

PERFECT SAMPLING FROM SPATIAL MIXING

WEIMING FENG, HENG GUO, AND YITONG YIN

ABSTRACT. We introduce a new perfect sampling technique that can be applied to general Gibbs distributions and runs in linear time if the correlation decays faster than the neighborhood growth. In particular, in graphs with sub-exponential neighborhood growth like \mathbb{Z}^d , our algorithm achieves linear running time as long as Gibbs sampling is rapidly mixing. As concrete applications, we obtain the currently best perfect samplers for colorings and for monomer-dimer models in such graphs.

1. INTRODUCTION

Spin systems model nearest neighbor interactions of complex systems. These models originated from statistical physics, and have found a wide range of applications in probability theory, machine learning, and theoretical computer science, often under different names such as *Markov random fields* or *Boltzmann machines*. Given an underlying graph $G = (V, E)$, a *configuration* σ is an assignment from vertices to a finite set of spins, usually denoted by $[q]$. The *weight* of a configuration is specified by the q -dimensional vector b_v assigned to each vertex $v \in V$ and the q -by- q symmetric interaction matrix A_e assigned to each edge $e \in E$, namely,

$$(1) \quad w(\sigma) = \prod_{v \in V} b_v(\sigma_v) \prod_{e = \{u, v\} \in E} A_e(\sigma_u, \sigma_v).$$

The equilibrium state of the system is described by the *Gibbs distribution* μ , where the probability of a configuration is proportional to its weight.

A central algorithmic problem related to spin systems is to sample from the Gibbs distribution. A canonical Markov chain for sampling approximately from the Gibbs distribution is the *Gibbs sampler* (a.k.a. *heat bath* or *Glauber dynamics*). The efficiency of the sampler is determined by its *mixing time*, namely how long it takes to converge to the desired distribution. One (conjectured) general criterion for rapid mixing of such Markov chains is the *spatial mixing* property [45], which roughly states that the correlation among variables decays rapidly in the system as their distances increase. It is widely believed that spatial mixing (in some form) implies the rapid mixing of the Gibbs sampler. However, rigorous implications have only been established for special classes of graphs or systems, such as for lattice graphs [33, 11], neighborhood amenable graphs [19], or for monotone systems [36], and very recently, for anti-ferromagnetic 2-spin systems [3, 7].

One main drawback for Gibbs samplers or Markov chains in general, is that one needs to know the mixing time in advance to implement the algorithms with provably small errors. The mixing time is usually hard to analyze and very pessimistic. The so-called *perfect samplers* are thus more desirable, which run in a Las Vegas fashion and return exact samples upon halting. There have been a number of techniques available to design perfect samplers, such as Coupling From The Past (CFTP) [37], including the monotone & anti-monotone CFTP [37, 25] and the bounding chains [27, 5], Randomness Recycler (RR) [16], and Partial Rejection Sampling (PRS) [24, 14]. Nevertheless, none of these techniques addresses general spin systems or relates to the important spatial mixing properties of the system.

In this paper, we introduce a new technique to perfectly sample from Gibbs distributions of spin systems. The correctness of our algorithm relies on only the conditional independence property of Gibbs distributions. Moreover, the expected running time is linear in the size of the system, when the correlation decays more rapidly than the growth of the neighborhood.

(Weiming Feng, Yitong Yin) STATE KEY LABORATORY FOR NOVEL SOFTWARE TECHNOLOGY, NANJING UNIVERSITY, 163 XIANLIN AVENUE, NANJING, JIANGSU PROVINCE, CHINA. E-mail: fengwm@smail.nju.edu.cn, yinyt@nju.edu.cn

(Heng Guo) SCHOOL OF INFORMATICS, UNIVERSITY OF EDINBURGH, INFORMATICS FORUM, EDINBURGH, EH8 9AB, UNITED KINGDOM. E-mail: hguo@inf.ed.ac.uk

Theorem 1.1 (informal). *For any spin system with bounded maximum degree,¹ if strong spatial mixing holds with a rate faster than the neighborhood growth of the underlying graph, then there exists a perfect sampler with running time $O(n)$ in expectation, where n represents the number of vertices of the graph.*

More details and undefined terms are explained in Section 2. Formal statements of our results are given in Theorem 2.4 for spin systems on sub-exponential neighborhood growth graphs, and in Theorem 2.6 for spin systems on general graphs. Applications on list colorings and on monomer-dimer models are given in Theorem 2.9 and Theorem 2.10.

Lattice graphs, such as \mathbb{Z}^d , are of special interests in statistical physics and combinatorics. These graphs have sub-exponential neighborhood growth, which implies that temporal mixing is equivalent to spatial mixing on them [11]. Therefore our sampler runs in linear time as long as the standard Glauber dynamics has $O(n \log n)$ mixing time. This is a direct strengthening of aforementioned results [33, 11, 19] from approximate to perfect sampling, with an improved running time.

Corollary 1.2 (informal). *For spin systems on graphs with sub-exponential neighborhood growth, if the Gibbs sampler has $O(n \log n)$ mixing time, where n is the number of vertices, or the system shows strong spatial mixing, then there exists a perfect sampler with running time $O(n)$ in expectation.*

It is worth noting that many traditional perfect sampling algorithms, especially those rely on CFTP [37, 27, 5], suffer from “non-interruptibility”. That is, early termination of the algorithm induces a bias on the sample. In contrast, our algorithm is interruptible in the following sense: conditioned on its termination at any particular step, the algorithm guarantees to return a correct sample. Therefore, had the algorithm been running for too long, one can simply stop it and restart. One can also run many independent copies in parallel and output the earliest returned sample without biasing the sample.

In addition, our algorithm can be used to solve the recently introduced dynamic sampling problem [14, 13], where the Gibbs distribution itself changes dynamically and the algorithm needs to efficiently maintain a sample from the current Gibbs distribution. The detail of this part is given in Section 7. Our perfect sampler also generalizes straightforwardly to Gibbs distributions with multi-body interactions (namely spin systems on hypergraphs / constraint satisfaction problems), and similar efficiency can be achieved when some appropriate variant of spatial mixing holds.

1.1. Algorithm overview. We give an overview of our main new idea. We briefly review the *Gibbs sampler*, which is a Markov chain on the state space $[q]^V$. At each step, a vertex $u \in V$ is picked uniformly at random and the current configuration $X \in [q]^V$ is updated by the simple rule:

- the spin X_u is redrawn according to the marginal distribution $\mu_u^{X_{\Gamma(u)}}$;

where $\Gamma(u) \triangleq \{v \in V \mid \{u, v\} \in E\}$ denotes the neighborhood of u , and $\mu_u^{X_{\Gamma(u)}}$ denotes the marginal distribution induced at vertex u by the Gibbs distribution μ , conditioned on the current spins of the neighborhood $\Gamma(u)$. It is a basic fact that this chain converges to the Gibbs distribution μ .

Our perfect sampler makes use of the same update rule. To expose our main idea, we first consider the single-site version of our perfect sampler, which works for systems with soft constraints (where all A_e and b_v are positive). The sampling algorithm is quite simple, described in Algorithm 1.

The algorithm starts from an arbitrary initial configuration $X \in [q]^V$, and gradually “repairs” X to a perfect sample of the Gibbs distribution μ . We maintain a set $\mathcal{R} \subseteq V$ of vertices that are currently “incorrect”, initially set as $\mathcal{R} = V$. At each step, a random vertex u is picked from \mathcal{R} , and we try to remove u from \mathcal{R} while maintaining the following invariant, denoted by the *conditional Gibbs property*:

- (★) $X_{\overline{\mathcal{R}}}$ always follows the law $\mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}}}$,
which is the marginal distribution induced by μ on $\overline{\mathcal{R}} \triangleq V \setminus \mathcal{R}$ conditioned on $X_{\mathcal{R}}$.

¹We remark that our algorithm remains in polynomial-time (but not in linear time) for graphs with unbounded degrees, as long as the degree does not grow too quickly. This requirement on the degree comes from the cost of updating a block of certain radius, similar to the cost of block dynamics, and the exact upper bound needed varies from problem to problem. See Theorem 2.6 and the discussion thereafter. For the most part, we state our results for bounded degree cases to keep the statements clean.

Algorithm 1: Perfect Gibbs sampler (single-site version)

```

1 Start from an arbitrary initial configuration  $X \in [q]^V$  and  $\mathcal{R} \leftarrow V$ ;
2 while  $\mathcal{R} \neq \emptyset$  do
3   pick a  $u \in \mathcal{R}$  uniformly at random;
4   let  $\mu_{\min}$  be the minimum value of  $\mu_u^\sigma(X_u)$  over all  $\sigma \in [q]^{\Gamma(u)}$  that  $\sigma_{\mathcal{R} \cap \Gamma(u)} = X_{\mathcal{R} \cap \Gamma(u)}$ ;
5   with probability  $\mu_{\min} / \mu_u^{X_{\Gamma(u)}}(X_u)$  do                                 $\triangleright$  Bayes filter
6     update  $X$  by redrawing  $X_u \sim \mu_u^{X_{\Gamma(u)}}$ ;                                 $\triangleright$  Gibbs sampler update
7      $\mathcal{R} \leftarrow \mathcal{R} \setminus \{u\}$ ;
8   else
9      $\mathcal{R} \leftarrow \mathcal{R} \cup \Gamma(u)$ ;
10 return  $X$ ;

```

This property ensures that the configuration on $\overline{\mathcal{R}}$ follows the correct distribution conditioned on the configuration on \mathcal{R} . In particular, when $\mathcal{R} = \emptyset$, $\overline{\mathcal{R}} = V$, $\mu_{\overline{\mathcal{R}}}^{X_{\overline{\mathcal{R}}}} = \mu$, and the sample X follows precisely the distribution μ . This is the goal of our algorithm: reduce \mathcal{R} to the empty set.

One first attempt would be to simply update the spin X_u according to its marginal distribution $\mu_u^{X_{\Gamma(u)}}$ as in the Gibbs sampler and then remove u from \mathcal{R} . This gives the transition $(X, \mathcal{R}) \rightarrow (X', \mathcal{R}')$, where $\mathcal{R}' = \mathcal{R} \setminus \{u\}$, and $X = X'$ except at vertex u , where $X'_u \sim \mu_u^{X_{\Gamma(u)}}$.

Ideally, if one had $X_{\overline{\mathcal{R}}} \sim \mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}'}}$, then after the spin of u being redrawn as $X'_u \sim \mu_u^{X_{\Gamma(u)}}$, the $X_{\overline{\mathcal{R}}}$ would be extended to a $X'_{\overline{\mathcal{R}}} \sim \mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}'}} \equiv \mu_{\overline{\mathcal{R}'}}^{X_{\mathcal{R}'}}$ as $X_{\mathcal{R}'} = X'_{\mathcal{R}'}$, satisfying the invariant (\star) on the new pair (X', \mathcal{R}') .

However, the invariant (\star) on the current (X, \mathcal{R}) only guarantees $X_{\overline{\mathcal{R}}} \sim \mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}}}$ rather than $X_{\overline{\mathcal{R}}} \sim \mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}'}}$. To remedy this, we construct a filter $\mathcal{F} = \mathcal{F}(u, X)$ that corrects $\mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}}}$ to $\mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}'}}$. We call \mathcal{F} the *Bayes filter*. Specifically, \mathcal{F} is determined by a biased coin depending on only part of X so that

$$(2) \quad \Pr[\mathcal{F} \text{ succeeds}] \propto \frac{\mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}'}}(X_{\overline{\mathcal{R}}})}{\mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}}}(X_{\overline{\mathcal{R}}})} = \frac{\mu_u^{X_{\mathcal{R}'}}(X_u)}{\mu_u^{X_{\Gamma(u)}}(X_u)},$$

where \propto is taken over all $X_{\overline{\mathcal{R}}}$ and the equality is due to Bayes' theorem:

$$\mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}}}(X_{\overline{\mathcal{R}}}) = \mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}' \wedge X_u}}(X_{\overline{\mathcal{R}}}) = \mu_u^{X_{\mathcal{R}' \wedge X_{\overline{\mathcal{R}}}}}(X_u) \cdot \mu_{\overline{\mathcal{R}}}^{X_{\mathcal{R}'}}(X_{\overline{\mathcal{R}}}) / \mu_u^{X_{\mathcal{R}'}}(X_u),$$

together with the fact that $\mu_u^{X_{\mathcal{R}' \wedge X_{\overline{\mathcal{R}}}}}(X_u) = \mu_u^{X_{\Gamma(u)}}(X_u)$, guaranteed by the *conditional independence property* (formally, Property 4.3) of the Gibbs distribution.

We observe that despite that the exact value of the marginal probability $\mu_u^{X_{\mathcal{R}'}}(X_u)$ in (2) is hard to compute, it does not depend on $X_{\overline{\mathcal{R}}}$. Therefore, $\mu_u^{X_{\mathcal{R}'}}(X_u)$ can be treated as a constant and (2) holds as long as $\Pr[\mathcal{F} \text{ succeeds}] \propto 1 / \mu_u^{X_{\Gamma(u)}}(X_u)$, which is satisfied precisely by our Bayes filter in Algorithm 1.

Now by (2), conditioned on the success of the filter \mathcal{F} , the new (X', \mathcal{R}') satisfies the invariant (\star) . Meanwhile, since the filter only reveals the neighborhood spins $X_{\Gamma(u)}$, upon the failure of \mathcal{F} , the invariant (\star) remains to hold as long as the revealed sites $\Gamma(u)$ are added into \mathcal{R} and X is unchanged.

This sampler is valid for general Gibbs distributions, since the only property we require is conditional independence. However for efficiency purposes, our general algorithm, Algorithm 2, uses *block* updates. A suitable block radius (chosen according to spatial mixing) is the key to efficiency in our analysis. Moreover, to make sure that marginal distributions are well-defined, we restrict our attention to *permissive* systems, which contain all soft constraint systems as well as all hard constraint systems of interest. The details are given in Section 2 and 3.

The algorithm is efficient as long as the size of \mathcal{R} shrinks in expectation in every step. For the more general Algorithm 2, this holds true when the correlation decays faster than the neighborhood growth. The details are in Section 5.

1.2. Related work. The conditional Gibbs property has been used implicitly or explicitly in previous works such as partial rejection sampling [24, 23, 22, 21, 20], dynamic sampling [14], and randomness recycler [16]. Furthermore, the invariance of the conditional Gibbs property ensures that the sampling algorithm is correct even when the input spin system is dynamically changing over time [14].

Before our work, all previous resampling algorithms [16, 24, 23, 22, 21, 20, 14] fall into the paradigm of rejection sampling: a new sample is generated, usually from modifying the old sample, and (part of) the new sample is rejected independently with some probability determined by the new sample. In our algorithm, the filtration is executed *before* the generation of the new sample, with a bias independent of the new sample.

Spatial mixing properties were known to imply rapid mixing of the Gibbs sampler for some particular systems. See for example [36]. For anti-ferromagnetic 2-spin systems, weak spatial mixing corresponds to the optimal threshold for efficient samplability [7, 3, 12, 17, 39, 38, 31, 45]. It remains an interesting open problem whether spatial mixing implies the existence of efficient samplers in general, and whether these samplers can be perfect.

2. OUR RESULTS

2.1. Model and definitions. Let $G = (V, E)$ be an undirected graph, and $[q] = \{1, 2, \dots, q\}$ a finite domain of $q \geq 2$ spins. An instance of q -state spin system is specified by a tuple $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$, where $\mathbf{b} = (b_v)_{v \in V}$ assigns every vertex $v \in V$ a vector $b_v \in \mathbb{R}_{\geq 0}^q$ and $\mathbf{A} = (A_e)_{e \in E}$ assigns every edge $e \in E$ a symmetric matrix $A_e \in \mathbb{R}_{\geq 0}^{q \times q}$. The *Gibbs distribution* $\mu_{\mathcal{I}}$ over $[q]^V$ is defined as

$$(3) \quad \forall \sigma \in [q]^V : \quad \mu_{\mathcal{I}}(\sigma) \triangleq \frac{w_{\mathcal{I}}(\sigma)}{Z_{\mathcal{I}}} = \frac{1}{Z_{\mathcal{I}}} \prod_{v \in V} b_v(\sigma_v) \prod_{e=\{u,v\} \in E} A_e(\sigma_u, \sigma_v),$$

where $w_{\mathcal{I}}(\sigma)$ is the *weight* defined in (1) and $Z_{\mathcal{I}} \triangleq \sum_{\sigma \in [q]^V} w_{\mathcal{I}}(\sigma)$ is the *partition function*.

We restrict our attention to the so-called “permissive” spin systems, where the marginal distributions are always well-defined. Let $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$ be an instance of spin system. A configuration on V is called *feasible* if its weight is positive, and a partial configuration is *feasible* if it can be extended to a feasible configuration. For any (possibly empty) subset $\Lambda \subseteq V$ and any (not necessarily feasible) partial configuration $\sigma \in [q]^{\Lambda}$, we use $w_{\mathcal{I}}^{\sigma}(\tau)$ to denote the weight of $\tau \in [q]^{V \setminus \Lambda}$ conditional on σ :

$$(4) \quad w_{\mathcal{I}}^{\sigma}(\tau) = \prod_{v \in V \setminus \Lambda} b_v(\tau_v) \prod_{\substack{e=\{u,v\} \in E \\ u,v \in V \setminus \Lambda}} A_e(\tau_u, \tau_v) \prod_{\substack{e=\{u,v\} \in E \\ u \in \Lambda, v \in V \setminus \Lambda}} A_e(\sigma_u, \tau_v).$$

Define the partition function $Z_{\mathcal{I}}^{\sigma}$ conditional on σ as $Z_{\mathcal{I}}^{\sigma} \triangleq \sum_{\tau \in [q]^{V \setminus \Lambda}} w_{\mathcal{I}}^{\sigma}(\tau)$.

Definition 2.1 (permissive). A spin system $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$, where $G = (V, E)$, is called *permissive* if $Z_{\mathcal{I}}^{\sigma} > 0$ for any partial configuration $\sigma \in [q]^{\Lambda}$ specified on any subset $\Lambda \subseteq V$.

Permissive systems are very common, including, for examples, uniform proper q -coloring when $q \geq \Delta + 1$, where Δ is the maximum degree, and spin systems with soft constraints, e.g. the Ising model, or with a “permissive” state that is compatible with all other states, e.g. the hardcore model.

For permissive systems, a feasible configuration is always easy to construct by greedy algorithm. More importantly, with permissiveness, marginal probabilities are always well defined, which is crucial for Gibbs sampler and spatial mixing property.

Formally, we use $\mu_{\mathcal{I}}^{\sigma}$ to denote the conditional distribution over $[q]^{V \setminus \Lambda}$ given $\sigma \in [q]^{\Lambda}$, that is,

$$(5) \quad \forall \tau \in [q]^{V \setminus \Lambda}, \quad \mu_{\mathcal{I}}^{\sigma}(\tau) \triangleq \frac{w_{\mathcal{I}}^{\sigma}(\tau)}{Z_{\mathcal{I}}^{\sigma}}.$$

And for any $v \in V \setminus \Lambda$, we use $\mu_{v, \mathcal{I}}^{\sigma}$ to denote the marginal distribution at v projected from $\mu_{\mathcal{I}}^{\sigma}$.

For any $u, v \in V$, we use $\text{dist}_G(u, v)$ to denote the shortest-path distance between v and u in G .

Definition 2.2 (strong spatial mixing [45, 44]). Let $\delta : \mathbb{N} \rightarrow \mathbb{R}^+$. A class \mathfrak{S} of permissive spin systems is said to exhibit *strong spatial mixing* with rate $\delta(\cdot)$ if for every instance $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A}) \in \mathfrak{S}$, where

$G = (V, E)$, for every $v \in V$, $\Lambda \subseteq V$, and any two partial configurations $\sigma, \tau \in [q]^\Lambda$,

$$(6) \quad d_{\text{TV}} \left(\mu_{v, \mathcal{I}}^\sigma, \mu_{v, \mathcal{I}}^\tau \right) \leq \delta(\ell),$$

where $\ell = \min\{\text{dist}_G(v, u) \mid u \in \Lambda, \sigma_u \neq \tau_u\}$, and $d_{\text{TV}} \left(\mu_{v, \mathcal{I}}^\sigma, \mu_{v, \mathcal{I}}^\tau \right) \triangleq \frac{1}{2} \sum_{a \in [q]} \left| \mu_{v, \mathcal{I}}^\sigma(a) - \mu_{v, \mathcal{I}}^\tau(a) \right|$ denotes the total variation distance between $\mu_{v, \mathcal{I}}^\sigma$ and $\mu_{v, \mathcal{I}}^\tau$. In particular, \mathfrak{S} exhibits *strong spatial mixing with exponential decay* if (6) is satisfied for $\delta(\ell) = \alpha \exp(-\beta\ell)$ for some constants $\alpha, \beta > 0$.

Our first result holds for spin systems on graphs with bounded neighborhood growth.

Definition 2.3 (sub-exponential neighborhood growth). A class \mathfrak{G} of graphs is said to have *sub-exponential neighborhood growth* if there is a function $s : \mathbb{N} \rightarrow \mathbb{N}$ such that $s(\ell) = \exp(o(\ell))$ and for every graph $G = (V, E) \in \mathfrak{G}$,

$$\forall v \in V, \forall \ell \geq 0, \quad |S_\ell(v)| \leq s(\ell),$$

where $S_\ell(v) \triangleq \{u \in V \mid \text{dist}_G(v, u) = \ell\}$ denotes the sphere of radius ℓ centered at v in G .

Note that graphs with sub-exponential neighborhood growth necessarily have bounded maximum degree because we can set $\ell = 1$ and get $s(1) = O(1)$.

2.2. Main results. Our first result shows that for spin systems on graphs with sub-exponential neighborhood growth, strong spatial mixing implies the existence of linear-time perfect sampler.

Theorem 2.4 (main theorem: bounded-growth graphs). *Let $q > 1$ be a finite integer and \mathfrak{S} a class of permissive q -state spin systems on graphs with sub-exponential neighborhood growth. If \mathfrak{S} exhibits strong spatial mixing with exponential decay, then there exists an algorithm such that given any instance $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A}) \in \mathfrak{S}$, the algorithm outputs a perfect sample from $\mu_{\mathcal{I}}$ within $O(n)$ time in expectation, where n is the number of vertices in G .*

The factor in $O(\cdot)$ is $q^{O(\Delta^{\ell_0})}$, which is the cost for an update in $(\ell_0$ -block) Gibbs sampler, where $\ell_0 = O(1)$ is the radius at which the decay catches up with neighborhood growth, as formally given in (33).

It is already known that for spin systems on sub-exponential neighborhood growth graphs, the strong spatial mixing with exponential decay implies $O(n \log n)$ mixing time for block Gibbs sampler [11], which only generates approximate samples. We give a perfect sampler with $O(n)$ expected running time under the same condition. The notion of sub-exponential neighborhood growth is related to, but should be distinguished from neighborhood-amenability (see e.g. [19]), which says that in an infinite graph, for any constant real $c > 0$, there is an ℓ such that $\frac{|S_{\ell+1}(v)|}{|B_\ell(v)|} \leq c$ holds everywhere.

Our main result on general graphs assumes the following strong spatial mixing condition.

Condition 2.5. *Let $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$ be a permissive spin system where $G = (V, E)$. There is an integer $\ell = \ell(q) \geq 2$ such that the following holds: for every $v \in V$, $\Lambda \subseteq V$, for any two partial configurations $\sigma, \tau \in [q]^\Lambda$ satisfying $\min\{\text{dist}_G(v, u) \mid u \in \Lambda, \sigma(u) \neq \tau(u)\} = \ell$,*

$$(7) \quad d_{\text{TV}} \left(\mu_{v, \mathcal{I}}^\sigma, \mu_{v, \mathcal{I}}^\tau \right) \leq \frac{\gamma}{5|S_\ell(v)|},$$

where $S_\ell(v)$ denotes the sphere of radius ℓ centered at v in G , and

$$(8) \quad \gamma = \gamma(v, \Lambda) \triangleq \min \left\{ \mu_{v, \mathcal{I}}^\rho(a) \mid \rho \in [q]^\Lambda, a \in [q] \text{ that } \mu_{v, \mathcal{I}}^\rho(a) > 0 \right\}$$

denotes the lower bound of positive marginal probabilities at v .

The above condition basically says that the spin systems exhibit strong spatial mixing with a decay rate faster than that of neighborhood growth, given that the marginal probabilities are appropriately lower bounded (which holds with $\gamma = \Omega(1)$ when entries of \mathbf{A} and \mathbf{b} are of finite precision and the maximum degree Δ is finitely bounded). Our result on general graphs is stated as the following theorem.

Theorem 2.6 (main theorem: general graphs). *Let \mathfrak{S} be a class of permissive spin systems satisfying Condition 2.5. There exists an algorithm which given any instance $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A}) \in \mathfrak{S}$, outputs a perfect sample from $\mu_{\mathcal{I}}$ within $n \cdot q^{O(\Delta^\ell)}$ time in expectation, where n is the number of vertices in G , Δ is the maximum degree of G , and $\ell = \ell(q)$ is determined by Condition 2.5.*

The $q^{O(\Delta^\ell)}$ factor in the time cost is contributed by the block Gibbs sampler update on ℓ -radius blocks. This extra cost could remain polynomial if $q = \omega(1)$, but the upper bound on q for that will vary from problem to problem. See e.g. the discussion after Theorem 2.9. The conditions of both Theorem 2.4 and Theorem 2.6 are special cases of a more technical condition (Condition 5.1), formally stated in Section 5.

2.3. Applications on specific systems. Our results can be applied on various spin systems. We consider two important examples:

- **Uniform list coloring:** A list coloring instance is specified by $\mathcal{I} = (G, [q], \mathcal{L})$, where $\mathcal{L} \triangleq \{L_v \subseteq [q] \mid v \in V\}$ assigns each vertex $v \in V$ a list of colors $L_v \subseteq [q]$. A proper list coloring of instance \mathcal{I} is a $\sigma \in [q]^V$ that $\sigma_v \in L_v$ for all $v \in V$ and $\sigma_u \neq \sigma_v$ for all $\{u, v\} \in E$. Let $\mu_{\mathcal{I}}$ denote the uniform distribution over all proper list colorings of \mathcal{I} .
- **The monomer-dimer model:** The monomer-dimer model defines a distribution over graph matchings. An instance is specified by $\mathcal{I} = (G, \lambda)$, where $G = (V, E)$ is a graph and $\lambda > 0$. Each matching $M \subseteq E$ in G is assigned a weight $w_{\mathcal{I}}(M) = \lambda^{|M|}$. Let $\mu_{\mathcal{I}}$ be the distribution over all matchings in G such that $\mu_{\mathcal{I}}(M) \propto w_{\mathcal{I}}(M)$.

We use $\deg_G(v)$ to denote the degree of v in G and $\Delta = \max_{v \in V} \deg_G(v)$ the maximum degree.

First, for the list coloring problem, we define the following two conditions for instance $\mathcal{I} = (G, [q], \mathcal{L})$. Let $\deg_G(v)$ denote the degree of $v \in V$ in graph G , and $\Delta \triangleq \max_{v \in V} \deg_G(v)$ the maximum degree.

Condition 2.7. For every $v \in V$, $|L(v)| \geq \alpha \deg_G(v) + \beta$, where either one of the followings holds:

- $\alpha = 2$ and $\beta = 0$;
- G is triangle-free, $\alpha > \alpha^*$ where $\alpha^* = 1.763 \dots$ is the positive root of $x^x = e$, and $\beta \geq \frac{\sqrt{2}}{\sqrt{2}-1}$ satisfies $(1 - 1/\beta)\alpha e^{\frac{1}{\alpha}(1-1/\beta)} > 1$.

Condition 2.8. For every $v \in V$, $|L(v)| \geq \Delta^2 - \Delta + 2$.

Our perfect sampler for list coloring runs in linear time in either above condition.

Theorem 2.9. Let \mathcal{Q} be a class of list coloring instances with at most q colors for a finite $q > 0$. If either of the two followings holds for all instances $\mathcal{I} = (G, [q], \mathcal{L}) \in \mathcal{Q}$:

- Condition 2.7 and G has sub-exponential neighborhood growth; or
- Condition 2.8,

then there exists a perfect sampler for $\mu_{\mathcal{I}}$ that runs in expected $O(n)$ time, where n is the number of vertices.

The constant factor in $O(\cdot)$, can be determined in the same way as in Theorem 2.4 in the case of Condition 2.7, but is much higher in the case of Condition 2.8 due to the arbitrary neighborhood growth. Nevertheless, even in this costly case, $\ell = O(q^2 \log q)$ and the overall overhead can be upper bounded by a rough estimate $\exp(\exp(\text{poly}(q)))$.

Sampling proper q -colorings (where the lists are identical for all vertices) has been extensively studied, especially using Markov chains. Approximate sampling received considerable attention [28, 9, 26, 35, 10]. The current best result [43, 6] is the $O(n \log n)$ mixing time of the flip chain if $q \geq (\frac{11}{6} - \epsilon_0)\Delta$ for some constant $\epsilon_0 > 0$. For perfect sampling q -colorings, Huber introduced a bounding chain [27] based on CFTP [46], which terminates within $O(n \log n)$ steps in expectation if $q \geq \Delta^2 + 2\Delta$. In a very recent breakthrough, Bhandari and Chakraborty [5] introduced a novel bounding chain that has expected running time $O(n \log^2 n)$ in a substantially broader regime $q > 3\Delta$. Another way to obtain perfect samplers is to use standard reductions between counting and sampling [30]. Using this technique, any FPTAS for the number of colorings can be turned into a polynomial-time perfect sampler. (It is important that the approximate counting algorithm is deterministic, or at least with errors that can be detected.) Currently, the FPTAS with the best regime for general graphs is due to Liu et al. [32], which requires $q \geq 2\Delta$, and it runs in time $n^{\text{EXP}(\Delta)}$.

Comparing to the results above, our algorithm draws perfect samples and achieves the $O(n)$ expected running time. In case of sub-exponential growth graphs such as \mathbb{Z}^d , it improves the result of [5] by requiring only $q > \alpha\Delta + O(1)$, where $\alpha > \alpha^* = 1.763 \dots$.

Next we consider the monomer-dimer model. It was proved in [4, 40] that instances $\mathcal{I} = (G, \lambda)$ on graphs G with maximum degree Δ exhibits strong spatial mixing with rate $(1 - \Omega(1/\sqrt{1 + \lambda\Delta}))^{-\ell}$. Applying our algorithm yields the following perfect sampling result.

Theorem 2.10. *Let \mathfrak{M} be a class of monomer-dimer instances $\mathcal{I} = (G, \lambda)$ on graphs G with sub-exponential neighborhood growth and constant λ . There exists an algorithm which given any instance $\mathcal{I} = (G, \lambda) \in \mathfrak{M}$, outputs a perfect sample from $\mu_{\mathcal{I}}$ within expected $O(n)$ time.*

Previously, Markov chains were the most successful techniques for sampling weighted matchings. The Jerrum-Sinclair chain [29] on a monomer-dimer model $\mathcal{I} = (G, \lambda)$, generates approximate samples from $\mu_{\mathcal{I}}$ within $\tilde{O}(n^2m)$ steps, where $m = |E|$. This chain also mixes in $O(n \log^2 n)$ time for finite subgraphs of the 2D lattice \mathbb{Z}^2 [42].

It is difficult to convert the Jerrum-Sinclair chain to perfect samplers. Before our work, the only perfect sampler for the monomer-dimer model we are aware of is the one obtained via standard reductions from sampling to counting [30] (similar to the case of colorings), together with deterministic approximate counting algorithms [4]. This is a perfect sampler with running time $n^{\text{Poly}(\Delta, \lambda)}$.

Our algorithm is the first linear-time perfect sampler for the monomer-dimer model on graphs with sub-exponential neighborhood growth, such as finite subgraphs of lattices \mathbb{Z}^d for any constant d and constant weight λ .

3. THE ALGORITHM

We now describe our general perfect Gibbs sampler. It generalizes the single-site version (Algorithm 1) by allowing block updates. This generalization allows us to bypass some pathological situations, and to greatly improve the efficiency of the algorithm. The pseudocode is given in Algorithm 2.

Let $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$ be a permissive spin system instance and $G = (V, E)$. For any $u \in V$ and integer $\ell \geq 0$, recall that $B_{\ell}(u) \triangleq \{v \in V \mid \text{dist}_G(u, v) \leq \ell\}$ denotes the ℓ -ball centered at u in G . And for any $B \subseteq V$, we use $\partial B \triangleq \{v \in V \setminus B \mid \{u, v\} \in E\}$ to denote the vertex boundary of B in G .

Algorithm 2: Perfect Gibbs sampler (general version)

Parameter: an integer $\ell \geq 0$;
1 Start from an arbitrary feasible configuration $X \in [q]^V$, i.e. $w_{\mathcal{I}}(X) > 0$;
2 $\mathcal{R} \leftarrow V$;
3 **while** $\mathcal{R} \neq \emptyset$ **do**
4 pick a $u \in \mathcal{R}$ uniformly at random and let $B \leftarrow (B_{\ell}(u) \setminus \mathcal{R}) \cup \{u\}$;
5 let μ_{\min} be the minimum value of $\mu_u^{\sigma}(X_u)$ over all $\sigma \in [q]^{\partial B}$ that $\sigma_{\mathcal{R} \cap \partial B} = X_{\mathcal{R} \cap \partial B}$;
6 **with probability** $\frac{\mu_{\min}}{\mu_u^{X_{\partial B}}(X_u)}$ **do** \triangleright Bayes filter
7 update X by redrawing $X_B \sim \mu_B^{X_{\partial B}}$; \triangleright block Gibbs sampler update
8 $\mathcal{R} \leftarrow \mathcal{R} \setminus \{u\}$;
9 **else**
10 $\mathcal{R} \leftarrow \mathcal{R} \cup \partial B$;
11 **return** X ;

The algorithm is parameterized by an integer $\ell \geq 1$, which is set rigorously in Section 5.

The initial $X \in [q]^V$ is an arbitrary feasible configuration, which is easy to construct by greedy algorithm since \mathcal{I} is permissive (Definition 2.1). After each iteration of the **while** loop, either X is unchanged or X_B is redrawn from $\mu_B^{X_{\partial B}}$, which is the marginal distribution of $\mu = \mu_{\mathcal{I}}$, conditioned on the current X_B . Thus, we have the following observation.

Observation 3.1. *In Algorithm 2, the configuration X is always feasible, i.e. $w_{\mathcal{I}}(X) > 0$.*

The observation implies that $\mu_u^{X_{\partial B}}(X_u) > 0$ all along. The Bayes filter in Line 5 is always well-defined. If the filter succeeds, X_B is resampled according to the correct marginal distribution $\mu_B^{X_{\partial B}}$ and u is

removed from \mathcal{R} (that is, u has been successfully “fixed”); otherwise, X is unchanged and \mathcal{R} is enlarged by ∂B (because variables in ∂B are revealed and no longer random).

The key to the correctness of Algorithm 2 is the conditional Gibbs property (\star): the law of X over $\bar{\mathcal{R}} \triangleq V \setminus \mathcal{R}$ is always the conditional distribution $\mu^{X_{\mathcal{R}}} = \mu_{\mathcal{I}}^{X_{\mathcal{R}}}$. By similar argument as in Section 1.1, just redrawing X_B from $\mu_B^{X_{\partial B}}$ will introduce a bias $\propto \mu_u^{X_{\partial B}}(X_u)$ to the sample X , relative to its target distribution $\mu^{X_{\mathcal{R}\setminus\{u\}}} = \mu_{\mathcal{I}}^{X_{\mathcal{R}\setminus\{u\}}}$. In the algorithm, we use the Bayes filter that succeeds with probability $\propto 1/\mu_u^{X_{\partial B}}(X_u)$ to cancel this bias, with the risk of enlarging \mathcal{R} by ∂B upon failure. Balancing the success probability and the size of ∂B is the key to getting an efficient algorithm, and this depends on choosing an appropriate ℓ according to the spatial mixing rate.

The efficiency of the algorithm, on the other hand, depends on the success rate for the filter at Line 6: if its success probability p is always close enough to 1, so that $(1-p)|\partial B| < p$, then set \mathcal{R} shrinks in expectation in the worst case, and the algorithm converges geometrically. Meanwhile, the success probability $p = \mu_{\min}/\mu_u^{X_{\partial B}}(X_u)$ is always close enough to 1 when the values of the marginal probability $\mu_u^\sigma(X_u)$ vary little for all $\sigma \in [q]^{\partial B}$ that $\sigma_{\mathcal{R}\cap\partial B} = X_{\mathcal{R}\cap\partial B}$, which is a decay of correlation property.

The correctness and efficiency of the algorithm are analyzed respectively in next two sections.

4. CORRECTNESS OF THE PERFECT SAMPLING

In this section, we prove the correctness of Algorithm 2, stated by the following theorem.

Theorem 4.1 (correctness theorem). *Given any permissive spin system $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$, Algorithm 2 with any parameter $\ell \geq 1$ terminates with probability 1, and outputs X that follows the law of $\mu_{\mathcal{I}}$.*

The theorem is implied by two key properties of the Gibbs distribution $\mu_{\mathcal{I}}$.

4.1. Key properties of Gibbs distributions. Note that the μ_{\min} in Algorithm 2 is determined by the set \mathcal{R} , the vertex $u \in \mathcal{R}$, and the partial feasible configuration $X_{\mathcal{R}}$. Formally, fixing the parameter $\ell \geq 0$ in Algorithm 2,

$$\mu_{\min}(\mathcal{R}, u, X_{\mathcal{R}}) \triangleq \min \left\{ \mu_{u, \mathcal{I}}^\sigma(X_u) \mid \sigma \in [q]^{\partial B} \text{ s.t. } \sigma_{\mathcal{R}\cap\partial B} = X_{\mathcal{R}\cap\partial B}, \text{ where } B \triangleq (B_\ell(u) \setminus \mathcal{R}) \cup \{u\} \right\}.$$

Property 4.2 (positive lower bound of μ_{\min}). *The lower bound $\gamma_{\mathcal{I}}$ of μ_{\min} is positive:*

$$(9) \quad \gamma_{\mathcal{I}} \triangleq \min \left\{ \mu_{\min}(\mathcal{R}, u, X_{\mathcal{R}}) \mid \mathcal{R} \subseteq V, u \in \mathcal{R}, X_{\mathcal{R}} \in [q]^{\mathcal{R}} \text{ s.t. } X_{\mathcal{R}} \text{ is feasible} \right\} > 0.$$

To state the next property, we introduce some notations: For any $\Lambda \subseteq V$, $\sigma \in [q]^\Lambda$ and $S \subseteq V \setminus \Lambda$, we use $\mu_{S, \mathcal{I}}^\sigma(\cdot)$ to denote the marginal distribution on S projected from $\mu_{\mathcal{I}}^\sigma$. For any disjoint sets $\Lambda, \Lambda' \subseteq V$, $\sigma \in [q]^\Lambda$ and $\sigma' \in [q]^{\Lambda'}$, we use $\sigma \uplus \sigma'$ to denote the configuration on $\Lambda \uplus \Lambda'$ that is consistent with σ on Λ and consistent with σ' on Λ' .

Property 4.3 (conditional independence). *Suppose $A, B, C \subset V$ are three disjoint non-empty subsets such that the removal of C disconnects A and B in G . For any $\sigma_A \in [q]^A$, $\sigma_B \in [q]^B$ and $\sigma_C \in [q]^C$,*

$$\mu_{B, \mathcal{I}}^{\sigma_A \uplus \sigma_C}(\sigma_B) = \mu_{B, \mathcal{I}}^{\sigma_C}(\sigma_B).$$

Theorem 4.1 is proved for general distribution μ over $[q]^V$ relying only on these two properties. Thus, Algorithm 2 is correct for general permissive Gibbs distributions.

In particular, we verify that all permissive spin systems satisfy these two properties. First, the conditional independence (Property 4.3) holds generally for Gibbs distributions [34]. Next, for the positive lower bound of μ_{\min} (Property 4.2): for spin systems with soft constraints, clearly Property 4.2 holds for all $\ell \geq 0$; and for general permissive spin systems \mathcal{I} , we need to verify that Property 4.2 holds if $\ell \geq 1$. Fix a tuple $(\mathcal{R}, u, X_{\mathcal{R}})$ in (9). The following fact follows from the definition of set B .

Fact 4.4. $\partial B \subseteq S_{\ell+1}(u) \cup \mathcal{R}$, where $B = (B_\ell(u) \setminus \mathcal{R}) \cup \{u\}$.

The fact implies $\partial B \setminus \mathcal{R} \subseteq S_{\ell+1}(u)$. Since $\ell \geq 1$, u is not adjacent to any vertex in $\partial B \setminus \mathcal{R}$. Since \mathcal{I} is permissive and $X_{\mathcal{R}}$ is feasible, $\mu_{u, \mathcal{I}}^\sigma(X_u) > 0$ for all $\sigma \in [q]^{\partial B}$ such that $\sigma_{\mathcal{R}\cap\partial B} = X_{\mathcal{R}\cap\partial B}$. This implies $\mu_{\min}(\mathcal{R}, u, X_{\mathcal{R}})$ is positive, thus the Property 4.2 holds.

We then prove Theorem 4.1 assuming only Property 4.2 and Property 4.3. More specifically, termination of the algorithm is guaranteed by Property 4.2, and correctness of the output upon termination is guaranteed by Property 4.3.

4.2. Termination of the algorithm. Denote by T the number of iterations of the **while** loop in Algorithm 2. To prove that the algorithm terminates with probability 1, we show that T is stochastically dominated by a geometric distribution. We use \mathcal{F} to denote the Bayes filter in Algorithm 2. Then,

$$\Pr[\mathcal{F} \text{ succeeds}] = \frac{\mu_{\min}(\mathcal{R}, u, X_{\mathcal{R}})}{\mu_{u, \mathcal{I}}^{X_{\partial B}}(X_u)} \geq \mu_{\min}(\mathcal{R}, u, X_{\mathcal{R}}).$$

If \mathcal{F} succeeds for $n = |V|$ consecutive iterations of the **while** loop, then the set \mathcal{R} must become empty and the algorithm terminates. By Property 4.2, we have

$$(10) \quad \forall k \geq 0, \quad \Pr[T \geq kn] \leq \gamma_{\mathcal{I}}^{kn}.$$

This implies T is stochastically dominated by a geometric distribution. Each iteration of the **while** loop terminates within finite number of steps. Thus, the algorithm terminates with probability 1.

4.3. Correctness upon termination. We show that upon termination, the output follows the correct distribution. Let $(X, \mathcal{R}) \in [q]^V \times 2^V$ be the random pair maintained by the algorithm. The following condition is the ‘‘loop invariant’’ of the random pair (X, \mathcal{R}) .

Condition 4.5 (conditional Gibbs property). *For any $R \subseteq V$ and $\sigma \in [q]^R$, conditioned on $\mathcal{R} = R$ and $X_R = \sigma$, the random configuration $X_{V \setminus R}$ follows the law $\mu_{\mathcal{I}}^{\sigma}$.*

Condition 4.5 is satisfied initially by the initial pair $(X, \mathcal{R}) = (X, V)$. Furthermore, consider the **while** loop that transforms

$$(X, \mathcal{R}) \rightarrow (X', \mathcal{R}').$$

Then next lemma shows that Condition 4.5 holds inductively assuming Property 4.3.

Lemma 4.6. *Suppose that $(X, \mathcal{R}) \in [q]^V \times 2^V$ is a random pair such that X is feasible and the pair (X, \mathcal{R}) satisfies Condition 4.5. Then, the random pair (X', \mathcal{R}') satisfies Condition 4.5.*

By Observation 3.1, the random configuration $X \in [q]^V$ maintained by algorithm is always feasible. Lemma 4.6 guarantees that Condition 4.5 holds all along for the random pair (X, \mathcal{R}) maintained by Algorithm 2. In particular, when the algorithm terminates, $\mathcal{R} = \emptyset$ and the output X follows the correct distribution $\mu_{\mathcal{I}}$. This proves Theorem 4.1.

Proof of Lemma 4.6. It is sufficient to prove that for any $R \subseteq V$, any feasible partial configuration $\rho \in [q]^R$ and any vertex $u \in R$, conditioned on $\mathcal{R} = R$, $X_R = \rho$, and the vertex picked in Line 4 being u , the new random pair (X', \mathcal{R}') after one iteration of the **while** loop satisfies Condition 4.5.

Fix $R \subseteq V$ and a feasible partial configuration $\rho \in [q]^R$. Let $\mathbf{u} \in R$ denote the uniform random vertex picked in Line 4. Fix a vertex $u \in R$. Let \mathcal{E} denote the event

$$\mathcal{E} : \quad X_R = \rho \wedge \mathcal{R} = R \wedge \mathbf{u} = u.$$

Since (X, \mathcal{R}) satisfies Condition 4.5 and given the set R , \mathbf{u} is independent from X , we have

$$(11) \quad \forall \tau \in [q]^{V \setminus R} : \quad \Pr[X_{V \setminus R} = \tau \mid \mathcal{E}] = \mu_{\mathcal{I}}^{\rho}(\tau).$$

Recall that \mathcal{F} is the Bayes filter. Depending on whether \mathcal{F} succeeds or not, we have two cases.

The easier case is when \mathcal{F} fails. Recall that the set B is fixed by R and u . In this case, $\mathcal{R}' = R \cup \partial B$ and $X' = X$. Conditioned on \mathcal{E} , we know that $X_u = \rho_u$ and $X_{\partial B \cap R} = \rho_{\partial B \cap R}$, the filter \mathcal{F} depends only on the partial configuration $X_{\partial B \setminus R}$. For any configuration $\sigma \in [q]^{\partial B \setminus R}$, conditioned on \mathcal{E} and $X_{\partial B \setminus R} = \sigma$, the failure of \mathcal{F} is independent from $X_{V \setminus (R \cup \partial B)} = X_{V \setminus \mathcal{R}'}$. Thus, by (11), conditioned on \mathcal{E} , $X_{\partial B \setminus R} = \sigma$ and the failure of \mathcal{F} , we have that $X'_{V \setminus \mathcal{R}'} = X_{V \setminus \mathcal{R}'} \sim \mu_{\mathcal{I}}^{\rho \cup \sigma}$, i.e. (X', \mathcal{R}') satisfies Condition 4.5.

Now we analyze the main case that \mathcal{F} succeeds. If this case does occur, we must have

$$(12) \quad \mu_{\min}(R, u, \rho) \triangleq \min\{\mu_{u, \mathcal{I}}^{\sigma}(\rho_u) \mid \sigma \in [q]^{\partial B} \text{ s.t. } \sigma_{R \cap \partial B} = \rho_{R \cap \partial B}\} > 0.$$

Define $R_u \triangleq R \setminus \{u\}$. The fact that \mathcal{F} succeeds means $\mathcal{R}' = R \setminus \{u\} = R_u$ and $X'_{R_u} = X_{R_u} = \rho_{R_u}$. Hence, we only need to show that

$$(13) \quad \forall \tau \in [q]^{V \setminus R_u} : \Pr \left[X'_{V \setminus R_u} = \tau \mid \mathcal{E} \wedge \mathcal{F} \text{ succeeds} \right] = \mu_{\mathcal{I}}^{\rho(R_u)}(\tau).$$

Recall $B = (B_\ell(u) \setminus R) \cup \{u\} = B_\ell(u) \setminus R_u$. We define the following set:

$$H \triangleq V \setminus \{R_u \cup B\}.$$

Namely, B is the set whose configuration is resampled, and H is the set whose configuration is untouched, i.e. $X'_H = X_H$. Note that $B \uplus H \uplus R_u = V$. By the chain rule:

$$(14) \quad \Pr \left[X'_{V \setminus R_u} = \tau \wedge \mathcal{F} \text{ succeeds} \mid \mathcal{E} \right] = \Pr \left[X'_H = \tau_H \wedge X'_B = \tau_B \wedge \mathcal{F} \text{ succeeds} \mid \mathcal{E} \right] \\ = \Pr[X'_H = \tau_H \mid \mathcal{E}] \cdot \Pr[\mathcal{F} \text{ succeeds} \mid \mathcal{E} \wedge X'_H = \tau_H] \cdot \Pr[X'_B = \tau_B \mid \mathcal{E} \wedge X'_H = \tau_H \wedge \mathcal{F} \text{ succeeds}].$$

As $X'_H = X_H$, (11) implies that

$$\Pr[X'_H = \tau_H \mid \mathcal{E}] = \mu_{H, \mathcal{I}}^\rho(\tau_H).$$

By Line 7 of Algorithm 2, X'_B is redrawn from the distribution $\mu_{B, \mathcal{I}}^{X_{\partial B}}(\cdot)$. By conditional independence property (Property 4.3), we have $\mu_{B, \mathcal{I}}^{X_{\partial B}}(\cdot) = \mu_{\mathcal{I}}^{X_{V \setminus B}}(\cdot)$. Note that $V \setminus B = R_u \uplus H$. Conditioned on $\mathcal{E} \wedge X'_H = \tau_H$, $X_{R_u} = \rho_{R_u}$ and $X_H = X'_H = \tau_H$, thus $\mu_{B, \mathcal{I}}^{X_{\partial B}}(\cdot) = \mu_{\mathcal{I}}^{\rho(R_u) \uplus \tau(H)}(\cdot)$. Hence,

$$(15) \quad (14) = \mu_{H, \mathcal{I}}^\rho(\tau_H) \cdot \mu_{\mathcal{I}}^{\rho(R_u) \uplus \tau(H)}(\tau_B) \cdot \Pr[\mathcal{F} \text{ succeeds} \mid \mathcal{E} \wedge X'_H = \tau_H].$$

To finish the proof, we need to calculate $\Pr[\mathcal{F} \text{ succeeds} \mid \mathcal{E} \wedge X'_H = \tau_H]$. This is done by the following claim, whose proof is deferred to the end of the section.

Claim 4.7. Assume (12). It holds that

$$(16) \quad \mu_{H, \mathcal{I}}^\rho(\tau_H) > 0 \iff \mu_{H, \mathcal{I}}^{\rho(R_u)}(\tau_H) > 0.$$

Furthermore, for τ_H such that $\Pr[X'_H = \tau_H \mid \mathcal{E}] = \mu_{H, \mathcal{I}}^\rho(\tau_H) > 0$,

$$(17) \quad \Pr[\mathcal{F} \text{ succeeds} \mid \mathcal{E} \wedge X'_H = \tau_H] = C \cdot \frac{\mu_{H, \mathcal{I}}^{\rho(R_u)}(\tau_H)}{\mu_{H, \mathcal{I}}^\rho(\tau_H)},$$

where $C = C(R, u, \rho) > 0$ is a constant depending only on R, u, ρ but not on τ .

Combining (15) and Claim 4.7, we have

$$(18) \quad \forall \tau \in [q]^{V \setminus R_u}, \Pr \left[X'_{V \setminus R_u} = \tau \wedge \mathcal{F} \text{ succeeds} \mid \mathcal{E} \right] = C \cdot \mu_{\mathcal{I}}^{\rho(R_u)}(\tau).$$

This equation can be verified in two cases:

- If $\mu_{H, \mathcal{I}}^\rho(\tau_H) = 0$, then by (16), $\mu_{H, \mathcal{I}}^{\rho(R_u)}(\tau_H) = 0$, thus LHS = RHS = 0.
- If $\mu_{H, \mathcal{I}}^\rho(\tau_H) > 0$, by (15) and (17), we have LHS = $C \cdot \mu_{H, \mathcal{I}}^{\rho(R_u)}(\tau_H) \cdot \mu_{\mathcal{I}}^{\rho(R_u) \uplus \tau(H)}(\tau_B)$ = RHS, where the last equation holds because $\tau \in [q]^{V \setminus R_u}$ and $V \setminus R_u = H \uplus B$.

Thus, the probability that \mathcal{F} succeeds is

$$(19) \quad \Pr[\mathcal{F} \text{ succeeds} \mid \mathcal{E}] = \sum_{\sigma \in [q]^{V \setminus R_u}} \Pr \left[X'_{V \setminus R_u} = \sigma \wedge \mathcal{F} \text{ succeeds} \mid \mathcal{E} \right] = \sum_{\sigma \in [q]^{V \setminus R_u}} C \cdot \mu_{\mathcal{I}}^{\rho(R_u)}(\sigma) = C,$$

where the last equation holds because $\sum_{\sigma \in [q]^{V \setminus R_u}} \mu_{\mathcal{I}}^{\rho(R_u)}(\sigma) = 1$ and $C = C(R, u, \rho) > 0$ is a constant depending only on R, u, ρ . Thus, for any $\tau \in [q]^{V \setminus R_u}$, combining (18) and (19), we have

$$\Pr \left[X'_{V \setminus R_u} = \tau \mid \mathcal{F} \text{ succeeds} \wedge \mathcal{E} \right] = \frac{\Pr \left[X'_{V \setminus R_u} = \tau \wedge \mathcal{F} \text{ succeeds} \mid \mathcal{E} \right]}{\Pr[\mathcal{F} \text{ succeeds} \mid \mathcal{E}]} = \frac{C \cdot \mu_{\mathcal{I}}^{\rho(R_u)}(\tau)}{C} = \mu_{\mathcal{I}}^{\rho(R_u)}(\tau),$$

where the last equation holds due to $C = C(R, u, \rho) > 0$. This proves (13). \square

Proof of Claim 4.7. We first introduce the following definitions. Recall that $R_u \uplus B \uplus H = V$. We further partition ∂B into two disjoint sets $\partial B \cap R_u$ and $\partial B \setminus R_u$. Define

$$(20) \quad \begin{aligned} S &\triangleq \partial B \setminus R_u = \partial B \cap H, \\ \Psi &\triangleq \partial B \cap R_u = \partial B \cap R. \end{aligned}$$

We now prove (16). Since $\rho = \rho(R_u) \uplus \rho(u)$, by the Bayes law, we have the following relation between $\mu_{H, \mathcal{I}}^{\rho(R_u)}(\tau_H)$ and $\mu_{H, \mathcal{I}}^\rho(\tau_H)$:

$$(21) \quad \mu_{H, \mathcal{I}}^\rho(\tau_H) = \mu_{H, \mathcal{I}}^{\rho(R_u) \uplus \rho(u)}(\tau_H) = \frac{\mu_{u, \mathcal{I}}^{\rho(R_u) \uplus \tau(H)}(\rho_u)}{\mu_{u, \mathcal{I}}^{\rho(R_u)}(\rho_u)} \cdot \mu_{H, \mathcal{I}}^{\rho(R_u)}(\tau_H).$$

Note that $\rho \in [q]^R$ is a feasible configuration, thus $\mu_{u, \mathcal{I}}^{\rho(R_u)}(\rho_u) > 0$ and the above ratio is well-defined. Note that $u \in B$ and $R_u \uplus B \uplus H = V$. The set ∂B separates u from $(R_u \uplus H) \setminus \partial B$. Note that $\partial B = S \uplus \Psi$, where S and Ψ is defined in (20). By the conditional independence property (Property 4.3), we have

$$(22) \quad \mu_{u, \mathcal{I}}^{\rho(R_u) \uplus \tau(H)}(\rho_u) = \mu_{u, \mathcal{I}}^{\rho(\Psi) \uplus \tau(S)}(\rho_u) \geq \mu_{\min}(R, u, \rho) > 0,$$

where the first inequality is because $\mu_{\min}(R, u, \rho)$ in (12) can be rewritten as $\min_{\eta \in [q]^S} \mu_{u, \mathcal{I}}^{\rho(\Psi) \uplus \eta}(\rho_u)$, and the second inequality is because $\mu_{\min}(R, u, \rho) > 0$ due to the lower bound in (12).

Next, we prove (17). Suppose $\mu_{H, \mathcal{I}}^\rho(\tau_H) > 0$. Combining (21) and (22), it remains to prove that

$$(23) \quad \Pr[\mathcal{F} \text{ succeeds} \mid \mathcal{E} \wedge X'_H = \tau_H] = C \cdot \frac{\mu_{H, \mathcal{I}}^{\rho(R_u)}(\tau_H)}{\mu_{H, \mathcal{I}}^\rho(\tau_H)} = C \cdot \frac{\mu_{u, \mathcal{I}}^{\rho(R_u)}(\rho_u)}{\mu_{u, \mathcal{I}}^{\rho(\Psi) \uplus \tau(S)}(\rho_u)}.$$

Conditional on \mathcal{E} , we have $X_\Psi = \rho_\Psi$ and $X_u = \rho_u$. Recall that $X'_H = X_H$, $S \subseteq H$ and $S \uplus \Psi = \partial B$. Conditional on $X'_H = \tau_H$, it holds that $X_S = \tau_S$. By the definition of the filter \mathcal{F} in Line 5 of Algorithm 2, we have that

$$(24) \quad \Pr[\mathcal{F} \text{ succeeds} \mid \mathcal{E} \wedge X'_H = \tau_H] = \frac{\mu_{\min}(\mathcal{R}, u, X_{\mathcal{R}})}{\mu_{u, \mathcal{I}}^{X_{\partial B}}(X_u)} = \frac{\mu_{\min}(R, u, \rho)}{\mu_{u, \mathcal{I}}^{\rho(\Psi) \uplus \tau(S)}(\rho_u)}.$$

Combining (24) and (23), we can set $C = C(R, u, \rho)$ in (23) as

$$(25) \quad C = C(R, u, \rho) \triangleq \frac{\mu_{\min}(R, u, \rho)}{\mu_{u, \mathcal{I}}^{\rho(R_u)}(\rho_u)} = \frac{1}{\mu_{u, \mathcal{I}}^{\rho(R_u)}(\rho_u)} \cdot \min_{\eta \in [q]^S} \mu_{u, \mathcal{I}}^{\rho(\Psi) \uplus \eta}(\rho_u).$$

Note that $\mu_{u, \mathcal{I}}^{\rho(R_u)}(\rho_u) > 0$ because ρ is feasible, and $\mu_{\min}(R, u, \rho) > 0$ due to the lower bound in (12). This implies $C(R, u, \rho) > 0$. Remark the sets S and Ψ are determined by R and u . This implies that the $C(R, u, \rho)$ defined as above depends only on R, u, ρ . This proves (17). \square

5. EFFICIENCY UNDER STRONG SPATIAL MIXING

In this section, we prove the following result.

Condition 5.1. Let $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$ be a permissive spin system where $G = (V, E)$. There is an integer $\ell = \ell(q) \geq 2$ such that the following holds: for every $v \in V$, $\Lambda \subseteq V$, and any two partial configurations $\sigma, \tau \in [q]^\Lambda$ satisfying $\min\{\text{dist}_G(v, u) \mid u \in \Lambda, \sigma_u \neq \tau_u\} = \ell$,

$$(26) \quad \forall a \in [q] : \left| \frac{\mu_{v, \mathcal{I}}^\sigma(a)}{\mu_{v, \mathcal{I}}^\tau(a)} - 1 \right| \leq \frac{1}{5 |S_\ell(v)|} \quad (\text{with the convention } 0/0 = 1),$$

where $S_\ell(v) \triangleq \{u \in V \mid \text{dist}_G(v, u) = \ell\}$ denotes the sphere of radius ℓ centered at v in G .

Theorem 5.2. Let \mathfrak{S} be a class of permissive spin systems satisfying Condition 5.1. Given any instance $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A}) \in \mathfrak{S}$, the Algorithm 2 with parameter $\ell = \ell^* - 1$ outputs a perfect sample from $\mu_{\mathcal{I}}$ within $O\left(n \cdot q^{2\Delta \ell^*}\right)$ time in expectation, where n is the number of vertices in G , Δ is the maximum degree of G , $\ell^* = \ell^*(q) \geq 2$ is determined by Condition 5.1, and $O(\cdot)$ hides only absolute constants.

The correctness part of Theorem 5.2 follows from Theorem 4.1, we focus on the efficiency part. The proof sketch is that if \mathfrak{I} satisfies Condition 5.1 with parameter ℓ^* , we set the parameter ℓ in Algorithm 2 as $\ell = \ell^* - 1$. We prove that after each iteration of the **while** loop, the size of \mathcal{R} decreases by at least $\frac{1}{5}$ in expectation. Note that the initial $\mathcal{R} = V$, thus the initial size of \mathcal{R} is n . By the optional stopping theorem, the number of iterations of the **while** loop is $O(n)$ in expectation. One can verify that the time complexity of the **while** loop is $O(q^{2\Delta^{\ell+1}}) = O(q^{2\Delta^{\ell^*}})$. This proves the running time result.

To analyze the expected size of \mathcal{R} after each iteration of the **while** loop, we analyze the Bayes filter \mathcal{F} in Line 5. The probability that \mathcal{F} fails is $1 - \mu_{\min}/\mu_u^{X_{\partial B}}(X_u)$. Suppose $\sigma \in [q]^{\partial B}$ achieves $\mu_{\min} = \mu_u^\sigma(X_u)$. By Fact 4.4, we can verify that σ and $X_{\partial B}$ can differ only at $\partial B \setminus \mathcal{R} \subseteq S_{\ell+1}(u) = S_{\ell^*}(u)$. By Condition 5.1, we know that $\Pr[\mathcal{F} \text{ fails}] \leq \frac{1}{5|S_{\ell^*}(u)|}$. In addition, we have

- if \mathcal{F} succeeds, the size of \mathcal{R} decreases by 1;
- if \mathcal{F} fails, the size of \mathcal{R} increases by $|\partial B \setminus \mathcal{R}|$, it easy to verify $\partial B \setminus \mathcal{R} \subseteq S_{\ell^*}(u)$ by Fact 4.4, thus, the size of \mathcal{R} increases by at most $|S_{\ell^*}(u)|$.

Thus, after each iteration of the **while** loop, the size of \mathcal{R} decreases by at least $\frac{1}{5}$ in expectation.

In the formal proof, we actually prove a stronger result. We first introduce the following condition.

Condition 5.3. Let $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$ be a permissive spin system where $G = (V, E)$. There is an integer $\ell = \ell(q) \geq 2$ such that the following holds: for every $v \in V$, any two disjoint sets $A, B \subseteq V$ with $\text{dist}_G(v, B) = \min\{\text{dist}_G(v, u) \mid u \in B\} = \ell$, and any partial configuration $\sigma \in [q]^A$,

$$(27) \quad \forall a \in [q], \quad 1 - \frac{\min_{\tau \in [q]^B} \mu_{v, \mathcal{I}}^{\sigma \cup \tau}(a)}{\mu_{v, \mathcal{I}}^\sigma(a)} \leq \frac{1}{5|S_\ell(v)|} \quad (\text{with the convention } 0/0 = 1),$$

where $S_\ell(v) \triangleq \{u \in V \mid \text{dist}_G(v, u) = \ell\}$ denotes the sphere of radius ℓ centered at v in G .

It is straightforward to verify that Condition 5.1 implies Condition 5.3. In the rest of this section, we prove that the efficiency result in Theorem 5.2 holds under Condition 5.3.

Let $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A}) \in \mathfrak{I}$ be the input instance satisfying Condition 5.3 with some $\ell^* \geq 2$. Set the parameter ℓ in Algorithm 2 as $\ell = \ell^* - 1$. Denote by T the number of iterations of the **while** loop in Algorithm 2. To prove the efficiency result in Theorem 5.2, we bound the maximum running time of the **while** loop and the expectation of T in the following two lemmas.

Lemma 5.4. The running time of each **while** loop is at most $O(q^{2\Delta^{\ell+1}}) = O(q^{2\Delta^{\ell^*}})$.

Lemma 5.5. $\mathbb{E}[T] \leq 5n$.

Since the input instance \mathcal{I} is permissive (Definition 2.1), the initial feasible configuration can be constructed by a simple greedy algorithm. The running time of the first two lines in Algorithm 2 is $O(n\Delta)$. Combining this with Lemma 5.4 and Lemma 5.5 proves the efficiency result in Theorem 5.2.

Proof of Lemma 5.4. We first show that $\mu_{u, \mathcal{I}}^{X_{\partial B}}(X_u)$ can be computed in time $O(q^{\Delta^{\ell+1}})$, where $O(\cdot)$ hides an absolute constant. Let $\tilde{G} = G[B \cup \partial B]$ and $\tilde{\mathcal{I}}$ be the instance restricted to \tilde{G} . By the conditional independence property (Property 4.3), it is straightforward to verify

$$(28) \quad \mu_{u, \mathcal{I}}^{X_{\partial B}}(X_u) = \mu_{u, \tilde{\mathcal{I}}}^{X_{\partial B}}(X_u)$$

since ∂B separates B from $V \setminus (B \cup \partial B)$ and $u \in B$. Since $|B| \leq |B_\ell(u)| \leq \frac{\Delta^{\ell+1}-1}{\Delta-1} \leq \Delta^{\ell+1}$, it takes at most $O(q^{\Delta^{\ell+1}})$ to enumerate all possibilities and compute $\mu_{u, \tilde{\mathcal{I}}}^{X_{\partial B}}(X_u)$ using (28). By Fact 4.4, $\partial B \subseteq S_{\ell+1}(u) \cup R$. Since $|\partial B \setminus R| \leq |S_{\ell+1}(u)| \leq \Delta^{\ell+1}$, we can enumerate all $[q]^{\partial B \setminus R}$ to compute μ_{\min} in time $O(q^{2\Delta^{\ell+1}})$. The total running time for the first three lines of the **while** loop is at most $O(q^{2\Delta^{\ell+1}})$.

Another non-trivial computation is to sample $X(B)$ from $\mu_{\mathcal{I}}^{X_{\partial B}}$. Similar to (28), conditional independence implies that this can be done by straightforward enumeration in time $O(q^{\Delta^{\ell+1}})$. The total running time of the **while** loop is thus $O(q^{2\Delta^{\ell+1}}) = O(q^{2\Delta^{\ell^*}})$. \square

Proof of Lemma 5.5. Define a sequence of random pairs $(X_0, \mathcal{R}_0), (X_1, \mathcal{R}_1), \dots, (X_T, \mathcal{R}_T)$, where each $(X_t, \mathcal{R}_t) \in [q]^V \times 2^V$. The initial (X_0, \mathcal{R}_0) is constructed by the first two lines of Algorithm 2. In t -th **while** loop, Algorithm 2 updates $(X_{t-1}, \mathcal{R}_{t-1})$ to (X_t, \mathcal{R}_t) . For any $t \geq 0$, we use a random variable $Y_t \triangleq |\mathcal{R}_t|$ to denote the size of \mathcal{R}_t . The stopping time T is the smallest integer such that $Y_t = 0$.

Define the *execution log* of Algorithm 2 up to time t as

$$\mathcal{L}_t \triangleq (X_0(\mathcal{R}_0), \mathcal{R}_0), (X_1(\mathcal{R}_1), \mathcal{R}_1), \dots, (X_t(\mathcal{R}_t), \mathcal{R}_t).$$

Note that the algorithm terminates at time T if and only if $\mathcal{R}_T = \emptyset$. In the t -th iteration of the **while** loop, we use \mathcal{F}_t to denote the Bayes filter and \mathbf{u}_t to denote the random vertex picked in Line 4. We have the following claim.

Claim 5.6. *Given any execution log \mathcal{L}_{t-1} created by Algorithm 2 such that $\mathcal{R}_{t-1} \neq \emptyset$, for any $u \in \mathcal{R}_{t-1}$,*

$$\Pr[\mathcal{F}_t \text{ succeeds} \mid \mathcal{L}_{t-1}, \mathbf{u}_t = u] \geq \begin{cases} 1 & \text{if } |S_t| = \emptyset; \\ 1 - \frac{2}{5|S_t|} & \text{if } |S_t| \neq \emptyset, \end{cases}$$

where $S_t = \partial B_t \setminus \mathcal{R}_{t-1}$ and $B_t = (B_t(u) \setminus \mathcal{R}_{t-1}) \cup \{u\}$ is the set B in the t -th iteration the **while** loop.

Note that given \mathcal{L}_{t-1} , the vertex $\mathbf{u}_t \in \mathcal{R}_{t-1}$ is sampled uniformly and independently. For any $1 \leq t \leq T$ and any execution log \mathcal{L}_{t-1} created by Algorithm 2, if $S_t = \emptyset$, by Claim 5.6, we have

$$\mathbb{E}[Y_t \mid \mathcal{L}_{t-1}, S_t = \emptyset] = Y_{t-1} - 1;$$

Suppose $S_t \neq \emptyset$. If \mathcal{F}_t fails, then $\mathcal{R}_t = \mathcal{R}_{t-1} \cup \partial B_t = \mathcal{R}_{t-1} \cup (\partial B_t \setminus \mathcal{R}_{t-1})$. In other words, $|S_t|$ new vertices will be added into \mathcal{R}_{t-1} if \mathcal{F}_t fails. We have

$$\begin{aligned} \mathbb{E}[Y_t \mid \mathcal{L}_{t-1}, S_t \neq \emptyset] &\leq Y_{t-1} - \Pr[\mathcal{F}_t \text{ succeeds} \mid \mathcal{L}_{t-1}, S_t \neq \emptyset] + \Pr[\mathcal{F}_t \text{ fails} \mid \mathcal{L}_{t-1}, S_t \neq \emptyset] \cdot |S_t| \\ \text{(by Claim 5.6)} \quad &\leq Y_{t-1} + \frac{2}{5|S_t|} - \frac{3}{5} \\ &\leq Y_{t-1} - \frac{1}{5}. \end{aligned}$$

Combining two cases together implies

$$\mathbb{E}[Y_t \mid \mathcal{L}_{t-1}] = \mathbb{E}[Y_t \mid (X_0(\mathcal{R}_0), \mathcal{R}_0), (X_1(\mathcal{R}_1), \mathcal{R}_1), \dots, (X_{t-1}(\mathcal{R}_{t-1}), \mathcal{R}_{t-1})] \leq Y_{t-1} - \frac{1}{5}.$$

We now define a sequence Y'_0, Y'_1, \dots, Y'_T where each $Y'_t = Y_t + \frac{t}{5}$. Thus, Y'_0, Y'_1, \dots, Y'_T is a supermartingale with respect to $(X_0(\mathcal{R}_0), \mathcal{R}_0), (X_1(\mathcal{R}_1), \mathcal{R}_1), \dots, (X_T(\mathcal{R}_T), \mathcal{R}_T)$ and T is a stopping time. Note that each $|Y'_t - Y'_{t-1}| \leq n + 1$ is bounded and $\mathbb{E}[T]$ is finite due to (10). Due to the optional stopping theorem [8, Chapter 5], we have $\mathbb{E}[Y'_T] \leq \mathbb{E}[Y'_0] = \mathbb{E}[Y_0]$. Hence

$$\mathbb{E}[T] \leq 5\mathbb{E}[Y_0] = 5n,$$

where the last equation is because $\mathbb{E}[Y_0] = \mathbb{E}[|\mathcal{R}_0|] = n$. \square

Proof of Claim 5.6. Suppose $S_t = \partial B_t \setminus \mathcal{R}_{t-1} = \emptyset$. This implies $\partial B_t \subseteq \mathcal{R}_{t-1}$. By the definition of μ_{\min} , we have $\mu_{u, \mathcal{I}}^{X_{t-1}(\partial B_t)}(X_{t-1}(u)) = \mu_{\min}$ and $\Pr[\mathcal{F}_t \text{ succeeds}] = 1$. In the following proof, we assume $S_t \neq \emptyset$.

We need the following property to prove the claim. Fix an execution log L_t up to time $t \geq 0$:

$$L_t = (\rho_0, R_0), (\rho_1, R_1), \dots, (\rho_t, R_t),$$

where each $R_i \subseteq V$, each $\rho_i \in [q]^{R_i}$. Assume $R_t \neq \emptyset$ and L_t is a feasible execution log, i.e. $\Pr[L_t = L_t] > 0$. We claim that given the log L_t , the random $X_t \in [q]^V$ satisfies $X_t(R_t) = \rho_t$ and

$$(29) \quad \forall \tau \in [q]^{V \setminus R_t}, \quad \Pr[X_t(V \setminus R_t) = \tau \mid L_t = L_t] = \mu_{\mathcal{I}}^{\rho_t}(\tau).$$

Equation (29) is proved by the induction on t . If $t = 0$, we have $R_0 = V$, Equation (29) holds trivially. Assume Equation (29) holds for all $t < k$. Fix any feasible execution log $L_k = (\rho_0, R_0), (\rho_1, R_1), \dots, (\rho_k, R_k)$ such that $R_k \neq \emptyset$. Since L_k is feasible, we have $R_{k-1} \neq \emptyset$. Consider the k -th iteration of the **while** loop. The k -th iteration exists because $R_{k-1} \neq \emptyset$. By induction hypothesis, conditioning on the execution log $\mathcal{L}_{k-1} = (\rho_0, R_0), (\rho_1, R_1), \dots, (\rho_{k-1}, R_{k-1})$, the random pair $(X_{k-1}, \mathcal{R}_{k-1})$ satisfies the Condition 4.5 and X_{k-1} is a feasible configuration (since L_k is a feasible execution log, thus ρ_{k-1} is feasible). By

Lemma 4.6, conditioning on the execution log $\mathcal{L}_{k-1} = (\rho_0, R_0), (\rho_1, R_1), \dots, (\rho_{k-1}, R_{k-1})$, the random pair (X_k, \mathcal{R}_k) satisfies the Condition 4.5. By Condition 4.5, assuming the further condition $\mathcal{R}_k = R_k$ and $X_k(R_k) = \rho_k$, Equation (29) holds for $t = k$. This proves Equation (29).

Consider a feasible execution log up to time $t - 1 \geq 0$:

$$L_{t-1} = (\rho_0, R_0), (\rho_1, R_1), \dots, (\rho_{t-1}, R_{t-1})$$

satisfying $R_{t-1} \neq \emptyset$, where each $R_i \subseteq V$ and each $\rho_i \in [q]^{R_i}$. Given the execution log $\mathcal{L}_{t-1} = L_{t-1}$, we fix a vertex $u \in R_{t-1}$ and assume $\mathbf{u}_t = u$. We analyze the t -th iteration of the **while** loop. To simplify the notation, we drop the index and denote

$$\mathbf{X} = X_{t-1}, \quad R = R_{t-1}, \quad \rho = \rho_{t-1}, \quad B = B_t = (B_{\ell}(u) \setminus R) \cup \{u\}, \quad S = S_t = \partial B \setminus R.$$

Note that the vertex \mathbf{u}_t is sampled from R uniformly and independently. By (29), given $\mathcal{L}_{t-1} = L_{t-1}$ and $\mathbf{u}_t = u$, it holds that $X(R) = \rho$ and (X, R) satisfies Condition 4.5. By Property 4.2, we know $\mu_{\min}(R, u, \rho) > 0$, thus the lower bound in (12) holds. According to the proof of Lemma 4.6, combining (19) and (25), we have

$$\Pr[\mathcal{F}_t \text{ succeeds} \mid \mathcal{L}_{t-1} = L_{t-1}, \mathbf{u}_t = u] = \frac{1}{\mu_{u, \mathcal{I}}^{\rho(R_u)}(\rho_u)} \cdot \min_{\eta \in [q]^S} \mu_{u, \mathcal{I}}^{\rho(\Psi) \uplus \eta}(\rho_u),$$

where $R_u = R \setminus \{u\}$, $S = \partial B \setminus R$ and $\Psi = \partial B \cap R$. Note that $\partial B = S \uplus \Psi$ and $u \in B$, the set ∂B separates u from $V \setminus (B \cup \partial B)$. Since $\Psi \subseteq R_u$ and two sets R_u and B are disjoint, by the conditional independence property (Property 4.3), we have $\mu_{u, \mathcal{I}}^{\rho(\Psi) \uplus \eta}(\rho_u) = \mu_{u, \mathcal{I}}^{\rho(R_u) \uplus \eta}(\rho_u)$. This implies

$$\Pr[\mathcal{F}_t \text{ succeeds} \mid \mathcal{L}_{t-1} = L_{t-1}, \mathbf{u}_t = u] = \frac{1}{\mu_{u, \mathcal{I}}^{\rho(R_u)}(\rho_u)} \cdot \min_{\eta \in [q]^S} \mu_{u, \mathcal{I}}^{\rho(R_u) \uplus \eta}(\rho_u).$$

By Fact 4.4, we have $S = \partial B \setminus R \subseteq S_{\ell+1}(u)$. We take $A = R_u, B = S, v = u, \sigma = \rho(R_u)$ and $a = \rho_u$ in Condition 5.3, since $\text{dist}_G(u, S) = \ell + 1 = \ell^*$ and $|S| \leq |S_{\ell^*}(u)|$, this proves that

$$\Pr[\mathcal{F}_t \text{ succeeds} \mid \mathcal{L}_{t-1} = L_{t-1}, \mathbf{u}_t = u] \geq 1 - \frac{1}{5|S_{\ell^*}(u)|} \geq 1 - \frac{1}{5|S|} \geq 1 - \frac{2}{5|S|}.$$

□

Remark 5.7. Suppose the input instance from the class \mathfrak{S} satisfies Condition 5.1 with some $\ell^* \geq 2$ and take $\ell = \ell^* - 1$ in Algorithm 2. We could tweak Algorithm 2 to reduce its running time to $O(n \cdot q^{\Delta \ell^*})$. Let $\Psi \triangleq \partial B \cap \mathcal{R}$ and $S \triangleq \partial B \setminus \mathcal{R}$. Note that $S \subseteq S_{\ell^*}(u)$ by Fact 4.4. The idea is that instead of calculating μ_{\min} , we may simply compute $\mu_{u, \mathcal{I}}^{X(\Psi) \uplus \sigma}(X_u)$ where $\sigma = \mathbf{1} \in [q]^S$ is a one-vector, then use (26) to get a lower bound μ_{low} of μ_{\min} as

$$\mu_{\text{low}} \triangleq \left(1 - \frac{1}{5|S_{\ell^*}(u)|}\right) \mu_{u, \mathcal{I}}^{X(\Psi) \uplus \sigma}(X_u).$$

By Condition 5.1, $\mu_{\text{low}} \leq \mu_{u, \mathcal{I}}^{X(\partial B)}(X_u)$. Then we use μ_{low} instead of μ_{\min} in the definition of \mathcal{F} . It is straightforward to check that Algorithm 2 is still correct with this tweak, and for each iteration of the **while** loop, given any \mathbf{X}, \mathcal{R} , it holds that

$$\begin{aligned} \Pr[\mathcal{F} \text{ succeeds} \mid \mathbf{X}, \mathcal{R}] &= \frac{\mu_{\text{low}}}{\mu_{u, \mathcal{I}}^{X(\partial B)}(X_u)} = \left(1 - \frac{1}{5|S_{\ell^*}(v)|}\right) \frac{\mu_{u, \mathcal{I}}^{X(\Psi) \uplus \sigma}(X_u)}{\mu_{u, \mathcal{I}}^{X(\partial B)}(X_u)} = \left(1 - \frac{1}{5|S_{\ell^*}(v)|}\right) \frac{\mu_{u, \mathcal{I}}^{X(\Psi) \uplus \sigma}(X_u)}{\mu_{u, \mathcal{I}}^{X(\Psi) \uplus X(S)}(X_u)} \\ &\geq \left(1 - \frac{1}{5|S_{\ell^*}(v)|}\right)^2 \quad (\text{by } S \subseteq S_{\ell^*}(u) \text{ and Condition 5.1}) \\ &\geq 1 - \frac{2}{5|S|}. \end{aligned}$$

This proves Claim 5.6. Besides, we do not need to enumerate all configurations in $[q]^S$ to compute μ_{\min} , the expected running time of Algorithm 2 can be reduced to $O(n \cdot q^{\Delta \ell^*})$.

In this section, we use Theorem 5.2 to prove other results mentioned in Section 2. We analyze the strong spatial mixing properties for the classes of spin systems mentioned in Section 2, so that we can use Theorem 5.2 to prove the existences of perfect samplers.

6.1. Spin systems on sub-exponential neighborhood growth graphs. In this section, we prove Theorem 2.4 using Theorem 5.2. We need the following proposition, which explains the relation between the multiplicative form of decay in (26) and the additive form of decay in (6). Similar results appeared in [1, 15, 41, 2].

Proposition 6.1. *Let $\delta : \mathbb{N} \rightarrow \mathbb{N}$ be a non-increasing function. Let \mathfrak{S} be a class of permissive spin system instances exhibiting strong spatial mixing with decay rate δ . For every instance $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A}) \in \mathfrak{S}$, where $G = (V, E)$, for every $v \in V$, $\Lambda \subseteq V$, and any two partial configurations $\sigma, \tau \in [q]^\Lambda$ with $\ell \geq 2$,*

$$(30) \quad \forall a \in [q], \quad \min \left(\left| \frac{\mu_{v, \mathcal{I}}^\sigma(a)}{\mu_{v, \mathcal{I}}^\tau(a)} - 1 \right|, 1 \right) \leq 10q \cdot |S_{\lfloor \ell/2 \rfloor}(v)| \cdot \delta(\lfloor \ell/2 \rfloor) \quad (\text{with the convention } 0/0 = 1),$$

where $\ell \triangleq \min\{\text{dist}_G(v, u) \mid u \in \Lambda, \sigma(u) \neq \tau(u)\} \geq 2$ and $S_{\lfloor \ell/2 \rfloor}(v) \triangleq \{u \in V \mid \text{dist}_G(v, u) = \lfloor \ell/2 \rfloor\}$ denotes the sphere of radius $\lfloor \ell/2 \rfloor$ centered at v in G .

The proof of Proposition 6.1 is deferred to the end of this section. We first use Proposition 6.1 to prove Theorem 2.4.

Proof of Theorem 2.4. Let q be a finite integer. Suppose \mathfrak{S} is a class of q -state spin systems that is defined on a class of graphs that have sub-exponential neighborhood growth in Definition 2.3 with function $s : \mathbb{N} \rightarrow \mathbb{N}$. Suppose \mathfrak{S} exhibits strong spatial mixing with exponential decay with some constants $\alpha > 0, \beta > 0$. Let δ be the function $\delta(x) = \alpha \exp(-\beta x)$.

Fix a instance $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A}) \in \mathfrak{S}$, where $G = (V, E)$. By Proposition 6.1, we have for any $\Lambda \subseteq V$, any $v \in V$, and any two partial configurations $\sigma, \tau \in [q]^\Lambda$ satisfying $\ell \triangleq \min\{\text{dist}_G(v, u) \mid u \in \Lambda, \sigma(u) \neq \tau(u)\} \geq 2$,

$$(31) \quad \forall a \in [q], \quad \min \left(\left| \frac{\mu_{v, \mathcal{I}}^\sigma(a)}{\mu_{v, \mathcal{I}}^\tau(a)} - 1 \right|, 1 \right) \leq 10q \cdot |S_{\lfloor \ell/2 \rfloor}(v)| \cdot \delta(\lfloor \ell/2 \rfloor) \\ (\text{by } |S_r(v)| \leq s(r)) \leq 10q \cdot s(\lfloor \ell/2 \rfloor) \cdot \delta(\lfloor \ell/2 \rfloor)$$

Note that $\delta(\ell) = \alpha \exp(-\beta \ell)$. We take $\ell_0 = \ell_0(q, \alpha, \beta, s)$ sufficiently large such that $\ell_0 \geq 2$ and

$$(32) \quad 10\alpha q \cdot s(\lfloor \ell_0/2 \rfloor) \exp(-\beta \lfloor \ell_0/2 \rfloor) \leq \frac{1}{5s(\ell_0)} \leq \frac{1}{5|S_{\ell_0}(v)|}.$$

Note that (32) is equivalent to

$$(33) \quad \delta(\lfloor \ell_0/2 \rfloor) = \alpha \exp(-\beta \lfloor \ell_0/2 \rfloor) \leq \frac{1}{50q \cdot s(\lfloor \ell_0/2 \rfloor) \cdot s(\ell_0)}.$$

Such $\ell_0 = \ell_0(q, \alpha, \beta, s)$ must exist because $s(r) = \exp(o(r))$ for all $r \geq 0$. Combining (31) and (32) implies that \mathcal{I} satisfies Condition 5.1 with $\ell_0 \geq 2$. By Theorem 5.2, if the parameter ℓ in Algorithm 2 is set so that $\ell = \ell_0 - 1$, given \mathcal{I} , the expected running time of Algorithm 2 is $n \cdot q^{O(\Delta^{\ell_0})}$. Since $\ell_0 = O(1)$, $q = O(1)$ and $\Delta \leq s(1) = O(1)$, the expected running time of Algorithm 2 is $O(n)$. \square

We now prove Proposition 6.1. Similar results are proved in [1, 15, 41, 2].

Proof of Proposition 6.1. Fix a instance $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A}) \in \mathfrak{S}$, where $G = (V, E)$. Fix two partial configurations $\sigma, \tau \in [q]^\Lambda$ with $\ell \triangleq \min\{\text{dist}_G(v, u) \mid u \in \Lambda, \sigma(u) \neq \tau(u)\}$. We use $D \triangleq \{v \in \Lambda \mid \sigma(v) \neq \tau(v)\}$ to denote the set at which σ and τ disagree. Fix a spin $a \in [q]$. Since \mathcal{I} is permissive (Definition 2.1),

we have

$$\begin{aligned}\mu_{v,I}^\sigma(a) = 0 &\iff b_v(a) \prod_{u \in \Gamma_G(v) \cap \Lambda} A_{uv}(a, \sigma(v)) = 0, \\ \mu_{v,I}^\tau(a) = 0 &\iff b_v(a) \prod_{u \in \Gamma_G(v) \cap \Lambda} A_{uv}(a, \tau(v)) = 0,\end{aligned}$$

where $\Gamma_G(v)$ is the neighborhood of v in G . Since $\ell \geq 2$, there is no edge between v and D , we have $\Gamma_G(v) \cap D = \emptyset$. This implies $\mu_{v,I}^\sigma(a) = 0$ if and only if $\mu_{v,I}^\tau(a) = 0$. If $\mu_{v,I}^\sigma(a) = \mu_{v,I}^\tau(a) = 0$, the proposition holds trivially. We assume

$$(34) \quad \mu_{v,I}^\sigma(a) > 0 \wedge \mu_{v,I}^\tau(a) > 0.$$

Define the set of vertices $H \triangleq S_{\lfloor \ell/2 \rfloor}(v) \setminus \Lambda$, where $S_{\lfloor \ell/2 \rfloor}(v) = \{u \in V \mid \text{dist}(u, v) = \lfloor \ell/2 \rfloor\}$ is the sphere of radius $\lfloor \ell/2 \rfloor$ centered at v in graph G . By the definitions, we have $H \cap D = \emptyset$. If $H = \emptyset$, then $S_{\lfloor \ell/2 \rfloor}(v) \subseteq \Lambda$, the proposition holds due to the conditional independence property. In the rest of the proof, we assume $H \neq \emptyset$.

For any two disjoint sets $S, S' \subseteq V$ and any partial configurations $\eta \in [q]^S, \eta' \in [q]^{S'}$, we use $\mu_I^{S \leftarrow \eta, S' \leftarrow \eta'}$ to denote the distribution $\mu_I^{\eta \uplus \eta'}$.

For any $\rho \in [q]^H$ satisfying $\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho) > 0$ and $\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho) > 0$ we have

$$\mu_{v,I}^\sigma(a) = \frac{\mu_{H,I}^\sigma(\rho) \cdot \mu_{v,I}^{\sigma \uplus \rho}(a)}{\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho)}, \quad \mu_{v,I}^\tau(a) = \frac{\mu_{H,I}^\tau(\rho) \cdot \mu_{v,I}^{\tau \uplus \rho}(a)}{\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho)}.$$

The first equation holds since $\mu_{H,I}^\sigma(\rho) \cdot \mu_{v,I}^{\sigma \uplus \rho}(a) = \mu_{v,I}^\sigma(a) \cdot \mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho)$ and $\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho) > 0$; the second equation holds similarly. Note that $\mu_{v,I}^\sigma(a) > 0$ and $\mu_{v,I}^\tau(a) > 0$. We have

$$\frac{\mu_{v,I}^\sigma(a)}{\mu_{v,I}^\tau(a)} = \left(\frac{\mu_{v,I}^{\sigma \uplus \rho}(a)}{\mu_{v,I}^{\tau \uplus \rho}(a)} \right) \left(\frac{\mu_{H,I}^\sigma(\rho)}{\mu_{H,I}^\tau(\rho)} \right) \left(\frac{\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho)}{\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho)} \right).$$

Note that $(\Lambda \setminus D) \cup H$ separates v from D in graph G , and the two configurations $\sigma \uplus \rho$ and $\tau \uplus \rho$ disagree only at D . By the conditional independence property, we have $\mu_{v,I}^{\sigma \uplus \rho}(a) = \mu_{v,I}^{\tau \uplus \rho}(a)$. Hence, we have

$$(35) \quad \frac{\mu_{v,I}^\sigma(a)}{\mu_{v,I}^\tau(a)} = \left(\frac{\mu_{H,I}^\sigma(\rho)}{\mu_{H,I}^\tau(\rho)} \right) \left(\frac{\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho)}{\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho)} \right).$$

Note that (35) holds for any $\rho \in [q]^H$ satisfying $\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho) > 0$ and $\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho) > 0$. Our goal is to choose a suitable ρ and bound the RHS. Let

$$\epsilon \triangleq \delta(\lfloor \ell/2 \rfloor).$$

Without loss of generality, we assume

$$(36) \quad 10q \cdot |S_{\lfloor \ell/2 \rfloor}(v)| \cdot \delta(\lfloor \ell/2 \rfloor) = 10q\epsilon \cdot |S_{\lfloor \ell/2 \rfloor}(v)| < 1.$$

If (36) does not hold, then the inequality (30) holds trivially. We have the following claim.

Claim 6.2. *Assume (36). There exists a configuration $\rho \in [q]^H$ satisfying $\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho) > 0$ and $\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho) > 0$ such that*

$$\left(1 - \frac{2q\epsilon}{1 + q\epsilon} \right)^{2m} \leq \left(\frac{\mu_{H,I}^\sigma(\rho)}{\mu_{H,I}^\tau(\rho)} \right) \left(\frac{\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho)}{\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho)} \right) \leq \left(1 + \frac{2q\epsilon}{1 - q\epsilon} \right)^{2m},$$

where $m \triangleq |S_{\lfloor \ell/2 \rfloor}(v)|$ and $\epsilon \triangleq \delta(\lfloor \ell/2 \rfloor)$.

The inequality (36) implies that

$$(37) \quad q\epsilon m \leq \frac{1}{10}.$$

Combining Claim 6.2 with the above, we have

$$\begin{aligned} & \left(\frac{\mu_{H,I}^\sigma(\rho)}{\mu_{H,I}^\tau(\rho)} \right) \left(\frac{\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho)}{\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho)} \right) \leq \exp\left(\frac{4q\epsilon m}{1 - q\epsilon}\right) \\ \text{(by (37))} & \leq \exp(5q\epsilon m) \\ \text{(by (37))} & \leq 1 + 10q\epsilon m. \end{aligned}$$

Similarly, we have

$$\begin{aligned} & \left(\frac{\mu_{H,I}^\sigma(\rho)}{\mu_{H,I}^\tau(\rho)} \right) \left(\frac{\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho)}{\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho)} \right) \geq (1 - 2q\epsilon)^{2m} \\ \text{(by (37))} & \geq \exp(-8q\epsilon m) \\ & \geq 1 - 10q\epsilon m. \end{aligned}$$

Recall $m \triangleq |S_{\lfloor \ell/2 \rfloor}(v)|$ and $\epsilon \triangleq \delta(\lfloor \ell/2 \rfloor)$. This proves the proposition. \square

Proof of Claim 6.2. Suppose $|H| = h \geq 1$. Let $H = \{v_1, v_2, \dots, v_h\}$. Define a sequence of subsets H_0, H_1, \dots, H_h as $H_i \triangleq \{v_j \mid 1 \leq j \leq i\}$. Note that $H_0 = \emptyset$ and $H_h = H$. We now construct the configuration $\rho \in [q]^H$ by the following h steps.

- initially, $\rho = \emptyset$ is an empty configuration;
- in i -th step, note that $\rho \in [q]^{H_{i-1}}$, choose $c_i \in [q]$ that maximizes $\mu_{v_i, I}^{\Lambda \leftarrow \sigma, v \leftarrow a, H_{i-1} \leftarrow \rho}(c_i)$ (break tie arbitrarily), extend ρ further at position v_i and set $\rho(v_i) = c_i$, thus $\rho \in [q]^{H_i}$ after the i -th step.

By the construction, we have

$$\forall 1 \leq i \leq h, \quad \mu_{v_i, I}^{\Lambda \leftarrow \sigma, v \leftarrow a, H_{i-1} \leftarrow \rho(H_{i-1})}(\rho(v_i)) \geq \frac{1}{q} > 0.$$

Recall $H \triangleq S_{\lfloor \ell/2 \rfloor}(v) \setminus \Lambda$. We have $\text{dist}_G(H, D) \geq \ell - \lfloor \ell/2 \rfloor \geq \lfloor \ell/2 \rfloor$, where D is the set at which σ and τ disagree, and $\text{dist}_G(H, D) \triangleq \min\{\text{dist}_G(u_1, u_2) \mid u_1 \in H \wedge u_2 \in D\}$. Recall $\epsilon \triangleq \delta(\lfloor \ell/2 \rfloor)$ and δ is a non-increasing function. By the strong spatial mixing property in Definition 2.2, we have

$$\forall 1 \leq i \leq h, \quad \mu_{v_i, I}^{\Lambda \leftarrow \tau, v \leftarrow a, H_{i-1} \leftarrow \rho(H_{i-1})}(\rho(v_i)) \geq \frac{1}{q} - \epsilon > 0,$$

where $\frac{1}{q} - \epsilon > 0$ holds due to (36). By the chain rule, we have $\mu_{H, I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho) > 0$ and $\mu_{H, I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho) > 0$.

We now prove that ρ satisfies the inequalities in Claim 6.2. For any $1 \leq i \leq h$, define

$$(38) \quad p_i \triangleq \mu_{v_i, I}^{\Lambda \leftarrow \sigma, v \leftarrow a, H_{i-1} \leftarrow \rho(H_{i-1})}(\rho(v_i)).$$

Recall that $\text{dist}_G(H, v) \geq \lfloor \ell/2 \rfloor$ and $\text{dist}_G(H, D) \geq \lfloor \ell/2 \rfloor$. Recall $\epsilon \triangleq \delta(\lfloor \ell/2 \rfloor)$ and δ is a non-increasing function. By the strong spatial mixing property in Definition 2.2, we have for any $c \in [q]$ and any $1 \leq i \leq h$,

$$(39) \quad 0 < p_i - \epsilon \leq \mu_{v_i, I}^{\Lambda \leftarrow \sigma, v \leftarrow c, H_{i-1} \leftarrow \rho(H_{i-1})}(\rho(v_i)) \leq p_i + \epsilon,$$

$$(40) \quad 0 < p_i - \epsilon \leq \mu_{v_i, I}^{\Lambda \leftarrow \tau, v \leftarrow c, H_{i-1} \leftarrow \rho(H_{i-1})}(\rho(v_i)) \leq p_i + \epsilon.$$

Note that $p_i - \epsilon \geq \frac{1}{q} - \epsilon > 0$ due to (36). Combining (39), (40) and the chain rule implies

$$\forall c, c' \in [q], \quad \prod_{i=1}^h \left(\frac{p_i - \epsilon}{p_i + \epsilon} \right) \leq \frac{\mu_{H, I}^{\Lambda \leftarrow \tau, v \leftarrow c}(\rho)}{\mu_{H, I}^{\Lambda \leftarrow \sigma, v \leftarrow c'}(\rho)} \leq \prod_{i=1}^h \left(\frac{p_i + \epsilon}{p_i - \epsilon} \right)$$

Note that $p_i \geq \frac{1}{q}$ for all $1 \leq i \leq h$ due the construction of ρ , and $q\epsilon < 1$ due to (36). We have

$$(41) \quad \forall c, c' \in [q], \quad \left(1 - \frac{2q\epsilon}{1+q\epsilon}\right)^h \leq \frac{\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow c}(\rho)}{\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow c'}(\rho)} \leq \left(1 + \frac{2q\epsilon}{1-q\epsilon}\right)^h,$$

and

$$(42) \quad \forall c, c' \in [q], \quad \left(1 - \frac{2q\epsilon}{1+q\epsilon}\right)^h \leq \frac{\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow c}(\rho)}{\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow c'}(\rho)} \leq \left(1 + \frac{2q\epsilon}{1-q\epsilon}\right)^h.$$

Note that

$$\mu_{H,I}^\sigma(\rho) = \sum_{c \in [q]} \mu_{v,I}^\sigma(c) \mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow c}(\rho), \quad \mu_{H,I}^\tau(\rho) = \sum_{c \in [q]} \mu_{v,I}^\tau(c) \mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow c}(\rho).$$

$\mu_{H,I}^\sigma(\rho)$ is a convex combination of $\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow c}(\rho)$, and $\mu_{H,I}^\tau(\rho)$ is a convex combination of $\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow c}(\rho)$. By (41) and (42), it holds that

$$\left(1 - \frac{2q\epsilon}{1+q\epsilon}\right)^{2h} \leq \left(\frac{\mu_{H,I}^\sigma(\rho)}{\mu_{H,I}^\tau(\rho)}\right) \left(\frac{\mu_{H,I}^{\Lambda \leftarrow \tau, v \leftarrow a}(\rho)}{\mu_{H,I}^{\Lambda \leftarrow \sigma, v \leftarrow a}(\rho)}\right) \leq \left(1 + \frac{2q\epsilon}{1-q\epsilon}\right)^{2h}.$$

Note that $h = |H|$ and $H \subseteq S_{\lfloor \ell/2 \rfloor}(v)$, then $m = |S_{\lfloor \ell/2 \rfloor}(v)| \geq h$. This proves the claim. \square

6.2. Spin systems on general graphs. In this section, we prove Theorem 2.6 by showing that Condition 2.5 implies Condition 5.1.

Proof of Theorem 2.6. Fix a instance $\mathcal{I} = (G, [q], \mathbf{A}, \mathbf{b}) \in \mathfrak{S}$ satisfying Condition 2.5 with parameter $\ell = \ell(q) \geq 2$. Fix subset $\Lambda \subseteq V$ and vertex $v \in V \setminus \Lambda$. For any two partial configurations $\sigma, \tau \in [q]^\Lambda$ satisfying $\min\{\text{dist}_G(v, u) \mid u \in \Lambda, \sigma(u) \neq \tau(u)\} = \ell \geq 2$, we claim

$$(43) \quad \forall a \in [q], \quad \mu_{v,\mathcal{I}}^\sigma(a) = 0 \iff \mu_{v,\mathcal{I}}^\tau(a) = 0.$$

Let $D \triangleq \{u \in \Lambda \mid \sigma(u) \neq \tau(u)\}$, $H \triangleq \Lambda \setminus D$ and $\rho \triangleq \sigma_H = \tau_H$. Since $\ell \geq 2$, $\Gamma_G(v) \cap \Lambda = \Gamma_G(v) \cap H$, where $\Gamma_G(v)$ is the neighborhood of v in G . Since \mathcal{I} is a permissive spin system (Definition 2.1), $\mu_{v,\mathcal{I}}^\sigma(a) = 0$ if and only if $b_v(a) \prod_{u \in \Gamma_G(v) \cap H} A_{uv}(a, \rho_u) = 0$; similarly, $\mu_{v,\mathcal{I}}^\tau(a) = 0$ if and only if $b_v(a) \prod_{u \in \Gamma_G(v) \cap H} A_{uv}(a, \rho_u) = 0$. This proves (43).

If $\mu_{v,\mathcal{I}}^\sigma(a) = \mu_{v,\mathcal{I}}^\tau(a) = 0$, then (26) holds trivially. Otherwise, by Condition 2.5, $\mu_{v,\mathcal{I}}^\sigma(a) \geq \gamma$ and $\mu_{v,\mathcal{I}}^\tau(a) \geq \gamma$, where $\gamma = \gamma(\Lambda, v) > 0$ is positive and depends only on Λ and v . By (8) and (7), we have

$$\left| \frac{\mu_{v,\mathcal{I}}^\sigma(a)}{\mu_{v,\mathcal{I}}^\tau(a)} - 1 \right| \leq \frac{\gamma + d_{\text{TV}}(\mu_{v,\mathcal{I}}^\sigma, \mu_{v,\mathcal{I}}^\tau)}{\gamma} - 1 \leq \frac{1}{5|S_\ell(v)|}.$$

This implies that any instance $\mathcal{I} = (G, [q], \mathbf{A}, \mathbf{b}) \in \mathfrak{S}$ satisfies Condition 5.1 with parameter $\ell = \ell(q) \geq 2$. Theorem 2.6 is a corollary of Theorem 5.2. \square

6.3. Uniform list coloring. We now prove Theorem 2.9. Let \mathfrak{Q} be a class of list coloring instances with at most q colors for a finite $q > 0$. Let $\alpha^* \approx 1.763\dots$ be the positive root of the equation $x^x = e$. Suppose there exist $\alpha > \alpha^*$ and $\beta \geq \frac{\sqrt{2}}{\sqrt{2}-1}$ satisfying $(1 - 1/\beta)\alpha e^{\frac{1}{\alpha}(1-1/\beta)} > 1$ such that for all $\mathcal{I} = (G = (V, E), [q], \mathcal{L}) \in \mathfrak{Q}$, the graph G is triangle-free and

$$\forall v \in V, \quad |L(v)| \geq \alpha \deg_G(v) + \beta.$$

Gamarnik, Katz, and Misra [18] proved that \mathfrak{Q} exhibits the strong spatial mixing with exponential decay. If \mathfrak{Q} is defined on sub-exponential neighborhood growth graphs, then by Theorem 2.4, the linear time perfect sampler exists for every instance in \mathfrak{Q} .

There are two remaining cases in Theorem 2.9. We now assume that the class \mathfrak{Q} of list coloring instances satisfies one of the following two conditions.

(I) there is an $s : \mathbb{N} \rightarrow \mathbb{N}$ with $s(\ell) = \exp(o(\ell))$ such that for any $\mathcal{I} = (G = (V, E), [q], \mathcal{L}) \in \mathfrak{Q}$,

$$\begin{aligned} \forall v \in V, \ell \geq 0, \quad |S_\ell(v)| &\leq s(\ell), \\ \forall v \in V, \quad |L(v)| &\geq 2 \deg_G(v); \end{aligned}$$

(II) for any $\mathcal{I} = (G = (V, E), [q], \mathcal{L}) \in \mathfrak{Q}$,

$$\forall v \in V, \quad |L(v)| \geq \Delta^2 - \Delta + 2.$$

Lemma 6.3. *Let \mathfrak{Q} be a class of list coloring instances with at most q colors for a finite $q > 0$. Suppose \mathfrak{Q} satisfies (I) or (II). There exist finite $A > 0$ and $\theta > 0$ such that for every $\mathcal{I} = (G, [q], \mathcal{L}) \in \mathfrak{Q}$, where $G = (V, E)$, for any $v \in V$, any $\Lambda \subseteq V$, and any $\sigma, \tau \in [q]^\Lambda$ with $\ell \triangleq \min\{\text{dist}_G(v, u) \mid u \in \Lambda, \sigma(u) \neq \tau(u)\} = \Omega(q \log q)$, it holds that*

$$\forall a \in [q] : \quad \left| \frac{\mu_{v, \mathcal{I}}^\sigma(a)}{\mu_{v, \mathcal{I}}^\tau(a)} - 1 \right| \leq \frac{Ae^{-\theta\ell}}{|S_\ell(v)|} \quad (\text{with the convention } 0/0 = 1),$$

where $A = A(q, s) > 0$ and $\theta = \frac{1}{2q} > 0$ if \mathfrak{Q} satisfies (I) with the function $s : \mathbb{N} \rightarrow \mathbb{N}$; or $A = \text{poly}(q)$ and $\theta = \frac{1}{2q^2} > 0$ if \mathfrak{Q} satisfies (II).

Theorem 5.2 together with Lemma 6.3 proves the remaining two cases in Theorem 2.9. We take a sufficiently large ℓ^* such that $\ell^* = \Omega(q \log q)$ and $Ae^{-\theta\ell^*} \leq \frac{1}{5}$. By Lemma 6.3, instances of \mathfrak{Q} satisfy Condition 5.1 with this $\ell^* \geq 2$. Thus the perfect sampler exists due to Theorem 5.2. Note that ℓ^* depends only on q and the function s , and for any instance $\mathcal{I} \in \mathfrak{Q}$, the maximum degree $\Delta \leq q$. Thus, the expected running time of our algorithm is $n \cdot q^{O(q^{\ell^*})} = O(n)$. Furthermore, if \mathfrak{Q} satisfies (II), then $\ell^* = \Theta(q^2 \log q)$, thus the expected running time is $n \cdot \exp(\exp(\text{poly}(q)))$.

6.3.1. *The multiplicative SSM of list coloring (proof of Lemma 6.3).* In [18, Theorem 3], Gamarnik, Katz, and Misra established the best known strong spatial mixing result for list colorings in bounded degree graphs. This is almost what we need, except that we want to control the decay rate under conditions (I) and (II). Going through the proof of [18, Theorem 3] and keeping track of the decay rate, we have the proposition below. The similar analysis technique are also used in [32].

Proposition 6.4 ([18]). *Let $\mathcal{I} = (G, [q], \mathcal{L})$ be a list coloring instance, where $G = (V, E)$. Assume that \mathcal{I} satisfies $|L(v)| \geq \deg_G(v) + 1$ for all $v \in V$. Suppose*

$$\max_{u \in V} \frac{\deg_G(u) - 1}{|L(u)| - \deg_G(u)} \leq \chi < 1.$$

Then for any $\Lambda \subseteq V$, any vertex $v \in V \setminus \Lambda$, and any two partial colorings $\sigma, \tau \in [q]^\Lambda$ satisfying $\ell \triangleq \min\{\text{dist}_G(v, u) \mid u \in \Lambda, \sigma(u) \neq \tau(u)\} = \Omega(\frac{\log q}{\log(1/\chi)})$, it holds that

$$\forall a \in [q] : \quad \left| \frac{\mu_{v, \mathcal{I}}^\sigma(a)}{\mu_{v, \mathcal{I}}^\tau(a)} - 1 \right| \leq B\chi^\ell, \quad (\text{with convention } 0/0 = 1)$$

where $B = \text{poly}(q/\chi)$ depends only on q and χ .

Proof of Lemma 6.3. Fix a instance $\mathcal{I} = (G, [q], \mathcal{L}) \in \mathfrak{Q}$, where $G = (V, E)$. Suppose \mathfrak{Q} satisfies Condition in (I). We have

$$\max_{u \in V} \frac{\deg_G(u) - 1}{|L(u)| - \deg_G(u)} \leq \max_{u \in V} \frac{\deg_G(u) - 1}{\deg_G(u)} = \frac{\Delta - 1}{\Delta} \leq \frac{q - 1}{q}.$$

The χ and B in Proposition 6.4 can be set as $\chi = \frac{q-1}{q}$ and $B = \text{poly}(q/\chi) \leq B_{\max} = \text{poly}(q)$. Then for any subset $\Lambda \subseteq V$, any vertex $v \in V \setminus \Lambda$, any two colorings $\sigma, \tau \in [q]^\Lambda$ that disagree on $D \subseteq \Lambda$ satisfying $\ell \triangleq \min\{\text{dist}_G(v, u) \mid u \in D\} = \Omega(\frac{\log q}{\log 1/\chi}) = \Omega(q \log q)$, it holds that

$$\forall a \in [q] : \quad \left| \frac{\mu_{v, \mathcal{I}}^\sigma(a)}{\mu_{v, \mathcal{I}}^\tau(a)} - 1 \right| \leq B_{\max} \chi^\ell \leq B_{\max} \cdot \frac{|S_\ell(v)|}{|S_\ell(v)|} \cdot \chi^\ell.$$

Since G has sub-exponential growth, we have that $|S_\ell(v)| \leq s(\ell) = \exp(o(\ell))$. Thus,

$$\forall a \in [q] : \left| \frac{\mu_{v, \mathcal{I}}^\sigma(a)}{\mu_{v, \mathcal{I}}^\tau(a)} - 1 \right| \leq B_{\max} \cdot \frac{s(\ell)}{|S_\ell(v)|} \cdot \left(\frac{q-1}{q} \right)^\ell \leq \frac{s(\ell) B_{\max}}{|S_\ell(v)|} \cdot e^{-\ell/q} \leq \frac{Ae^{-\theta\ell}}{|S_\ell(v)|},$$

for some $A = A(q, s) > 0$ and $\theta = \frac{1}{2q} > 0$.

Suppose \mathfrak{L} satisfies (II). Recall that Δ is the maximum degree of graph G . we have

$$\max_{u \in V} \frac{\deg_G(u) - 1}{|L(u)| - \deg_G(u)} \leq \frac{\Delta - 1}{(\Delta - 1)^2 + 1}.$$

The χ and B in Proposition 6.4 can be set as $\chi = \frac{\Delta-1}{(\Delta-1)^2+1}$ and $B = \text{poly}(q/\chi)$. Thus $1/\chi \leq \Delta^2 \leq q^2$. We have $B = \text{poly}(q/\chi) \leq B_{\max} = \text{poly}(q)$. For any subset $\Lambda \subseteq V$, any vertex $v \in V \setminus \Lambda$, any two colorings $\sigma, \tau \in [q]^\Lambda$ that disagree on $D \subseteq \Lambda$ satisfying $\ell \triangleq \min\{\text{dist}_G(u, v) \mid u \in D\} = \Omega(\frac{\log q}{\log 1/\chi}) = \Omega(\log q)$, it holds that

$$\forall a \in [q] : \left| \frac{\mu_{v, \mathcal{I}}^\sigma(a)}{\mu_{v, \mathcal{I}}^\tau(a)} - 1 \right| \leq B_{\max} \chi^\ell \leq B_{\max} \cdot \frac{\Delta(\Delta-1)^{\ell-1}}{|S_\ell(v)|} \cdot \chi^\ell,$$

where the last inequality due to $|S_\ell(v)| \leq \Delta(\Delta-1)^{\ell-1}$. Since $\chi = \frac{\Delta-1}{(\Delta-1)^2+1}$, we have

$$\begin{aligned} \forall a \in [q] : \left| \frac{\mu_{v, \mathcal{I}}^\sigma(a)}{\mu_{v, \mathcal{I}}^\tau(a)} - 1 \right| &\leq \frac{B_{\max} \Delta}{\Delta-1} \cdot \frac{1}{|S_\ell(v)|} \cdot \left(\frac{(\Delta-1)^2}{(\Delta-1)^2+1} \right)^\ell \\ \text{(by } \Delta \leq q) &\leq \frac{2B_{\max}}{|S_\ell(v)|} \cdot \left(\frac{(q-1)^2}{(q-1)^2+1} \right)^\ell \\ &\leq \frac{2B_{\max}}{|S_\ell(v)|} \cdot e^{-\frac{\ell}{2q^2}} = \frac{Ae^{-\theta\ell}}{|S_\ell(v)|}, \end{aligned}$$

where $A = 2B_{\max} = \text{poly}(q)$ and $\theta = \frac{1}{2q^2} > 0$. □

6.4. The monomer-dimer model. We now prove Theorem 2.10. We first present the monomer-dimer model instance as a spin system instance, then we use Theorem 2.4 to prove Theorem 2.10.

Given a graph $G = (V, E)$, we use $G^* = (V^*, E^*) = \text{Lin}(G)$ to denote the *line graph* of G . Each vertex $v_e \in V^*$ in line graph G^* represents an edge $e \in E$ in the original graph G , and two vertices $v_e, v_{e'}$ in G^* are adjacent if and only if e and e' share a vertex in G . We call $S \subseteq V^*$ an *independent set* in G^* if no two vertices in S are adjacent in G^* . It is easy to verify that there is a one-to-one correspondence between the matchings in G and the independent sets in G^* .

Given a monomer-dimer model instance $\mathcal{I} = (G, \lambda)$, we define a *hardcore model* instance $\mathcal{I}^* = (G^*, \lambda)$ in the line graph $G^* = \text{Lin}(G)$. Each independent set S in G^* is assigned a weight $w_{\mathcal{I}^*}(S) = \lambda^{|S|}$. Let $\mu_{\mathcal{I}^*}$ be a distribution over all independent sets in G^* such that $\mu_{\mathcal{I}^*}(S) \propto w_{\mathcal{I}^*}(S)$. Hence, \mathcal{I}^* is a spin system instance and \mathcal{I}^* is permissive. Besides, if we can sample independent sets from $\mu_{\mathcal{I}^*}$, then we can sample matchings from $\mu_{\mathcal{I}}$.

Suppose the class of monomer-dimer model instances \mathfrak{M} satisfies the condition in Theorem 2.10. Then, there exist a constant C and a function $s : \mathbb{N} \rightarrow \mathbb{N}$ with $s(\ell) = \exp(o(\ell))$ such that for all $\mathcal{I} = (G, \lambda) \in \mathfrak{M}$, $\lambda \leq C = O(1)$, $|S_\ell(v)| \leq s(\ell) = \exp(o(\ell))$ for all $v \in V$ and $\ell \geq 0$, and $\Delta_G \leq s(1) = O(1)$. Thus, \mathfrak{M} exhibits strong spatial mixing with exponential decay with constants $\alpha = \alpha(C, s) > 0$ and $\beta = \beta(C, s) > 0$ [4, 40]. Observe that if e_1, e_2, \dots, e_ℓ is a path of edges in G , then $v_{e_1}, v_{e_2}, \dots, v_{e_\ell}$ is a path of vertices in G^* , and vice versa. Hence, the following results hold for the class of hardcore instances $\mathfrak{H} = \{\mathcal{I}^* = (G^*, \lambda) \mid \mathcal{I} \in \mathfrak{M}\}$.

- The class of hardcore instances \mathfrak{H} exhibits strong spatial mixing with exponential decay with constants $\alpha' = \alpha'(C, s) > 0$ and $\beta = \beta(C, s) > 0$.
- for any instance $(G^*, \lambda) \in \mathfrak{H}$, the graph G^* has sub-exponential growth. Suppose $G^* = (V^*, E^*)$ is the line graph of $G = (V, E)$. For all $e = \{u, v\} \in E$, $\ell \geq 1$, it holds that $|S_\ell^*(v_e)| \leq \Delta_G (|S_{\ell-1}(u)| + |S_{\ell-1}(v)|) \leq 2s(1)s(\ell-1) = \exp(o(\ell))$, where $v_e \in V^*$ represents the edge e , $S_\ell^*(v_e)$ is the sphere of radius ℓ centered at v_e in G^* , $S_{\ell-1}(v)$ is the sphere of radius $\ell-1$ centered at v in G .

Note that the number of vertices in G^* is at most $n\Delta_G = O(n)$, where n is number of vertices in G . Theorem 2.10 is a corollary of Theorem 2.4.

7. DYNAMIC SAMPLING

In this section, we use our algorithm to solve the dynamic sampling problem [14, 13]. In this problem, the Gibbs distribution itself changes dynamically and the algorithm needs to maintain a random sample efficiently with respect to the current Gibbs distribution.

We first define the update for the spin system instance. Let $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$ be a spin system instance, where $G = (V, E)$.

- updates for vertices: modifying the vector b_v of vertex $v \in V$;
- updates for edges: modifying the matrix A_e of edge $e \in E$; or adding new edge $e \notin E$.

We use (D_V, D_E, C) to denote the update for instance \mathcal{I} , where $D_V \subseteq V$, $D_E \subseteq \{\{u, v\} \mid u, v \in V, u \neq v\}$, and $C = (b_v)_{v \in D_V} \cup (A_e)_{e \in D_E}$. For each $v \in D_V$, we modify its vector to $b_v \in C$, and for each $e \in D_E$, we either add the new edge e with matrix $A_e \in C$ (if $e \notin E$), or modify its matrix to $A_e \in C$ (if $e \in E$).

Definition 7.1 (dynamic sampling problem). Given a spin system instance $\mathcal{I} = (G, [q], \mathbf{b}, \mathbf{A})$ where $G = (V, E)$, a random sample $X \in [q]^V$ such that $X \sim \mu_{\mathcal{I}}$, and an update (D_V, D_E, C) that modifies the instance \mathcal{I} to an updated instance $\mathcal{I}' = (G', [q], \mathbf{b}', \mathbf{A}')$ where $G' = (V, E')$, the algorithm updates X to a new sample $X' \in [q]^V$ such that $X' \sim \mu_{\mathcal{I}'}$.

Theorem 7.2. *Let \mathfrak{S} be a class of permissive spin systems satisfying Condition 5.1. There exists an algorithm such that if the updated instance $\mathcal{I}' = (G', [q], \mathbf{b}', \mathbf{A}') \in \mathfrak{S}$, then the algorithm solves the dynamic sampling problem within $\Delta(|D_V| + |D_E|)q^{O(\Delta')}$ time in expectation, where Δ is the maximum degree of G' and $\ell = \ell(q) \geq 2$ is determined by Condition 5.1.*

Suppose $q, \Delta, \ell = O(1)$. By Theorem 7.2, the running time of our algorithm is linear in the size of the update. Hence, the efficient dynamic sampling algorithm exists if strong spatial mixing holds with a rate faster than the neighborhood growth. The relation between the spatial mixing property and the static sampling is well studied, we extend such relation further to the dynamic setting.

The dynamic sampling algorithm is given in Algorithm 3.

Algorithm 3: Dynamic perfect Gibbs sampler

Input: a spin system instance $\mathcal{I} = (G = (V, E), [q], \mathbf{b}, \mathbf{A})$, a random sample $X \sim \mu_{\mathcal{I}}$, and an update (D_V, D_E, C) that modifies \mathcal{I} to $\mathcal{I}' = (G' = (V, E'), [q], \mathbf{b}', \mathbf{A}')$.

- 1 $\mathcal{D} \leftarrow D_V \cup (\bigcup_{e \in D_E} e)$ and $\partial\mathcal{D} \leftarrow \{v \in V \setminus \mathcal{D} \mid \exists u \in \mathcal{D} \text{ s.t. } \{u, v\} \in E'\}$;
- 2 based on $X_{\partial\mathcal{D}}$, modify the partial configuration $X_{\mathcal{D}}$ so that $w_{\mathcal{I}'}(X) > 0$;
- 3 $\mathcal{R} \leftarrow \mathcal{D} \cup \partial\mathcal{D}$;
- 4 **while** $\mathcal{R} \neq \emptyset$ **do**
- 5 $(X, \mathcal{R}) \leftarrow \text{Fix}(\mathcal{I}', X, \mathcal{R})$;
- 6 **return** X ;

In Algorithm 3, the set $\mathcal{D} \subseteq V$ contains all the vertices incident to the update. Note that the input X can be an infeasible configuration for \mathcal{I}' , i.e. $w_{\mathcal{I}'}(X) = 0$, because the configuration $X_{\mathcal{D}}$ may violate the new constraints in C . Hence, in Line 2, we modify the configuration $X_{\mathcal{D}}$ so that $w_{\mathcal{I}'}(X) > 0$. Given the $X_{\partial\mathcal{D}}$, this step can be achieved by a simple greedy algorithm since \mathcal{I}' is permissive. Then, we construct the initial \mathcal{R} as $\mathcal{D} \cup \partial\mathcal{D}$. In Line 5, we call the subroutine Fix on the updated instance \mathcal{I}' .

Note that $\mathcal{R} = \mathcal{D} \cup \partial\mathcal{D}$ and $\partial\mathcal{D}$ separates \mathcal{D} from $V \setminus \mathcal{R}$ in both G and G' . In Line 2, we only modify the partial configuration $X_{\mathcal{D}}$. Such modification only reviews the information of X in $\mathcal{D} \cup \partial\mathcal{D}$. Thus, after the modification, the $X_{V \setminus \mathcal{R}}$ follows the distribution $\mu_{\mathcal{I}'}^{X_{\mathcal{R}}} = \mu_{V \setminus \mathcal{R}, \mathcal{I}'}^{X_{\partial\mathcal{D}}}$ due to the conditional independence property. Since two instances \mathcal{I} and \mathcal{I}' differ only at the subset \mathcal{D} , due to the conditional independence property, two distributions $\mu_{\mathcal{I}}^{X_{\mathcal{R}}}$ and $\mu_{\mathcal{I}'}^{X_{\mathcal{R}}}$ are identical. Thus, the initial X, \mathcal{R} satisfies Condition 4.5 with respect to \mathcal{I}' , and X is a feasible configuration for \mathcal{I}' . In each iteration of the **while** loop, we call the subroutine Fix on \mathcal{I}' . By the identical proof in Section 4, the output $X \sim \mu_{\mathcal{I}'}$.

Algorithm 4: $\text{Fix}(\mathcal{I}, X, \mathcal{R})$

Input: a spin system instance $\mathcal{I} = (G = (V, E), [q], \mathbf{b}, \mathbf{A})$, a configuration $X \in [q]^V$, a non-empty subset $\mathcal{R} \subseteq V$, and an integer parameter $\ell \geq 0$;

- 1 pick a $u \in \mathcal{R}$ uniformly at random and let $B \leftarrow (B_\ell(u) \setminus \mathcal{R}) \cup \{u\}$;
- 2 let μ_{\min} be the minimum value of $\mu_{u, \mathcal{I}}^\sigma(X_u)$ over all $\sigma \in [q]^{\partial B}$ that $\sigma_{\mathcal{R} \cap \partial B} = X_{\mathcal{R} \cap \partial B}$;
- 3 **with probability** $\frac{\mu_{\min}}{\mu_{u, \mathcal{I}}^{X_{\partial B}}(X_u)}$ **do**
- 4 update X by redrawing $X_B \sim \mu_{B, \mathcal{I}}^{X_{\partial B}}$;
- 5 $\mathcal{R} \leftarrow \mathcal{R} \setminus \{u\}$;
- 6 **else**
- 7 $\mathcal{R} \leftarrow \mathcal{R} \cup \partial B$;
- 8 **return** (X, \mathcal{R})

Let Δ denote the maximum degree of graph G' . Note that $|\mathcal{D}| = O(|D_V| + |D_E|)$. The time complexity of the first three lines of Algorithm 3 is $O(\Delta|\mathcal{D}|)$. Note that the size of the initial \mathcal{R} is $O(\Delta|\mathcal{D}|)$. The efficiency result in Theorem 7.2 follows from the identical proof in Section 5.

8. CONCLUSION AND OPEN PROBLEMS

The connection between efficient algorithms and spatial mixing (decay of correlation) has been a long-lasting theme in the study of sampling and approximate counting algorithms. In this work, we introduce a new approach for perfect sampling, that relates efficient perfect sampling to strong spatial mixing, for Gibbs distributions on graphs with sub-exponential neighborhood growth.

Our perfect sampling approach is generic. It is based on the classic Gibbs sampler, while previous perfect sampling techniques were designed for specific systems or subclasses of systems. It is surprising to us that the Gibbs sampler, studied for decades as the go-to algorithm for approximate sampling, can be turned into a perfect sampler by simply adding a filter that accesses only local information.

One key insight in designing our algorithm is to preserve the so called conditional Gibbs property (\star), a strong invariant property that implies interruptible perfect sampling as well as dynamic perfect sampling. An important open problem is to establish the same implication from spatial mixing to efficient perfect sampling on general graphs. This is interesting even for special systems, e.g. the hardcore or monomer-dimer models.

We want to point out that our current algorithm preserves the conditional Gibbs invariant (\star) in a quite pessimistic way: the invariant holds for any conditioning of the random configuration of the “incorrect” variables. We see that even such a straightforward implementation of the invariant is sufficient to give efficient perfect sampling under spatial mixing on graphs of bounded neighborhood growth. It is promising to have cleverer algorithms that exploit the true power of the conditional Gibbs property by implementing this invariant in a more global fashion, on average over randomness.

In general, do efficient perfect sampling algorithms exist whenever efficient approximate samplers exist? Our results provide evidence towards a positive answer, especially for spin systems on graphs with bounded neighborhood growth. However, in general the gap between approximate and perfect sampling persists (e.g. for sampling proper graph colorings [5]). Designing efficient perfect sampling algorithms matching their approximate counterparts remains an interesting research direction.

ACKNOWLEDGEMENT

The research was supported by the National Key R&D Program of China 2018YFB1003202 and NSFC under Grant Nos. 61722207 and 61672275. We thank Jingcheng Liu for helpful comments. Heng Guo and Yitong Yin want to thank the hospitality of the Simons Institute for the Theory of Computing at UC Berkeley, where part of the work was done.

REFERENCES

- [1] Kenneth S. Alexander. On weak mixing in lattice models. *Probab. Theory Related Fields*, 110(4):441–471, 1998.
- [2] Kenneth S. Alexander. Mixing properties and exponential decay for lattice systems in finite volumes. *Ann. Probab.*, 32(1A):441–487, 2004.
- [3] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *arXiv preprint arXiv:2001.00303*, 2020. FOCS 2020, to appear.
- [4] Mohsen Bayati, David Gamarnik, Dimitriy Katz, Chandra Nair, and Prasad Tetali. Simple deterministic approximation algorithms for counting matchings. In *STOC*, pages 122–127. ACM, 2007.
- [5] Siddharth Bhandari and Sayantan Chakraborty. Improved bounds for perfect sampling of k -colorings in graphs. In *STOC*, pages 631–642. ACM, 2020.
- [6] Sitan Chen, Michelle Delcourt, Ankur Moitra, Guillem Perarnau, and Luke Postle. Improved bounds for randomly sampling colorings via linear programming. In *SODA*, pages 2216–2234. SIAM, 2019.
- [7] Zongchen Chen, Kuikui Liu, and Eric Vigoda. Rapid mixing of Glauber dynamics up to uniqueness via contraction. *arXiv preprint arXiv:2004.09083*, 2020. FOCS 2020, to appear.
- [8] Rick Durrett. *Probability: theory and examples*. Cambridge University Press, fifth edition, 2019.
- [9] Martin E. Dyer and Alan M. Frieze. Randomly coloring graphs with lower bounds on girth and maximum degree. *Random Struct. Algorithms*, 23(2):167–179, 2003.
- [10] Martin E. Dyer, Alan M. Frieze, Thomas P. Hayes, and Eric Vigoda. Randomly coloring constant degree graphs. *Random Struct. Algorithms*, 43(2):181–200, 2013.
- [11] Martin E. Dyer, Alistair Sinclair, Eric Vigoda, and Dror Weitz. Mixing in time and space for lattice spin systems: A combinatorial view. *Random Struct. Algorithms*, 24(4):461–479, 2004.
- [12] Charilaos Efthymiou, Thomas P. Hayes, Daniel Štefanković, Eric Vigoda, and Yitong Yin. Convergence of MCMC and loopy BP in the tree uniqueness region for the hard-core model. In *FOCS*, pages 704–713. IEEE, 2016.
- [13] Weiming Feng, Kun He, Xiaoming Sun, and Yitong Yin. Dynamic MCMC sampling. *arXiv preprint arXiv:1904.11807*, 2019.
- [14] Weiming Feng, Nisheeth K. Vishnoi, and Yitong Yin. Dynamic sampling from graphical models. In *STOC*, pages 1070–1081. ACM, 2019.
- [15] Weiming Feng and Yitong Yin. On local distributed sampling and counting. In *PODC*, pages 189–198. ACM, 2018.
- [16] James A. Fill and Mark Huber. The randomness recycler: a new technique for perfect sampling. In *FOCS*, pages 503–511. IEEE, 2000.
- [17] Andreas Galanis, Daniel Štefanković, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combin. Probab. Comput.*, 25(4):500–559, 2016.
- [18] David Gamarnik, Dmitriy Katz, and Sidhant Misra. Strong spatial mixing of list coloring of graphs. *Random Struct. Algorithms*, 46(4):599–613, 2015.
- [19] Leslie Ann Goldberg, Russell A. Martin, and Mike Paterson. Strong spatial mixing with fewer colors for lattice graphs. *SIAM J. Comput.*, 35(2):486–517, 2005.
- [20] Heng Guo and Kun He. Tight bounds for popping algorithms. *Random Struct. Algorithms*, 2020. to appear.
- [21] Heng Guo and Mark Jerrum. Perfect simulation of the hard disks model by partial rejection sampling. *Ann. Inst. Henri Poincaré Comb. Phys. Interact.*, 2019. to appear.
- [22] Heng Guo and Mark Jerrum. A polynomial-time approximation algorithm for all-terminal network reliability. *SIAM J. Comput.*, 48(3):964–978, 2019.
- [23] Heng Guo and Mark Jerrum. Approximately counting bases of bicircular matroids. *Combin. Probab. Comput.*, 2020. to appear.
- [24] Heng Guo, Mark Jerrum, and Jingcheng Liu. Uniform sampling through the Lovász local lemma. *J. ACM*, 66(3):18:1–18:31, 2019.

- [25] Olle Häggström and Karin Nelander. Exact sampling from anti-monotone systems. *Statist. Neerlandica*, 52(3):360–380, 1998.
- [26] Thomas P. Hayes and Eric Vigoda. A non-Markovian coupling for randomly sampling colorings. In *FOCS*, pages 618–627. IEEE, 2003.
- [27] Mark Huber. Perfect sampling using bounding chains. *Ann. Appl. Probab.*, 14(2):734–753, 2004.
- [28] Mark Jerrum. A very simple algorithm for estimating the number of k -colorings of a low-degree graph. *Random Struct. Algorithms*, 7(2):157–165, 1995.
- [29] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, 1989.
- [30] Mark Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoret. Comput. Sci.*, 43:169–188, 1986.
- [31] Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *SODA*, pages 67–84. SIAM, 2013.
- [32] Jingcheng Liu, Alistair Sinclair, and Piyush Srivastava. A deterministic algorithm for counting colorings with 2Δ colors. In *FOCS*, pages 1380–1404. IEEE Computer Society, 2019.
- [33] Fabio Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on probability theory and statistics (Saint-Flour, 1997)*, volume 1717 of *Lecture Notes in Math.*, pages 93–191. Springer, Berlin, 1999.
- [34] Marc Mezard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [35] Michael Molloy. The glauber dynamics on colorings of a graph with high girth and maximum degree. *SIAM J. Comput.*, 33(3):721–737, 2004.
- [36] Elchanan Mossel and Allan Sly. Exact thresholds for Ising-Gibbs samplers on general graphs. *Ann. Probab.*, 41(1):294–328, 2013.
- [37] James G. Propp and David B. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct. Algorithms*, 9(1-2):223–252, 1996.
- [38] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *J. Stat. Phys.*, 155(4):666–686, 2014.
- [39] Allan Sly and Nike Sun. Counting in two-spin models on d -regular graphs. *Ann. Probab.*, 42(6):2383–2416, 2014.
- [40] Renjie Song, Yitong Yin, and Jinman Zhao. Counting hypergraph matchings up to uniqueness threshold. *Inform. and Comput.*, 266:75–96, 2019.
- [41] Yinon Spinka. Finitary codings for spatial mixing Markov random fields. *Ann. Probab.*, 48(3):1557–1591, 2020.
- [42] Jacob van den Berg and Rachel Brouwer. Random sampling for the monomer-dimer model on a lattice. *J. Math. Phys.*, 41(3):1585–1597, 2000.
- [43] Eric Vigoda. Improved bounds for sampling colorings. *J. Math. Phys.*, 41(3):1555–1569, 2000.
- [44] Dror Weitz. *Mixing in time and space for discrete spin systems*. ProQuest LLC, Ann Arbor, MI, 2004. Thesis (Ph.D.)—University of California, Berkeley.
- [45] Dror Weitz. Counting independent sets up to the tree threshold. In *STOC*, pages 140–149. ACM, 2006.
- [46] David B. Wilson. Generating random spanning trees more quickly than the cover time. In *STOC*, pages 296–303. ACM, 1996.