentially in t to stationarity, so the deviation of the sampling distribution from the desired one can be made very small at modest computational expense. More significant is having a definite termination condition, in contrast to Markov chain simulation which has to be carried on for ‘sufficiently many’ steps. This requires an a priori bound on the ‘mixing time’ of a Markov chain, which must be analytically derived and may be much larger than actually necessary. In contrast, perfect samplers can respond to characteristics of a problem instance, and produce results more quickly on favourable instances.

Initially, it is not obvious that perfect sampling is possible in complex settings. The first demonstration came with the Coupling From The Past (CFTP) approach of Propp and Wilson [PW98]. There have been other recent approaches to the problem of efficient perfect sampling which move away from Markov chains, such as the partial rejection sampling of Guo, Jerrum and Liu [GJL19], the randomness recycler of Fill and Huber [FH00], and the work on perfect simulation by Huber [Hub16]. Our contribution is in a similar direction.

Our particular goal in this work is to produce perfect samples from a spin system, even when the underlying graph G may be infinite. Although the specification of a spin system is local, it may, if the conditions are right, define a unique Gibbs measure on an infinite graph G . Classic examples of this phenomenon include the Ising and hardcore models on $\mathbb{Z}^d$, in the subcritical regime. Since we cannot write down an infinite spin configurations, we need to explain what we mean by perfect sampling in this situation. Consider a spin system on $\mathbb{Z}^2$ that has a unique Gibbs measure. As an example of what we want to achieve, we might cite the problem of sampling (perfectly) from the the distribution of spins on an $L \times L$ square fragment of $\mathbb{Z}^2$ induced by the infinite Gibbs measure.

Van den Berg and Steif [BS99] pioneered perfect sampling from infinite Gibbs measures. They showed that perfect sampling is possible in the case of the ferromagnetic Ising model on $\mathbb{Z}^2$ in the uniqueness regime, i.e., when there is a unique Gibbs measure. They also gave a precise negative result that essentially rules out perfect simulation in non-uniqueness. Their approach is through CFTP. Spinka [Spi20], in a wide-ranging investigation, has taken forward this approach. In contrast we depart from CFTP, and take our inspiration from partial rejection sampling and related work.

The closest prior work to ours is that of Feng, Guo and Yin [FGY19]. Like them, our key assumptions are that the spin system displays ‘strong spatial mixing’, and that the underlying graph has ‘subexponential growth’. Definitions of these terms can be found in Section 4. Roughly speaking, we require correlations between spins to decay rapidly with distance, and local neighbourhoods of G to grow slowly enough with increasing radius. Taken together, these conditions allow us to largely discount the effect of spins at far away vertices, which is key to the efficiency of our algorithm. The conditions of strong spatial mixing and subexponential growth are quite standard and have been widely studied [GMP05; Wei06]. Various authors have sought to leverage the property of strong spatial

mixing to algorithmic ends [FGY19; Spi20; SSY13; SSvY17; DSVW04]. Intriguingly, it has been suggested [CLV20a] that strong spatial mixing may be equivalent to spectral independence, a concept that is currently the subject of intensive study. In support of this idea is that in the case of subexponential growth, strong spatial mixing with exponential rate implies spectral independence: these two conditions bound the total influence on a vertex, which in turn bounds the rows of the influence matrix, and thus the eigenvalues, giving spectral independence. As for the second condition, subexponential growth currently seems important for perfect sampling, even though some approaches to approximate sampling dispense with it [Liu21].

The main difference between our approach and that of Feng et al. is the following. Their strategy is to start with a spin configuration on the whole of G , and gradually ‘repair’ the configuration until it becomes a perfect sample. As this approach seems inapplicable to infinite graphs, we instead ‘grow’ the perfect sample one vertex at a time. Our recursive algorithm tries to guess the spin at a vertex v without looking at spins that live on a sphere of radius ℓ centred at v ; when it cannot guess, the algorithm recursively determines the spins on the sphere before returning to assign a spin to v .

As it runs, algorithm traces out something akin to a branching process, whose expected size is finite with probability 1 when sufficiently strong spatial mixing holds. When successful, the algorithm is linear time: for each vertex whose spin we sample, the algorithm takes constant time. Another view is that the algorithm calculates marginal probabilities in sublinear time. An advantage to the algorithm is that we do not need a configuration on an entire graph to determine the spin at a single vertex. This gives the possibility of drastically more efficient sampling than any algorithm requiring an entire configuration, and allows us to sample from infinite graphs as well. The algorithm provides a novel perspective on the use of strong spatial mixing on spin systems. Indeed, the algorithm could be viewed as an effective proof that strong spatial mixing together with subexponential growth implies uniqueness of the Gibbs measure.

2 Spin systems

For the time being, we restrict attention to finite graphs, and defer the discussion of infinite graphs until later. Given a finite graph $G = (V, E)$ and $q \in \mathbb{N}$ we work with the set of configurations $\Omega_V = [q]^V$. When there is no room for confusion, we will write Ω instead of Ω_V . Naturally, when $W \subset V$ we write $\Omega_W = [q]^W$ for the partial configurations restricted to W .

For $\sigma \in \Omega$ and $W \subset V$ we denote its restriction to W by σ_W . For simplicity, when $W = \{v\}$ we will write σ_v instead of $\sigma_{\{v\}}$. In this same situation, we denote the spin at v by i , and write $\sigma_v = i$ rather than $\sigma_v = \{(v, i)\}$.

Given a field $b : [q] \rightarrow \mathbb{R}$ and a symmetric interaction weight $A : [q] \times [q] \rightarrow \mathbb{R}$, the Gibbs distribution gives to each configuration σ the weight

$$\prod_{v \in V} b(\sigma_v) \prod_{(u, v) \in E} A(\sigma_u, \sigma_v).$$

Define the partition function $Z(G)$ to be the sum of all possible weights

$$Z(G) := \sum_{\sigma \in \Omega} \prod_{v \in V} b(\sigma_v) \prod_{(u,v) \in E} A(\sigma_u, \sigma_v).$$

Then we define the Gibbs measure to be the measure $\mu : \Omega \rightarrow [0, 1]$ where

$$\mu(\sigma) = \frac{\prod_{v \in V} b(\sigma_v) \prod_{(u,v) \in E} A(\sigma_u, \sigma_v)}{Z(G)}. \quad (2.1)$$

Frequently in this paper we will want to look at the marginal distribution induced by fixing the configuration on some vertices. For any W , define the marginal distribution of μ on W by

$$\mu_W(\sigma) = \sum_{\sigma' \in \Omega : \sigma'_W = \sigma} \mu(\sigma'),$$

for all $\sigma \in \Omega_W$. Also, for any $\Lambda \subset V$ and $\tau \in \Omega_\Lambda$, define the marginal distribution with boundary condition (Λ, τ) by

$$\mu_W^{(\Lambda, \tau)}(\sigma) := \frac{1}{\mu_\Lambda(\tau)} \sum_{\substack{\sigma' \in \Omega : \\ \sigma'_\Lambda = \tau, \sigma'_W = \sigma}} \mu(\sigma'), \quad (2.2)$$

assuming $\mu_\Lambda(\tau) \neq 0$. We abbreviate $\mu_W^{(\Lambda, \tau)}$ to μ_W^τ when the set Λ is clear from the context. With this definition of the Gibbs measure, we have a fixed probability space determined by a few parameters. We call this a spin system $\mathcal{S} = (G, q, b, A)$.

As usual, the *Gibbs property* of μ is important to us. Suppose $W \subset V$ and that the configuration $\tau \in \Omega_{V \setminus W}$ is *feasible*, i.e., $\mu_{V \setminus W}(\tau) > 0$. Define the *boundary* of W by

$$\partial W = \{v \in V \setminus W : \{u, v\} \in E \text{ for some } u \in W\},$$

and let $\tau' \in \Omega_{\partial W}$ be the restriction of τ to ∂W . Then $\mu_W^{(V \setminus W, \tau)} = \mu_W^{(\partial W, \tau')}$; in other words, conditioning on the complement of W is equivalent to conditioning on its boundary.

3 Sketch of the basic algorithm

Before describing the full algorithm and the properties it requires of the spin system, we pause to describe a basic version of the algorithm, the ‘radius 1’ version of the full algorithm. After this initial consideration, we will expand the algorithm to a radius determined by spatial mixing.

Suppose we start with a graph $G = (V, E)$ and we want to know the spin at a single vertex $v \in V$. We give this vertex a realisation y of a $U[0, 1]$ (i.e., uniform in the unit interval) random variable. If we knew the spins of the neighbours of v it would be a simple task to partition the unit interval according to the probabilities of each possible spin at v , and then use y to determine the spin at v . Unfortunately, we do not and so we need some way around this.

Instead of just partitioning $[0, 1]$ according to the respective spin probabilities, we partition it according to the minimum probability for each spin given any configuration on v 's neighbours. Let $W = \{w_1, \dots, w_d\}$ be the neighbours of v . For $i \in [q]$ define

$$p_v^i = \min_{\tau \in \Omega_W} \mu_v^\tau(i).$$

So p_v^i is the minimum, over all assignments of spins to neighbours of v , of the probability that vertex v should be assigned spin i . For the purposes of the algorithm we will partition $[0, 1]$ into intervals $I_1, \dots, I_q$ of sizes $p_v^1, \dots, p_v^q$ with a remaining interval which we call the ‘zone of indecision’ I_0 of size $p_v^0 = 1 - \sum_{i=1}^q p_v^i$.

If the point $y \in [0, 1]$ selected earlier falls in $\cup_{i=1}^q I_i$ then we can assign a spin to v . If not, we need to know the spins of the neighbours in order to split I_0 into different spin regions. We will call the above process on each of the neighbours recursively in turn.

How we make these recursive calls is important—the outcome for W must follow the marginal distribution μ_W , not the product of the marginal distributions at each neighbour, that is, for $\sigma \in \Omega_W$ we want

$$\mu_W(\sigma) = \mu_{w_1}(\sigma_{w_1}) \cdot \mu_{w_2}^{\sigma_{w_1}}(\sigma_{w_2}) \cdots \mu_{w_d}^{\sigma_{W-w_d}}(\sigma_{w_d}).$$

To this end, when we call the algorithm on w_j , we condition on the spins of $w_1, \dots, w_{j-1}$ to ensure the correctness of the probabilities. Once we have the spins at each of the neighbours, we subdivide the zone of indecision I_0 into q intervals corresponding to the different spins. We now use y to determine the spin at v and discard the intermediate working (including any spins applied to vertices other than v) that we have used during the algorithm.

Why should this strategy work? Well, if we chose the spins σ_W first, then subdivided the interval $[0, 1]$ appropriately, and then chose a uniform random point y in the interval, we would certainly assign a spin to v with the correct probability. We merely observe that, if we choose the random point y first, then at least some of the time we do not need to know σ_W and we can omit the first step. If the probability of omitting the first step is large enough, we can avoid an infinite recursion.

3.1 Example: the hardcore model

A classic example of a spin system is the hardcore model. The hardcore model is a distribution over independent sets of a graph, which we model as a spin system with $q = 2$ spins. For consistency with the notation already introduced, we label these spins 1 and 2 rather than the usual 0 and 1. The one constraint is that neighbouring vertices are not allowed to both have spin 2. The hardcore distribution gives independent sets a weight depending on their size. To be explicit, the associated field and interaction weights are:

$$b = \begin{pmatrix} 1 \\ \lambda \end{pmatrix}, \quad A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}.$$

Now, to see the basic algorithm in action, we consider its behaviour on the hardcore model. Calling the algorithm at a vertex v , we see that the $p_v^1 = \frac{1}{1+\lambda}$ and $p_v^2 = 0$. (If one of the neighbours of v has spin 2, then we cannot assign spin 2 to vertex v . So $p_v^2 = 0$. On the other hand, if all the neighbours have spin 1 there is just a probability $\frac{1}{1+\lambda}$ that we will

assign spin 1 to v . So $p_v^1 = \frac{1}{1+\lambda}$.) Thus we set the spin to 1 immediately with probability $\frac{1}{1+\lambda}$ and call the algorithm recursively on the neighbours with probability $\frac{\lambda}{1+\lambda}$.

Consider the recursive calls on the neighbours, $w_1, \dots, w_d$. If the algorithm sets the spin of any w_i to 2 then we set the spin of v to 1 and terminate the algorithm. Otherwise, the neighbouring vertices all have spin 1. Conditioned on this event, v has probability $\frac{1}{1+\lambda}$ of being 1 and probability $\frac{\lambda}{1+\lambda}$ of being 2. Thus the entire zone of indecision is moved to the case of the spin being 2 and we terminate with this assignment.

We can view the recursive calls of the algorithm as a random tree T_A . By comparing T_A to a branching process, we will see that for graphs with maximum degree Δ , when $\lambda < \frac{1}{\Delta-1}$ the expected time to set the spin of one vertex is $O(1)$.

Each vertex in T_A will have no children when $y \in I_1 \cup I_2$ and will have a child for each neighbour when $y \in I_0$. The number of neighbours is bounded by Δ and the probability of having children is $p_v^0 = \frac{\lambda}{1+\lambda}$, so we can bound the size of T_A by a branching process with the offspring distribution ξ where

$$\mathbb{P}[\xi = 0] = \frac{1}{1+\lambda}, \quad \mathbb{P}[\xi = \Delta] = \frac{\lambda}{1+\lambda}.$$

The expected size of such a branching process is finite when

$$\Delta \cdot \frac{\lambda}{1+\lambda} < 1 \iff \lambda < \frac{1}{\Delta-1}.$$

To be exact, the expected size of this branching process—and the expected run-time of the algorithm—is

$$\frac{1}{1 - \Delta \cdot \frac{\lambda}{1+\lambda}} = \frac{1+\lambda}{1 - (\Delta-1)\lambda}.$$

Thus, the size of T_A is finite when $\lambda < \frac{1}{\Delta-1}$ and in this case the expected run-time of the algorithm is bounded by $\frac{1+\lambda}{1-(\Delta-1)\lambda}$.

4 Spin systems on infinite graphs: our requirements

Up to this point, our graphs have been finite, but we now want to extend the discussion to infinite graphs. Although we now allow the vertex set to be countably infinite, we do assume that graphs are locally finite, i.e., that all vertex degrees are finite. For perfect sampling to make sense, we need the system we are dealing with to have unique infinite Gibbs measure. Briefly, a Gibbs measure μ for an infinite graph G , with interaction weights A and field b , is a probability measure that has the following properties:

- Events that are determined by a finite set of spin variables have a well defined probability of occurrence. Technically, cylinder events are measurable.
- For any finite vertex subset $W \subset V$, the probability distribution of σ_W (i.e., the marginal distribution on W) is consistent with that given by equation (2.1) for the

finite graph $G[W]$. Specifically, suppose $W \subset V$ is any finite vertex subset, and define the boundary of W by

$$\partial W = \{v \in V \setminus W : \{u, v\} \in E \text{ for some } u \in W\}.$$

Then, for all $\tau \in \Omega_{\partial W}$, the probability distribution $\mu_W^{(\partial W, \tau)}$ (which is well defined as it is determined by a finite set of spin variables) agrees with equation (2.2), where $\Omega = [q]^{W \cup \partial W}$ and $\Lambda = \partial W$.

When there is a unique infinite Gibbs measure satisfying the above conditions we say that the system is in the ‘uniqueness regime’. For a thorough treatment of these ideas see, e.g., Friedli and Velenik [FV18].

Luckily, we will require strong spatial mixing for the algorithm, which itself implies that there is a unique infinite measure μ with the above properties. A perfect sampling algorithm for μ is required to open up a finite window onto a perfect sample from μ .

The basic algorithm, which was sketched in the context of the hardcore model, made recursive calls on the immediate neighbours of vertex v . It is correct, but works for only small values of λ . We therefore analyse a more general, radius ℓ , version of the algorithm, which recurses on vertices at distance ℓ from v , rather than its immediate neighbours. Let d_G denote graph distance and

$$S_\ell = S_\ell(v) := \{w \in V : d_G(v, w) = \ell\}$$

denote the sphere of radius ℓ about vertex v . Let (Λ, σ) be the partial assignment of existing spins. Then we redefine probabilities with which we partition the unit interval. Let

$$p_v^i := \min_{\tau \in \Omega_{S_\ell \setminus \Lambda}} \mu_v^{(\Lambda, \sigma) \oplus (S_\ell \setminus \Lambda, \tau)}(i) \quad \forall i = 1, \dots, q,$$

$$p_v^0 := 1 - \sum_{i \in [q]} p_v^i.$$

Here, $\oplus$ denotes ‘concatenation’ of spins. Thus $(\Lambda, \sigma) \oplus (S_\ell \setminus \Lambda, \tau)$ is an assignment of spins to $\Lambda \cup S_\ell$ that agrees with σ on Λ and τ on $S_\ell \setminus \Lambda$. We will simplify this notation to $\sigma \oplus \tau$ if the domains of σ and τ are clear from the context.

To analyse the efficiency of the algorithm, we consider its behaviour as a branching process where at a given vertex v the probability of no children is $1 - p_v^0$ and the probability of $|S_\ell \setminus \Lambda|$ children is p_v^0 . We want to limit this branching, and to do so we leverage strong spatial mixing. Say that a partial assignment (W, τ) is *feasible* if $\mu_W(\tau) > 0$. Also denote total variation distance between two distributions by d_{TV} .

Definition 4.1. *Let $\mathcal{S}$ be a spin system and $f : \mathbb{N} \rightarrow \mathbb{R}^+$ a function. Given $W \subset V, v \in V \setminus W$, and $\tau^1, \tau^2 \in \Omega_W$, let $\ell = \min\{d_G(v, w) : w \in W, \tau_w^1 \neq \tau_w^2\}$. If for all W with $V \setminus W$ finite, all $v \in V \setminus W$, and all feasible $\tau^1, \tau^2 \in \Omega_W$ with $\tau^1 \neq \tau^2$,*

$$d_{\text{TV}}(\mu_v^{\tau^1}, \mu_v^{\tau^2}) \leq f(\ell),$$

then we say that $\mathcal{S}$ exhibits strong spatial mixing with rate f .

Note that there is no problem giving a meaning to the conditional distributions $\mu_v^{\tau^1}$ and $\mu_v^{\tau^2}$: since $V \setminus W$ finite, the Gibbs property of μ ensures that $\mu_v^{\tau^1}$ and $\mu_v^{\tau^2}$ depend on a finite set of spins. We want to use this bound to control the size of the zone of indecision which branches the algorithm.

Lemma 4.2. *For any $\Lambda \subset V$, $v \in V \setminus \Lambda$, and feasible partial configuration (Λ, σ) ,*

$$p_v^0 \leq q \cdot \max_{\tau^1, \tau^2 \in \Omega_{S_\ell(v) \setminus \Lambda}} d_{\text{TV}} \left(\mu_v^{\sigma \oplus \tau^1}, \mu_v^{\sigma \oplus \tau^2} \right),$$

where the maximisation is over all feasible $\sigma \oplus \tau^1$ and $\sigma \oplus \tau^2$.

Proof. For convenience, set $W = S_\ell(v) \setminus \Lambda$. Let $\tau' \in \Omega_W$ be an arbitrary configuration with $\sigma \oplus \tau'$ feasible. Then

$$p_v^0 = 1 - \sum_{i \in [q]} p_v^i = \sum_{i \in [q]} \left(\mu_v^{\sigma \oplus \tau'}(i) - p_v^i \right) \leq \sum_{i \in [q]} \max_{\tau \in \Omega_W} \left(\mu_v^{\sigma \oplus \tau}(i) - p_v^i \right).$$

Clearly, for every $i \in [q]$ it holds that

$$\max_{\tau \in \Omega_W} \left(\mu_v^{\sigma \oplus \tau}(i) - p_v^i \right) \leq \max_{\tau^1, \tau^2 \in \Omega_W} d_{\text{TV}} \left(\mu_v^{\sigma \oplus \tau^1}, \mu_v^{\sigma \oplus \tau^2} \right),$$

so it follows that

$$p_v^0 \leq \sum_{i \in [q]} \max_{\tau \in \Omega_W} \left(\mu_v^{\sigma \oplus \tau}(i) - p_v^i \right) \leq q \cdot \max_{\tau^1, \tau^2 \in \Omega_W} d_{\text{TV}} \left(\mu_v^{\sigma \oplus \tau^1}, \mu_v^{\sigma \oplus \tau^2} \right).$$

□

It is worth pointing out that we cannot do much better than this bound. Consider the family of distributions $\{\mu_j : j \in [q]\}$ over $[q]$ such that

$$\mu_j(i) = \begin{cases} 0 & : \text{if } i = j, \\ \frac{1}{q-1} & : \text{otherwise.} \end{cases}$$

Then $p_v^0 = 1$ and for all $j, k \in [q], j \neq k$

$$d_{\text{TV}}(\mu_j, \mu_k) = \frac{1}{q-1}.$$

We will leverage this bound and strong spatial mixing for the full algorithm.

Finally, a useful property of a graph is *sub-exponential neighbourhood growth*.

Definition 4.3. *Let $G = (V, E)$. We say that G has sub-exponential neighbourhood growth if for all $c > 1$ there exists an $N \in \mathbb{N}$ such that for all $v \in V$ and all $\ell > N$*

$$|S_\ell(v)| < c^\ell.$$

This property provides a restriction on how fast the number of recursive calls in our algorithm grows.

5 The algorithm in detail

The general algorithm is quite similar to the basic algorithm, but rather than conditioning on neighbours we condition on S_ℓ , the sphere of radius ℓ around our chosen vertex.

The algorithm takes a spin system on a graph G with a partial configuration σ on some subset Λ of the vertices. (Informally, the partial configuration (Λ, σ) represents the spins that have already been determined.) The algorithm is then called on some unassigned vertex v with a distance parameter ℓ . Recall that $S_\ell(v) := \{w \in V : d_G(v, w) = \ell\}$ is the sphere of radius ℓ centred at vertex v . Recall also the probabilities with which we split the unit interval used to decide the spin at v :

$$p_v^i := \min_{\tau \in \Omega_{S_\ell \setminus \Lambda}} \mu_v^{(\Lambda, \sigma) \oplus (S_\ell \setminus \Lambda, \tau)}(i) \quad \forall i = 1, \dots, q,$$

$$p_v^0 := 1 - \sum_{i \in [q]} p_v^i.$$

We draw a sample $y \sim U[0, 1]$ from the uniform distribution on $[0, 1]$. If y is in a section of the interval corresponding to p_v^i , with $i \neq 0$, then we set the spin to i —if not, we recursively call the algorithm on the vertices in $S_\ell \setminus \Lambda$ in turn, each time conditioning on the assignments to the previous vertices in this set. (See Algorithm 1.)

Algorithm 1 $\text{SSMS}(\mathcal{S} = (G = (V, E), q, b, A), (\Lambda, \sigma), v, \ell)$

- 1: **Input:** The algorithm takes a spin system $\mathcal{S}$, a set of known vertices $\Lambda \subset V$ with a configuration $\sigma \in \Omega_\Lambda$, a vertex to sample $v \notin \Lambda$, and a distance $\ell \in \mathbb{N}$.
 - 2: **Output:** The algorithm returns the partial configuration passed in with a spin at v as well: $(\Lambda, \sigma) \oplus (v, i)$ for some $i \in [q]$.
 - 3: $S_\ell \leftarrow \{w \in V : d_G(v, w) = \ell\}$
 - 4: **for** $i \in [q]$ **do**
 - 5: $p_v^i \leftarrow \min_{\tau \in \Omega_{S_\ell \setminus \Lambda}} \mu_v^{\sigma \oplus \tau}(i)$
 - 6: $I_i \leftarrow \left[\sum_{j=1}^{i-1} p_v^j, \sum_{j=1}^i p_v^j \right)$
 - 7: **end for**
 - 8: $p_v^0 \leftarrow 1 - \sum_{i \in [q]} p_v^i$
 - 9: $I_0 \leftarrow [1 - p_v^0, 1]$
 - 10: Sample $y \sim U[0, 1]$, a realisation of a uniform $[0, 1]$ random variable.
 - 11: **if** $y \in I_0$ **then**
 - 12: $(J_1, \dots, J_q) \leftarrow \text{BD-SPLIT}(\mathcal{S}, (\Lambda, \sigma), v, \ell, (p_v^1, \dots, p_v^q))$
 - 13: Find $i \in [q]$ such that $y \in J_i$
 - 14: **return** $((\Lambda, \sigma) \oplus (v, i))$
 - 15: **else**
 - 16: Find i such that $y \in I_i$
 - 17: **return** $((\Lambda, \sigma) \oplus (v, i))$
 - 18: **end if**
-

Once the spins on $S_\ell \setminus \Lambda$ (and hence $S_\ell \cup \Lambda$) are decided, we subdivide the interval I_0 (the zone of indecision) into q subintervals, whose lengths can now be calculated precisely

in subroutine BD-SPLIT, by brute force. (See Algorithm 2.) The spin at the original vertex v is now determined using the original sample y from $U[0, 1]$.

Algorithm 2 BD-SPLIT $(\mathcal{S}, (\Lambda, \sigma), v, \ell, (p_v^1, \dots, p_v^q))$

```

1: Give  $S_\ell(v) \setminus \Lambda$  an ordering  $S_\ell(v) \setminus \Lambda = \{w_1, \dots, w_m\}$ .
2:  $(\Lambda', \sigma') \leftarrow (\Lambda, \sigma)$ 
3: for  $j \in [1, \dots, m]$  do
4:    $(\Lambda', \sigma') \leftarrow \text{SSMS}(\mathcal{S}, (\Lambda', \sigma'), w_j, \ell)$ 
5: end for
6: for  $i \in [q]$  do
7:    $\rho_v^i \leftarrow \mu_v^{\sigma'}(i) - p_v^i$ 
8:    $I_i \leftarrow \left[ \sum_{j=1}^q p_v^j + \sum_{k=1}^{i-1} \rho_v^k, \sum_{j=1}^q p_v^j + \sum_{k=1}^i \rho_v^k \right)$ 
9: end for
10: return  $(I_1, \dots, I_q)$ 

```

As we will see more concretely, there is a trade-off in this algorithm. As the radius ℓ grows, the calculations to determine the subdivision of I_0 rapidly increase, but the expected number of recursive calls will decrease.

An important thing to notice about the algorithm is that it does not terminate for all choices of $\mathcal{S}$ and ℓ , even if G is finite. We will see that the correctness of the algorithm essentially depends on whether it terminates with probability 1.

It is worth noting that the algorithm can be made more efficient in the BD-SPLIT. Rather than run every recursive call, we can update the $I_0, \dots, I_q$ after each individual call and terminate BD-SPLIT once I_0 no longer contains the uniform random variable.

5.1 Run-time

Our run-time analysis is in terms of the number of calls of SSMS. For fixed ℓ , the amount of time each call takes will be $O(1)$.

Our viewpoint is that the algorithm behaves similarly to a branching process where each recursive call of the algorithm is a child of the vertex which initiated the call. This comparison is not exact of course, as the calls at different vertices will not branch according to one distribution, but it is enough to give us a broad class of situations where the algorithm terminates.

Theorem 5.1. *Let $\mathcal{S} = (G = (V, E), q, b, A)$ be a spin system which exhibits strong spatial mixing with rate f . Suppose that we have a growth bound for G of the form*

$$|S_\ell(v)| \leq g(\ell),$$

for all $v \in V$. Let $\ell \in \mathbb{N}$ be such that

$$q \cdot f(\ell) \cdot g(\ell) \leq \alpha < 1,$$

Then for all $v \in V$, $\Lambda \subset V$, and $\sigma \in \Omega_\Lambda$ such that

$$\mu_\Lambda(\sigma) > 0,$$

the expected number of times $\text{SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell)$ calls SSMS (in total) is bounded by $\frac{1}{1-\alpha}$.

Proof. Let us construct a random tree T_A out of our recursive calls. Our root will be our initial call. Recursively, if at some node w of the tree we fall in the zone of indecision, we add a child to the node for every recursive call that is initiated by BD-SPLIT. If we do not fall in the zone of indecision, then w has no children.

Let T_B be the branching process with offspring distribution ξ where

$$\mathbb{P}[\xi = 0] = 1 - q \cdot f(\ell), \quad \mathbb{P}[\xi = g(\ell)] = q \cdot f(\ell).$$

Then T_B stochastically dominates T_A . Indeed, at every node $w \in T_A$, the probability of falling in the zone of indecision is bounded by $q \cdot f(\ell)$ due to the spin system's strong spatial mixing. Further, when w does have children, the number of children is bounded above by $|S_\ell(w)| \leq g(\ell)$.

Now we can leverage the theory of branching processes to finish off our bound. We know that

$$\mathbb{E}[\xi] \leq \alpha < 1,$$

so by the fundamental theorem of branching processes, we know that

$$\mathbb{E}[|T_A|] \leq \mathbb{E}[|T_B|] \leq \frac{1}{1 - \alpha}.$$

□

From this, we see that under appropriate conditions we can sample sections of graphs of size n in time $O(n)$. It is also worth noticing that the above conditions give termination of the algorithm with probability 1. Note that, unless ℓ is very small, the computation required to compute the subdivision of I_0 in BD-SPLIT will be huge. So although the expected time to compute one spin is constant, that constant may be impractically large. This will be a particular issue when the system is close to non-uniqueness.

The qualitative nature of the main result in this section is summarised in the following corollary.

Corollary 5.2. *Let $G = (V, E)$ be a graph with sub-exponential growth bounded by g and let $\mathcal{S} = (G, q, b, A)$ be a spin system which exhibits strong spatial mixing with rate f which is inversely exponential. Consider any choices of $v \in V$, $\Lambda \subset V$, and $\sigma \in \Omega_\Lambda$. Then the running time τ of $\text{SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell)$ satisfies $\mathbb{E}[\tau] = O(1)$.*

Proof. In Theorem 5.1, choose ℓ large enough to satisfy the inequality $q \cdot f(\ell) \cdot g(\ell) \leq \alpha < 1$. □

5.2 Correctness

Our understanding of the correctness of the algorithm depends on our view of the algorithm as a branching process.

Theorem 5.3. *Let $\mathcal{S} = (G = (V, E), q, b, A)$ be a spin system which exhibits strong spatial mixing with rate f . Suppose for all $v \in V$, all finite $\Lambda \subset V$, and all $\sigma \in \Omega_\Lambda$ such that*

$$\mu_\Lambda(\sigma) > 0,$$

$\text{SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell)$ terminates with probability 1. Then

$$\mathbb{P}[\text{SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell) = i] = \mu_v^\sigma(i).$$

Now rather than directly relating the output distribution of the algorithm to the true probability, we want to pass between the two via a similar algorithm — which is just a depth-bounded version of the original one — that is more readily seen to be correct.¹

Assume we have an oracle $\mathcal{O}$ which takes the input of SSMS and a spin i and produces the true probability $\mu_v^\sigma(i)$. (Since we have strong spatial mixing, this probability is well defined.) For some height $h \in \mathbb{N}$, our new bounded algorithm will exactly match SSMS until the recursion depth reaches h . At that point, the algorithm will use the oracle on all nodes at the deepest level. This is Algorithm 3.

Algorithm 3 BOUNDED-SSMS($\mathcal{S}, (\Lambda, \sigma), v, \ell, h$)

```

1: Input: The algorithm takes a spin system  $\mathcal{S}$ , a set of known vertices  $\Lambda \subset V$  with a
   configuration  $\sigma \in \Omega_\Lambda$ , a vertex to sample  $v \notin \Lambda$ , a distance  $\ell \in \mathbb{N}$ , and a height  $h \in \mathbb{N}$ .
2: Output: The algorithm returns the partial configuration passed in with a spin at  $v$  as
   well:  $(\Lambda, \sigma) \oplus (v, i)$  for some  $i \in [q]$ .
3: if  $h = 0$  then
4:   Use  $\mathcal{O}$  to sample  $i \sim \mu_v^\sigma$ 
5:   return  $(\Lambda, \sigma) \oplus (v, i)$ 
6: end if
7:  $S_\ell \leftarrow \{w \in V : d_G(v, w) = \ell\}$ 
8: for  $i \in [q]$  do
9:    $p_v^i \leftarrow \min_{\tau \in \Omega_{S_\ell \setminus \Lambda}} \mu_v^{\sigma \oplus \tau}(i)$ 
10:   $I_i \leftarrow \left[ \sum_{j=1}^{i-1} p_v^j, \sum_{j=1}^i p_v^j \right)$ 
11: end for
12:  $p_v^0 \leftarrow 1 - \sum_{i \in [q]} p_v^i$ 
13:  $I_0 \leftarrow [1 - p_v^0, 1]$ 
14: Sample  $y \sim U[0, 1]$ , a realisation of a uniform  $[0, 1]$  random variable.
15: if  $y \in I_0$  then
16:    $(J_1, \dots, J_q) \leftarrow \text{BOUNDED-BD-SPLIT}(\mathcal{S}, (\Lambda, \sigma), v, \ell, (p_v^1, \dots, p_v^q), S_\ell \setminus \Lambda, h - 1)$ 
17:   Find  $i \in [q]$  such that  $y \in J_i$ 
18:   return  $((\Lambda, \sigma) \oplus (v, i))$ 
19: else
20:   Find  $i$  such that  $y \in I_i$ 
21:   return  $((\Lambda, \sigma) \oplus (v, i))$ 
22: end if

```

Now we see the correctness:

Lemma 5.4. *Under the conditions of Theorem 5.3, for all $h \in \mathbb{N}$ and for all $i \in [q]$*

$$\mathbb{P}[\text{BOUNDED-SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell, h) = i] = \mu_v^\sigma(i).$$

¹The issue with the original version of the algorithm is that the recursion it employs is not well founded, so we cannot directly prove correctness by induction, as we should like to do.

Algorithm 4 BOUNDED-BD-SPLIT $(\mathcal{S}, (\Lambda, \sigma), v, \ell, (p_v^1, \dots, p_v^q), h)$

```

1: Give  $S_\ell(v) \setminus \Lambda$  an ordering  $S_\ell(v) \setminus \Lambda = \{w_1, \dots, w_m\}$ .
2:  $(\Lambda', \sigma') \leftarrow (\Lambda, \sigma)$ 
3: for  $j \in [1, \dots, m]$  do
4:    $(\Lambda', \sigma') \leftarrow \text{BOUNDED-SSMS}(\mathcal{S}, (\Lambda', \sigma'), w_j, \ell, h)$ 
5: end for
6: for  $i \in [q]$  do
7:    $\rho_v^i \leftarrow \mu_v^{\sigma'}(i) - p_v^i$ 
8:    $I_i \leftarrow \left[ \sum_{j=1}^q p_v^j + \sum_{k=1}^{i-1} \rho_v^k, \sum_{j=1}^q p_v^j + \sum_{k=1}^i \rho_v^k \right]$ 
9: end for
10: return  $(I_1, \dots, I_q)$ 

```

Proof. We prove this by induction on h . For the base case of $h = 0$ the lemma is true due to the use of the oracle $\mathcal{O}$. Now suppose it holds for $h = k$. Given the order $S_\ell(v) \setminus \Lambda = \{w_1, \dots, w_m\}$ that the algorithm processes the vertices in, we have

$$\begin{aligned}
& \mathbb{P}[\text{BOUNDED-SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell, k+1) = i] \\
&= p_v^i + \sum_{\tau \in \Omega_{S_\ell(v) \setminus \Lambda}} (\mu_v^{\sigma \oplus \tau}(i) - p_v^i) \\
&\quad \times \prod_{j=1}^m \mathbb{P} \left[\text{BOUNDED-SSMS} \left(\mathcal{S}, (\Lambda, \sigma) \oplus \bigoplus_{k=1}^{j-1} (w_k, \tau_{w_k}), w_j, \ell, k \right) = \tau_{w_j} \right] \\
&= p_v^i + \sum_{\tau \in \Omega_{S_\ell(v) \setminus \Lambda}} (\mu_v^{\sigma \oplus \tau}(i) - p_v^i) \prod_{j=1}^m \mu_{w_j}^{(\Lambda, \sigma) \oplus \bigoplus_{k=1}^{j-1} (w_k, \tau_{w_k})}(\tau_{w_j}),
\end{aligned}$$

where the second inequality comes from the induction hypothesis. Finally, observing that

$$p_v^i = \sum_{\tau \in \Omega_{S_\ell(v) \setminus \Lambda}} p_v^i \cdot \prod_{j=1}^m \mu_{w_j}^{(\Lambda, \sigma) \oplus \bigoplus_{k=1}^{j-1} (w_k, \tau_{w_k})}(\tau_{w_j}),$$

we see that

$$\begin{aligned}
& \mathbb{P}[\text{BOUNDED-SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell, k+1) = i] \\
&= p_v^i + \sum_{\tau \in \Omega_{S_\ell(v) \setminus \Lambda}} (\mu_v^{\sigma \oplus \tau}(i) - p_v^i) \prod_{j=1}^m \mu_{w_j}^{(\Lambda, \sigma) \oplus \bigoplus_{k=1}^{j-1} (w_k, \tau_{w_k})}(\tau_{w_j}) \\
&= \sum_{\tau \in \Omega_{S_\ell(v) \setminus \Lambda}} \mu_v^{\sigma \oplus \tau}(i) \prod_{j=1}^m \mu_{w_j}^{(\Lambda, \sigma) \oplus \bigoplus_{k=1}^{j-1} (w_k, \tau_{w_k})}(\tau_{w_j}) \\
&= \mu_v^\sigma(i).
\end{aligned}$$

□

Now we use this lemma to bridge the gap between the algorithm and the true probability.

Proof of Theorem 5.3. Let ν_h and ν be the distributions of $\text{BOUNDED-SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell, h)$ and $\text{SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell)$ respectively. Recall that

$$d_{\text{TV}}(\nu_h, \nu) = \inf_{\substack{X \sim \nu, \\ Y \sim \nu_h}} \mathbb{P}[X \neq Y],$$

where the random variables X and Y are on a joint sample space. While the depth of recursion remains below h , we can couple the random choices made by the bounded and unbounded versions of the algorithm. Thus, using T_A and T_B from the run-time argument,

$$\begin{aligned} d_{\text{TV}}(\nu, \nu_h) &\leq \mathbb{P}[\text{height}(T_A) \geq h] \\ &\leq \mathbb{P}[\text{height}(T_B) \geq h] \\ &\xrightarrow{h \rightarrow \infty} 0. \end{aligned}$$

Now since ν_h is also the correct distribution, we have that for all h

$$|\mathbb{P}[\text{SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell) = i] - \mu_v^\sigma(i)| \leq d_{\text{TV}}(\nu, \nu_h)$$

which implies $\mathbb{P}[\text{SSMS}(\mathcal{S}, (\Lambda, \sigma), v, \ell) = i] = \mu_v^\sigma(i)$, as required. $\square$

6 Examples

6.1 The Hardcore model

In Section 3.1 we informally analysed the radius 1 version of the algorithm. We return to this case as an example application of Theorem 5.3. Setting $\ell = 1$, we have $f(1) = \lambda/(1+\lambda)$, $g(1) = \Delta$ and $q = 2$. Then the algorithm has constant expected run-time provided $q f(\ell) g(\ell) < 1$, which holds when $\lambda < 1/(2\Delta - 1)$. This is worse than the condition that arose in our informal analysis, owing to some slack in Lemma 4.2. We remark that for non-constant Δ , the radius 1 algorithm remains efficient for small enough λ , evading the requirement of constant Δ .

Now we use Theorem 5.2 to find significantly more flexibility. On general graphs of maximum degree Δ we know from Weitz [Wei06] that the hardcore model exhibits strong spatial mixing whenever

$$\lambda < \lambda_c := \frac{(\Delta - 1)^{\Delta-1}}{(\Delta - 2)^\Delta}.$$

Let $\mathcal{G}$ be a class of graphs of maximum degree Δ and subexponential growth. It follows from Theorem 5.2 that for all $\lambda < \lambda_c$ and large enough ℓ (depending on λ) the expected run-time of the algorithm is constant per assigned spin. It is worth noting that for the case of a general graph of constant maximum degree Δ it was recently shown, by Chen, Liu and Vigoda [CLV20a], that Glauber dynamics mixes in time $O(n \log n)$ when $\lambda < \lambda_c$. Thus, for general graphs our algorithm is efficient in the same regime as the Glauber dynamics, but we do require the additional condition of subexponential growth. On the plus side, our algorithm is linear in the number of assigned spins, works for infinite graphs, and produces a perfect sample. Spinka [Spi20, Cor. 1.6] also shows how to produce perfect samples when

$\lambda < \lambda_c$, though his proofs are restricted to the square lattices $\mathbb{Z}^d$. (It could well be the case that Spinka’s approach works in a more general setting.)

For a more concrete example, consider the integer lattice $\mathbb{Z}^2$. In the radius 1 case we have efficiency for $\lambda < \frac{1}{3}$, but by leveraging strong spatial mixing our range increases to at least $\lambda < 2.538$, a bound that was derived by Sinclair, Srivastava, Štefankovič and Yin [SSvY17]. In fact, the true threshold of strong spatial mixing on $\mathbb{Z}^2$ is believed to be even higher: approximately 3.7962 [BET80; VVY15]. Note that the graphs in this example have subexponential (indeed, polynomial) growth. In the case of $\mathbb{Z}^2$, Spinka’s approach [Spi20] has the same range of validity as ours.

6.2 The Monomer-Dimer model

Another widely studied model which can be viewed as a spin system is the *monomer-dimer* model. When we restrict ourselves to a finite graph $G = (V, E)$, the monomer-dimer model with dimer activity γ is the distribution over matchings on the graph with the probability of a matching M being proportional to $\gamma^{|M|}$. To view this model as a spin system we simply consider the line-graph of G .

Definition 6.1. *Let $G = (V, E)$ be a finite graph. The line-graph $L(G)$ is the graph with one vertex corresponding to each edge in G and an edge between any $e, e' \in E$ which are incident.*

The monomer-model can then be represented as the hardcore model on $L(G)$ with $\lambda = \gamma$. The monomer-dimer model exhibits strong spatial mixing when G has sub-exponential neighbourhood growth — a fact essentially established by van den Berg and Brouwer [BB00] and extended by Bayati et al. [BGK+07] — so returning to Theorem 5.2 we see that SSMS allows us to sample partial configurations of the monomer-dimer model at any fixed λ in linear time, at least for graphs of subexponential growth.

Under the same conditions, there already exists a sampling algorithm, based on Glauber dynamics, that runs in polynomial time [JS89]. Indeed, recent work of Chen et al. [CLV20a] shows that Glauber dynamics mixes in time $O(n \log n)$. The advantage of the current algorithm is that it is exact and we can sample just a partial configuration on the graph.

6.3 The q -colour Model

Another well-known spin system model is the q -colour model where each spin corresponds to a colour. Here we define for every $i, j \in [q]$

$$b(i) = 1, \quad A(i, j) = \mathbf{1}_{[i \neq j]},$$

i.e., a colour cannot be adjacent to itself but every other configuration is equally weighted. By Theorem 5.2, whenever we have a graph G (more precisely, a class of graphs) with sub-exponential neighbourhood growth on which the q -colour model exhibits strong spatial mixing, SSMS will allow us to sample partial configurations of the q -colour model on the graph in linear time.

One such regime is revealed by Goldberg, Martin and Paterson [GMP05]. In addition to our neighbourhood growth condition, if G is triangle-free with maximum degree $\Delta \geq 3$

then we have strong spatial mixing when

$$q > \alpha\Delta - \gamma \approx 1.76\Delta - 0.47,$$

where $\alpha^\alpha = e$ and $\gamma = \frac{4\alpha^3 - 6\alpha^2 - 3\alpha + 4}{2(\alpha^2 - 1)}$.

Another strong spatial mixing regime is described by Efthymiou et al. [EGH+19]. Here, for the $(d + 1)$ -regular tree, strong spatial mixing holds when $q > 1.59d$. This result raises hope that the dependence of q on Δ can be improved in due course.

6.4 The ferromagnetic Ising model on the 2-dimensional square lattice

The ferromagnetic Ising model (in the absence of an external field) is a spin model characterised by the following interactions:

$$b = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad A = \begin{pmatrix} \lambda & 1 \\ 1 & \lambda \end{pmatrix},$$

with $\lambda \geq 1$. On the 2-dimensional square lattice, it is known that this model undergoes a phase transition at $\lambda_c = 1 + \sqrt{2}$. When $\lambda < \lambda_c$ there is a unique Gibbs measure, while for $\lambda > \lambda_c$ there are two. For general graphs with degree bound Δ we know — see Theorem 1 of Zhang, Liang and Bai [ZLB11] or Theorem 1.1 of Shao and Sun [SS21] — that strong spatial mixing holds for $\lambda < \Delta/(\Delta - 2)$, which in our case equates to $\lambda < 2$. It is reasonable to expect that strong spatial mixing holds up to λ_c , but this has not been proved. Currently, the best we can say is that approach works when $\lambda < 2$.

Martinelli, Olivieri and Schonmann [MOS94] show how to bootstrap weak spatial mixing to a form of strong spatial mixing in 2-dimensional systems. However, their strong mixing condition is restricted to rectangular regions $V \setminus W$. (Refer to Definition 4.1.)

Note added in proof: In addition to these results, in a recent paper Ding, Song, and Sun [DSS22] show that strong spatial mixing extends further on the Ising model. Fixing d and letting β_c be the critical inverse temperature of the Ising model on $\mathbb{Z}^d$, they show that for any finite graph $G \subset \mathbb{Z}^d$ and any $\lambda \in [0, \beta_c)$, the spin system exhibits strong spatial mixing. Due to this result and Theorem 5.1, for any d , $\lambda \in [0, \beta_c)$, and any $G \subset \mathbb{Z}^d$, ssms will allow us to sample partial configurations of the Ising model with inverse temperature λ on G in linear time.

Note that efficient perfect samplers are known which do work up to λ_c , the first being due to van den Berg and Steif [BS99] using an algorithm inspired by Propp and Wilson’s Coupling From The Past (CFTP). Spinka [Spi20] also covers this situation in the context of a more general treatment.

7 Future Work

While strong spatial mixing is sufficient for efficiency of the algorithm, we have not proven that it is necessary. Indeed, in the analysis there remains the possibility that strong spatial mixing is overkill for ensuring efficiency—strong spatial mixing does not tell us

about the average proportion of bad events in the algorithm, but just gives us the worst case proportion of bad events.

The question then is in what ways we can relax this condition. A few possible directions include

- Is there a condition which exactly captures this average case of bad events?
- Can we modify the algorithm so that it requires only the relaxed definition of strong spatial mixing used by van den Berg and Steif? If so, our approach would apply quite generally to 2-dimensional spin systems in the uniqueness phase.
- What do conditions on properties like the connective constant tell us about the behaviour of the algorithm?
- What properties of a spin system do we require in general for linear time algorithms?

References

- [ALGV18] N. Anari, K. Liu, S. O. Gharan, and C. Vinzant, “Log-concave polynomials II: high-dimensional walks and an FPRAS for counting bases of a matroid,” *CoRR*, vol. abs/1811.01816, 2018. arXiv: 1811 . 01816. [Online]. Available: <http://arxiv.org/abs/1811.01816>.
- [BB00] J. van den Berg and R. Brouwer, “Random sampling for the monomer-dimer model on a lattice,” in 3, vol. 41, Probabilistic techniques in equilibrium and nonequilibrium statistical physics, 2000, pp. 1585–1597. DOI: 10 . 1063 / 1 . 533198.
- [BET80] R. J. Baxter, I. G. Enting, and S. K. Tsang, “Hard-square lattice gas,” *J. Statist. Phys.*, vol. 22, no. 4, pp. 465–489, 1980, ISSN: 0022-4715. DOI: 10 . 1007/BF01012867.
- [BGK+07] M. Bayati, D. Gamarnik, D. Katz, C. Nair, and P. Tetali, “Simple deterministic approximation algorithms for counting matchings,” in *STOC’07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, ACM, New York, 2007, pp. 122–127. DOI: 10 . 1145/1250790 . 1250809.
- [BS17] A. Barvinok and P. Soberón, “Computing the partition function for graph homomorphisms,” *Combinatorica*, vol. 37, no. 4, pp. 633–650, 2017, ISSN: 0209-9683. DOI: 10 . 1007/s00493-016-3357-2.
- [BS99] J. van den Berg and J. E. Steif, “On the existence and nonexistence of finitary codings for a class of random fields,” *Ann. Probab.*, vol. 27, no. 3, pp. 1501–1522, 1999, ISSN: 0091-1798. DOI: 10 . 1214/aop/1022677456.
- [CLV20a] Z. Chen, K. Liu, and E. Vigoda, “Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion,” *ArXiv*, vol. abs/2011.02075, 2020.
- [CLV20b] —, “Rapid mixing of Glauber dynamics up to uniqueness via contraction,” *arXiv preprint arXiv:2004.09083*, 2020.

- [DSS22] J. Ding, J. Song, and R. Sun, “A new correlation inequality for ising models with external fields,” *Probability Theory and Related Fields*, pp. 1–16, 2022. DOI: 10.1007/s00440-022-01132-1.
- [DSVW04] M. Dyer, A. Sinclair, E. Vigoda, and D. Weitz, “Mixing in time and space for lattice spin systems: A combinatorial view,” *Random Structures Algorithms*, vol. 24, no. 4, pp. 461–479, 2004, ISSN: 1042-9832. DOI: 10.1002/rsa.20004. [Online]. Available: <https://doi-org.ezproxy.library.qmul.ac.uk/10.1002/rsa.20004>.
- [EGH+19] C. Efthymiou, A. Galanis, T. P. Hayes, D. Štefankovič, and E. Vigoda, “Improved strong spatial mixing for colorings on trees,” in *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*, ser. LIPIcs. Leibniz Int. Proc. Inform. Vol. 145, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019, Art. No. 48, 16.
- [FGY19] W. Feng, H. Guo, and Y. Yin, “Perfect sampling from spatial mixing,” *ArXiv*, vol. abs/1907.06033, 2019.
- [FH00] J. A. Fill and M. Huber, “The randomness recycler: A new technique for perfect sampling,” in *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, IEEE Comput. Soc. Press, Los Alamitos, CA, 2000, pp. 503–511. DOI: 10.1109/SFCS.2000.892138.
- [FV18] S. Friedli and Y. Velenik, *Statistical mechanics of lattice systems*. Cambridge University Press, Cambridge, 2018, pp. xix+622, A concrete mathematical introduction, ISBN: 978-1-107-18482-4.
- [GJL19] H. Guo, M. Jerrum, and J. Liu, “Uniform sampling through the Lovász local lemma,” *J. ACM*, vol. 66, no. 3, Art. 18, 31, 2019, ISSN: 0004-5411. DOI: 10.1145/3310131.
- [Gla63] R. J. Glauber, “Time-dependent statistics of the Ising model,” *J. Mathematical Phys.*, vol. 4, pp. 294–307, 1963, ISSN: 0022-2488. DOI: 10.1063/1.1703954.
- [GMP05] L. A. Goldberg, R. Martin, and M. Paterson, “Strong spatial mixing with fewer colors for lattice graphs,” *SIAM J. Comput.*, vol. 35, no. 2, pp. 486–517, 2005, ISSN: 0097-5397. DOI: 10.1137/S0097539704445470.
- [Hub16] M. L. Huber, *Perfect simulation*, ser. Monographs on Statistics and Applied Probability. CRC Press, Boca Raton, FL, 2016, vol. 148, pp. xxii+228, ISBN: 978-1-4822-3244-8.
- [JS89] M. Jerrum and A. Sinclair, “Approximating the permanent,” *SIAM J. Comput.*, vol. 18, no. 6, pp. 1149–1178, 1989, ISSN: 0097-5397. DOI: 10.1137/0218077.
- [KO20] T. Kaufman and I. Oppenheim, “High order random walks: Beyond spectral gap,” *Comb.*, vol. 40, no. 2, pp. 245–281, 2020. DOI: 10.1007/s00493-019-3847-0.

- [Liu21] K. Liu, “From coupling to spectral independence and blackbox comparison with the down-up walk,” in *Approximation, randomization, and combinatorial optimization. algorithms and techniques*, ser. LIPIcs. Leibniz Int. Proc. Inform. Vol. 207, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2021, Art. No. 32, 21.
- [MOS94] F. Martinelli, E. Olivieri, and R. H. Schonmann, “For 2-D lattice spin systems weak mixing implies strong mixing,” *Comm. Math. Phys.*, vol. 165, no. 1, pp. 33–47, 1994, ISSN: 0010-3616.
- [PR17] V. Patel and G. Regts, “Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials,” *SIAM J. Comput.*, vol. 46, no. 6, pp. 1893–1919, 2017, ISSN: 0097-5397. DOI: 10.1137/16M1101003.
- [PW98] J. G. Propp and D. B. Wilson, “How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph,” in 2, vol. 27, 7th Annual ACM-SIAM Symposium on Discrete Algorithms (Atlanta, GA, 1996), 1998, pp. 170–217. DOI: 10.1006/jagm.1997.0917.
- [Spi20] Y. Spinka, “Finitary codings for spatial mixing Markov random fields,” *Ann. Probab.*, vol. 48, no. 3, pp. 1557–1591, 2020, ISSN: 0091-1798. DOI: 10.1214/19-AOP1405.
- [SS21] S. Shao and Y. Sun, “Contraction: A unified perspective of correlation decay and zero-freeness of 2-spin systems,” *J. Stat. Phys.*, vol. 185, no. 2, Paper No. 12, 25, 2021, ISSN: 0022-4715. DOI: 10.1007/s10955-021-02831-0. [Online]. Available: <https://doi.org/10.1007/s10955-021-02831-0>.
- [SSvY17] A. Sinclair, P. Srivastava, D. Štefankovič, and Y. Yin, “Spatial mixing and the connective constant: Optimal bounds,” *Probab. Theory Related Fields*, vol. 168, no. 1-2, pp. 153–197, 2017, ISSN: 0178-8051. DOI: 10.1007/s00440-016-0708-2.
- [SSY13] A. Sinclair, P. Srivastava, and Y. Yin, “Spatial mixing and approximation algorithms for graphs with bounded connective constant (extended abstract),” in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science—FOCS 2013*, IEEE Computer Soc., Los Alamitos, CA, 2013, pp. 300–309. DOI: 10.1109/FOCS.2013.40.
- [VVY15] J. C. Vera, E. Vigoda, and L. Yang, “Improved bounds on the phase transition for the hard-core model in 2 dimensions,” *SIAM J. Discrete Math.*, vol. 29, no. 4, pp. 1895–1915, 2015, ISSN: 0895-4801. DOI: 10.1137/140976923.
- [Wei06] D. Weitz, “Counting independent sets up to the tree threshold,” in *STOC’06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, ACM, New York, 2006, pp. 140–149. DOI: 10.1145/1132516.1132538.
- [ZLB11] J. Zhang, H. Liang, and F. Bai, “Approximating partition functions of the two-state spin system,” *Inform. Process. Lett.*, vol. 111, no. 14, pp. 702–710, 2011, ISSN: 0020-0190. DOI: 10.1016/j.ipl.2011.04.012. [Online]. Available: <https://doi.org/10.1016/j.ipl.2011.04.012>.